Computers and Chemical Engineering 25 (2001) 1055–1074

Computers & Chemical Engineering

www.elsevier.com/locate/compchemeng

Ecological interface design for petrochemical applications: supporting operator adaptation, continuous learning, and distributed, collaborative work

Greg A. Jamieson, Kim J. Vicente *

Department of Mechanical and Industrial Engineering, Cognitive Engineering Laboratory [1], *University of Toronto, 5 King's College Road, Toronto, Ontario, Canada M5S 3G8*

Received 10 January 2000; received in revised form 9 January 2001; accepted 18 January 2001

Abstract

Future support systems for operators of petrochemical refineries will have to support operator adaptation to unanticipated events, foster continuous learning, and facilitate distributed, collaborative work. This paper describes Ecological Interface Design, a candidate framework for human–computer interface design that has the potential to fulfill these diverse demands. Support for adaptation and continuous learning is demonstrated though the design of a novel operator interface for a fluid catalytic cracking unit. While the framework forms a basis upon which a distributed, collaborative support system may be built, no such design is presented here. The process of the application of the framework is described in detail, including the domain modelling activity and a description of the resulting graphical user interface. Limitations to applying the design approach to operational plants are discussed. © 2001 Elsevier Science Ltd. All rights reserved.

Keywords: Ecological interface design; Cognitive engineering; Adaptation; Learning; Collaboration

1. Introduction

The purpose of this article is to describe and demonstrate the application of a novel framework for the design of human–computer interfaces for use in the petrochemical industry. As national and international economic competition has increased, petrochemical plants are continually trying to 'do more with less'. The need for continuous improvement in quality and productivity has increased drastically. This trend has caused plants to become more complex than in the past, a trend that is sure to increase in the future (Pool, 1997; Vicente, 1999).

This increase in complexity has changed the demands on operators of process control plants. There are three requirements in particular that are worth highlighting; adaptation to novelty, continuous learning, and distributed, collaborative work.

An increase in complexity has led to a concomitant increased need to support *adaptation to novelty*, particularly to events that are unfamiliar to operators and that have not been anticipated by designers. Accident analyses have repeatedly shown that unanticipated events pose the greatest threat to system safety (see Vicente and Rasmussen, 1992 for a review). In these situations, operators cannot rely on procedures or automation provided by the designers because the event was not anticipated. Thus, operators must adapt to the new and unfamiliar situation and improvise a solution themselves. This is a demanding task, especially when we consider that the threat to equipment, the environment, the worker, and public safety is usually at its peak during these unfamiliar and unanticipated events. Thus, a framework for designing human–computer interfaces for monitoring and control of petrochemical processes must provide a way of identifying what type of information support operators need in order to be able to deal successfully with unanticipated events. We refer to such frameworks as being *event-independent* (Vicente and Tanabe, 1993).

* Corresponding author.
E-mail address: benfica@mie.utoronto.ca (K.J. Vicente).
[1] www.mie.utoronto.ca/labs/cel/.

0098-1354/01/$ - see front matter
PII: S0098-1354(01)00678-0

A second challenge that has emerged in the petrochemical industry is an increased requirement for *continuous learning*. Process plants are continuously striving to improve their operations by increasing the efficiency of existing processes, and by developing new processes that have valuable functionality. As a result, contemporary plants change much more frequently than in the past. These work environments are teeming with change and novelty, as continual attempts at quality improvement are being made. This trend means that there is an increased burden on operators to continually learn about the structure and behavior of the plant. Therefore, a framework for designing human–computer interfaces for petrochemical process control systems must foster continuous learning so that operators can keep pace with the turbulent change that characterizes their work environments.

A third challenge that has emerged in the petrochemical process industries is the increased requirement to support *distributed*, *collaborative work*. As plant-wide control systems have become increasingly complex, there is a greater need to integrate different kinds of specialized knowledge (e.g. plant operations vs. plant management) in making decisions about control. More and more, these different types of knowledge will be held by different individuals, thereby creating an increased need for collaboration. Moreover, as control rooms become increasingly centralized due to the power of information technology, the distance from the control room to the operating unit tends to increase. While this tends to bring some plant personnel closer together (e.g. panel operators, engineers, and managers), it also creates a greater need to coordinate work activities across individuals who are distributed spatially (e.g. control room operators vs. field operators).

Collectively, these challenges constitute a new set of design requirements for operator support systems. Future designs will have to support adaptation to unanticipated events, foster continuous learning, and facilitate distributed, collaborative work if they are to meet the challenges imposed by today's competitive marketplace. In the following section, we describe a candidate framework for human–computer interface design that has the potential to fulfill these diverse demands. Following the theoretical description, we provide an example application of the framework to a fluid catalytic cracking unit. We demonstrate how the resultant design tries to foster adaptation and continuous learning and we discuss how the same design framework might also support distributed, collaborative activity.

2. The ecological interface design framework

Ecological Interface Design (EID) is a theoretical framework for interface design for complex sociotechnical systems (Vicente and Rasmussen, 1990, 1992). It 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 work domain 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, parallel, perceptual-motor interaction with the world; 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 serial, analytical problem solving. These ways of processing information lead to qualitatively distinct behaviors; skill-based behavior (SBB), rule-based behavior (RBB), and knowledge-based behavior (KBB). The systems engineer can view this taxonomy as a concise organization of the wealth of knowledge that has been gained about human cognition in a form than can guide systems design. The SRK framework is used in EID to guide the design of the visual *form in* which information should be displayed in an interface. The idea is to take advantage of operators' powerful pattern recognition and psychomotor abilities, allowing people to deploy everyday skills that have been honed 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 plant structure and dynamics that 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 the SRK framework, as follows:

1. *To support knowledge-based behavior* — represent the work domain in the form of an AH to serve as an externalized mental model that will support problem solving;
2. *To support rule-based behavior* — provide a consistent one-to-one mapping between the work domain constraints and the cues provided by the interface;
3. *To support skill-based behavior* — support interaction via time-space signals by allowing the operator to act directly on the display.

For a detailed justification of these principles, see Vicente and Rasmussen (1990, 1992). Specific examples of their application to the present interface are given later in this article.

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 language for modeling the same underlying work domain. For process control plants, five levels of constraint have been found to be of use (Rasmussen, 1985): the purposes for which the work domain 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 work domain 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 work domain purposes, whereas lower levels *represent physical* information about how those purposes are realized by equipment.

2.1. Contemporary interfaces for unit operations

Computer systems in process control rooms are continually evolving to provide the unit operations team with more, and more integrated, information[2]. A contemporary control room typically includes an information system that provides graphical representations of process data, trending capabilities, alarm management tools, and often event analyses, advanced control capabilities, and optimisation software. There are many diverse systems of this type, but it is important to differentiate EID from existing information support systems. We approach this question by returning to the three design principles introduced above.

2.1.1. Support SBB

Allowing the operator to act directly on the display through direct manipulation is an important part of supporting SBB. Touch-screen interfaces, which are almost ubiquitous in the process industries, can provide this functionality. This feature has generally been well maintained by the mouse as PCs have become more widely employed. However, both of these input devices are used almost exclusively to navigate through the displays and to select points of interest. Most interaction with the data takes place through keyboard text entry. The principle of supporting SBB encourages us to design interfaces that allow operators to interact directly with the data. For example, changing setpoints, adjusting alarm limits, and switching between control modes are all actions that could be taken through direct interaction with the display.

2.1.2. Support RBB

Many of the computer-based tools in contemporary control rooms support rule-based execution by providing a direct association between recognizable events and appropriate responses. Alarms often prompt pre-planned responses, diagnostic tools recommend control actions, and computerized procedures direct the operator to take certain actions. Supporting these associations is a critical role for an interface and contemporary displays perform this role fairly well. With the increasing use of graphical information displays comes greater opportunity to extend this rule-based support. For example, a graphical element for indicating the level in a vessel might also indicate options for reducing that level when it exceeds a high limit. In this way, the graphical element provides a cue to the operator as to what action to take based on the relevant functional constraint.

2.1.3. Support KBB

While contemporary computer tools are fairly effective at supporting the execution of anticipated tasks, they are typically poor at supporting KBB. They represent the physical constraints of the process (e.g. by way of a mimic display), but they do not comprehensively convey the functional constraints that must be considered during fault management. EID seeks to overcome this shortcoming by making all of the goal-relevant constraints visible in the interface, thereby providing a visualization of a functional model of the process that the operator can use in problem solving.

2.1.4. Summary

Contemporary interfaces for process unit operations are fairly effective at supporting SBB and RBB, although there is still room for improvement. However, these interfaces generally lack explicit support for KBB. This is exactly the sort of support that the operations team needs when managing unanticipated events. In fact, the ability of human operators to adapt to novel situations is the central justification for keeping them in the control room. Examples of how all three of these design principles are at work in the present interface are given later.

[2] Recent software applications in the petrochemical industry are geared towards total plant, or enterprise resource, management. However, to date, EID has been applied exclusively to the design of interfaces for unit operations. EID may have contributions to make to total plant management, but our present discussion should be viewed only in the context of unit operations.

2.2. Prior applications and findings

EID has been employed to create user interfaces for a number of process-related work domains. These include thermal-hydraulic process simulations (Vicente and Rasmussen, 1990; Reising and Sanderson, 1998), a power plant feedwater sub-system (Dinadis and Vicente, 1996), aviation engineering systems (Dinadis and Vicente, 1999), a conventional power generation simulation (Burns, 2000a), and full-scope simulators for a marine nuclear propulsion system (Yamaguchi and Tanabe, 2000) and a nuclear power plant (Itoh, Sakuma and Monta, 1995).

Vicente and his colleagues have reported a series of studies on EID using a thermal–hydraulic process simulation (Vicente, Christoffersen and Pereklita, 1995; Christoffersen, Hunter and Vicente, 1996; Pawlak and Vicente, 1996; Christoffersen, Hunter and Vicente, 1997, 1998). Results from these studies show that ecological interfaces lead to more consistent control performance, improved fault detection times, more accurate fault diagnoses, better adaptation to tight operating constraints, and deeper understanding of the controlled process (see Vicente, 1997 for a review). Reising and Sanderson (2000a,b), working with a different process simulation, reported results of evaluations of control performance, fault detection, and fault diagnosis that reinforce those summarized by Vicente. Burns (2000a,b) reported similar results on a larger-scale application of EID to the monitoring of a power plant simulation. Thus, a body of empirical research is forming that shows repeatable performance advantages for ecological interfaces across applications, domains, and levels of fidelity. This accumulated work has also shown that operators can effectively learn to employ non-traditional graphical representations in reasoning about the control of complex systems.

Despite these applications and encouraging results, EID cannot be introduced arbitrarily to new work domains for two reasons. First, experience in domains outside of petrochemical processing shows that scaling up from generic laboratory process simulations to larger systems is difficult (Burns, 2000a,b). Second, new application domains may have unique aspects that pose novel challenges. For example, in petrochemical processing, it is not immediately clear how chemical reactions, the behavior of catalysts, or strong pressure dependencies can be incorporated into the AH. Accordingly, a proof of concept study is required because it cannot be taken for granted that the initial successes of the design framework will scale up, and transfer over, to unit operations in the petrochemical industry. In the following section, we undertake this proof of concept by first presenting the results of an AH analysis of a petrochemical process and then introducing a novel user interface design concept for that process.

3. Ecological interface design applied to petrochemical processing

The application of EID to fluid catalytic cracking (FCC) can be summarized in two steps. First, an AH analysis of the work domain must be conducted. Second, the information content and structure specified by the AH analysis must be mapped onto a visual form. Following a brief introduction of the work domain, each of these steps is described in turn.

3.1. Work domain introduction

We employed a simulation of a fluid catalytic cracking unit (FCCU) originally detailed by McFarlane, Reineman, Bartee and Georgakis (1993). That article contains all of the governing equations for the simulation and a complete description of the model parameters. The figures in the McFarlane et al. paper have been assimilated into a pseudo piping and instrumentation diagram (P&ID) shown in Fig. 1. The simulation is a simplification of a hypothetical plant. The work domain modelling effort described here is based solely on the model description provided by McFarlane et al. and not on a general treatment of FCC processes. Thus, assumptions and simplications built into the McFarlane simulation will thereby be subsumed by the work domain description.

3.2. Work domain analysis

Fig. 2 presents an overview of the work domain analysis for the FCCU. This figure consists of one abstraction dimension (vertical) and one decomposition dimension (horizontal). Taken together, the five levels of abstraction constitute the AH discussed above. The three levels of decomposition constitute an aggregation-decomposition hierarchy. Unlike the AH, this latter hierarchy is characterized by part-whole relations between the levels. While orthogonal to the AH, it is often useful to include the aggregation-decomposition dimension because higher levels of abstraction often make more sense when described in terms of units or systems as opposed to individual components. This can be seen in Fig. 2 by noting that not all cells of the space are completed. Rather, descriptions at higher levels of abstraction tend to be most informative at higher levels of aggregation. Whereas descriptions of the FCCU are, in principle, possible at all combinations of the abstraction and aggregation levels, we chose to populate only those cells that are useful for our analysis. Each of these cells represents a different, but complete, model of the same work domain at varying levels of abstraction and aggregation. In the following paragraphs, we offer examples of the contents of the populated cells.

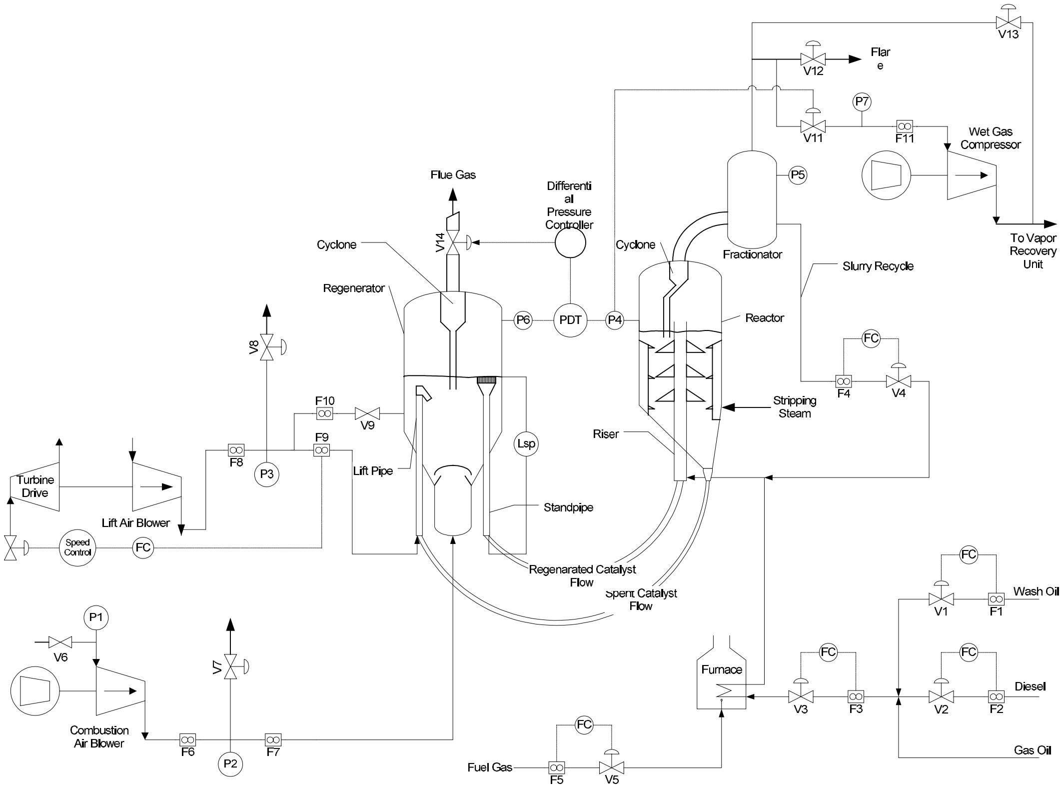

Fig. 1. The McFarlane FCCU (adapted from McFarlane et al., 1993).

3.2.1. Part-whole decomposition

Fig. 3 shows the range of representations along the part-whole dimension of Fig. 2. Each of the three levels represents a different degree of resolution of the FCCU. At the lowest level are blowers, valves, and heaters. The next levels aggregate these components into meaningful sub-systems and then the sub-systems into a whole system. The arrows at these levels represent physical inputs, outputs, or connections to, from, or between entities. Note that all three descriptions refer to the same work domain.

3.2.2. Physical form

No cells at the Physical Form (PFo) level are defined in Fig. 2 because our FCCU model does not represent physical location or appearance. In an operational plant, this level would be filled with representations of the location and appearance of the physical structures. For example, chemical plants frequently provide video images of flare towers so that operators can monitor the flare to make sure that it is not burning black. The flare image would constitute a PFo representation of one plant component. Other examples of PFo representations include photographs and other multi-media access to the appearance and location of the equipment.

Decomposition

Abstraction	System (S)	Unit (U)	Component (C)
Functional Purpose (FP)	FP-S		
Abstract Function (AF)		AF-U	AF-C (mass) AF-C (energy)
Generalized Function (GF)		GF-U	GF-C
Physical Function (PFn)			↕ PFn-C
Physical Form (PFo)			

Fig. 2. An overview of the abstraction hierarchy for the McFarlane et al. (1993) FCCU.

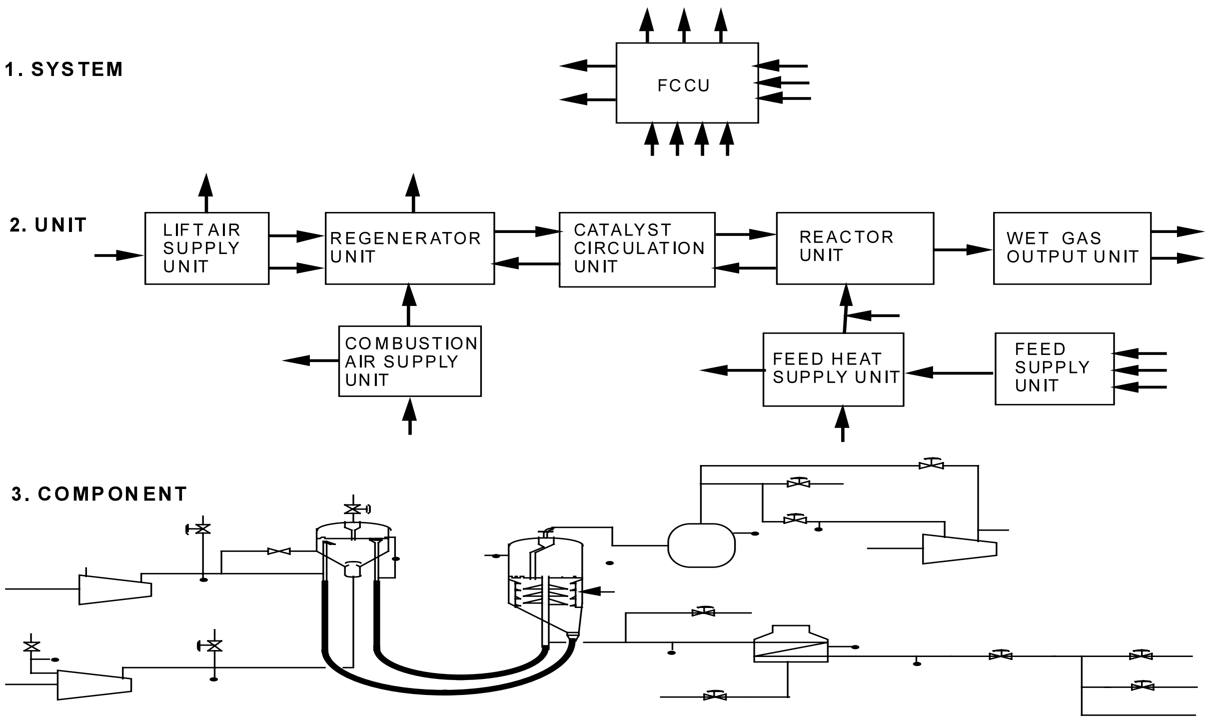

Fig. 3. The FCCU at three levels of part-whole decomposition.

3.2.3. Physical function

The Physical Function (PFn) representation abstracts from the PFo level by assigning names to the equipment and showing how they are connected. It also contains information related to the state of manipulable components. For example, a photograph of a butterfly valve would constitute a PFo representation, whereas the PFn designation 'Flow Valve V5' assigns a name (Flow Valve) and an indicator (V5) of a characteristic of that component. In this case, the indicator refers to the current state of the valve in terms of its position. The PFn-Component (PFn-C) cell representation of the FCCU is shown in Fig. 4.

The PFn representation resembles a traditional piping and instrumentation diagram (P&ID) except that instrumentation is not specified by the AH (compare with Fig. 1). Also absent from this model are the various controllers employed in the McFarlane et al. (1993) simulation. In constructing an AH, particular attention is paid to restricting descriptions to the means-ends structure of the plant elements. While control systems are crucial to the successful operation of a modern petrochemical plant, they do not lend themselves to characterisation by structural means-ends descriptions. Other analysis techniques can be employed to model the behavior of control systems (Vicente, 1999).

The lines connecting the nodes at the PFn level represent physical relationships between the structural components. For example, the two U-bend lines connecting the reactor and regenerator are meant to imply that there is a direct physical connection between these units. While this statement may strike the reader as being obvious, higher levels of abstraction do not follow this rule.

3.2.4. Generalized function

The Generalized Function (GF) level for the FCCU reveals information about heat transfers and flows of commodities (e.g. catalyst, feed products). It is the first level of abstraction that is divorced from the physical characteristics of the equipment. Here, the process is described according to the terminology of traditional disciplines of engineering. For example, descriptions of heat transfers, thermodynamics, and fluid flows are common to GF representations. Chemical reactions are also presented at this level of abstraction.

Fig. 5 shows the transition between the cells of the PFn and GF levels at the Component level of decomposition. The figure is the manifestation of the arrow between the two cells on the lower right in Fig. 2. This representation provides a very detailed explanation of how each node at the PFn level is related to its associated structural end(s) at the GF level. Conversely, each node at the GF level is connected to its structural means at the PFn level. Thus, a given node can have a single or multiple means and ends, emphasising the homomorphic (many-to-one) nature of the AH representation. Although we have not chosen to do so, similar diagrams could be created for other adjacent pairs of abstraction levels.

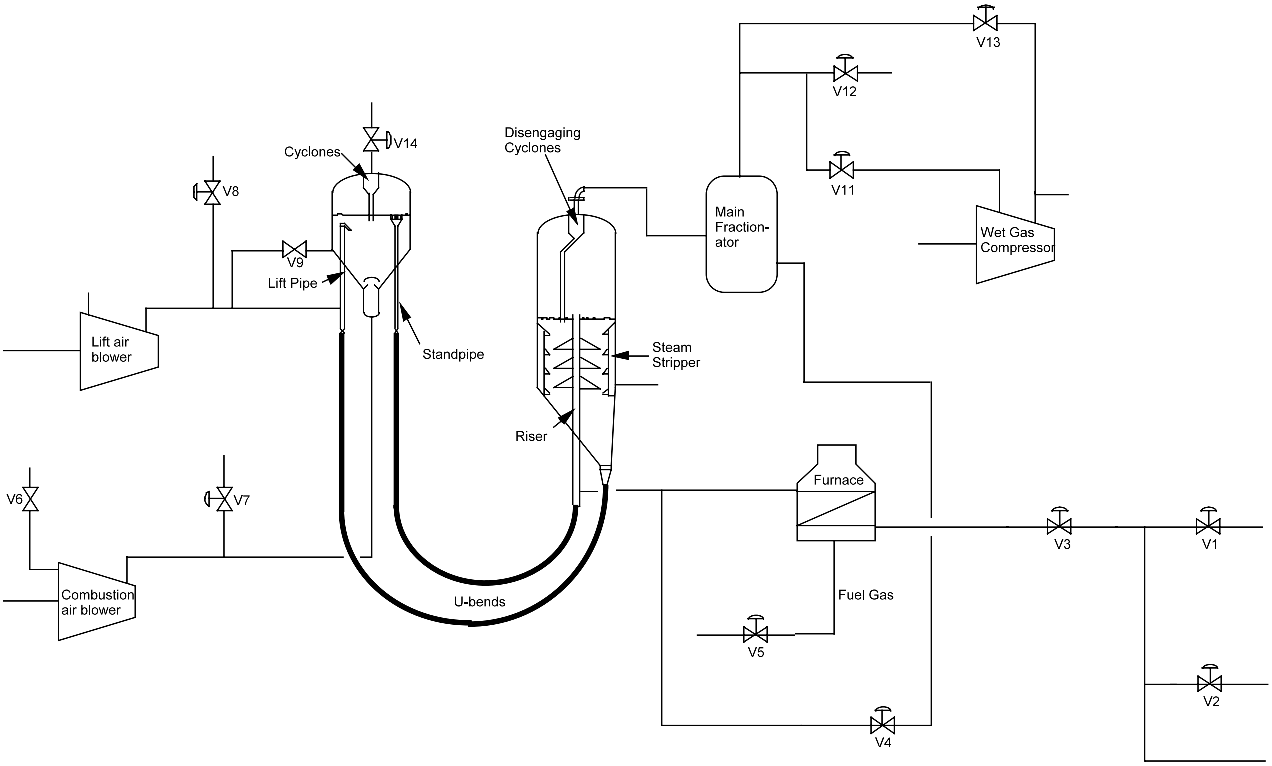

Fig. 4. The physical function level at the component level of decomposition.

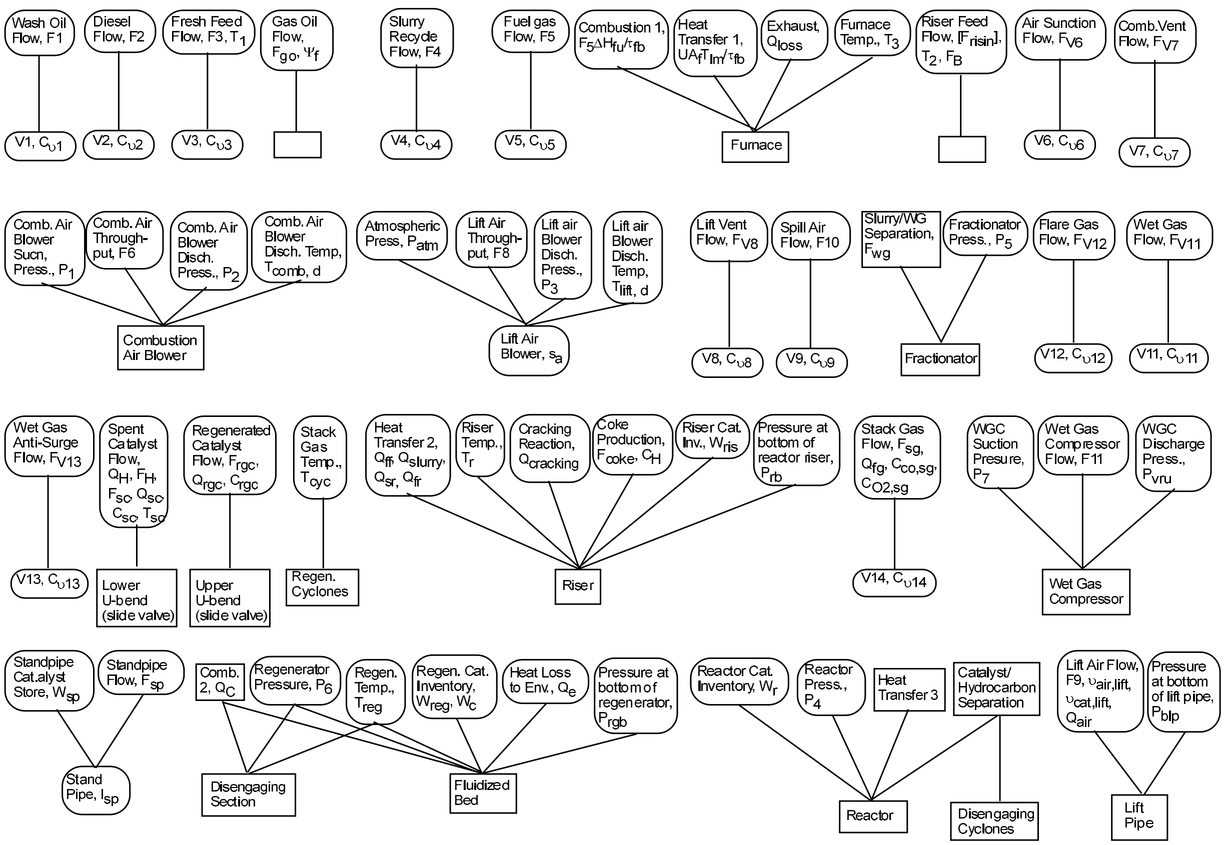

Fig. 5. Detail of transition from the physical function to generalized function level.

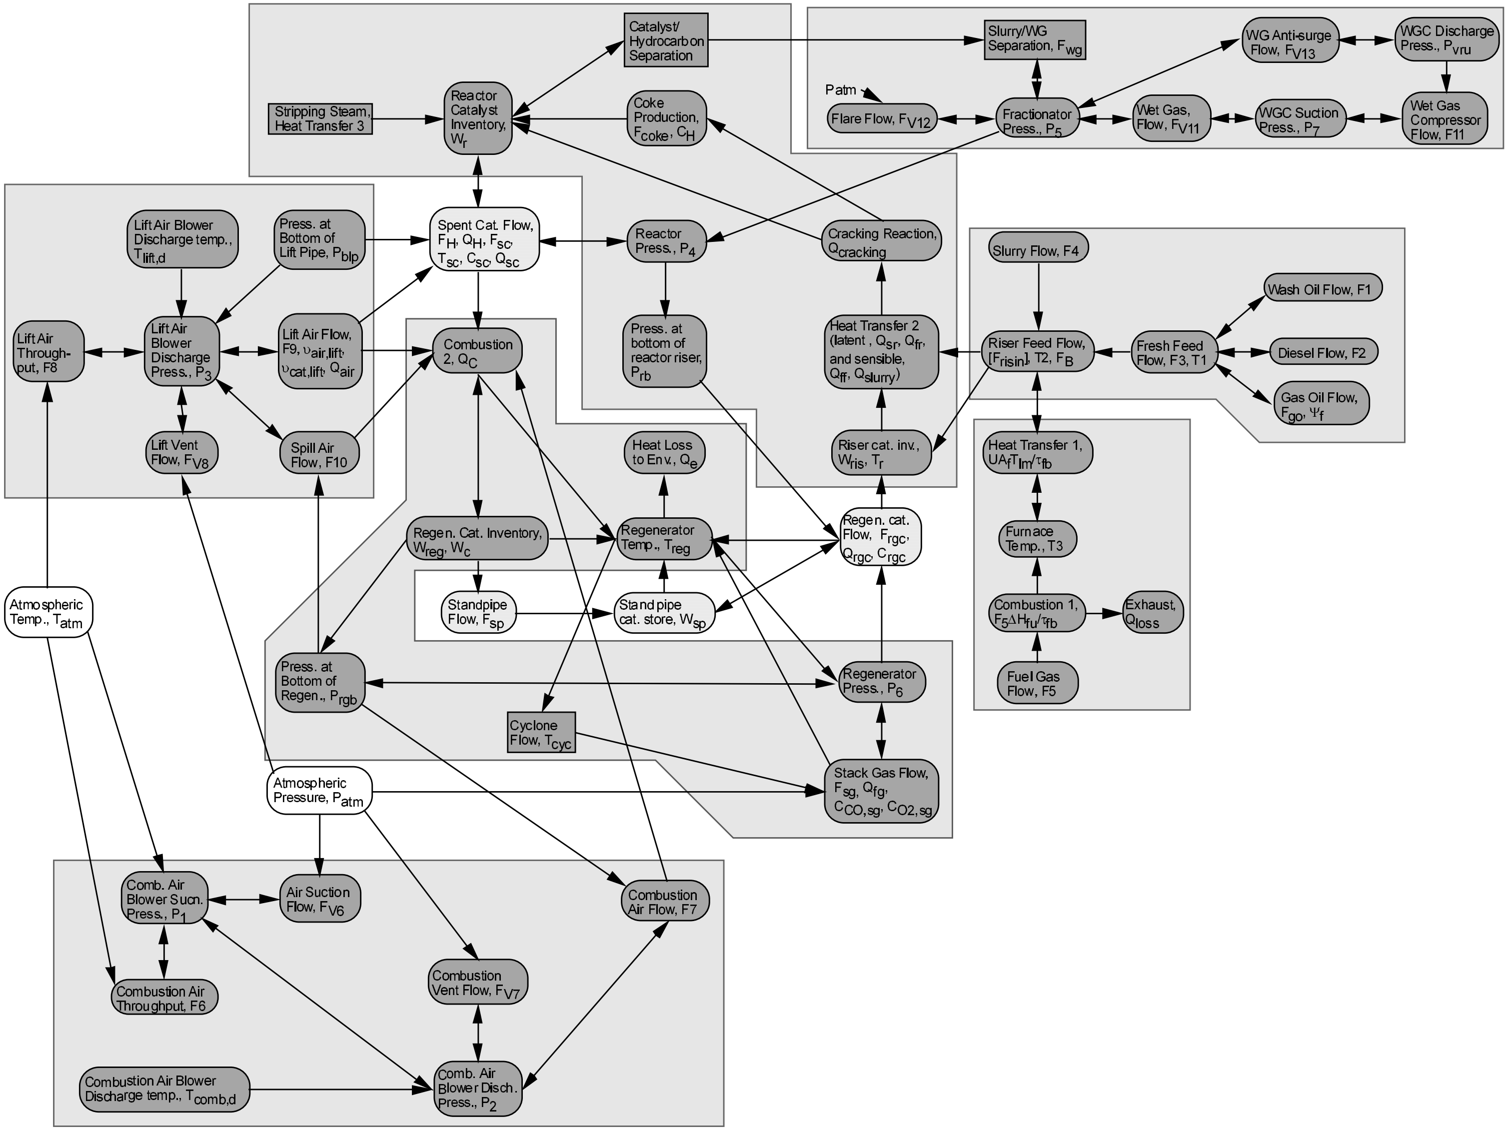

Fig. 6. The generalized function level at the component level of decomposition.

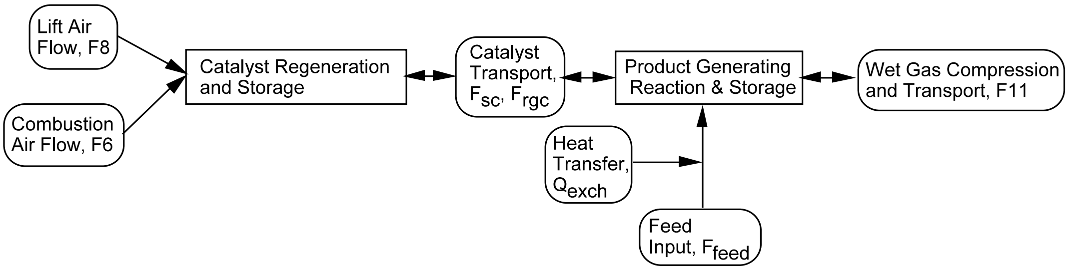

Fig. 7. The generalized function level at the unit level of decomposition.

The GF-Component (GF-C) cell representation is shown in Fig. 6. The nodes at this level represent general functions of the plant, e.g. flows, heat transfers. These nodes are the same as those that appear on the top of each of the four rows of Fig. 5. The connections between nodes in Fig. 6 represent causal relationships between functions, whereas the connections between nodes at the PFn level emphasised physical relationships between components.

The GF-C cell representation is rather complicated. The extensive connections between the nodes reflect the high degree of coupling between plant functions. The influence of pressure propagations is a major driver of this inter-dependency. Note that the complexity increases around those nodes related to the reactor and regenerator units. Whereas the functionality around the Feed Input and Heat Transfer Units is essentially sequential, the functionality in the reactor and regenerator is circular.

The GF level at the Unit (GF-U) level of decomposition is shown in Fig. 7. This representation is an aggregation of the functional relations described in Fig. 6. Note that, despite transitioning levels of decomposition, the same language is employed to describe the functions because the level of abstraction is the same. Each level of the AH employs the same language, regardless of the level of decomposition. Note that this representation is simpler than that at the GF-C level. Moving up a level of decomposition allows operators to think about the plant in fewer terms, exploiting hierarchy to reduce memory demands.

3.2.5. Abstract function

Fig. 8 shows the Abstract Function (AF) level at the unit level of decomposition (AF-U). The AF representations for the component level are too large to be included in this article (see Jamieson, 1998 for additional detail). However, Fig. 8 should be sufficient to describe this level of abstraction. The AF level reveals information about the first principles that govern mass and energy relationships in the plant. Connections between nodes at this level again reflect causality. Note that, at this level of abstraction, the various commodities are no longer distinguishable, as they are all represented as masses. Further, different modes of heat transfer are all treated as energy exchanges. Representing the functionality of the plant in terms of its mass and energy relations allows the analyst to exploit physical conservation laws. These laws cannot be expressed at the GF level because they do not hold in that context. For example, there is no such thing as a 'law of conservation of volume'.

3.2.6. Functional purpose

The overall purpose of an FCCU is to contribute to the financial viability of the installation as a whole. This function is specified as the sole node in the Functional Purpose (FP) level of abstraction (and is therefore not accompanied by a figure). This level is only described for the system level of decomposition (FP-S). In the McFarlane et al. (1993) model, this purpose can be described by a single variable, the flow rate of wet gas out of the wet gas compressor (F11 in Fig. 1). In practice, a collection of variables would have to be considered to make this assessment about an operational FCCU, including relative proportions of each product, their respective qualities, and their values on the market. However, the McFarlane model is substantially simplified in this regard.

3.2.7. Summary

Our AH contains six different representations of the FCCU (one of which is not shown here) and one detail of the connections between the representations (see Fig. 5). All of these representations must appear in our ecological interface in order to support operator adaptation, to facilitate continuous learning, and to support distributed, collaborative work among plant personnel. In the following section, we describe a novel interface design for the FCCU developed in accordance with the principles derived from the SRK taxonomy.

3.3. Prototype interface description

Woods (in preparation) has established terminology to describe the organization of space in a graphical interface. We have employed the following terms in the interface description.

1. Viewport: Any screen real estate where a process view can appear. The number of viewports limits the number of process views that can appear in parallel.

2. Process view: A coherent unit of representation of a portion of the underlying process that can be displayed in a viewport.
3. Workspace: The set of viewports and classes of process views that can be seen in parallel or in sequence.

Fig. 9 presents a workspace overview of the FCCU prototype EID. The workspace is divided into eight fixed viewports that will be described in detail in the following paragraphs. Each viewport can present one or more process views or representations that convey the physical and functional information specified in the AH.

In order to make explicit how the semantic mapping has been performed, Fig. 9 has been supplemented with an iconic representation of the FCCU AH for each of the process views that can appear in each of the eight viewports. The icons show which cells of the AH are represented by each process view and where they appear in the workspace. For example, the Paulsen displays (described below) which appear in Viewport 1 in the upper, left-hand corner of the interface are comprised of information solely from the GF-U cell of the AH (see Fig. 7). Other process views integrate content from multiple cells of the AH.

Fig. 10 shows a full screen view of the prototype ecological interface that we developed for the McFarlane et al. (1993) FCCU. The process views are laid out in the workspace roughly to correspond with iconic plant representations in two overview displays described below. Thus, the two air supply units are situated to the left (viewports 2 and 3) and the feed and feed preheat units to the right (viewport 8). The wet gas output unit occupies viewport 7 in the upper-right corner and the middle of the interface is reserved for conveying the strong interactions between the reactor and regenerator units (viewports 5 and 6). This spatial consistency between process view location and the appearance of the iconic representations in the overview displays is intended to support the operator in locating information across displays.

3.3.1. Viewport 1: Paulsen overview displays

Viewport 1 in the upper left corner of the workspace contains a Paulsen (1992) overview display of either pressure or temperature relationships. That is, the operator has a choice between viewing the Paulsen temperature process view (as shown in Fig. 10) or the Paulsen pressure process view (not shown). Each of these views is comprised of a sequential layout of the process

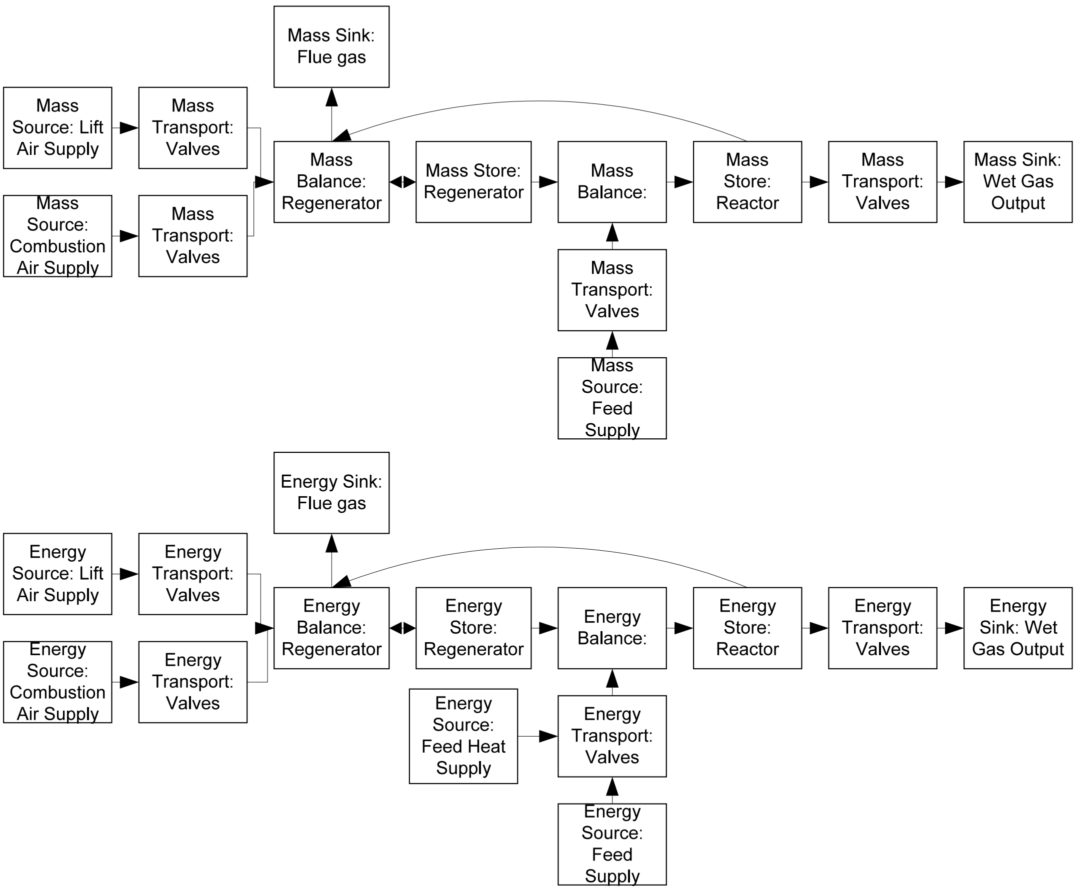

Fig. 8. The abstract function level at the unit level of decomposition (mass and energy treated separately).

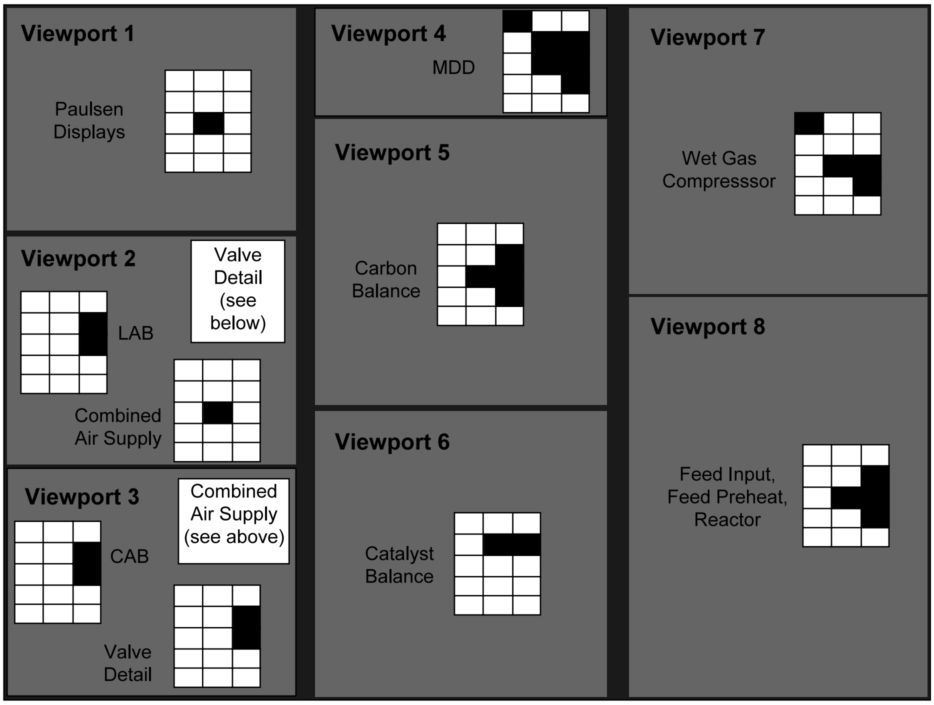

Fig. 9. An overview of the AH mappings of the ecological interface for the McFarlane et al. (1993) FCCU.

components along the horizontal axis and the parameter of interest (pressure or temperature) along the vertical axis. Thus, the graphics describe pressure or temperature profiles in spatial terms relative to the components of the FCCU. Also visible on both Paulsen displays are constraints imposed by physical or relational limitations. For example, the temperature of the furnace (T3) must remain below 1350 F. This constraint is conveyed in the Paulsen temperature view by the red arrow and horizontal line intended to suggest an upper limit to the viewer.

The role of the Paulsen display in the interface is to provide one type of overview of the process information. Given that the temperature and pressure distributions across a unit will be relatively consistent over a period of normal operations, the operator will learn that a certain profile is acceptable. Should the operational state start to change, however, the operator will be presented with a perceptual cue (i.e. an unusual profile) to that change as opposed to having to deduce it from individual temperature or pressure points on a mimic display.

3.3.2. *Viewports 2 and 3: air supply unit views and valve detail displays*

Viewports 2 and 3 communicate information about the performance of the two air blowers. The two viewports are usually populated by process views of the Lift Air Blower (LAB) and Combustion Air Blower (CAB). Both views include emergent-feature polar star displays (Coekin, 1969) depicting compressor performance (described in Viewport 7).

Another process view that can appear in viewports 2 and 3 is the Air Supply Combined Unit process view (not shown). This view provides an aggregated form of the component level air blower process views. This aggregated display allows the operator to maintain an awareness of the state of the two air supply units while opening up a viewport to display other information.

A third process view available for display in viewports 2 and 3 is the Valve Detail View (again not shown). These graphics employ a detailed three-axis representation to show explicitly how the valve setting relates to the percentage of flow, given the valve characteristic curve and the root pressure difference across the valve.

3.3.3. *Viewport 4: the mass data display*

A Mass Data Display (MDD) process view (Beuthel, Boussoffara, Elzer, Zinser and Tiβen, 1995) is located in a dedicated viewport in the top-centre of the workspace in Fig. 10 (viewport 4). The MDD is given this prominent location because it allows the operator to maintain a broad overview of the complete plant. Moreover, the MDD accomplishes this objective in a very small amount of space.

In the MDD, the state of each function and equipment has been mapped onto the orientation of a single line. These lines are grouped by unit and displayed inside iconic line diagrams of each unit. When the state of the function or equipment that is in a normal range, the corresponding line remains horizontal. However, as the state strays from normal, the line rotates counterclockwise. This rotation cannot convey precise quantitative information, but it is very effective at showing deviations from normal. Moreover, in each iconic field, process disturbances tend to appear as fields of non-parallel lines, a perceptually salient indicator of a potential problem. As process disturbances propagate

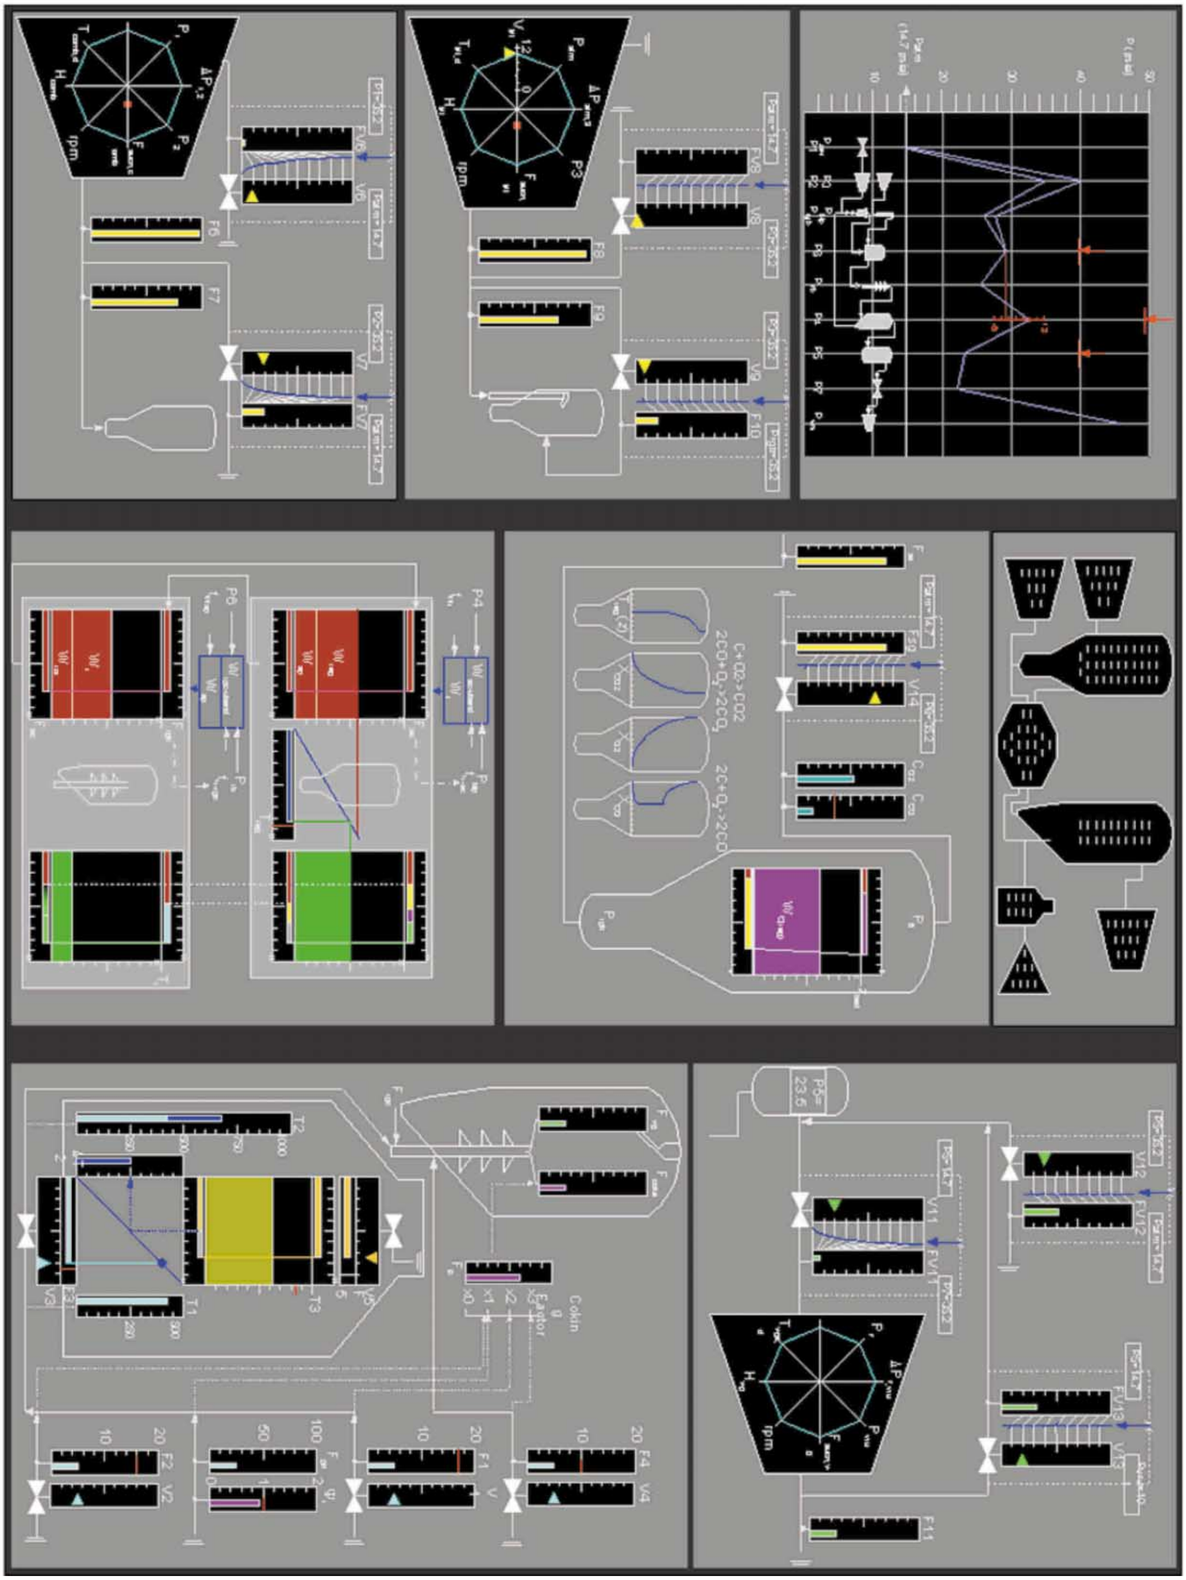

Fig. 10. The ecological interface for the McFarlane et al. (1993) FCCU.

through a plant, the MDD can also allow operators to see where the disturbance originated and how it is affecting other regions by observing the propagation of off-horizontal lines.

3.3.4. Viewports 5 and 6: balance displays

Viewports 5 and 6 are in Fig. 10 are reserved for crucial process balance information. The two views that can appear here are the Carbon Balance View (viewport 5) and the Catalyst Balance View (viewport 6). The Carbon Balance View details the critical regenerator function of burning coke off of the spent catalyst. The supply of air is related to the inventory of catalyst in the regenerator and the concentrations of coke on the spent and regenerated catalyst.

An iconic referent for the regenerator vessel contains a balance graphic first introduced by Vicente and Rasmussen (1990) and described in detail by Pawlak and Vicente (1996). Stacked bars at the top of the balance graphic denote incoming flow of mass (in this case coke). Stacked bars at the bottom of the graphic denote mass outflow. In between these stacked bars is a rectangle that increases in height to denote increasing mass inventory contained in the vessel (in this case, the regenerator). A line is drawn from the end of the stacked inflow bars to the end of the stacked outflow bars. The line depicts the emergent relationship of expected mass level gradient in the vessel. Taken as a whole, the graphic reflects a funnel metaphor. When the inflow of mass exceeds the outflow, the line is positively sloped and provides a visual cue that suggests that mass inventory should be increasing. When the outflow exceeds the inflow, the top of the funnel is larger than the bottom and mass inventory should be decreasing. This balance graphic is employed multiple times in the interface for both mass and energy relationships. The balance graphic in the Carbon Balance View is used to show the inventory of carbon in the regenerator. Carbon flows into the regenerator on spent catalyst from the reactor. Following the combustion reaction in the regenerator, a majority of the carbon leaves the regenerator as stack gas. However, depending on the extent of combustion, some carbon may remain entrained in the regenerated catalyst as it is transported back to the riser. The two mass outflows are distinguished by color; yellow for the stack gas (which matches the airflow) and maroon for catalyst (which is employed in the Catalyst Balance View described below). The relationship between these flows results in a carbon level gradient that is conveyed by the line connecting the endpoints of the stacked flow bars.

Also included in this process view are spatial temperature and gas concentration profiles. The profiles appear in four iconic representations of the regenerator located at the bottom of the process view next to the large regenerator icon. These plots are intended to allow the operator to perceive the profiles of these parameters in the regenerator. A visualisation of these profiles can assist the operator in understanding the progress of the chemical reactions that take place in the regenerator.

The Catalyst Balance View (viewport 6) is a graphical representation of the catalyst cycle in the FCCU. This view allows the operator to perceive where the catalyst is concentrated in the plant and how effectively it is being transported between the reactor and the regenerator. In addition, this display conveys the critical energy relations between these two vessels.

Four balance graphics are employed in the Catalyst Balance process view; two for the mass and energy inventories in the regenerator and two for the mass and energy inventories in the reactor (Vicente and Rasmussen, 1990). The relationship between the inventory of mass and energy in the regenerator vessel (the top two balance graphics) is shown to dictate the vessel temperature (Pawlak and Vicente, 1996). The mass level on the left side of the process view determines the slope of the rotating bar (fixed to a scale placed between the balance graphics) that reflects an incident beam determined by the energy level to the right. The reflected beam falls on the scale to denote a predicted temperature value in the vessel. A temperature bar shows the measured regenerator temperature, allowing the operator to compare the measured and predicted values. A red line on the temperature scale between the two balance graphics also depicts a physical constraint on the temperature in the regenerator.

The model equations for the reactor temperature suggest a different relationship from that shown by the regenerator. The temperature is treated as a product of the energy inflows and outflows, independent of the mass inventory in the vessel. The result of this simplification is a less coupled graphic for the regenerator mass and energy functions. The two balance graphics for reactor mass and energy do not have a rotating reflecting bar. Each of these functions is still important, but their relationship does not determine the temperature in the riser.

The Catalyst Balance process view also provides information relating to the force balance on the catalyst in the two U-bend lines. A free-body diagram is used above each pair of balance displays to specify the forces acting on each mass and the resultant flow of regenerated catalyst into the reactor and spent catalyst into the regenerator. A box located on top of the reactor and regenerator mass balance graphics represents the relative magnitude of mass for each flow. The boxes are divided according to the two mass contributions and arrows are drawn to represent forces acting on them. Following the free-body diagram metaphor, if the forces on one side of the box exceed those on the other, the box will move with the greater force. An arrow extending down from the box points to the predicted mass flow associated with that mass/force combination. The operator can compare the predicted value to a measured value depicted by the mass inflow bar.

3.3.5. Viewport 7: wet gas output unit view

The upper-right hand corner of the workspace is occupied by the wet gas compressor (WGC) process view. The WGC process view contains the sole representation of the plant's FP level of abstraction. The wet gas flow to the vapor recovery unit VRU (F11) allows operators to ascertain whether or not they are meeting their production goals.

Included in each of the LAB, CAB, and WGC process views is a polar star display (described above) depicting the compressor operating variables. This display consists of 8 scales for variables related to compressor performance. The scales are arranged in a symmetrical spoke pattern. The current value for each of the variables is noted along these scales. In addition, a thin line is drawn to connect each value to its two neighbors. The result is an emergent geometric figure that characterizes the performance of the compressor. When the scales are normalized to the steady state operating characteristics of the compressor, a symmetric octagon indicates a compressor that is operating normally. When the measured variable values deviate from the expected, the connecting lines will no longer form a regular octagon. This disfiguring of the display allows operators to rely on their perceptual skills to alert them to abnormal situations (see Dinadis and Vicente, 1999 for an example).

The polar star configuration was selected for two reasons. First, air compressors demonstrate tight local coupling between a set of variables and a relatively unidimensional impact (gas flow rate) on neighbouring components. The polar star allows for a representation of that coupling by connecting normalized values of those local parameters. Second, the relative size of the emergent octagon conveys a perceptual cue as to the volume of gas that is being moved by the compressor.

3.4. Viewport 8: semantic mapping in the feed input, feed preheat, and reactor unit views

How does this prototype ecological satisfy the three design principles set out by the EID framework? Fig. 11, an expanded image of viewport 8, is provided to answer this question.

3.4.1. Supporting SBB: the operator should be able to act directly on the display

The right side of the process view includes three valve and flow representations. Each is depicted using a scale to indicate the flow rate through the valve and an accompanying scale for the valve setting, indicated by a triangle. This design allows the operator to manipulate the valve setpoint directly by moving the triangle along the scale. This direct manipulation of controls is often preferable to entering settings using numeric keypads because it reduces the opportunity for input errors. Note how this sort of direct interaction with the interface differs from navigation or point selection using a touch-screen or mouse. The operator is able to interact directly with the components.

3.4.2. Supporting RBB: provide a consistent one-to-one mapping between the work domain constraints and the cues provided by the interface

An example of the support for RBB is the furnace temperature balance graphic in the lower, left-hand corner of Fig. 11. The temperature of the flow into the furnace is shown on a vertical scale to the right in cyan. The product flow through the heater is shown on a horizontal scale below, also in cyan. A vertical line connects this flow to a blue rotating bar that is fixed at the intersection of the product flow scale and a delta-T scale (at left, in blue). The orange portion of the orange/grey line on top shows the amount of energy transferred to the fluid in the tubes. The relevant constraint to be displayed is that the change in temperature of the product stream across the furnace is a function of the inlet temperature, flow rate, and rate of energy transfer. The dashed line that descends from the heat transfer line down to the rotating bar and across to the delta-T scales depicts this relationship. By mapping this constraint onto a graphical element, an operator can visualize the two means of affecting the outlet temperature; by decreasing the flow rate, which will cause the blue line to rotate left and thereby project the dashed line at a higher point on the delta-T scale; or by increasing the heat transfer in the furnace (i.e. by increasing the length of the orange line). A different graphical element is provided to show what constraints are associated with increasing the length of the orange line.

3.4.3. Supporting KBB: represent the work domain in the form of an AH to serve as an externalized mental model that will support problem solving

The furnace in Fig. 11 is presented as a graphical model to support an understanding of how the manipulation of feed flows and fuel gas combine to yield temperature of feed entering the reactor riser. An energy balance graphic is employed to relate how energy released from fuel gas combustion is either transferred to the feed or lost to the environment. The portion that is transferred to the feed results in a temperature increase across the heater. The dynamics of this relationship are shown by a pivoting reflecting arm fixed on the inflow meter whose angle is dictated by the rate of feed

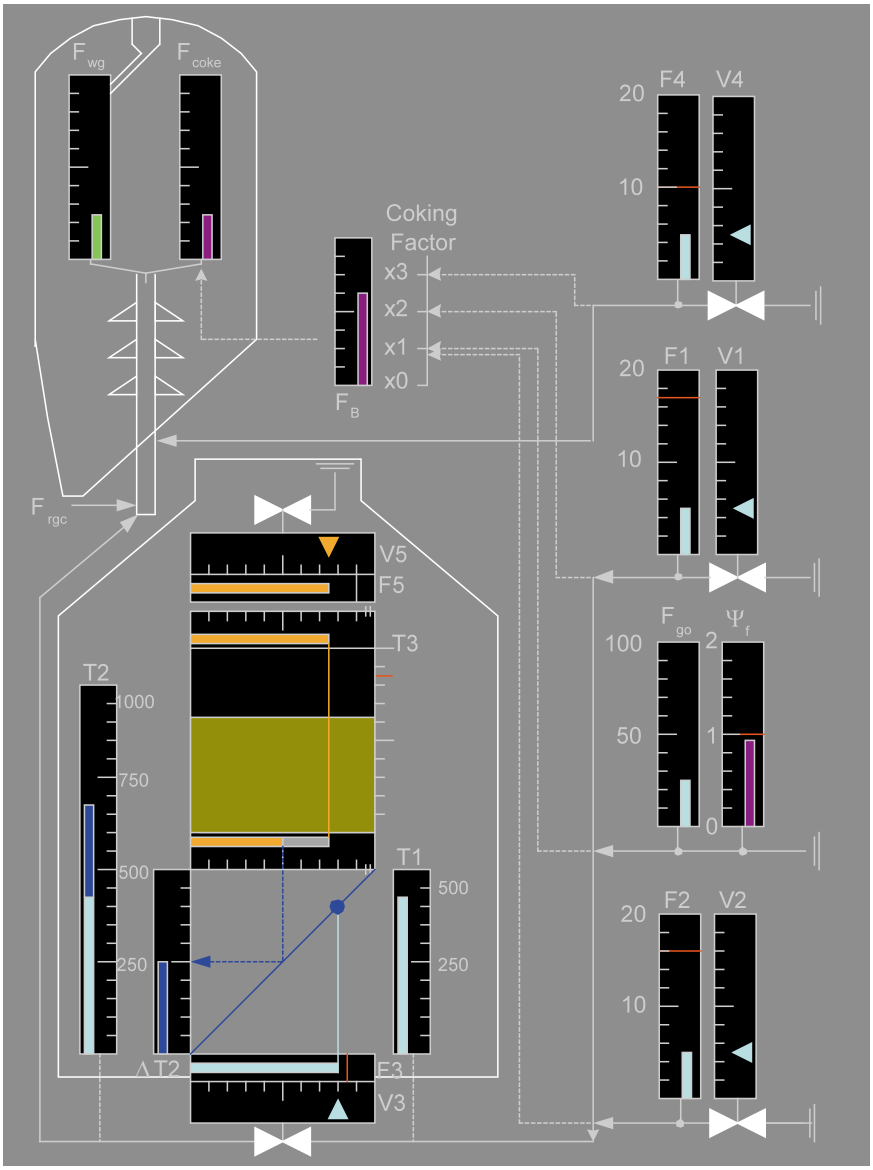

Fig. 11. Feed input, feed preheat, and reactor unit views.

flow. A line is extended down from the end of the energy outflow bar that depicts energy transferred to the feed. This line reflects off the reflecting bar and falls on a scale that indicates temperature change across the heater. This amount is added to the temperature of the incoming feed and displayed on a tall temperature scale to the left of the balance graphic.

3.5. Limitations

This research is a first time proof of concept that EID can be effectively applied to the petrochemical domain. In drawing this conclusion, however, we should note four limitations to the work presented here. First, the ecological interface for the McFarlane et al. (1993) FCCU has neither been implemented nor tested. However, valuable such findings would be, we did not set out to confirm that the performance advantages observed in other applications of EID would also manifest for this design. Implementing and testing an ecological interface for an operational petrochemical unit remains as an important area of future research.

Second, most FCCUs in service are unlikely to have the sensor set necessary to provide all of the data required by the prototype interface. Two factors play heavily in this shortage; technology and cost limitations. The former pertains primarily to the mass and energy functions specified at the AF level of the AH. Because no measurement technology exists to determine these values directly, they would have to be derived.

Other functions may not be measurable or derivable. The cost limitation pertains to the capital and support costs associated with sensors. Adding additional flow meters, for example, would require an investment in the hardware, would increase instrumentation support costs, and would add opportunities for instrumentation-based faults. Reising (1999) has examined the implications of limited sensor sets and their impact in the implementation of ecological interfaces in process control.

A third limitation of this work is that control information is not included because it is deliberately not modeled by the AH framework. Clearly, automation technology plays a crucial role in petrochemical processing from the regulatory level, to model-based predictive control, to enterprise optimization. To be considered for regular use in production facilities, EID must be expanded to include all of these types of control information. This represents a critical area for future research.

Finally, EID in general (and this work in particular) does not offer a systematic methodology for mapping the constraints identified in the AH onto the visual form. No existing interface design methodology for complex systems that offers such a method. Although the SRK framework provides a scaffold to guide the design of a visual form, the designer must fill in many details. Reising and Sanderson (1998) discuss the possibility of making this process of semantic mapping more systematic.

4. Alternative approaches

For many readers, using the AH framework to model the physical and functional relationships of a petrochemical process will be unfamiliar. Therefore, it may be helpful to compare and contrast the characteristics of the AH with those of a modeling framework that is more familiar to the chemical engineer. The signed directed graph (SDG) framework (Iri, Aoki, O'Shima and Matsuyama, 1979; Umeda, Kuriyama, O'Shima and Matsuyama, 1980; Wilcox and Himmelblau, 1994) employs a graph-theoretic approach to model the cause-effect relationships of a target process by representing process variables and parameters as graph nodes. Immediate causal relations are depicted as directed arcs between those nodes with values of (+) and (−) indicating whether the connected nodes change in the same or opposite direction, respectively. Current states for the parameters are described relative to a nominal steady state and reduced to values of (0), (+), or (−). Vedam and Venkatasubramanian (1997, 1999) have constructed a SDG for the same simulation of a Model IV FCCU that is employed in this article.

4.1. A comparison of models

The most telling distinction between the AH and SDG frameworks is the purposes for which they are employed. Whereas the AH is intended as an aid to operator problem solving, the SDG framework has generally been employed in support of automated fault diagnosis (e.g. Iri et al., 1979; Umeda et al., 1980; Wilcox and Himmelblau, 1994; Vedam and Venkatasubramanian, 1997, 1999). One of the advantages of its use in this capacity is that, when employed in an executable form, the SDG will provide an exhaustive evaluation of possible root causes for faults. However, this presents a problem in that the method lacks resolution, often returning multiple possible root causes for single faults. In response to this difficulty, "A knowledge base consisting of knowledge about reliability of equipment, infeasible root nodes and information about equipment maintenance is used to improve the resolution of diagnosis," (Vedam and Venkatasubramanian, (1999) p. 909). Using this sort of knowledge base introduces event-specific knowledge that will constrain the fault diagnosis algorithm in a non-causal, probabilistic manner. In many cases, this proves to be an efficient way to zero in on a correct fault diagnosis. However, in the case of unanticipated faults, such information can lead to inaccurate diagnoses that subsequently mislead the operations team.

In contrast, the AH retains its event-independence by excluding potentially relevant information about known faults and predicted likelihoods of failure. When employed in EID, the role of the AH is to identify the content and structure of the information to be contained in an operator interface. The responsibility for detecting an abnormality and resolving the location of a fault is assigned to the human operator. The interface designer supports this activity by ensuring that the relevant process constraints are made visible in the interface. The operator is free to incorporate probabilistic information.

It is important to emphasize that this key difference lies in the purpose of the application of the frameworks, as opposed to differences in their structure. In fact, both frameworks have the potential to yield event independent models because they both distinguish process constraints from system states. While the need to distinguish between constraints and states may appear obvious to the builder of an SDG or AH, some fault propagation models rely exclusively on state-based representations. Once a state-dependent component is introduced to the model, it loses its objectivity in examining the sorts of unanticipated events that can be of greatest danger to the process. In the SDG, the constraint/state distinction is notationally explicit with constraints being represented as nodes and links, while states are represented by categorical values. In the AH, constraints are also represented by nodes and directional links. State variables are identified and appended to nodes as descriptors.

A second distinguishing feature between the frameworks is the treatment of controls. The SDG framework employs specific methods for including the functions of controllers in the causal modeling. In contrast, the AH intentionally excludes the activity of automatic controls in the construction of the model for reasons discussed previously.

These crucial differences notwithstanding, there are a number of similarities between the frameworks. First, both frameworks include cause-effect models, although with differing levels of prominence. For the SDG, cause-effect is the only relation used to describe the process and the analysis is performed at a single level of aggregation (i.e. the process variable level). In contrast, the cause-effect model in the AH framework is of secondary importance; appearing exclusively in the topological relations described in the GF and AF levels (see Figs. 6 and 8). (In fact, at the GF-C level, the representations formed by the AH and SDG frameworks are very similar.) Because it is intended to support goal-directed operator problem solving, the AH framework focuses primarily on the structural means-ends relations and maintains these relations consistently across five levels of analysis. Thus, cause-effect models are a subset of the models produced by the AH framework. In comparison, the SDG does not describe the structural means-ends or part-whole relations that are captured by the AH. Thus, it does not provide as solid a basis for supporting human problem solving activities, nor was it originally intended to.

Finally, models in both frameworks may be derived from a combination of characteristic equations of the process and a combination of operating and engineering expertise. In both cases, the former is preferred and was afforded in the case of the Model IV FCCU because McFarlane et al. (1993) provide explicit process equations in their description. However, in many practical applications, such detailed mathematical descriptions are not available, and both cause-effect and functional relations must be extracted from the knowledge of plant engineers and operators. However, it should be emphasized that if prior operating experience is used to introduce relations between process states to the model, the model can no longer be deemed event independent.

4.2. Complementary frameworks

This comparison is not intended to suggest that the SDG could not be used in support of operator display design. Nor does it preclude the use of the AH in the design of a machine-based constraint violation detection algorithm (see Bisantz and Vicente, 1994). Rather, the two modeling frameworks exhibit characteristics that make them better suited for certain applications, as do all modeling frameworks. In fact, this explains why many applications of SDG are coupled with additional modeling and analytical frameworks that offer complementary strengths (e.g. Mylaraswamy and Venkatasubramanian, 1997; Vedam and Venkatasubramanian, 1997, 1999). Moreover, it also suggests that the AH can be meaningfully coupled with additional modeling techniques to assist in automated fault detection. Integrating automated fault diagnosis algorithms into EID is an outstanding issue for future research.

5. Achieving effective operator support through EID

This research represents the first comprehensive application of EID to petrochemical processing. In the final section of this article, we argue that EID can address the three requirements for operator support tools that we laid out in the introduction.

5.1. Support for unanticipated events

EID provides a basis for helping operators cope with unanticipated events by using the AH to support problem solving. The rationale is very similar to that behind analytical redundancy techniques in control theory (Frank, 1990). The AH represents the various layers of goal-relevant constraints that govern a process control plant when it is operating normally. When a fault occurs, one or more of those constraints will be violated, even if the fault is unfamiliar to operators and unanticipated by designers. Because it is based on an AH representation, the interface will show the variables that enter into the goal-relevant constraints, thereby allowing operators to compare the behavior that is expected from the plant given the constraints with the current behavior of the plant. Any mismatch (i.e. a residual) is an indication that the work domain is not behaving as it should be (see Vicente, 1999 for more details).

5.2. Support for continuous learning

EID also provides a basis for continuous learning by mapping the goal-relevant constraints identified by the AH representation onto the perceptual features of the interface. By creating a transparent, visualization of otherwise opaque, abstract properties (e.g. mass and energy flows; chemical reactions), we are providing a rich source of feedback to operators about the structure and behavior of the plant. This feedback serves as an important input to learning processes. In effect, operators can learn every day on the job because the feedback they obtain from the interface provides them with a salient indication of how the plant responds to their or the automation's actions. In addition, feedback is also provided about the relationships that hold between variables under various operating regimes. Thus, as changes are made to the plant in the name of continuous improvement, an interface based on EID should foster continuous learning because operators should be able to update their understanding of the work domain based on the rich and salient feedback provided by the interface.

Consider, for example, the mass and energy balance graphics in viewport 6 (see Fig. 10). The constraint relationships described by this display are amongst the most complicated and critical in the operation of the unit. While operators will be skilled at monitoring and controlling this process, they are unlikely to have the abstract, technical expertise required to understand the mass and energy relationships at play. The advantage of the ecological interface is that these relationships are made accessible in graphical terms. Through the course of their daily monitoring and control activities, operators will see how changes in catalyst mass and energy will impact regenerator temperature. They will also be able to see how the forces acting on the catalyst result in its circulation between the reacting and regenerating vessels. This exposure to relevant properties of their work environment affords an opportunity for operators to continually improve their understanding of the relationships that make the unit work effectively.

5.3. Support for distributed, collaborative work

EID also provides a basis for supporting distributed, collaborative work by using the multiple levels of the AH to represent the varying information that is required by different stakeholders in the work domain. Each level of the AH provides a different language for modeling the FCCU, and thus reveals different insights into the structure and behavior of the plant. This conceptual diversity provides a basis for accommodating the well-known fact that different individuals with different responsibilities tend to think about the plant in different ways. The AH contains these models *as well as* the linkages between them so that people can coordinate and communicate effectively.

For example, managers tend to be concerned with meeting productivity and environmental regulation constraints, both of which would be represented in the Functional Purpose level of the AH. In the present case, this information is constrained to the wet gas flow rate in viewport 7. The Plant engineers, on the other hand, tend to be more concerned with optimizing the production of the plant by thinking in terms of mass and energy balances, which are represented at the Abstract Function level of the AH. They will find this information in mass and energy balances in viewports 5, 6, and 8. In contrast, control room operators tend to be more concerned with the status and configuration of plant functions and components, which are represented at the Generalized Function and Physical Function levels of the AH, respectively. Along with the overview in viewport 4, this information is contained in all of the viewports (as the interface is primarily intended for panel operators). Finally, field operators will be more concerned with where particular components are located and what their appearance is, properties that are represented at the Physical Form level of the AH (and not contained in the interface).

These examples show that EID provides a global, shared representation-in the form of the AH — that can support communication and coordination across individuals that are geographically distributed and that have diverse concerns, and thus, information needs. The interface discussed here does not provide an explicit architecture for supporting this collaboration, however. Nor does it provide a mechanism for coordinating the contributions of persons with varying types of expertise. What it does do is present the control problem through an integrated series of representations that cut across the various stakeholders, placing the interests of those stakeholders in a common context.

6. Conclusions

In this paper, we have undertaken a proof of concept to show that EID can be applied to petrochemical processing. This represents the first application of EID to this challenging domain. In building an AH representation for a prototype FCCU, we specified the content and structure of the information that had to appear in an interface to support effective operator interaction. In mapping that information to a graphical representation through the SRK taxonomy, we grant operators the ability to interact with the information using all three of the ways that they have for processing information.

The EID framework is a promising design framework for coping with the challenging trends that are being observed in the petrochemical processing industry. We have argued why there are good reasons to believe that this approach to human–computer interface design can support operator adaptation, foster continuous learning, and facilitate distributed, collaborative work. Most of these arguments have been supported by data collected under representative laboratory conditions in other application domains. However, additional research is required to determine whether the anticipated benefits of EID can be realized in operational settings in the petrochemical industry.

Acknowledgements

This research was sponsored by research grants from the Natural Sciences and Engineering Research Council of Canada, NOVA Chemicals, and Honeywell Technology Center. We are grateful to the members of the Abnormal Situation Management Consortium for their insightful comments throughout the progress of the research, particularly Jamie Errington of Nova Chemicals and Peter Bullemer of the Honeywell Technology Center. We also wish to thank two anonymous reviewers for their insightful comments on previous drafts.

References

Beuthel, C., Boussoffara, P., Elzer, P., Zinser, K., & Tiβen, A. (1995). Advantages of mass-data-displays in process s&c. In: Proceedings of the 6th IFAC/IFIP/IFORS/IEA Symposium on Analysis, Design, and Evaluation of Man-machine Systems, Cambridge, MA:IFAC, pp. 439–444.

Bisantz, A. M., & Vicente, K. J. (1994). Making the abstraction hierarchy concrete. *International Journal of Human-Computer Studies*, *40*, 83–117.

Burns, C. M. (2000a). Putting it all together: improving display integration in ecological displays. *Human Factors*, *42*, 226–241.

Burns, C. M. (2000b). Navigation strategies with ecological displays. *International Journal of Human-Computer Studies*, *52*, 111–129.

Christoffersen, K., Hunter, C. N., & Vicente, K. J. (1996). A longitudinal study of the effects of ecological interface design on skill acquisition. *Human Factors*, *38*, 523–541.

Christoffersen, K., Hunter, C. N., & Vicente, K. J. (1997). A longitudinal study of the effects of ecological interface design on fault management performance. *International Journal of Cognitive Ergonomics*, *1*, 1–24.

Christoffersen, K., Hunter, C. N., & Vicente, K. J. (1998). A longitudinal study of the impact of ecological interface design on deep knowledge. *International Journal of Human-Computer Studies*, *48*, 729–762.

Coekin, J.A. (1969). A versatile presentation of parameters for rapid recognition of total state. In: *Proceedings of the IEE international symposium on man-machine systems* (no pagination). Cambridge, England: IEE.

Dinadis, N., & Vicente, K. J. (1996). Ecological interface design for a power plant feedwater subsystem. *IEEE Transactions on Nuclear Science*, *43*, 266–278.

Dinadis, N., & Vicente, K. J. (1999). Designing functional visualizations for aircraft systems status displays. *International Journal of Aviation Psychology*, *9*, 241–269.

Frank, P. M. (1990). Fault diagnosis in dynamic systems using analytical and knowledge-based redundancy — a survey and some new results. *Automatica*, *26*, 459–474.

Iri, M., Aoki, K., O'Shima, E., & Matsuyama, H. (1979). An algorithm for diagnosis of system failures in the chemical process. *Computers and Chemical Engineering*, *3*, 489–493.

Itoh, J., Sakuma, A., & Monta, K. (1995). An ecological interface for supervisory control of BWR nuclear power plants. *Control Engineering Practice*, *3*, 231–239.

Jamieson, G.A. (1998). Ecological interface design for petrochemical processing applications. Unpublished M.A.Sc. Thesis, University of Toronto, Toronto, Canada.

McFarlane, R. C., Reineman, R. C., Bartee, J. F., & Georgakis, C. (1993). Dynamic simulator for a model IV fluid catalytic cracking unit. *Computers and Chemical Engineering*, *17*, 275–300.

Mylaraswamy, D., & Venkatasubramanian, V. (1997). A hybrid framework for large scale process fault diagnosis. *Computers and Chemical Engineering*, *21*(Suppl.), S935–S940.

Paulsen, J.L. (1992). *A pipe reactor design and the interface system* (RISØ-M-2931). Roskilde, Denmark: Systems Analysis Department, Risø National Laboratory.

Pawlak, W. S., & Vicente, K. J. (1996). Inducing effective operator control through ecological interface design. *International Journal of Human-Computer Studies*, *44*, 653–688.

Pool, R. (1997). *Beyond engineering*: *how society shapes technology*. New York: Oxford University Press.

Rasmussen, J. (1985). The role of hierarchical knowledge representation in decision making and system management. *IEEE Transactions on Systems*, *Man*, *and Cybernetics*, *SMC-15*, 234–243.

Rasmussen, J. (1986). *Information processing and human-machine interaction: an approach to cognitive engineering*. New York: North Holland.

Reising, D.V.C., & Sanderson, P.M. (1998). Designing displays under ecological interface design: towards operationalizing semantic mapping. In: *Proceedings of the 42 annual meeting of the human factors and ergonomics society*, Santa Monica, CA:HFES, pp. 372–376.

Reising, D.V.C., & Sanderson, P.M. (2000a). Testing the impact of instrumentation location and reliability on ecological interface design: fault diagnosis performance. In: *Proceedings of the IEA 2000/HFES 2000 congress*, *Volume 3*, Santa Monica, CA: HFES, pp. 591–594.

Reising, D.V.C., & Sanderson, P.M. (2000b). Testing the impact of instrumentation location and reliability on ecological interface design: control performance. In: *Proceedings of the IEA 2000/HFES 2000 congress*, *Volume 1*, Santa Monica, CA: HFES, pp. 124–127.

Umeda, T., Kuriyama, T., O'Shima, E., & Matsuyama, H. (1980). A graphical approach to cause and effect analysis of chemical processing systems. *Chemical Engineering Science*, *35*, 2379–2388.

Vedam, H., & Venkatasubramanian, V. (1997). Signed digraph based multiple fault diagnosis. *Computers and Chemical Engineering*, *21*(Suppl.), S655–S660.

Vedam, H., & Venkatasubramanian, V. (1999). PCA-SDG based process monitoring and fault diagnosis. *Control Engineering Practice*, *7*, 903–917.

Vicente, K. J. (1997). Operator adaptation in process control: a three-year research program. *Control Engineering Practice*, *5*, 407–416.

Vicente, K. J. (1999). *Cognitive work analysis*: *toward safe*, *productive*, *and healthy computerbased work*. Mahwah, NJ: Erlbaum.

Vicente, K. J., Christoffersen, K., & Pereklita, A. (1995). Supporting operator problem solving through ecological interface design. *IEEE Transactions on Systems*, *Manand Cybernetics*, *SMC-25*, 529–545.

Vicente, K. J., & Rasmussen, J. (1990). The ecology of human-machine systems II: mediating 'direct perception' in complex work domains. *Ecological Psychology*, *2*, 207–250.

Vicente, K. J., & Rasmussen, J. (1992). Ecological interface design: theoretical foundations. *IEEE Transactions on Systems*, *Man*, *and Cybernetics*, *SMC-22*, 589–606.

Vicente, K.J., & Tanabe, F. (1993). Event-independent assessment of operator information requirements: Providing support for unanticipated events. In: *Proceedings of the topical meeting on nuclear plant instrumentation*, *control and man-machine interface technologies*, LaGrange Park, IL: American Nuclear Society, pp. 389–393.

Wilcox, N. A., & Himmelblau, D. M. (1994). The possible cause and effect graphs (PCEG) model for fault diagnosis — 1. Methodology. *Computers and Chemical Engineering*, *18*, 103–116.

Woods, D.D. (in preparation). Visualizing function: the theory and practice of representation design in the computer medium.

Yamaguchi, Y., & Tanabe, F. (2000). Creation of interface system for nuclear reactor operation: practical implication of implementing EID concept on large complex system. In: *Proceedings of the IEA 2000/HFES 2000 congress*, *Volume 3*, Santa Monica, CA: HFES, pp. 571–574.