Cognitive
Engineering
Laboratory

ECOLOGICAL INTERFACE DESIGN FOR PETROCHEMICAL PROCESS CONTROL: INTEGRATING TASK- AND SYSTEM-BASED APPROACHES

Greg A. Jamieson

CEL 02-01

Cognitive Engineering Laboratory, Department of Mechanical & Industrial Engineering, University of Toronto
5 King's College Road, Toronto, Ontario, Canada M5S 3G8
Phone: +1 (416) 978-7399 Fax: +1 (416) 978-3453
Email: benfica@mie.utoronto.ca URL: www.mie.utoronto.ca/labs/cel/

Cognitive Engineering Laboratory

Director: Kim J. Vicente, B.A.Sc., M.S., Ph.D., P. Eng.

The Cognitive Engineering Laboratory (CEL) at the University of Toronto (U of T) is located in the Department of Mechanical & Industrial Engineering, and is one of three laboratories that comprise the U of T Human Factors Research Group. CEL was founded in 1992 and is primarily concerned with conducting basic and applied research on how to introduce information technology into complex work environments, with a particular emphasis on power plant control rooms. Professor Vicente's areas of expertise include advanced interface design principles, the study of expertise, and cognitive work analysis. Thus, the general mission of CEL is to conduct principled investigations of the impact of information technology on human work so as to develop research findings that are both relevant and useful to industries in which such issues arise.

Current CEL Research Topics

CEL has been funded by Atomic Energy Control Board of Canada, AECL Research, Alias|Wavefront, Asea Brown Boveri Corporate Research - Heidelberg, Defense and Civil Institute for Environmental Medicine, Honeywell Technology Center, Japan Atomic Energy Research Institute, Microsoft Corporation, Natural Sciences and Engineering Research Council of Canada, Nortel Networks, Nova Chemicals, Rotoflex International, Westinghouse Science & Technology Center, and Wright-Patterson Air Force Base. CEL also has collaborations and close contacts with the Mitsubishi Heavy Industries and Toshiba Nuclear Energy Laboratory. Recent CEL projects include:

- Developing advanced human-computer interfaces for the petrochemical industry to enhance plant safety and productivity.
- Understanding control strategy differences between people of various levels of expertise within the context of process control systems.
- Developing safer and more efficient interfaces for computer-based medical devices.
- Creating novel measure of human performance and adaptation that can be used in experimentation with interactive, real-time, dynamic systems.
- Investigating human-machine system coordination from a dynamical systems perspective.

CEL Technical Reports

For more information about CEL, CEL technical reports, or graduate school at the University of Toronto, please contact Dr. Kim J. Vicente at the address printed on the front of this technical report.

ECOLOGICAL INTERFACE DESIGN FOR PETROCHEMICAL PROCESS CONTROL: INTEGRATING TASK- AND SYSTEM-BASED APPROACHES

by

Greg A. Jamieson

A thesis submitted in conformity with the requirements for the degree of Doctor of Philosophy

Graduate Department of Mechanical and Industrial Engineering

The University of Toronto

Ecological interface design for petrochemical process control: Integrating task- and system-based approaches. Greg A. Jamieson, Doctor of Philosophy, 2003, Department of Mechanical and Industrial Engineering, University of Toronto

ABSTRACT

Abnormal events in petrochemical plants cost the industry billions of dollars annually. In part, these events are difficult to deal with because current interfaces do not adequately inform operators about the state of the process. Ecological human-machine interfaces use a system-based analysis to identify higher-level process functions that should be presented in the operator interface. Several laboratory simulator studies have shown that, in comparison with contemporary interfaces, ecological interfaces can lead to faster fault detection, more accurate root-cause diagnosis, and more effective control responses. However, ecological interfaces differ from more traditional human centered interfaces that use a task-based analysis to inform the design process. In this dissertation, the first ecological interface for an industrial system to integrate both system- and task-based information was designed and implemented. A second ecological interface was created, drawing exclusively from the traditional system-based analyses. The process of integrating the two types of analyses and using them to inform the designs is discussed. Professional operators used the novel interfaces in an industrial petrochemical process simulator to monitor for, diagnose, and respond to several types of process events. Operators using the traditional ecological interfaces completed trials more quickly and executed fewer control actions than their counterparts using current process displays. Operators using the integrated (task- and system-based) ecological interface also showed these benefits, and in addition, showed improved fault diagnoses and better performance scores. The implications and opportunities for introducing ecological interfaces into industrial control rooms are discussed.

TABLE OF CONTENTS

TABLE OF TABLES

TABLE OF FIGURES

LIST OF APPENDICES

ACKNOWLDGEMENTS

This research was funded by cash and in kind grants from the Natural Science and Engineering Research Council, NOVA Chemicals, Ltd., and Honeywell, Inc. It was conducted in close association with engineers, operators, and managers at the Nova Chemicals, Ltd. Ethylene 1 plant in Joffre, AB. The author is indebted to Jeff Cole, Joanne Cundict, Keath Doe, Jamie Errington, Doug Lutz, and Matthew Sevenpifer for the many hours of technical, engineering, and organizational assistance that they provided. The interface design was conducted in collaboration with Wayne Ho, an M. Eng. candidate in MIE. Robert Chang, a University of Toronto undergraduate Computer Engineering student, developed the interface prototype that was used for the formative evaluation. Ron Blair of Blair Consulting implemented the displays. Mandy Li, a University of Toronto undergraduate Industrial Engineering student, assisted with the testing of the display implementation. Tamara Arenovich served as a statistical consultant

My deepest thanks go to Kim Vicente for creating the ultimate environment for learning about the philosophy and practice of our science. His ideas and contributions permeate this dissertation. Thanks are also due to Mark Chignel and John Shaw for their insightful contributions as committee members. Finally, Alex Kirlik provided a valuable external appraisal of this work and Paul Milgram served as a voting member of the final examination committee.

No words of thanks can repay Adeline for the love, support, and patience that she gave to me throughout.

GLOSSARY

ADS: Abstraction-Decomposition Space

AF: Abstract Function

AF-C: Abstract Function - Component

AHR: Acetylene Hydrogenation Reactor

C2: Hydrocarbons with two carbons, e.g. acetylene (C2H2), ethylene (C2H4)

CTA: Control Task Analysis

DCS: Distributed Control System

DL: Decision Ladder

E1/E2: Ethylene 1 and Ethylene 2 production units

GF: Generalized Function

GF-C: Generalized Function - Component

GUS: Global User Station

EID: Ecological Interface Design

H2/CO: Hydrogen and carbon-monoxide flow

HTA: Hierarchical Task Analysis

IR: Information requirement

MOV: Motor-operated valve

P+F: Physical and Function ecological interface

P+F+T: Physical, Functional, and Task ecological interface

PFm: Physical Form

PFn: Physical Function

TCT: Trial Completion Time

WRC: Weight Ratio Controller

INTRODUCTION

Purpose

This dissertation is situated at the intersection of a practical problem in process control and a theoretical problem in cognitive engineering. The practical problem is the management of abnormal situations by process operations personnel. Cochran & Bullemer (1996) defined an abnormal situation as any process disturbance that requires operator action to restore the plant to a normal operating condition. Their estimate places the annual cost of abnormal situations in the petrochemical industry at USD $10B in preventable losses. One challenge in recovering those losses is overcoming the failure of user interface technologies to match advances in automation technology and plant complexity (Bullemer & Nimmo, 1994).

Having identified user interface technologies as an area for improvement in addressing losses to abnormal situations, we encounter a theoretical problem in cognitive engineering. That is, a) what work analysis methods should be used to generate the design requirements for these interfaces, and b) how can those requirements be translated into effective visual forms? Vicente (1999) pointed out that work analysis methods for complex systems have typically fallen into two general categories; task-based and system-based[1] approaches. Whereas these approaches have usually been applied independently, Vicente (1999) demonstrated that the

---

[1] Readers familiar with Cognitive Work Analysis (Rasmussen, Pejtersen, & Goodstein, 1994; Vicente, 1999) will note that we are using the term "system-based analysis" in place of "work domain analysis". There are two reasons for this. First, it establishes a parallelism between work domain analysis as an instance of system-based analysis and hierarchical task analysis as an instance of task analysis. Second, it is consistent with the prior efforts of Miller and Vicente (1999). For the purposes of this dissertation, the terms "system-based analysis" and "work domain analysis" are more or less equivalent.

advantages of each approach are complementary and that additive benefits could be derived from integrating task- and system-based approaches. Miller and Vicente (1999) initiated a study to determine whether these methods could be integrated in the context of a representative petrochemical process. While their efforts again showed that the methods yielded complementary information requirements, the results of the analyses were neither integrated nor used to develop novel interfaces. Thus, it remains to be demonstrated that the combined results of task- and system-based approaches to work analysis can be effectively translated into a single, integrated interface form that is understandable to process operators and compatible with their sensory and cognitive limitations (Vicente, 1990). The primary objective of this dissertation is to determine whether, and under what conditions, there is an advantage to adding task-based information requirements to system-based ecological interface design for petrochemical process control.

Abnormal Situations. A fundamental distinction for the research described in this dissertation is the classification of events that occur in process operations. Vicente and Rasmussen (1992) classified events according to two dimensions: 1) whether the event is familiar or unfamiliar to an operator, and 2) whether an event is anticipated or unanticipated by designers of the process. Based on these two distinctions, they identified three event classes: 1) familiar and anticipated events (often addressed by standard operating procedures), 2) unfamiliar but anticipated events (often addressed by emergency operating procedures), and 3) unfamiliar and unanticipated events (not addressed by procedures). (The fourth category, events that are familiar to operators but unanticipated by designers might include known deviations from proscribed operating practice. This category of events is not discussed by Vicente and Rasmussen (1992) and is not addressed here). Vicente and Rasmussen (1992) note that the most severe accidents in process operations have historically fallen into the category of unfamiliar and unanticipated events.

While Vicente and Rasmussen (1992) spoke of events in terms of whether they are Familiar or Unfamiliar to operators, empirical studies of EID have distinguished between events from a systems perspective as either Normal or Fault events. Industry shares this system-centric view, using the terms Normal and Abnormal, wherein an abnormal situation is defined as any process disturbance that requires operator action to return the process to a normal operating condition (Cochran & Bullemer, 1994). While acknowledging that an abnormal situation could conceivably be familiar to operators (however undesirable that might be) or vice versa, the similarity between the terminology in prior empirical work and in industry leads suggests that the terms are largely consistent. In this work, the Normal/Abnormal terms are employed as opposed to the Familiar/Unfamiliar. The difference appears to be minor and not highly relevant in this case.

Two Approaches to Content Specification. A second fundamental distinction in this dissertation involves the task- and system-based approaches to work analysis discussed above. Task-based methods are more prevalent in human factors engineering. They include methods such as Plan Goal Graph (Geddes, 1989) and Hierarchical Task Analysis (HTA; Shepherd, 1989). System-based analysis methods are less common, although they have enjoyed increasing attention in the past decade (e.g., Vicente, 1999). These include the Abstraction Decomposition Space (ADS; Rasmussen, 1985) and Multi-level Flow Modeling (MFM; Lind, 1994). The particular concern of this dissertation is determining the advantage of integrating task- and system-based approaches as opposed to comparing or selecting between them.

Vicente (1999) has characterized the choice between task- and system-based methods as a trade-off between efficiency and robustness. The object of analysis in task analysis is known events, and the tasks, activities, or actions that must be performed in order to meet higher level goals in the context of those events. By focusing on what the analyst knows to be critical to successful task completion, the process of analysis and design allows for efficient handling of known events. This efficiency comes from the analyst embedding her acquired knowledge in interfaces and/or procedures that the operator can then use in response to those events.

In contrast, the object of analysis in system-based methods is the process for which an information system is to be designed. The analysis is conducted independent of the events, tasks, activities, and actions that might be relevant. The advantage of this approach is that the results of the analysis are more robust. That is, they are applicable to all events, even those that the analyst is unable to anticipate (Vicente & Rasmussen, 1992). Comparatively speaking, it can be said that task-based approaches trade-off robustness in favor of efficiency whereas system-based methods trade-off efficiency in favor of robustness.

Interface design can be generically described as having two stages: 1) the specification of information content and structure via work analysis, and 2) the subsequent semantic mapping of those requirements onto visual forms. Task-based approaches to content specification can be coupled with various techniques for developing form (e.g., Direct Manipulation Interfaces (Hutchins, Hollan, & Norman, 1986; Schneiderman, 1983), emergent feature displays (Buttigieg & Sanderson, 1991; Carswell & Wickens, 1987), the Proximity Compatibility Principle (Wickens & Andre, 1990; Wickens & Carswell, 1995), and perceptual organization techniques (Woods, in preparation)). In contrast, the Ecological Interface Design (EID) framework couples system-based analysis using ADS [2] with semantic mapping according

[2] The term "abstraction hierarchy" is more commonly used in the cognitive engineering literature than the more accurate "abstraction decomposition space". However, the terms are usually redundant and the reader can treat them as being so for the purposes of this dissertation.

to Rasmussen's (1983) Skills, Rules, Knowledge (SRK) taxonomy (Vicente & Rasmussen, 1992). Thus, while EID also takes advantage of the various form-developing techniques noted above, the SRK taxonomy provides a theory-based structure for the semantic mapping.

While their paper on the EID framework focuses on the use of the ADS, Vicente and Rasmussen (1992) anticipated the eventual need to expand the framework to include the products of other work analyses to direct interface content. However, EID has predominantly been attempted using ADS exclusively (although the Bookhouse project represents a notable counter-example; see Rasmussen, Pejtersen, & Goodstein, 1994). Thus, adding task-based content to EID represents an extension of the framework as commonly practiced. It is important to identify whether and how this integration can be accomplished. Also, an investigation of the differential advantages (if any) of integrating task-based information into the framework should be undertaken.

Conclusion. Taken together, the two fundamental issues noted above form the foundation of the work described here. The purpose is to identify the relationship between three different categories of process events and two different approaches to interface content specification. Vicente and Rasmussen (1992) predicted that EID would be most beneficial in supporting operators in managing events that are both abnormal (hence, unfamiliar to them) and unanticipated by designers. Several studies have confirmed that ecological interfaces lead to better performance in such events as opposed to normal events (see Vicente, 2002). However, although both single and multiple fault scenarios have been compared, the scenarios employed in those studies have not explicitly distinguished abnormal events that are anticipated by designers (Anticipated/Abnormal) from abnormal events that are unanticipated (Unanticipated/Abnormal). Further, no study of EID has used scenarios for which operating procedures were provided to the participants. In this study, the three event categories

(Anticipated/Normal, Anticipated/Abnormal, Unanticipated/Abnormal) were distinguished and evaluated experimentally between interfaces that either included or excluded task-based information. Thus, an important theoretical prediction of EID (i.e., that EID is most beneficial in supporting unanticipated abnormal events) was tested.

Corollary Issues

In addition to the two core issues of event classification and design approaches, this work addresses several corollary issues of importance to cognitive engineering research.

State of the art in operator support tools in practice. Current interfaces in the petrochemical process industry generally fail to provide higher-level information to operators in a form that is both effective and usable (Jamieson & Vicente, 2001). The existing interface for the target process tested here is typical of contemporary process displays. The display scheme is based primarily on a process mimic graphic that places some setting and flow values in their physical context. Trending capability and tabular alarm summary pages are also provided to support the mimic displays. Very little higher order information is made available and human factors standards for interface design are often not employed. In most cases, experienced operators designed these displays based on reviews of piping and instrumentation diagrams and drawing from their prior experience as users.

This research represents an alternative approach to this interface design practice. In addition to the hypotheses mentioned in the previous sections, we anticipate that professional operators using the novel ecological interfaces will perform better on a variety of event scenarios than their counterparts working with the conventional interface. The researchthereby furthers the state of the art.

Fidelity and scale. Much of the EID research completed to date has been conducted with a single representative microworld simulation with several ecological interface variations (Vicente, 1997). Reising (1999) also constructed a novel ecological interface for a laboratory simulation of similar complexity and extended the research by investigating the impact of sensor reliability and availability. Burns (2000) designed an ecological display (i.e., with no operator control capability) for a more complex desktop process simulation and evaluated different approaches to display integration. All of these studies employed undergraduate students as participants. Sharp and Helmicki (1998) tested ecological displays for neonatal intensive care monitoring with both medical students and resident physicians. Ham and Woon (2001a, 2001b) conducted an empirical evaluation of ecological interface content in a nuclear plant simulator, but again with students and without employing the configural graphics that have typified the visual form of other ecological displays. Finally, Toshiba has also developed an ecological interface for a full-scale nuclear power simulation (Itoh, Monta, Sakuma, & Makino, 1993). However, no empirical evaluation of this interface has been reported. Similarly, Yamaguchi and Tanabe (2000) describe an interface for a shipboard nuclear power simulator. Again, no empirical evaluation of this interface has been published.

This dissertation builds upon these efforts in terms of fidelity and scale. In terms of number of process variables, the target system is of a complexity comparable to that employed by Burns (2000). However, the use of a full-scale industrial simulator allowed for the development of interactive interfaces (i.e., with both display and control capabilities). The use of professional operators as opposed to students gives the study face validity in generalizing results to industry. In a further enhancement to generalizability to industry, this study looked at abnormal events that were both anticipated and unanticipated. Prior studies have not included proceduralized events, either normal or abnormal. Finally, the simulator used in the present study runs on the same digital instrumentation and control system that is employed in the actual plant. Thus, the challenges of implementing an ecological interface in practice have been realistically addressed and overcome. This scaling up brings EID closer to industrial practice and operations than prior investigations (see Vicente, 2002).

New Challenges for Content Integration and Parsing. The experience of creating an ecological interface with both system- and task-based content yields insight into a final corollary issue, namely the process of consolidating the content called for by the two analyses and then parsing it into graphical views. The integration process has not been a central problem in previous ecological interface designs because only system-based methods of content specification were used. The qualitatively different characteristics of the system- and task-based content may act as a barrier to effective integration. The parsing challenge is an inherent part of interface design and is faced by all designers. However, the introduction of information that was particularly relevant to specific tasks complicates the parsing decision because the two approaches lend themselves to different parsing techniques. System-based approaches have usually been matched with parsing along functional dimensions (see Burns, 2000; or Reising, 1999) whereas task-based approaches are more compatible with parsing based on events or activities.

Starting Points: Review of Miller & Vicente

This dissertation was based on a series of prior analytical studies by Miller and Vicente (1998a, 1998b, 1998c, 1998d, 1999a, 1999b, 2001). A brief review of this line of research will help to place the present study in context. Miller and Vicente (1998a, 1998b) set out to determine whether, and to what advantage, task- and system-based modeling and interface design techniques could be integrated. Their work started with a task-based modeling activity, with the object of evaluation being the DURESS II thermal hydraulic process simulator and testbed (Miller & Vicente, 1998a). A system-based modeling of DURESS II had already been completed (Bisantz & Vicente, 1994). Subsequently, Miller and Vicente (1998c, 1999a, 2001) provided a comparison of the types of information identified by each of the analyses. They concluded that the two analysis techniques each produced unique information requirements that had complementary strengths and weaknesses.

As a next step, the comparison was scaled up to a moderately complex, real-world system. In consultation with an industry sponsor, an acetylene hydrogenation reactor (AHR) was selected as a representative petrochemical process for analysis. The AHR is a critical sub-section of the ethylene refining process, whereby acetylene is removed from the ethylene product stream through a catalytic chemical reaction. A system-based analysis of the AHR in NOVA Chemicals Ltd.'s Ethylene 1 (E1) unit was performed, and information requirements were generated (Miller & Vicente, 1998d). The analysis was conducted using the abstraction-decomposition space (ADS) framework (Rasmussen, 1985). The ADS is a two-dimensional hierarchical modeling framework that supports system-based analysis. The two dimensions are the means-ends (i.e., abstraction) and part-whole (i.e., decomposition) hierarchies. Moving through the levels of abstraction means moving from more concrete descriptions of the system to more abstract descriptions, although each description is complete at its given level. Similarly, moving through the levels of aggregation means moving from a component-wise description of the system to a description of the system as a whole unit. Taken together, the two dimensions describe a physical and functional constraint space that governs opportunities for operator behavior.

To complete the comparison, Miller and Vicente (1999b) then performed a task-based analysis of the same process, and compared the resulting information requirements generated from each analysis. The method selected for this analysis was Hierarchical Task Analysis (HTA: Shepherd, 1989). HTA breaks down tasks into a part-whole hierarchy of stem and leaf plans of action. It expresses sequence constraints between the plans and denotes the role of various actors in carrying them out.

With these two analyses for the AHR completed, Miller and Vicente (1999b) once again concluded that, even when completed sequentially, system- and task-based analysis techniques each contributed unique and complementary display requirements.

The products of the previous line of research included two separate sets of information requirements for interfaces to support monitoring and control of the NOVA E1 AHR. NOVA's entire E1 process (i.e., including the AHR) was also modeled in a high fidelity process simulation that was used for training operators. Taken together, these resources formed the basis for the dissertation described here. In the following chapter, the process of transforming those two sets of information requirements into a consolidated list of interface content requirements is described. Because of its extent and detail, the actual process and product of the interface design activity is documented in a technical report (Jamieson & Ho, 2001).

The resulting ecological interfaces and the current E1 AHR interface were subjected to an empirical evaluation with professional operators. A description of the method of that study follows the information requirements analysis. Subsequently, the quantitative and qualitative results of the study are reported. A final chapter draws conclusions from the empirical study and comments on their bearing on the research objectives.

INFORMATION REQUIREMENTS ANALYSIS

The process of extending, modifying, and integrating the information requirements generated by Miller and Vicente (1999b) is described in this chapter. This process culminates in a comprehensive list of information requirements that inform the subsequent interface design step. Before proceeding with this discussion, however, a description of the AHR process is required.

Acetylene Hydrogenation Reactor

System description: The target system for this prototype ecological interface is an Acetylene Hydrogenation Reactor (AHR), a sub-system of the NOVA Chemicals Ltd. E1 Ethylene refining process (see Figure 1). Nova operations personnel refer to the AHR as 'the reactor section'. The terms 'AHR' and 'reactor section' are used interchangeably in this dissertation. The AHR is a moderately complex domain by contemporary chemical engineering standards. It is, however, arguably the most critical part of the ethylene refining process. Upsets in the AHR are relatively common and costly, being the single most frequent cause of upsets in the overall ethylene process and with downtime costing roughly $1000/minute (Miller & Vicente, 1998).

The ethylene refinery converts a stream composed mainly of ethane (C_2H_6) into ethylene (C_2H_4). The natural gas is first heated in pyrolytic furnaces where some of the ethane is separated into ethylene, hydrogen (H_2), and trace substances including acetylene (C_2H_2) and carbon monoxide (CO). Subsequent sub-systems treat the process flow and extract impurities. The reactor section is located in the latter half of the refining process. The AHR receives a C2 stream consisting of ethane, ethylene, and various trace elements such as acetylene. Subsystems downstream of the AHR separate the ethane and ethylene from the trace elements,

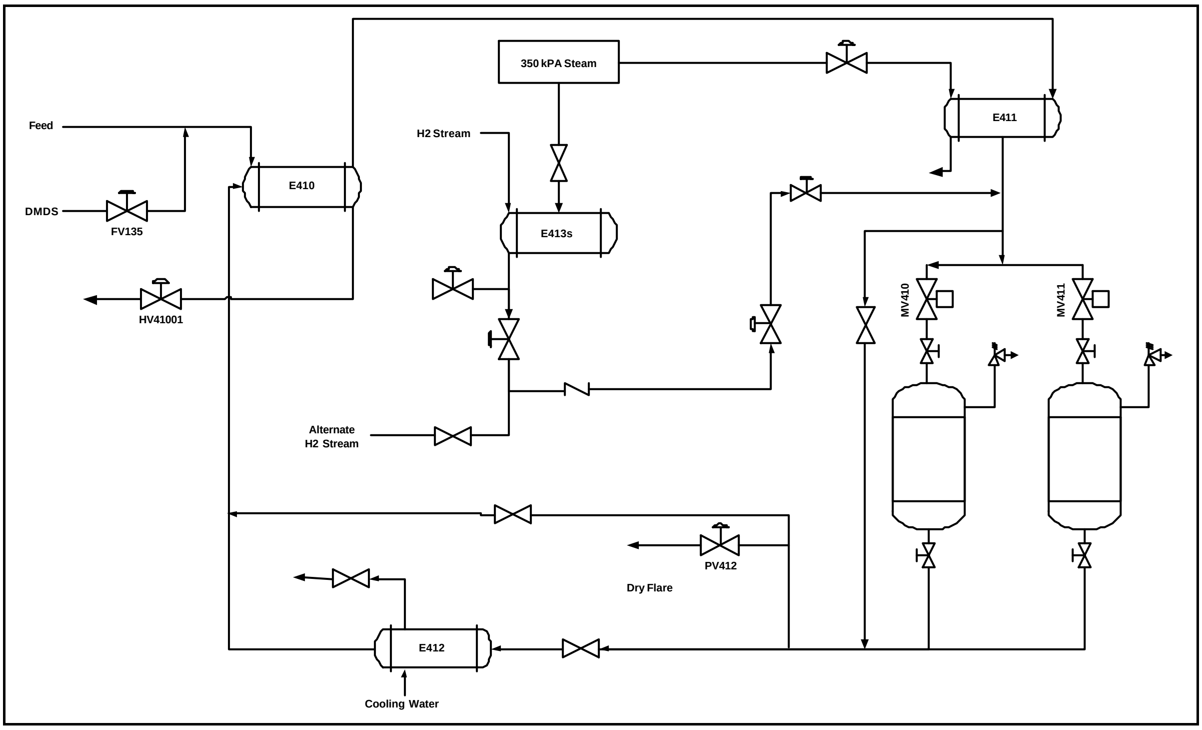

Figure 1: NOVA's E1 acetylene hydrogenation reactor.

but these processes are particularly sensitive to acetylene in the process flow. The purpose of the AHR is to remove this acetylene. To accomplish this, the C2 stream is reacted with hydrogen to convert the acetylene (C2H2) into ethylene (C2H4) in a process called hydrogenation.

<u>Process Description:</u> The process flow for the AHR is described next. The reader is encouraged to follow along with Figure 1.

1. Raw natural gas feed enters the E1 facility and undergoes pyrolysis in a number of furnaces (not shown in Figure 1). The AHR operator controls the addition of DiMethyle DiSulfide (DMDS) in these furnaces. DMDS reduces CO production, which impacts the reaction (see below).
2. The gases enter the AHR system in two streams. The feed stream consists primarily of ethylene (C2H4) and ethane (C2H6) with trace amounts of acetylene (C2H2). The H2 stream consists primarily of H2 and CO.

3. The H2 stream is heated in E413s and then routed to an intersection with the C2 stream after E411.
4. An alternate H2 stream is available if needed.
5. Prior to its intersection with the H2 stream, the C2 stream is heated twice. The first time is via E410, which uses hot effluent from the reactor (see below) to heat incoming C2 feed. The second is the steam heater E411.
6. The two streams intersect and flow into one or both of the two reactor vessels.
7. The reactor vessels are filled with a fixed catalyst bed which promotes the following reactions:
   - $C_2H_2 + H_2 \rightarrow C_2H_4$ + heat = "acetylene conversion to ethylene"
   - $C_2H_4 + H_2 \rightarrow C_2H_6$ + heat = "ethylene conversion to ethane"
   - (with a great deal of heat and/or pressure) $C_2H_4 \rightarrow C + CH_4$ + an even greater deal of heat = "ethylene decomposition"
8. Acetylene conversion is desired. Ethylene conversion is undesirable, but tolerable in small quantities. Ethylene decomposition is highly undesirable and dangerous.
9. The catalyst has many weak and a few strong sites. Precedence for a reactant being adsorbed on catalyst sites is as follows (assuming adequate H2):
   1. CO on strong
   2. CO on weak
   3. Ethylene on strong
   4. Acetylene on strong
   5. Acetylene on weak
   6. Ethylene on weak
10. Thus, the reactor should be managed as follows:

- Ensure that there is enough CO in the reactor to occupy all of the strong sites:
    - otherwise, ethylene will occupy strong sites and be converted to ethane. This is both inefficient (youare trying to maximize ethylene content) AND dangerous because excess ethylene conversion can use up available H2 leaving none for acetylene conversion.
- Minimize excess CO so as to avoid occupying weak sites.
    - CO on weak sites can mean not enough sites available for the acetylene reaction.
- Control the ratio of H2 feed to C2 feed (and the heat of both) to minimize ethylene conversion while sustaining acetylene conversion:
    - too little H2 (and/or too little heat) and there will notbe enough for total acetylene conversion,
    - too much H2 (and/or too much heat) and, after all acetylene conversion, ethylene conversion will occur producing undesirable ethane.

11. Increased heat increases the rate of all reactions. This increases the overall activity of the catalyst, but reduces selectivity. Heat in the reactors can be increased by increasing the heat of the incoming gas streams, which, in turn, can be accomplished by increasing heat transfer in E410 and E411.
12. Other reactions are possible given the presence of trace elements in the feed that have the effect of making the catalyst unreactive.
13. After reaction, the effluent flows downstream to E412. This cooler can be bypassed.
14. After E412, the reacted product stream is routed through E410 where it heats the incoming C2 stream as described above. After E412, the reacted, cooled product stream proceeds out of the AHR subsystem via HV41001.

15. Once the two input streams are mixed, they can be diverted to flare manually or automatically. Flare paths include manual valves attached to the reactors, and a control valve after the reactors (PV412). The mixed stream can also be bypassed around the reactors, and the H2/CO stream can be vented to atmosphere before it is mixed with the C2 stream and enters the reactor by a set of automatically controlled, pressure sensitive block and bleed valves.

There are several differences between the AHR process described here and the process described by (Miller & Vicente, 1998) because the analyses described in that work were performed on the actual process whereas the target system for the prototype displays is the AHR section of the E1 process *simulator*. The simulator is a first principles model of the E1 process tuned with time constants so that the simulated process dynamics are highly representative of the actual process behavior. While the simulated E1 process is very similar to the actual process, some equipment and secondary functions are not modeled. However, the simulation captures the most critical process paths and all of the control options available to the panel operator. NOVA uses the simulator as both a training platform and as a testbed for process instrumentation and control system modifications. Therefore, its faithfulness to the actual process is sufficient for industrial applications. The analyses were edited for content based on the simulator model because the goal of the research was to test the interface in that environment.

Information Requirements

At several points in the preceding discussion, the term “information requirement” has been used. An information requirement (IR) is a statement of a specific piece of information, or a relationship, that must be incorporated into an interface intended to support effective control of the target process. It is assumed that a comprehensive set of IRs forms a necessary body of information for this task. The aim of integrating both task- and system-based analysis methods can be cast as an effort to move closer to identifying a sufficient set of information requirements that, if met in an information system, would allow for effective adaptation under the broadest range of events (see Vicente, 1999, for a more comprehensive discussion).

Extension: Control Task Analysis for Automatic Controls

The first step in the treatment of the IRs was extending the analyses conducted by Miller and Vicente. A second task-based analysis was required because the HTA conducted by Miller and Vicente (1999) did not address the behavior of the automated controllers used in the process. The primary knowledge source of the HTA was the procedures and these procedures assume that the behavior of the controllers is; a) known to the operator, and b) reliable. The ADS had also not captured the activity of these controllers because control is assumed to be an object of task analysis, which the ADS does not address. Thus, the role of the automation had not been translated into requirements for the novel displays. This was clearly a problem as these controllers are an important part of the operations activity.

To fill in this gap, a Control Task Analysis (CTA[3]) was conducted to supplement the HTA. Another form of task analysis, CTA aims to "identify the requirements associated with known classes of events in a particular application domain" (Vicente, 1999, p. 182). In this particular application, the events are the control actions required to direct the process in a

---

[3] The acronym CTA is commonly associated with the term Cognitive Task Analysis, a class of work analysis methods loosely related by their focus on eliciting cognitive tasks that cannot be directly observed. This is potentially confusion because Control Task Analysis could conceivably be employed as a Cognitive Task Analysis method. However, CTA is used here exclusively to refer to Cognitive Task Analysis as a phase of Word Domain Analysis (Vicemte, 1999).

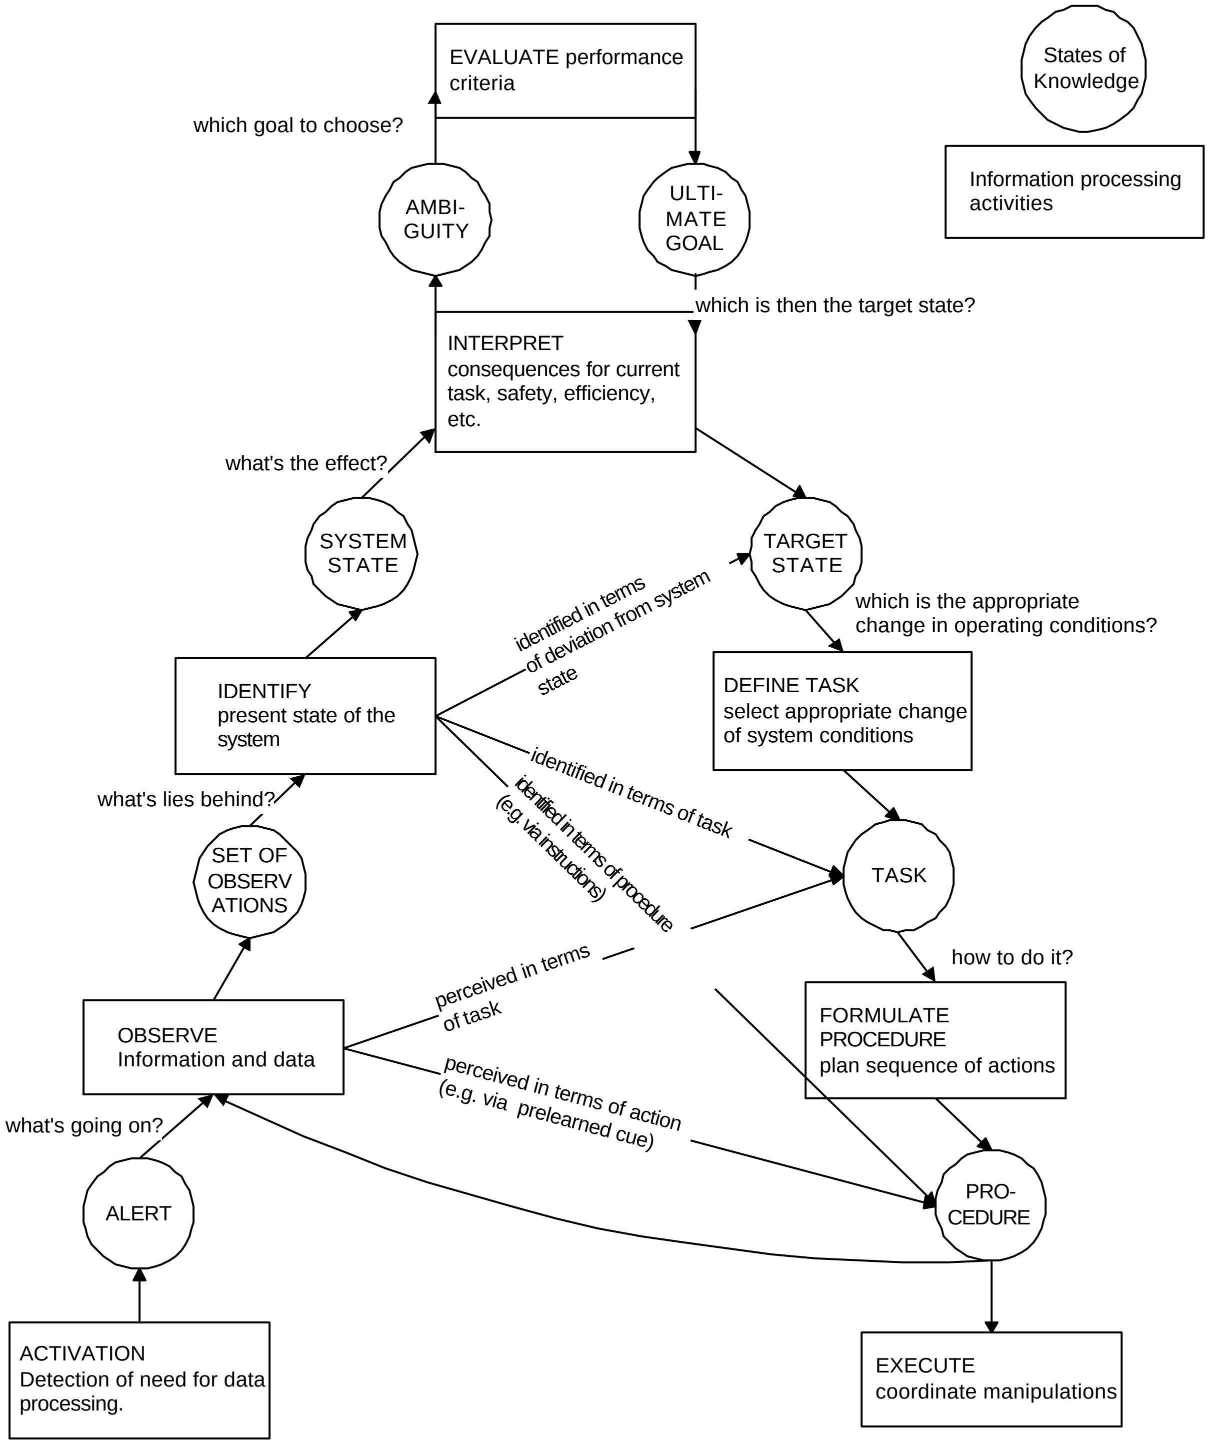

Figure 2: Decision ladder template for control task analysis (Rasmussen, 1976).

productive way. The CTA method was used in place of the HTA method because it does not distinguish between human or machine actors that complete the control tasks (see below)

A non-linear information processing framework called the Decision Ladder (DL; Rasmussen, 1976) is used to perform the CTA (see Figure 2). The DL is a sequence of information processing steps (the boxes) and intermediate states of knowledge (the circles). The framework is non-linear, because the component boxes and circles do not have to be completed in a specific sequence. Rather, the framework allows an actor (a generic term that incorporates the operator or the automation) to take advantage of shortcuts in association between the steps (the arrows). For example, in most cases neither an operator nor a machine needs to reason about the meaning of an error signal; it is merely a signal that implies an immediate corrective action. Therefore, the Interpret and Evaluate steps can be bypassed.

It is important to note that CTA does not differentiate between automated and manual execution of these actions. The information requirements for completing the task are identical whether a human operator or a machine carries out the task. Moreover, it is assumed that, in order to support effective monitoring of the automation, all of these information requirements must be met in the display.

The Weight Ratio Controller (WRC) serves as an example of this concept. The WRC is a controller that receives flow rate and composition inputs from the ethylene stream, calculates an ethylene flow rate to match a setpoint ratio, and outputs a control signal to a valve on the hydrogen stream to achieve this ratio. Figure 3 summarizes a DL for the WRC. The task is cyclical, with changes in Target State acting as an input to the cycle. The Target State defines the ratio of the weight of flow of H2/CO to the weight of flow of C2s in the incoming streams. These flow values are two of the data points that the controller must gather in the Observe process. Also collected in the Observe process are the analyzer readings that specify the composition of each of these flows. These data form a Set of Observations. The Identify process calculates the current weight ratio, compares it to the setpoint target, and uses the

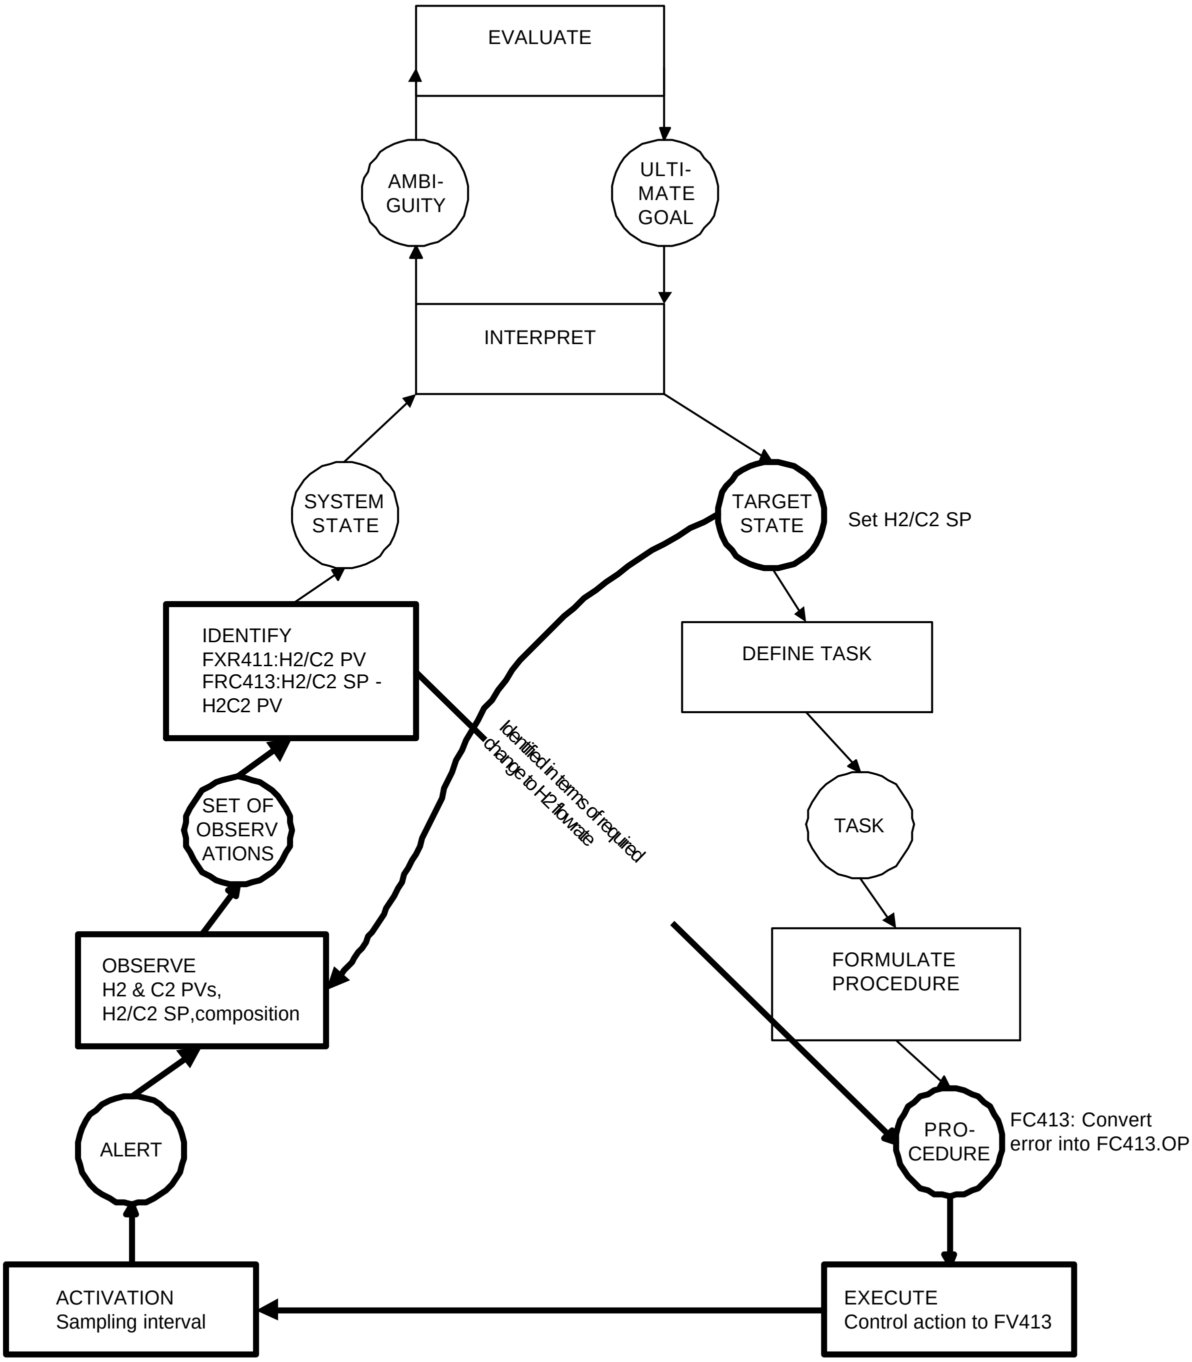

Figure 3: Decision ladder for the H2/C2 weight ratio control.

difference to determine the necessary change to the output of the control valve on the H2/CO stream. In this control task, once the process state has been identified, the control proceeds directly to a pre-defined Procedure because the Define Task and Formulate Procedure steps are either built into the automatic controller or judged based on skill and experience of the human operator. In this case, the procedure is simply to determine a valve output to move the process to its desired state (i.e., the setpoint Target State). The output can either be calculated by a controller algorithm or judged by the human operator. The next step in the task is to return to the Activation process, signaling the controller (again, human or machine) to complete another cycle.

There are several other controllers employed in the AHR. DLs for each of them are included in Appendix A. Some of them are simple regulatory controllers while others reflect more complicated control strategies.

Adding CTA to the previously completed analyses rounds out the sources of information requirements for the AHR. The next section turns to the challenge of integrating these requirements, a necessary step in informing the design of an ecological interface consisting of both system- and task-based models.

Modification: Adjusting for System Boundary and Simulator Fidelity

The HTA performed by Miller and Vicente (1999) used the existing written procedures as the sole knowledge source. They acknowledge a crucial limitation to this approach in that actual operational behavior can differ substantially from what a procedure dictates. While a separate study leads us to believe that this is less of an issue for the NOVA E1 facility than for other plants (Jamieson and Miller, 2000), it is more important to recall the purpose of the Miller and Vicente (1999) work, which was to compare the types of information produced by two analysis techniques. Further, our present purpose is to determine whether the additional information from a task-analysis can be effectively integrated into an ecological interface. Both of those purposes can be served by an analysis without extensive field verification of the procedure. We have taken to heart the cautions issued by Miller and Vicente (1999) in using their results for interface design.

One of the limitations of basing the HTA on the procedures is that the procedures are not written with the same system boundary that was employed in the ADS. The boundaries used for the ADS were confined to a sub-unit of the process (i.e., the AHR), whereas the procedures address equipment throughout the ethylene unit. Therefore, most of the procedures substantially cut across the ADS system boundaries. This difficulty arose because Miller and Vicente (1998d) had sought to isolate a manageable section of the entire E1 process to model with the ADS.

The result of the system boundary discrepancy was that the HTA identified some information requirements that were outside the scope of the ADS. Therefore, the information requirements from the HTA that were outside of the system boundaries were removed from the analysis results[4]. The removed items can be divided into two categories; 1) equipment and functions addressed in the procedures that were either upstream or downstream of the reactor section, or 2) equipment and functions addressed in the procedures that were within the reactor section, but not included in the ADS because they were not directly relevant to process operations. For example, some lines used in reactor swings and catalyst regeneration were not included in the ADS, but are mentioned in the HTA.

A second filtering process was required when the focus of analysis shifted from the actual process to the simulator model. This shift in focus took place for the purposes of this dissertation and was required because an empirical evaluation of the displays was to take place in the simulator. Equipment not included in the simulator model of the reactor section was removed from the ADS and not included in the display design process. The original AHR

[4] This modification to the IRs was performed by the author, subsequent to Miller and Vicente's (1999b) initial analysis.

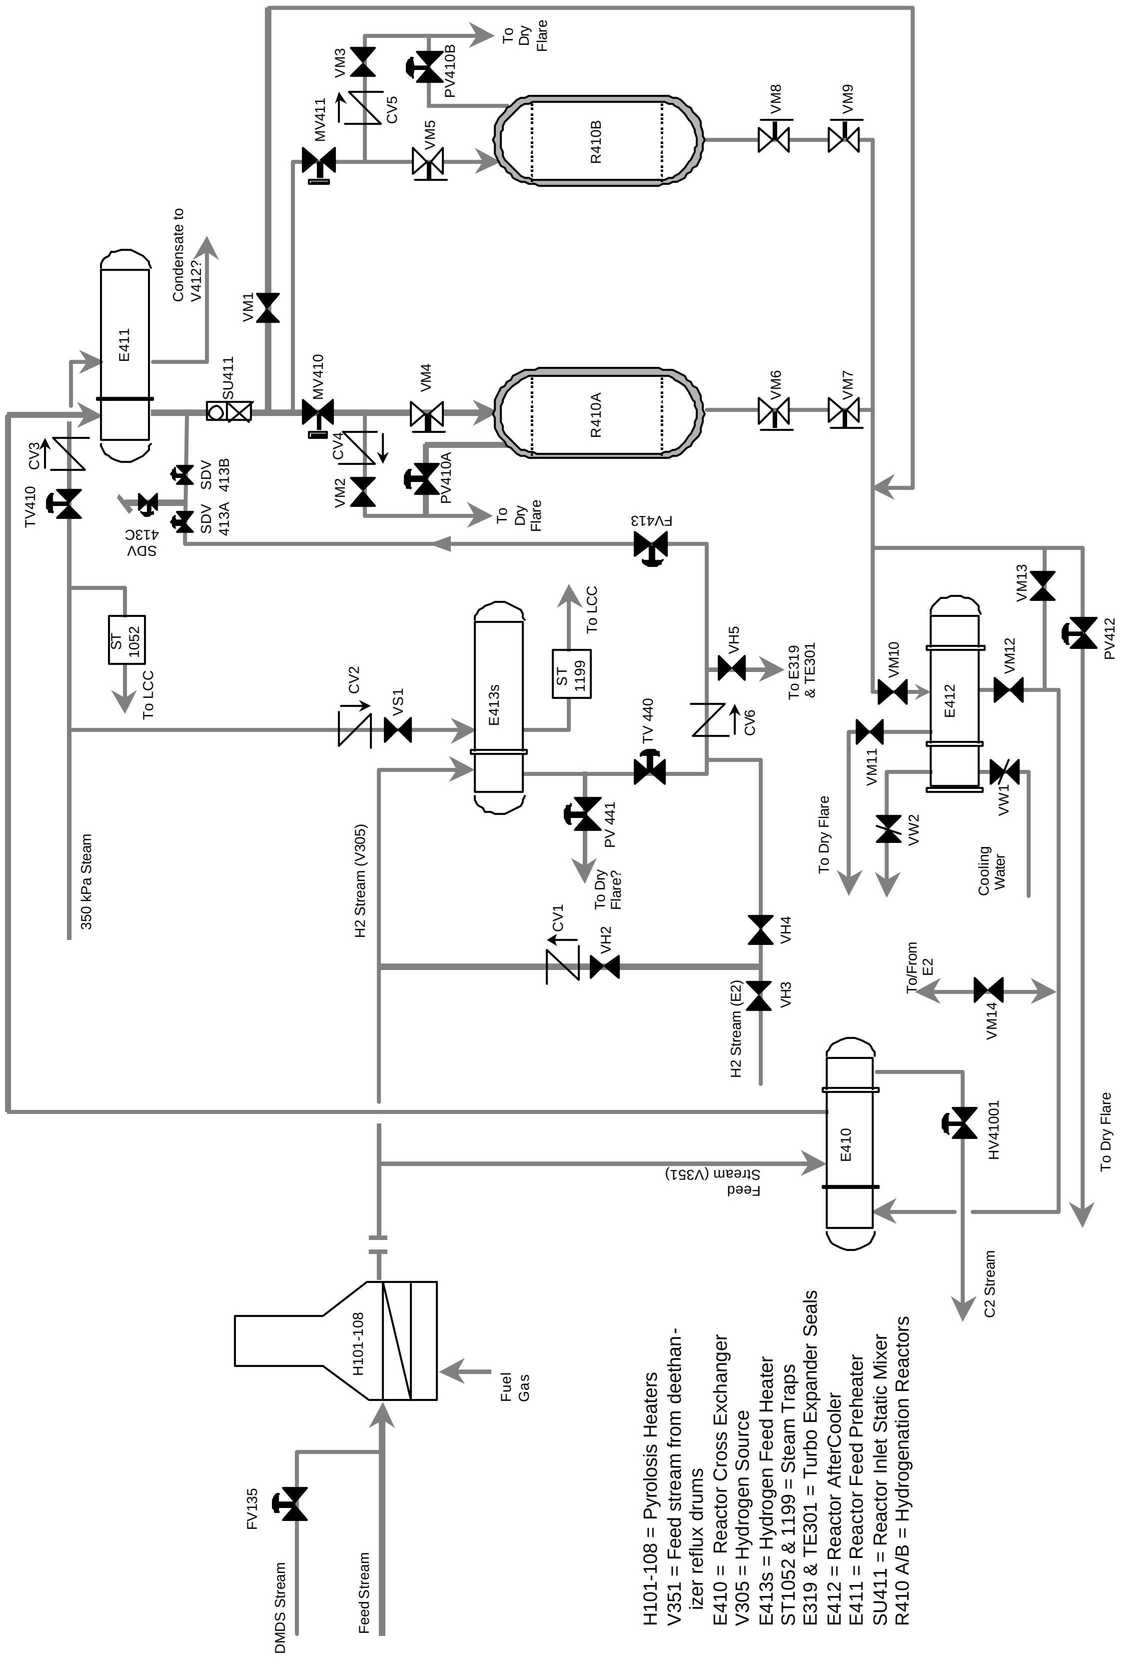

Figure 4: The system description used for the Miller and Vicente (1998d) ADS analysis.

model is shown in Figure 4. Using the notation employed by Miller and Vicente (1998), the following equipment was removed:

CV1, VH2, and line to H2/CO inlet stream

VM2, CV4, and line to flare

VM3, CV5, and line to flare

CV2 CV3 VH4 VH5 VM7 VM9 VW1 VM12

SDV413B, SDV413C, SU411

VM11 and line to flare

Fuel gas flow to furnaces

ST1052 and line

It should be noted that, although this is a substantial amount of equipment, many of the components combined with remaining equipment to form functions that are retained in the simulated process. In fact, while both of these filtering processes had an effect on the level of complexity in the target system (i.e., by reducing the number of components), neither had a substantial impact on the control tasks or strategies employed by the operators.

Integration: The IR Matrix

Once the three lists of information requirements (ADS, HTA, and CTA) had been modified based on the system boundary and simulator fidelity criteria, they were integrated. An information requirements (IR) matrix was formed listing the requirements from each of the three analyses (ADS, HTA, CTA). Then, each requirement was cross-checked against the requirements drawn from the other two analyses. Duplicate IRs were removed. Frequently, requirements from one analysis had to be assimilated or parsed to reconcile them with requirements from another analysis (see below). The final IR Matrix is given in Appendix B.

For each resultant IR, the analysis or analyses that identified it are noted in the matrix. Thus, the matrix contains a column for each analysis and checkmarks in the appropriate cell when the analysis identified the IR (see Table 1 for an example). In many cases, whether an analysis had identified an IR was not easy to assess, primarily because the language and specificity of the IRs differed according to the original analysis source. For example, the ADS would identify an IR indicating that the state of a valve must appear in the interface. The HTA might include a statement to close the same valve. It seems reasonable to conclude that a designer using the results of the HTA would conclude that a representation of the state of the valve in question was called for. In general, the analysts employed a liberal decision criterion by affirming that an analysis had identified an IR if a reasonably skilled designer could deduce from the analysis that that IR was present.

Table 1: A section of the IR matrix.

Identify state, and range of possible states, for Physical Function-Component functions	**ADS**	**HTA**	**CTA**
E410 shell side INLET temperature	✔		
E410 shell side OUTLET temperature	✔		
E410 tube side INLET temperature	✔		
E410 tube side OUTLET temperature	✔		
E411 shell side INLET temperature	✔	✔	
E411 shell side OUTLET temperature	✔	✔	✔
E411 tube side INLET temperature	✔	✔	
E411 tube side OUTLET temperature	✔	✔	

<u>Differences in level of specificity.</u> One of the early problems encountered was that the IRs differed substantially in their respective specificity. The requirements generated by the ADS and the CTA were very precise while those from the HTA ranged from precise to vague. For example, both the ADS and CTA identified this very specific IR: “The effect of the setting of

SDV413a on the flow of H2 to the C2 stream". Example HTA IRs include specificrequirements such as "Acetylene outlet analyzer update schedule" as well as more vague requirements such as "Sequence constraints on procedure steps".

The ADS requirements were precise because they appeared in groups based on level of abstraction and aggregation (see Table 1 header). For example, one group falls under the header, "Identify values, and range of possible values for the Generalized Function-Component functions" and then goes on to name those functions, such as, "The flow of H2/CO entering the AHR." Each group heading created a local context for those IRs that included the level of abstraction (i.e., Generalized Function), level of aggregation (i.e., Component), and the type of information in that category (i.e., values and range of possible values). Thus, each requirement drew on the context created by its group, allowing for more specific requirements within that context. The CTA requirements were also related by group, with the individual control tasks creating the shared context. The specificity of these IRs also benefited from the precise nature of feedback control algorithms. In contrast, the HTA requirements were identified based on their order of appearance in the procedure sequence. Removed from the procedure, these requirements lost their context and were often vaguely stated. Moreover, Jamieson and Miller (2000) showed that NOVA has an operations culture that allows operators a great deal of freedom when following their defined procedures. Thus, many of the details cited by the ADS as IR's are 'assumed' in the procedures that were referenced by the HTA. Had the details been available in the procedure, it is reasonable to assume that the HTA would have captured them.

Analysis overlap. A Venn diagram (see Figure 5) shows the proportion of the total IRs accounted for by each analysis and the overlap between the various analyses. Take, for example, the proportions for the ADS. The ADS alone identified 51% of the total IRs. However, the HTA identified a proportion of the same IRs, namely 23% of the total set. All three analyses (ADS, HTA, and CTA) identified the same 10% of the total IRs. However, there was no overlap between ADS and CTA that was not also shared with the HTA. Generally speaking, there is a large overlap in the information requirements generated by each analysis. However, there is also substantial uniqueness in each of them. Both of these observations conform to the observations of Miller and Vicente (2001).

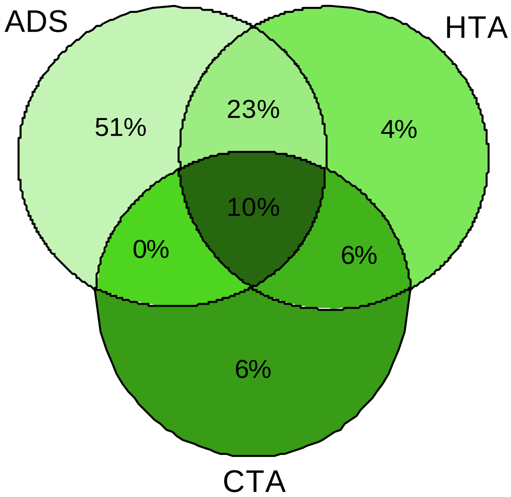

Figure 5: Proportion of IRs captured by each analysis (circles not to scale).

There are several difficulties with the counts that make up the diagram. One of them is the level of specificity problem discussed above. The count treats all requirements as having the same value. This can be misleading. For example, the ADS is credited for identifying the IR corresponding to the location and appearance of a particular in-line valve. This rather trivial IR is counted as one. In comparison, the IR corresponding to the order requirements in procedure steps is contained as a single IR. This critical requirement is clearly more important than the location and appearance of an in-line value. Another analyst could have chosen to identify each sequential order requirement between any two procedure steps as a single IR, thereby increasing the number of IRs cited by the HTA that were not captured by the ADS.

Given these problems, how is the IR matrix to be meaningfully interpreted? The answer is that it gives qualitative insight into the value of each analysis. The large overlap confirms that system- and task-based analyses are both productive methods for building a knowledge base for interface design. However, each analysis clearly identifies requirements not cited by the others. Thus, the analyses are mutually re-enforcing.

Observations for Future Applications

In retrospect, the process of integrating the information requirements benefited from having the foundation established by the ADS. The two-dimensional hierarchical structure of this analysis provided a framework within which the requirements gleaned from both of the task analyses could be inserted. Had the analysts known this in advance, the HTA-derived requirements could have been initially stated in relation to the ADS structure. This would have simplified the reconciliation process described above. In contrast, the ADS and HTA analyses were conducted quite separately using different knowledge sources. The prior knowledge obtained from the ADS analysis, while certainly valuable in making sense of the procedures, was not fully exploited in performing the task analysis.

This situation may be reversed for analysts who are highly familiar with the operation of the plant. They may find it beneficial to start with a task-based analysis and use that as a foundation to integrate ADS results. There is, however, an argument against this. Because the ADS identifies functionality that exists independent of operational tasks, it should be broader and more inclusive than any task-based analysis. Thus, from an analytical perspective, it may always hold an edge over task analyses as a foundation-forming analysis.

Summary

The process of integrating the information requirements was, in at least one way, much easier than expected because the IR was a useful concept in all three analyses. The ADS, HTA, and CTA were all fairly adept at leading to clearly defined data elements that need to appear in an effective interface. Notwithstanding the level of specificity challenge noted above, the IRs from each analysis were fairly compatible with those of the others. More difficulty in assimilating the results of the analyses had been anticipated. In the next two sections, we use the IR Matrix to evaluate the coverage of the existing E1 sensor set and interface.

Existing E1 AHR Interface and Information Availability

We next turn to an assessment of the proportion on the IRs appear in the existing interface. This first requires a description of the current interface.

Description. The existing interface suite for the E1 process is based on a mimic display format supplemented by textual information pertaining to control parameters and sensed process values. An example of a typical process view from this suite is shown in Figure 6. The display suite for the entire E1 plan comprises several dozen of these types of screens, supported by scores of others that provide rarely accessed details about controller logic and alarm configurations (because this information is largely static). Some operators have also created various summaries of certain types of information such as a listing of the setpoint, output and process variable for all controllers in a certain plant area.

The screen in Figure 6 shows the major pieces of equipment of the E1 AHR connected by arrows representing the designed direction of flow. Process variables are placed near the lines or equipment to provide information about measured variables at those locations. On the left-hand side of the screen is a table of controller parameters that can be swapped out with a table showing analyzer values. A second mimic display shows the reactor section in greater detail (Figure 7) along with the controller or analyzer information. These two displays make up the complete set of mimic displays for the AHR. Not included in either figure is the Change Detail zone at the bottom of the screen where operators get detailed information about a particular data point and make changes to controlled variables.

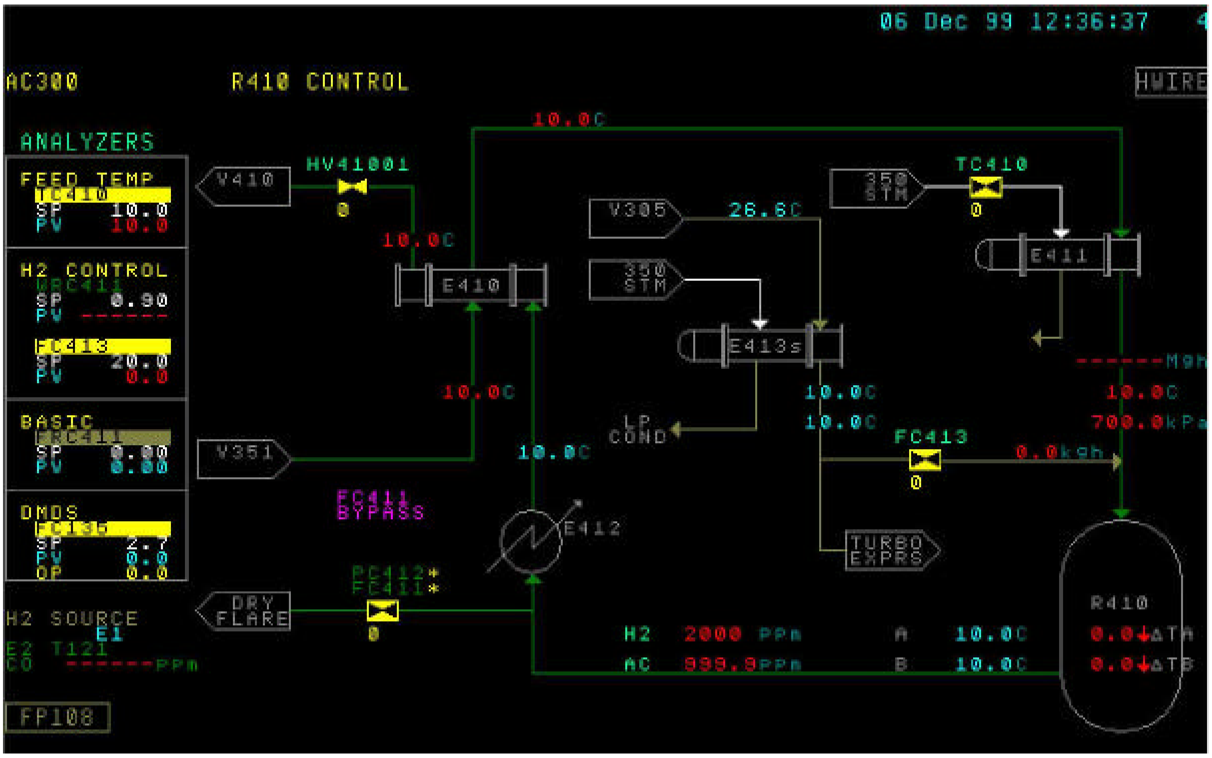

Figure 6: Existing mimic diagram for the AHR.

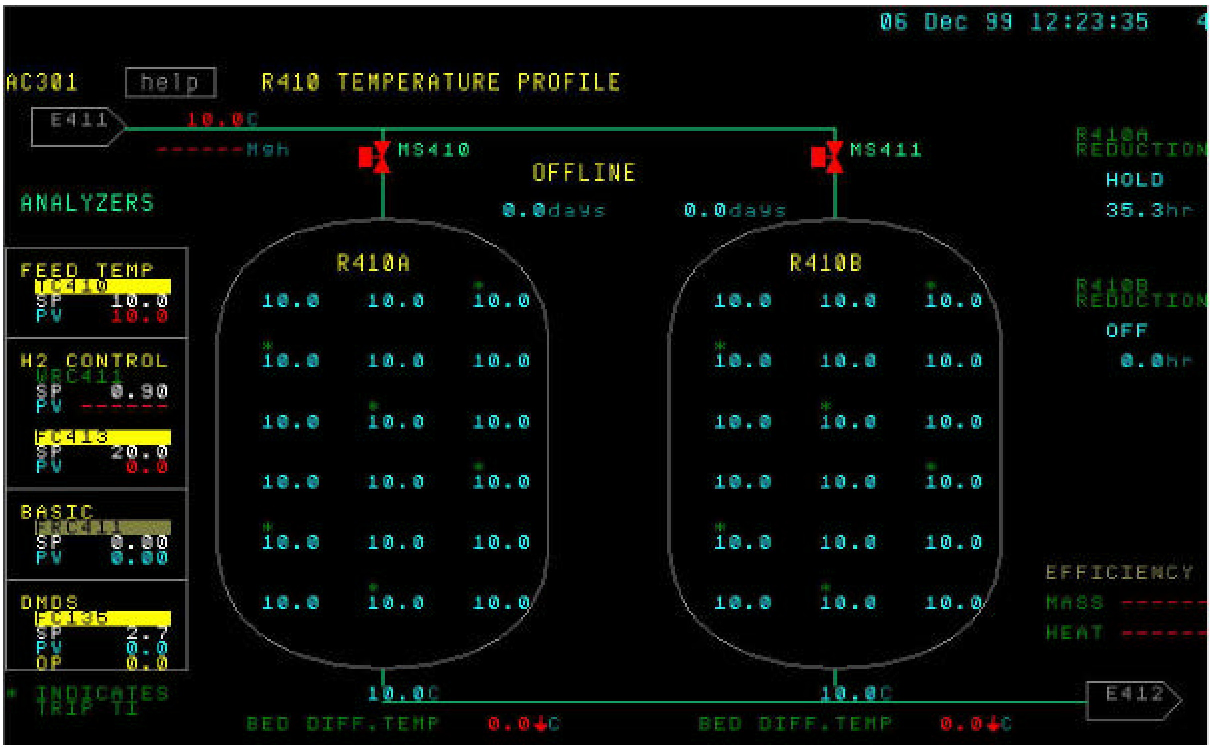

Figure 7: AHR mimic diagram showing reactor detail.

In addition to the mimic diagrams, several other software applications provide interfaces that support operator monitoring and control. The two most important of these are the alarm summary page and the data historian. The alarm summary page is the key interface to the alarms implemented on the instrumentation and control software. Examples of these ‘soft’ alarms include limits on equipment temperature, process pressure, and control configurations. A summary of the alarms currently in alarm state is provided to the operator in tabular format with information about the alarm priority, the time that it came in, a description, and the type of alarm (e.g., high limit, low limit, etc.). These software alarms are distinguished from the hard-wired panel alarms that are also in use for the most critical process parameters.

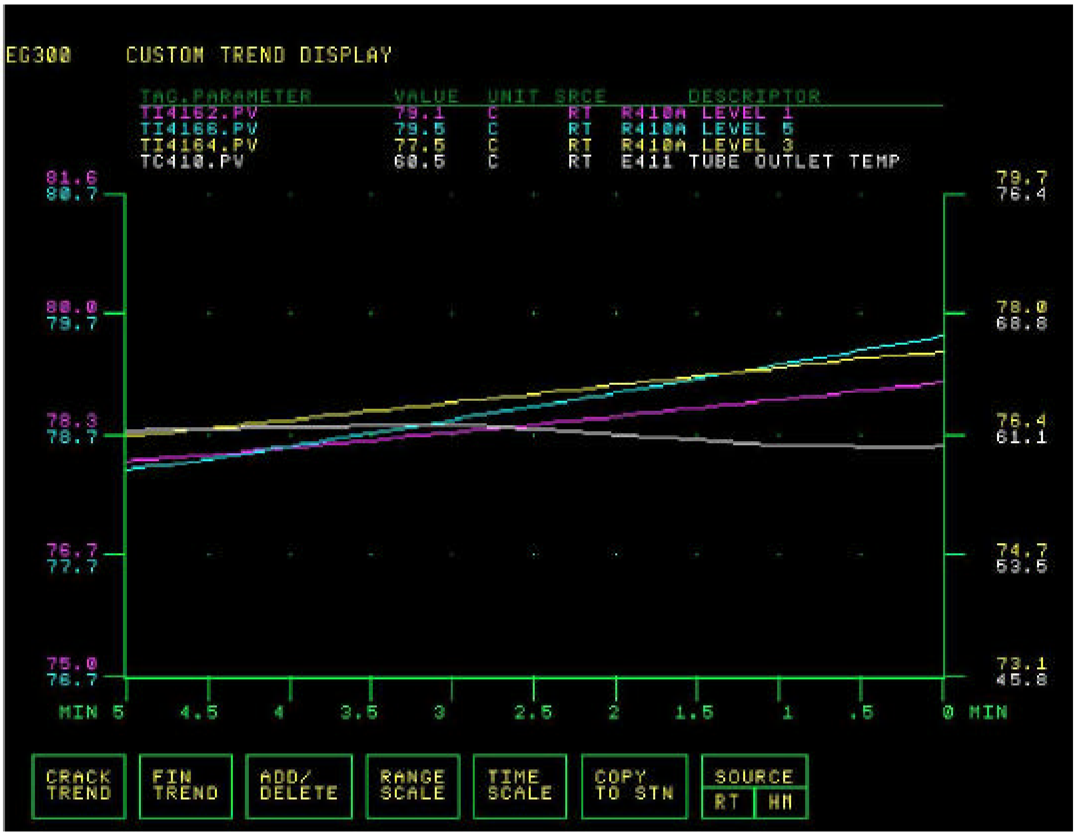

Figure 8: Trend view employed in current E1 plant.

The other important software application is the data historian. The historian keeps a record of the process variables for at least several days for the operator to access. However, the operator will typically only look at a history of values over the duration of a shift (12 hours). A trending package (Figure 8) allows the operator to plot up to 4 of these variables simultaneously. Many operators consider history information to be their most important diagnostic tool for investigating disturbances. It is also a useful monitoring tool when the time scale is set to a short duration.

Evaluation of how well the current interface suite meets IRs. The IR Matrix contains a column with the heading "E1" that includes a check mark for each IR that is contained in the existing interface. Figure 9 provides a comparison of the combined number of IRs generated from all sources to the number of IRs met by the existing interface suite. There is a clear shortfall in the existing displays, although the reasons for this are not consistent across the categories. At low levels, such at the Physical Form (PFm) and the components-wise topological Physical Function (PFn) levels, the E1 displays exclude many IRs to reduce display clutter. Rarely used valves, for example, are not included. Thus, the displays are optimized for normal operations. Many Physical Function and Generalized Function IRs at the component level are not measured, although they could be. This shortage of information propagates, limiting the availability of aggregate and abstract function (AF) information. An exception comes with the requirements generated by the CTA. A majority of these are met in the existing interface, largely because the DCS must use this information for automatic control. Notably, none of the task-only IRs are met. This can be attributed to the common practice of using procedures that provide these requirements.

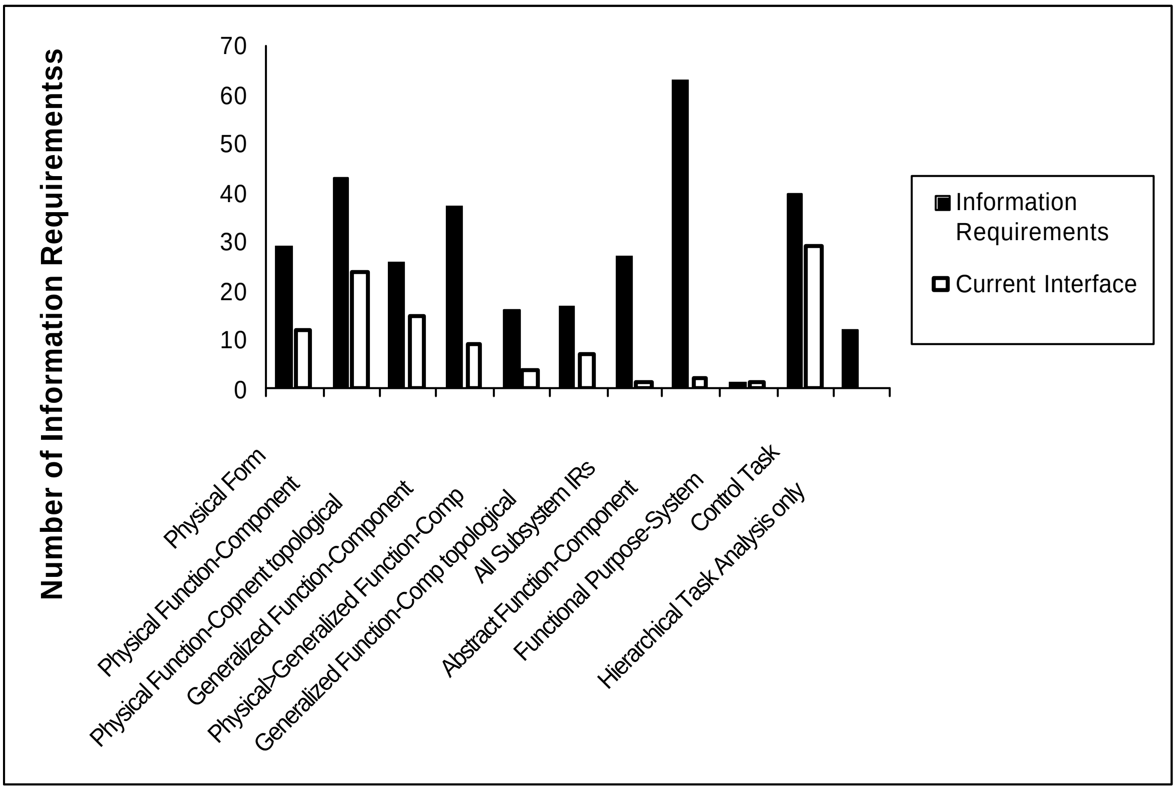

Figure 9: Comparison of IRs identified versus IRs included in the current E1 AHR interface.

Evaluation of the IR coverage of the existing sensor set. Beyond identifying the proportion of the IRs that are met in the existing interface, we can assess what proportion of the IRs could be met with the existing sensor set if the necessary derivations were performed. The check-list for this question appears in the column with the heading DCS? DCS stands for Distributed Control System and is the name given to the instrumentation and control system that communicates sensed and derived information to the operator. A summary of this evaluation appears in Figure 10.

Comparing Figure 10 to Figure 9, it is clear that there is not a great deal of information that could be supported by the existing instrumentation and control system that is not already available in the E1 displays. There is some GF-C and AF-C information that could be derived and added, but not a great deal. In addition, there is a component of CTA and HTA

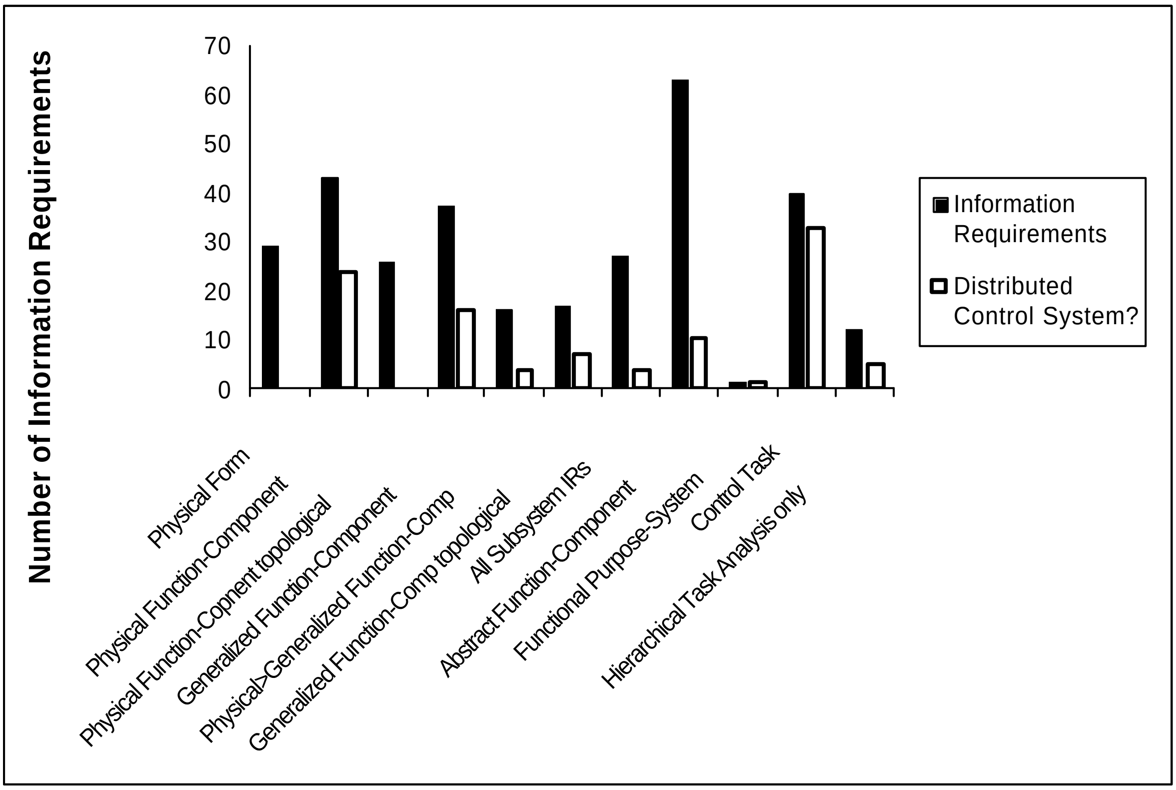

Figure 10: Comparison of IRs generated versus IRs that could be met given existing sensors.

requirements that could also be added using the existing sensor set. Overall, however, the existing interface suite appears to be exploiting much of the available information.

One caveat to this and the following analysis is that some information requirements are static relationships that do not require sensing. For example, the IR that states, "The two reactors (R410 A and B) are located side by side, but there is no direct connection between them," can be met with a drawing that requires no connection to an information system. Thus, there are several cases where an IR is marked as "N/A" in a coverage column. In Figure 10, this shows up as missing columns for the PFm and PFn-C topological categories.

<u>Evaluation of the IR coverage of the simulated process.</u> The simulator, because it is based on a first principles model of the process, can contain more information than the current instrumentation and control system. This provides an opportunity to design interfaces that contain a greater proportion of the IRs than would be possible given the existing sensor set in

the E1 plant simply because the simulator does not include a limiting model of the sensor set. The matrix column headed “Sim?” includes a check mark for all of the IRs that are available or can be derived from the simulator model. A summary of this column is contained in Figure 11.

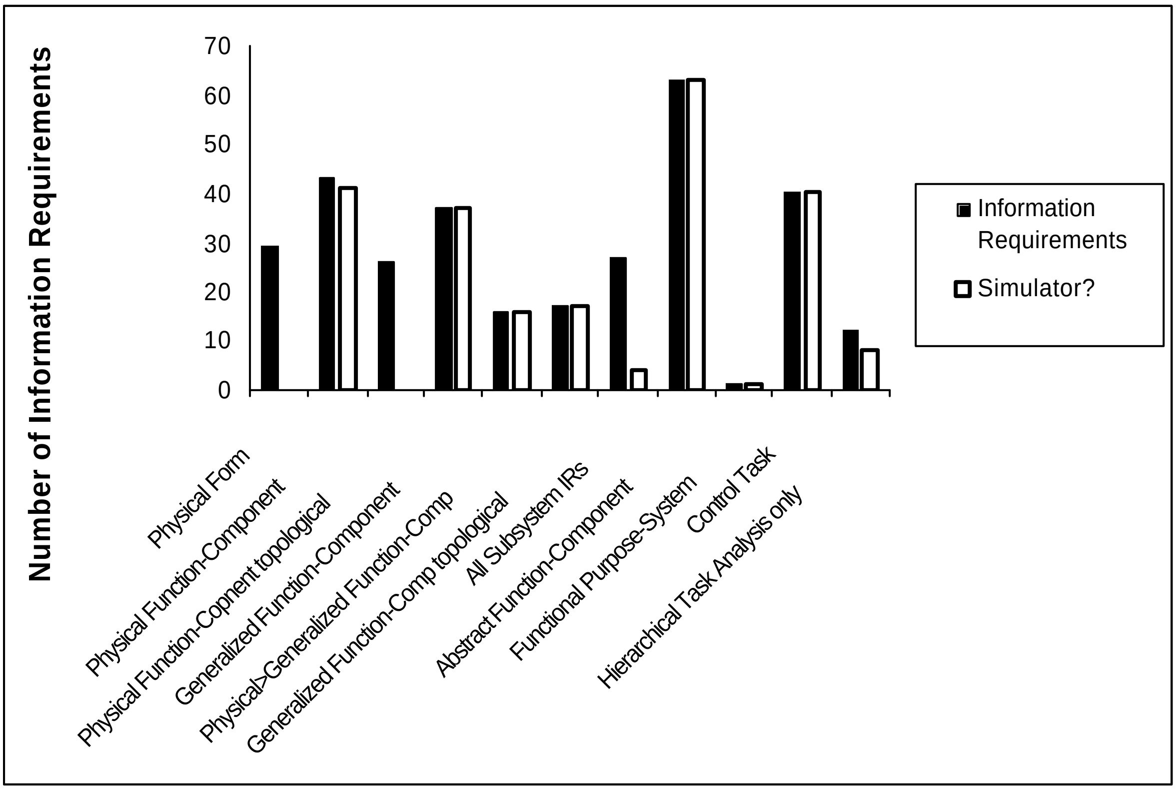

Figure 11: Comparison of IRs generated to IRs supported by the simulator model.

The difference between Figure 11 and Figure 10 is a stark one. The simulator model provides much greater access to the IRs, either through model parameters or through derivation of those parameters to determine highe r order relationships.

Summary

The three analyses described in the preceding sections indicate what proportion of the information requirements are met by the existing information system, how many more could be supported by that system, and the proportion supported by the process simulation. While there is some room for expanding the information available to operators of the current process, the bulk of the shortfall in the existing interfaces cannot be made up with the existing sensor suite. In contrast, the simulator provides an opportunity to design and implement a new set of interfaces that are supported by information that would not normally be available to operators. The following chapter describes such an interface design and implementation activity.

INTERFACE DESIGN

In the next two sections, I give a brief overview of two novel interfaces that were designed using the information requirements described in the previous chapter. One of the interfaces is based only on IRs that were identified by the ADS analysis. Because the product of this analysis is physical and functional information, that interface is designated P+F. The other interface is based on IRs from all three analyses. Since the HTA and CTA provide task-based information, this display is designated P+F+T. A detailed description of the interfaces is provided in Appendix C.

Physical+Functional+Task (P+F+T) Interface

The P+F+T interface (see Figure 12) contains the full suite of graphical elements and views that were developed based on all three work analyses. Ten process views are distributed across a workspace that spans two stacked 21"monitors, each at a resolution of 1280x1024 (for a total screen size of 1280x2048). The ten views are described below and mapped onto Figure 13.

<u>Schematic view.</u> The Schematic View provides primarily physical and low-level functional information about the process. It is based on a mimic display annotated with digital values for critical variables and iconic representations of valve states. It is similar to the mimic display available in the Current interface.

<u>Trend view</u>. The Trend View allows the operator to view a history of several process values in parallel. Any data element that is recognized by the information system (i.e., the DCS or the simulator) can be trended. A trend capability is incorporated in the Current interface

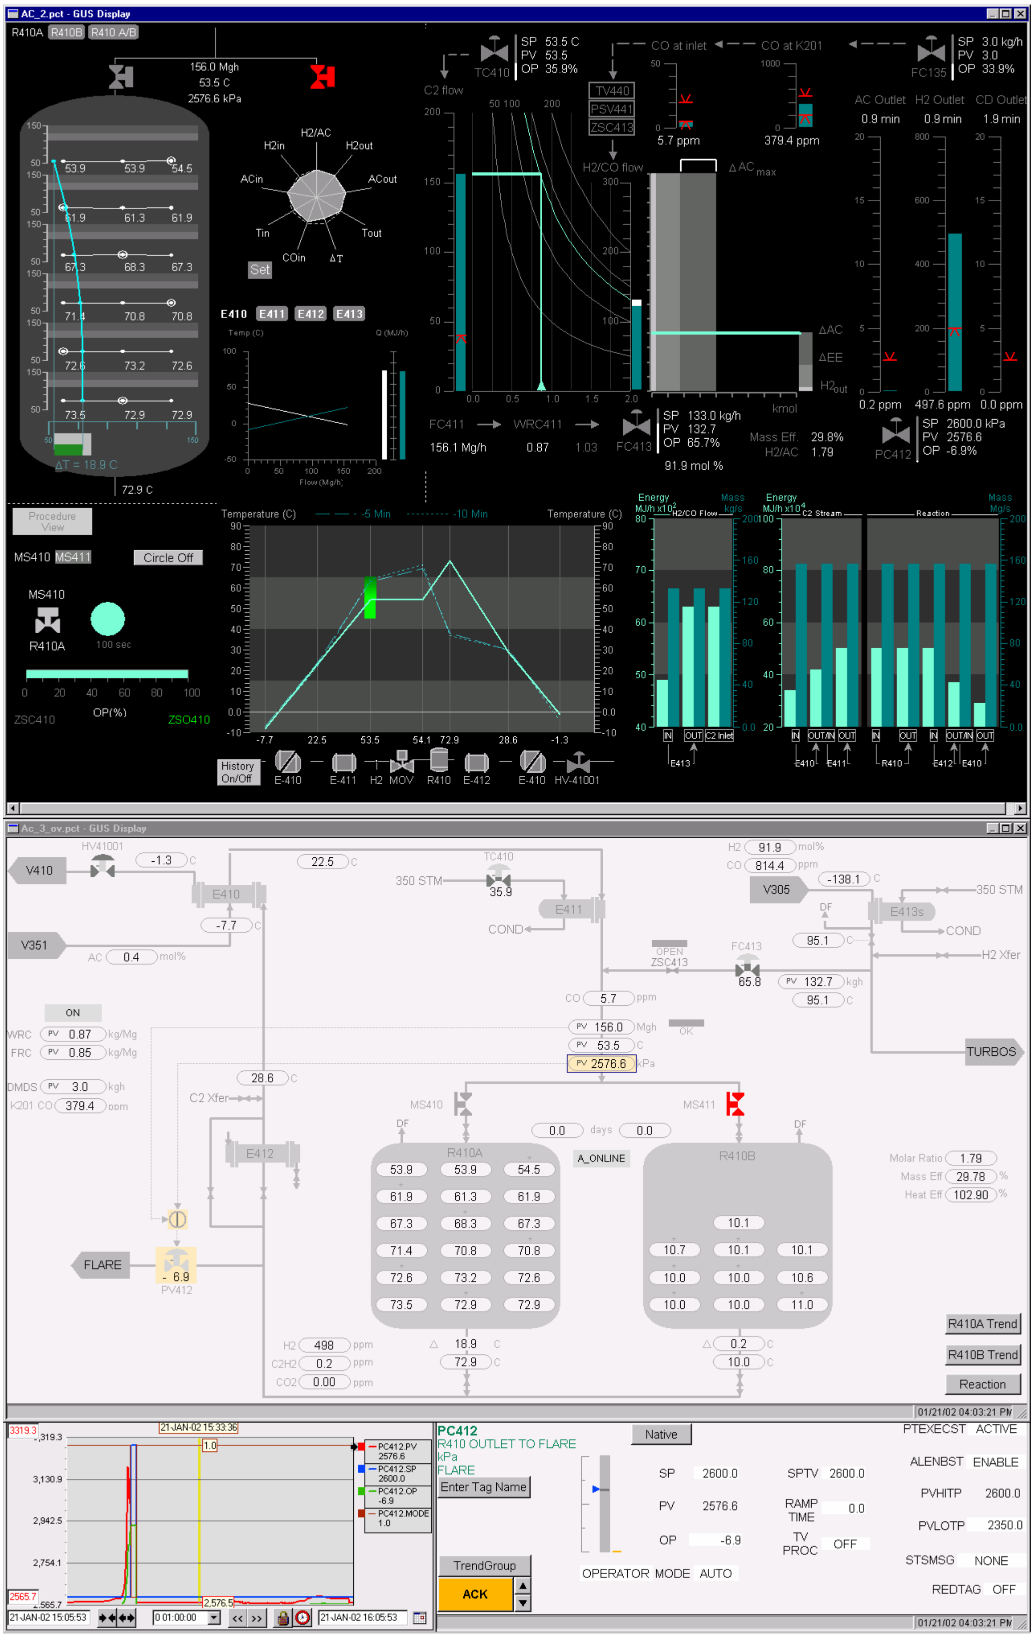

Figure 12: The P+F+T interface.

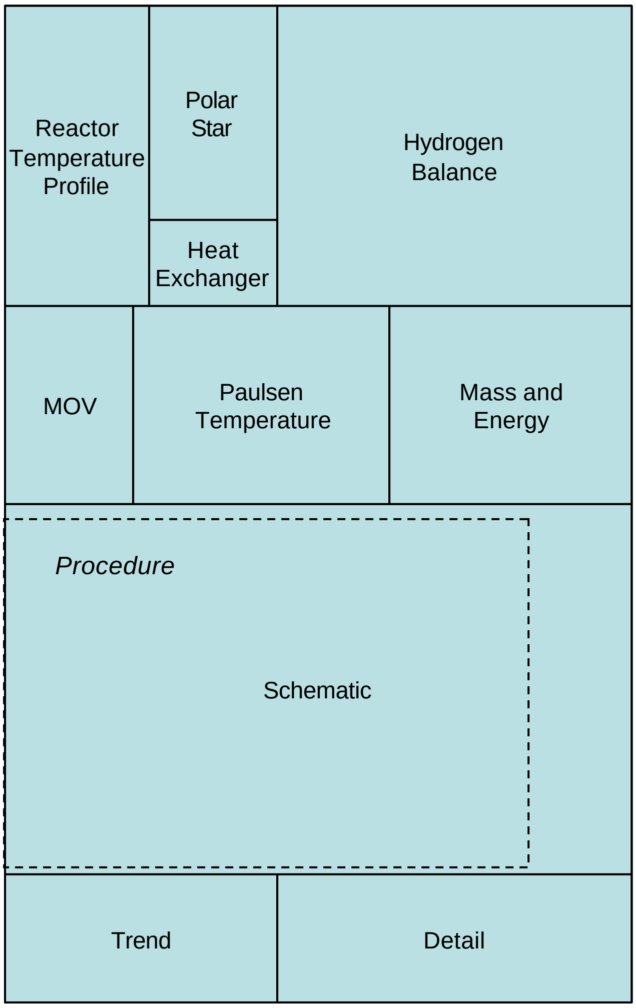

Figure 13: View key for P+F+T interface.

Detail view. All changes to controlled process variables are made in the Detail View. It provides the detailed process parameters associated with any data point. A view similar to the Detail View is available in the Current interface.

Reactor temperature profile view. This view shows the horizontal and vertical temperature distributions across the reactor. A combination of configural graphics and digital values is used. Indications for high temperature alarms are also included.

Polar star view. The polar star (Wolff, 1967) supports monitoring of critical reactor variables though a concise configural display. Parameter values on the graphic can be ‘reset’ by the operator to increase its context specificity.

Paulsen temperature view. A graphical profile of the equipment temperatures across the process is provided in the Paulsen (1992) View along with digital values. Process constraints are shown where they are available.

Hydrogen balance view. The hydrogen balance provides detailed graphical representations of the reactor functions and purpose. The view can be used to monitor the reactor and diagnose anomalies in its performance.

Mass and energy view. A graphical profile for mass and energy flow through the process is provided in this display of high level functional information.

Heat exchanger view. The first law of thermodynamics for heat exchangers is shown in a combination of configural and non-configural graphics. The view allows for trouble-shooting of heat exchanger problems.

MOV view. A predictive display depicts the process of opening and closing motor-operated valves (MOV). The view reduces the need for communication between panel and field operators.

Procedure view. Step-by-step procedure support for the reactor swing and temperature runaway procedures is provided here. Text, configural, and non-configural graphics are employed. The Procedure view is the only view that is not always visible. It appears in the location in Figure 13 denoted by the dashed box.

P+F Interface Description

The P+F interface is formed by removing two process views and three graphical elements from the P+F+T interface. The upper monitor for this interface is shown in Figure 14; the lower monitor is identical to that of the P+F+T interface. The graphics removed are those associated with information requirements that were identified only by the task-based analyses. The following graphical elements and views are not drawn:

- Procedure View
- MOV View
- H2 Balance View: Weight Ratio Control graphic
- H2 Balance View: Analyzer update timers
- H2 Balance View: Setpoint and Output data on control valves

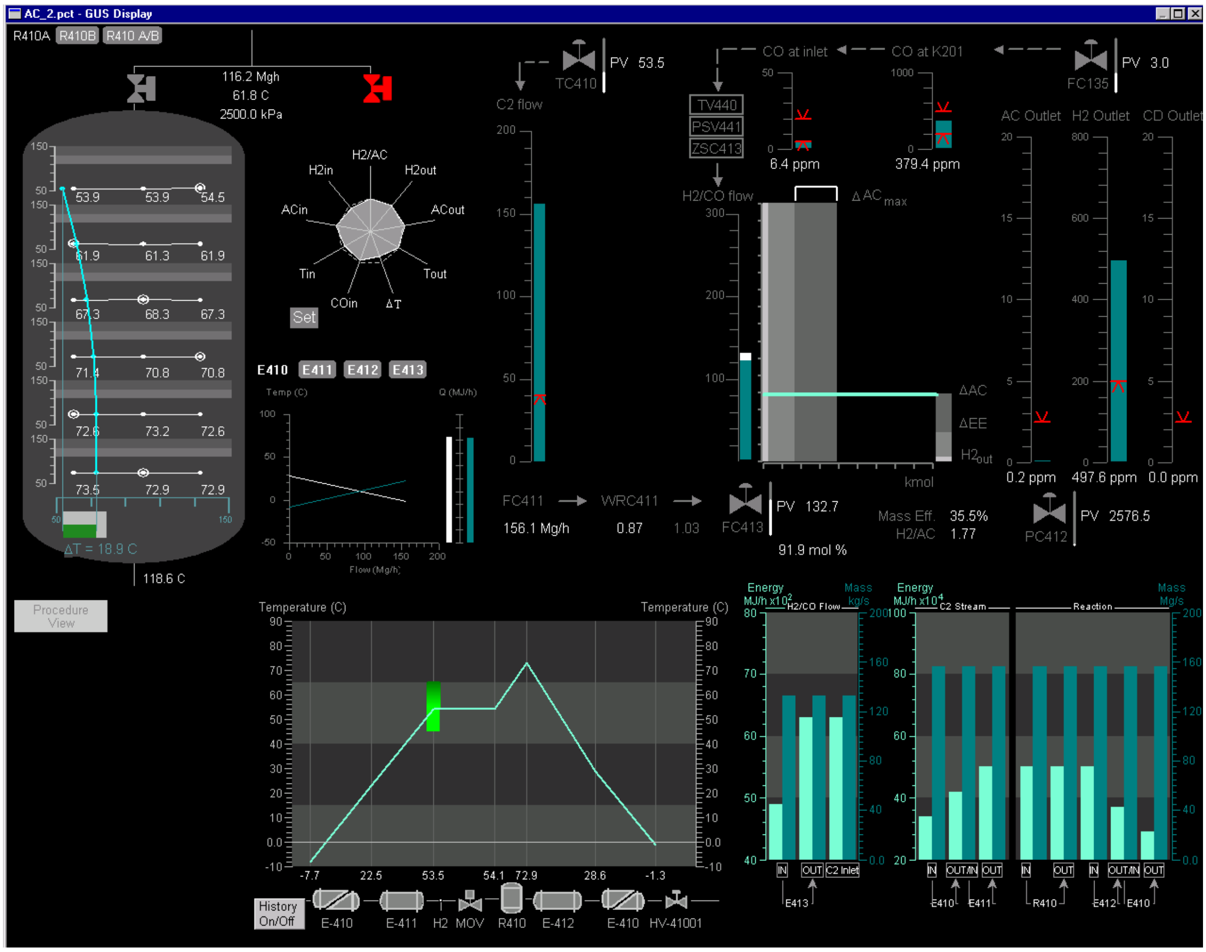

Figure 14: The P+F interface upper monitor.

Design Example

In this section, the graphical elements that are comprised of IRs for the WRC and surrounding process are examined in detail to serve as an example of how the entire set of IRs was parsed into process views. First, the IRs upon which the graphic was based are spelled out (in Table 2). Then, an annotated description of the appearance and behavior of the graphic (which appears in Figure 15) exemplifies how the IRs were instantiated. A similar description could be offered for any of the graphics in the P+F or P+F+T displays.

Table 2: Information Requirements for the weight ratio controller and surrounding process.

# on Figure 15	IR	ADS	HTA	CTA
1.	FV413 (appearance and location)	✓	✓	✓
2.	FV413 setting	✓	✓	✓
3.	The effect of FV413 setting on flow through the valve	✓	✓	✓
4.	The flow of CO into the reactor	✓	✓	
5.	The flow of H2 into the reactor	✓	✓	
6.	The flow of C2H2 into the reactor	✓	✓	
7.	The flow of C2H4 into the reactor	✓	✓	
8.	The flow of C2H6 into the reactor	✓	✓	
9.	The effect of the setting of SDV413a on the flow of H2 to the C2 stream	✓		✓
10.	The effect of the state of PV441 on flow of H2/CO through TV440 (to the C2 stream)	✓		
11.	The effect of the setting of TV440 on flow of H2/CO through CV6 and (to the C2 stream)	✓	✓	✓
12.	H2/CO Heat and Supply Input	✓		
13.	C2 Heat and Supply Input	✓		
14.	H2/C2 weight ratio SP		✓	✓
15.	H2/C2 weight ratio PV		✓	✓
16.	H2/C2 weight ratio error (SP-PV)			✓
17.	H2/C2 weight ratio OP (to FC413 SP)			✓
18.	FC413.SP			✓
19.	FC413.OP			✓
20.	Error (FC413.SP -FC413.OP)			✓
21.	ZSC-413 A status			✓
22.	TIC440.OP			✓
23.	Low flow limit on FC411.PV of 40 Mg/hr		✓	

H2 flow control information requirements. In Table 2, the information requirements that relate to the WRC and the surrounding process are consolidated from the IR Matrix. Note that each of the analyses contributed requirements for this graphic. The left-most column of the table contains a number to facilitate the following discussion.

The set of IRs in Table 2 defines the information that the designer must communicate through the graphical elements. Identifying a graphical form that communicates these requirements according to the design principles EID is an intuitive, creative process. No effort is made here to formalize that process. The next section describes how the IRs were mapped onto a graphical form.

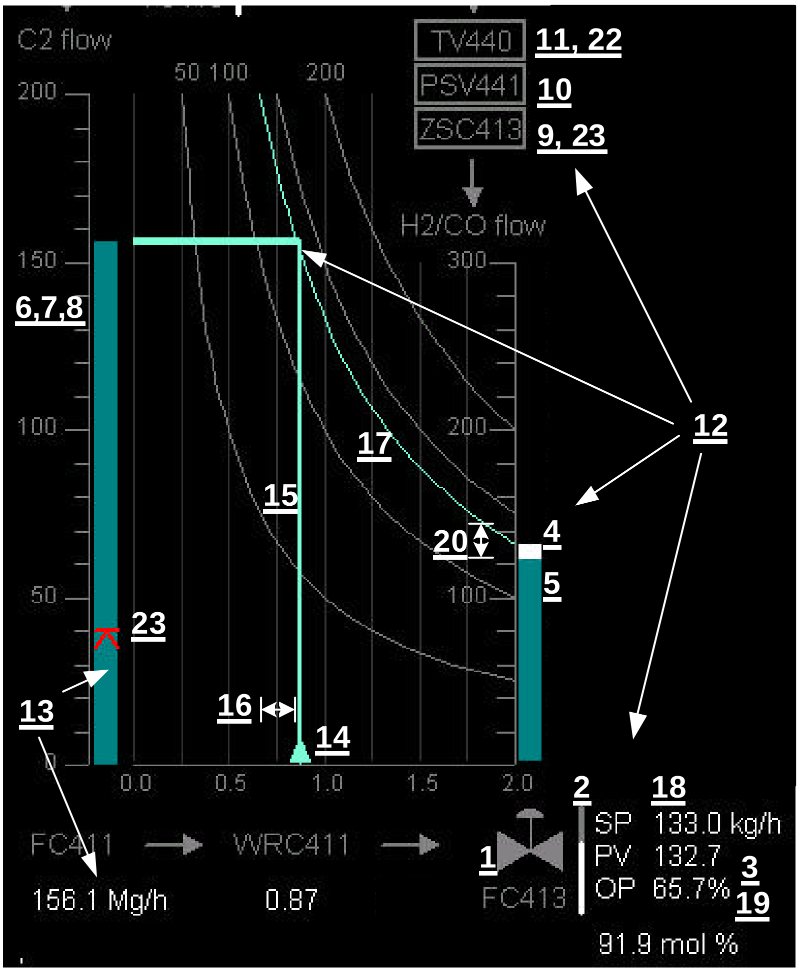

Figure 15: Hydrogen balance view: Ratio controller. Numbers correspond to IRs in Table 2.

<u>Ratio controller graphic.</u> The Ratio Controller Graphic allows the operator to monitor the control of the H2/CO stream (see Figure 15). Two flow-rate graphics are used to display the flow-rate of the C2 stream [13] and the H2/CO stream [4, 5]. Each graphic has a gray fixed scale, ranging from 0 to 200 in Mg/hr for the C2 flow (FC411) and 0 to 300 in kg/hr for the

H2/CO flow (FC413). Above each of the scales, the type of flow is denoted. A minimum limit for the FC411.PV is set at 40 Mg/hr [23]. The limit is indicated by red limit arrow formed by a chevron below a horizontal line segment.

For the C2 flow, a cyan bar indicates the flow rate [13]. A divided cyan [5] and white [4] bar is used to describe the flow rate of the H2/CO stream. The cyan portion of the bar represents the amount of H2, the white portion represents the amount of CO and impurities.

The behavior of the ratio controller is represented by a graphic that lies between the two flow rate bars [15, 17]. The baseline of the graphic reflects the value of the weight ratio controller (WRC) setpoint. It is drawn in a medium gray and has a range of 0.0 to 2.0. Darker vertical gray lines extend from the baseline at 0.25 unit increments, extending to a point equal to the maximum values of the C2 flow scale. Four curved constant value lines are drawn in medium gray on the plot at values of 50, 100, 150, and 200 kg/hr. At every point along each of the lines, the product of the value along the baseline and the value along the C2 flow scale is a constant. The lines are static. A fifth constant value line is drawn in green [17]. The value of this line is determined by the current value of FC413.PV [4,5]. A small green triangle [14] rests on the baseline at the present value of the setpoint of the weight ratio controller. A green line [15] extends vertically from the horizontal axis at the present value of the controller, up to the present constant value line. The difference between the position of this vertical green line and the green triangle is the error signal [16]. A second green line is drawn horizontally from the top of the C2 flowrate column across to the current constant value line. The two lines meet at the current constant value line defined by FC413.PV [17].

<u>Support for automation monitoring or manual setpoint control.</u> The graphic, as it is described above, can be used to monitor the performance of the weight ratio controller. When the WRC is keeping the ratio near the setpoint, the triangle and the vertical green line will

align. However, under some operating conditions, the operator will elect to turn off the WRC and set the flowrate FC413.PV manually. The graphic assists the operator by continuing to present the multiplication operation that is required to choose a hydrogen flow rate. In this case, the vertical green bar becomes a marker for the desired ratio, as opposed to the PV of the controller. The operator's task is to adjust the setpoint of the hydrogen flow controller (FC413.SP) so that the curved green live meets the intersection of the vertical and horizontal green lines. In effect, this is a direct replication of the automated behavior of the controller. The example shows that the same graphic can be used to support both automation monitoring and manual control.

Integrating system- and task-based information in EID. In addition to serving as a good example of how manual and automated control tasks can be supported by a single graphic, the graphic in Figure 15 also conveys some of the challenges to integrating task- and system-based information. It should be emphasized that, while this question was an important one for the design activity, the approach that we took here is one that may not be appropriate for other applications. One of the unique design challenges for this dissertation was to create a set of graphics that satisfied all of the IRs while simultaneously allowing us to extract the IRs that originated in the task-based analyses. This was necessary to create the experimental conditions (namely the P+F and P+F+T interfaces) that would allow an empirical evaluation of the theoretical question. Other applications of an integrated approach to IR generation and interface design will not likely assume this requirement.

Figure 16 shows how the graphic from Figure 15 appears when the task-based IRs are stripped away. Clearly this is not a particularly elegant design. However, it is the one that appears in the P+F interface. What was lost in this case is not only the information from the task-based analyses, but also the advantage of the configural graphic that brings the relationship between the C2 and H2 flows into view. It would have been possible to design a replacement for Figure 16 that better reflected the relationship between these two flows. However, when extended to the scale of the whole design activity, this approach would have been prohibitively time consuming.

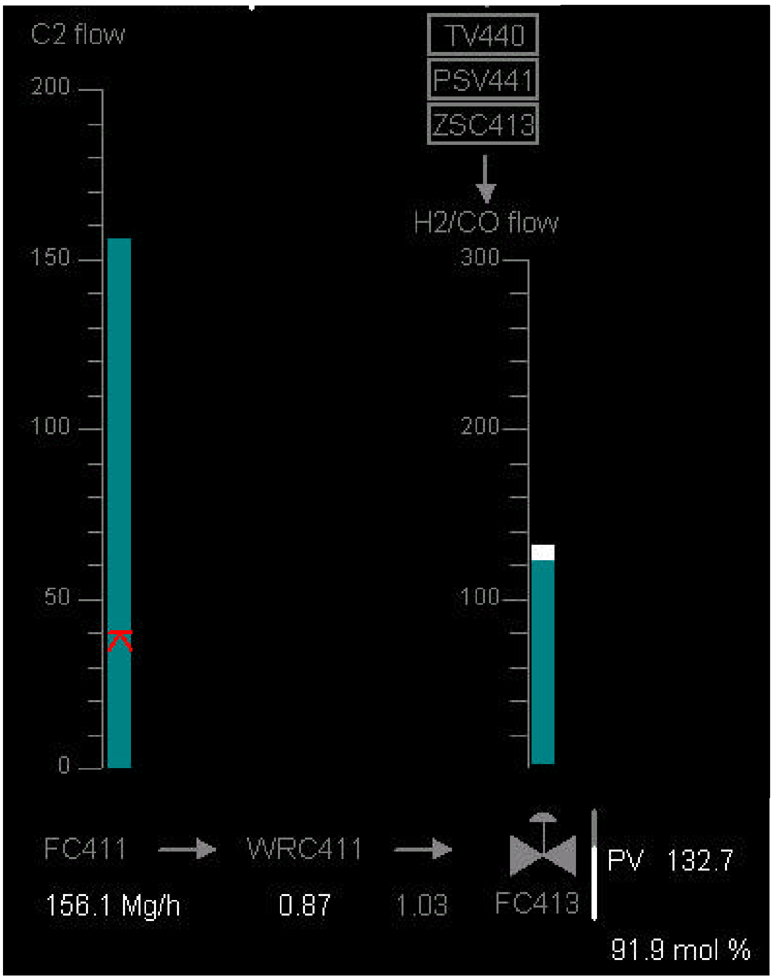

Figure 16: The C2 and H2 flow graphics based solely on system IRs.

There are alternative approaches to integrating the information requirements from the analyses. Just as prior applications of EID have parsed displays based on levels of the abstraction hierarchy (e.g., Burns, 1998; Reising, 1999; Ham and Woon, 2001b), it would have been possible to parse the displays so that the task-based IRs were all contained in the same process view. However, while this would have made it easier to extract the task-based IRs to form a P+F interface, it would have mapped IRs that are spatially and functionally related onto different, separated displays.

Implementation

One of the corollary issues for this dissertation was identifying and overcoming the challenges of designing and implementing an ecological interface in an environment that reasonably replicates the challenges that would be faced in an industry application. The following paragraphs contain a description of how this implementation was completed. Thereafter, we extract some take-away lessons from the experience.

Architecture. A set of ActiveX controls written in Microsoft Visual Basic 6.0 comprises a library of visual components that provide the final interface between the user and the DCS (see Figure 17). The DCS interfaces with the simulator via the Honeywell Total Plant Solution system (TPS); a newer DCS technology. The ActiveX controls are written to be compatible with the Honeywell TPS, Global User Station (GUS) Display Builder. This approach allows the plant/process engineer to incorporate the visual components in the same way that other GUS graphic objects are used.

The ActiveX visual components are passive. The GUS interface detects any change in variables assigned to the ActiveX components and sends the values to the components via component properties. When the ActiveX visual component needs to send information to the DCS (e.g., when an operator changes a setpoint), an event is triggered which activates GUS BASIC statements associated with the ActiveX components (written by NOVA personnel). These statements send data to the DCS or to other graphic elements.

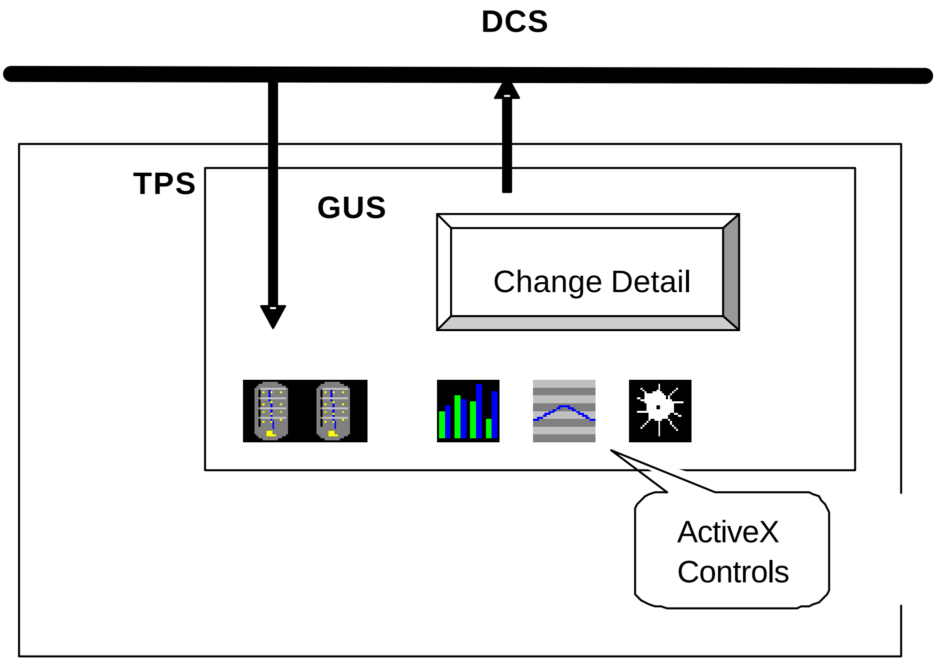

Figure 17: Implementation architecture.

The ActiveX visual components have a set of property pages that allow the plant/process engineer to assign preliminary values to properties of the components. This is the same way properties are initialized with any other ActiveX component. In addition, these ActiveX visual components have an initialization file associated with them to provide access to constants that facilitate fine tuning of the visualizations such as axis ranges, line widths, colors, separation of graphic elements, font sizes, etc. Another file common to these components is a Log file, which may be used to accumulate diagnostic information as to the performance of the controls.

As part of the deployment, each ActiveX visual component is packaged as a Visual Basic (VB) Application. Installation of the test application registers the controls in the Windows NT registry, creates a directory specific to each control where supporting files such as the constants file are housed, and provides a convenient mechanism to observe the operation of the ActiveX component under simulated but static conditions. Each installation conforms to the Microsoft Un-Install conventions.

Extractions from the implementation experience. The ecological interfaces were implemented by a professional software contractor with substantial experience in building software for complex graphics and prior experience with the TPS system. Technicians at NOVA provided GUS support to facilitate the data connection. Honeywell's own software developers had predicted that we would run into substantial difficulty in using ActiveX controls within the GUS environment. This did not turn out to be the case. However, we were convinced that it would have been substantially more difficult to try to implement the displays in GUS directly. Visual Basic affords the programmer much greater control over the appearance and behavior of the graphical elements than is possible with GUS. That we were able to code that behavior, wrap it up as an ActiveX control, and embed it in the GUS environment as a passive object was a key implementation finding. It suggests that other graphics environments (i.e., competitor products to Honeywell's GUS) might also be capable of hosting such controls. This would markedly expand the range of possible applications for a single control.

Implementing these displays in VB requires a great deal of programming skill and ingenuity. VB makes it very easy to create generic graphical interfaces base on standard components such as scroll bars and pull-down menus. An implementer who is adept at using the application for that purpose might be overwhelmed with the design requirements of displays of the complexity used in this and other applications of EID. Thus, it is imperative that the implementer be both knowledgeable about constructing ActiveX controls and creative. Moreover, he must have a grasp of the design intent of the graphics so that he can recommend alternatives that make a concept more feasible or less costly.

Conclusion

This chapter provides a detailed design example to demonstrate how a set of information requirements was instantiated in a graphical form. It also shows how the task-based elements were designed into P+F+T graphics in such a way that they could be removed to form the P+F graphics. Repeated for all of the other graphics in the interfaces, this process led to creation of two novel interfaces for the AHR (described in Appendix C). The following chapter outlines an experiment to compare the performance of professional plant operators using these interfaces to respond to normal, anticipated abnormal, and unanticipated abnormal process events.

EXPERIMENTAL METHOD

Participants

30 male professional operators participated in the study. The target population was qualified panel operators from two nearly identical production units (i.e., NOVA's E1 and E2 ethylene refining units). Both units contain the same equipment, similar levels and types of automation, and largely comparable procedures. The primary contingent of participants was current operators from each of these two units. Participants were also solicited from a third unit (NOVA's E3 ethylene unit), who were former qualified operators of E1 or E2. Table 3 details the demographics and operating experience of the participants.

Table 3: Participant demographics.

	Mean (years)	Std. Dev. (years)	Range (years)
Age	44.5	6.7	(31,57)
Industry Experience	19.7	5.9	(9,36)
Panel operating experience	6.6	2.9	(0.75,12.83)
NOVA Tenure	16.4	5.4	(3.7,24)
E1/E2 operating experience	3.6	2.4	(0.75,10)

26 of the participants had corrected or uncorrected 20/20 vision with no color deficiency. Three reported that they exhibited corrected or uncorrected 20/20 vision with a color vision deficiency and one had corrected vision of less that 20/20, but no color deficiency. On a scale from 1 to 5 with 1 being Novice and 5 being Expert, the mean self-rating of level of expertise with Microsoft Windows applications was 2.73 (s=0.85), with a range of 1 to 4.

Experimental Design

A 3x3 split-plot factorial design was employed, with Interface Group (Current, P+F, P+F+T) as a between participants factor and Scenario (Anticipated Normal, Anticipated Abnormal, Unanticipated Abnormal) as a within participants factor. The order of presentation of the scenarios was randomized. An equal number of participants (n=10) from each of three units were randomly assigned to each of the three interface condition (although participant cancellations lead to a slight upset of this balance). Experimental materials are shown in Appendix D.

Hypotheses

Hypotheses for performance as a function of Interface type and Scenario are shown in Figure 18. The hypotheses are stated as follows:

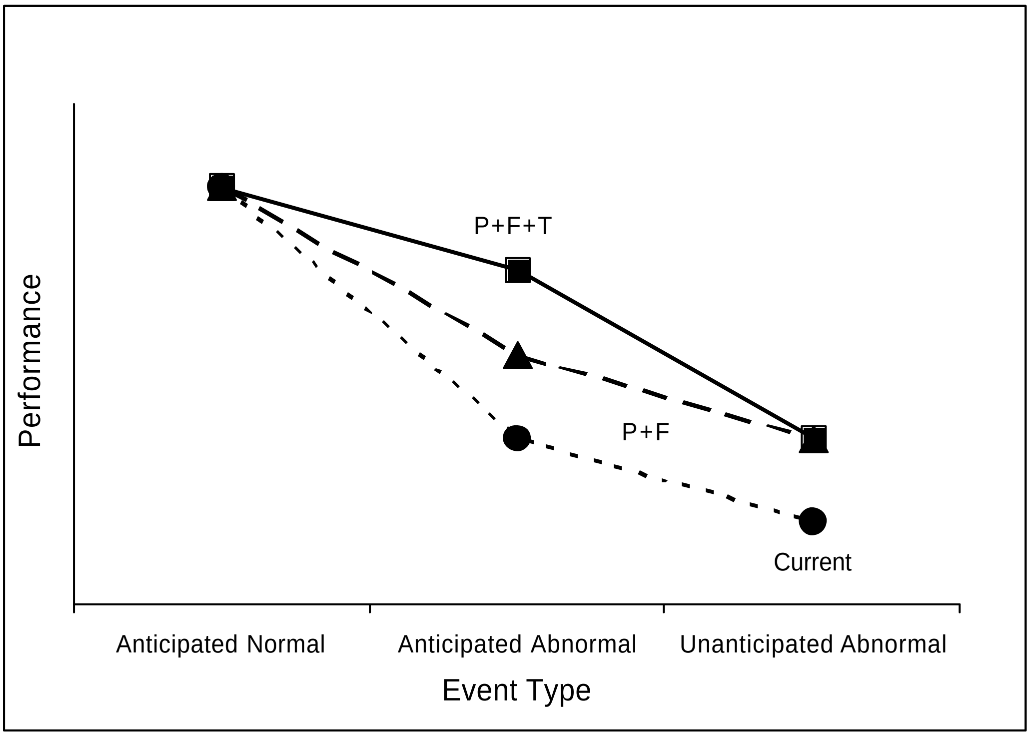

Figure 18: Hypothesized performance effects of Interface x Scenario interaction.

H1: An interaction between Interface type and Scenario classification was predicted. Operators using the P+F+T interface should perform better than the Current and P+F conditions in the Anticipated/Abnormal scenario. The rationale for this prediction is that the additional task information in the P+F+T interface should provide comparatively more support for anticipated events than either the Current or P+F interfaces. In addition, the P+F and P+F+T users should perform better than the Current interface for the Unanticipated/Abnormal condition. In this case, the P+F and P+F+T interfaces should equally support abnormal event management because both include functional information (and task information will not be relevant), whereas the Current interface lacks the functional information required to support abnormal events.

H2: A Scenario effect is predicted. The two Abnormal conditions should induce poorer performance than the Normal condition because they should require problem solving. Also, the Unanticipated/Abnormal condition should yield the worst performance because no procedure exists to direct the operator's response. Thus, the operator will need to generate a response as opposed to following an instruction.

H3: An Interface effect is predicted. The ecological interfaces should lead to better performance for the abnormal process events. This would be consistent with previous research in both laboratory and full-scope simulator studies.

Event Scenarios

For the purpose of this study, the dimensions of Normality and Anticipatedness were defined as follows. Any event for which a procedure had been written was deemed Anticipated. Events for which no procedure was written were considered to be Unanticipated. Normal events were those that were scheduled in the plant, regularly practiced in the simulator, and would likely have been executed by an operator in practice. Abnormal events were those that were unscheduled, rarely practiced, and most operators would never have encountered in operations. A description of each of the events is given below.

Anticipated/Normal. A reactor swing serves as the event for this category. Swinging the reactors involves redirecting process flow from one reactor to a parallel reactor. In actual operations, a reactor swing takes place on either a three-month or a 6-12 month cycle, depending on the type of catalyst contained in the reactor. The procedure is also regularly practiced in simulated operations. The scenario requires several control actions executed at the panel, and requires continual monitoring of a dynamic plant state. All of the activities are localized to the reactor section and are reviewed in advance by the operator.

Anticipated/Abnormal. A reactor temperature runaway best fits this categorization of event. The procedure has been employed only a few times in the history of the plant. It is practiced by operators, but infrequently. Some interaction with the outside operator is advisable (although not required), several timely control actions are called for, and continual monitoring of plant state is critical. Of the eligible procedures, the runaway had the most activities concentrated in the reactor section of the process. The event was instigated by a dual equipment failure caused by a localized fire in the unit. Ten minutes of steady state operation preceded the onset of the fault.

Unanticipated/Abnormal. A leak in the feed pre-heat exchanger serves as the event in this category. The leak raises the pressure in the shell side of the exchanger and inhibits the heating of the process stream. The cooler stream is less reactive, causing the reactor temperatures to fall. Keen monitoring, quick decision-making, prompt control action, and interaction with the outside operator are all required. Most of this activity should be confined to the AHR section of the process. Seven minutes of steady state operation preceded the onset of the fault.

Selection of Experimental Scenarios

Four criteria were used for selecting the two Anticipated events. First, they had to be events for which procedures were already written. Second, as far as possible, their causes and mitigations had to be constrained to the AHR section of the process. Third, several required control manipulations were desired. Fourth, the events had to be capable of being modeled in the simulator. Selecting events for the Anticipated Normal and Anticipated Abnormal conditions was a simple process once the criteria were specified because only one procedure satisfied all of the criteria. Using these four criteria, and with the aid of two domain experts, all of the procedures for the E1 unit were evaluated for use.

Of the five qualifying Anticipated Normal procedures, only the reactor swing had several control actions required within the reactor section of the process. Of the nine Anticipated Abnormal procedures, four were signaled by hard-wired alarms (which would have made detection a trivial task) and of the remaining five, the temperature runaway stood out in terms of prescribed control actions within the reactor section boundaries.

Selecting an Unanticipated Abnormal event and an instigating cause for the Anticipated Abnormal event was a more difficult challenge. In the case of the Unanticipated Abnormal scenario, the event had to follow the second and third criteria listed above (i.e., several steps confined to the reactor section). In addition, it had to be as close to novel as possible, as well as functionally meaningful. The tube leak was a good compromise in this case. There was no existing procedure for dealing with it and no procedure for dealing with the functional impact.

The functional meaningfulness criterion made it particularly challenging to create an initiating cause for the Anticipated Abnormal event because most of the plausible causes for a runaway have either been targeted by hard-wired alarms (e.g., compressor trips) or are familiar to the operator (e.g., certain failures that can lead to temperature runaway). (Note: The hard-wired alarms are distinct from the software driven alarm summary, which was disabled in the simulator. The hard-wired alarms appear on a fixed alarm panel and could not be universally silenced.) We eventually managed to generate a plausible root cause that could be modeled in the simulator. However, it required a pair of component failures.

In sum, the effective engineering of the information systems and operating procedures in use in the current plant made it difficult to select events that truly fit the theoretical categories.

Experimental Platform: Process Simulation

An industrial process simulation of the entire E1 unit served as the experimental platform. The system boundaries of the simulation are much broader than the boundaries of the AHR (for which the novel displays had been designed). For the purposes of this study, the operators were instructed to monitor and control only the AHR section. Tasks were selected that did not require control outside of the AHR. The simulation is model-based and has been tuned to closely match the dynamic behavior of the real unit. The operator controls the simulator through an emulated instrumentation and control system that is effectively identical to that used in the actual unit. A suite of eight 21" monitors and five keypad stations replicates the display and control equipment in the E1 unit (see Figure 19). In sum, the experimental platform closely approximates the appearance, behavior, and capabilities of the actual working environment.

Two artificial modifications were made to this environment to elicit monitoring behavior and collect data. First, the auditory indications from the alarm summary software were disabled. This forced the operators to rely on the graphical displays (new or old) to detect and diagnose upsets. This artificial manipulation was necessary because a key performance variable is event detection time. Auditory alarms would make many event detection tasks trivial and thereby subvert our evaluation of the displays. There is also an alarm panel in the simulator display suite that was visible to all participants (see two upper center monitors in Figure 19) and could not be silenced or covered. One or both of the panel alarms for the high reactor temperature and the automatic hydrogen trip sounded during the Anticipated/Abnormal

Figure 19: NOVA E1 simulator facility.

scenario, but never before the operator had detected the induced disturbance or pulled the manual trip button.

Second, a desk-mounted microphone was added to the operator station to facilitate data collection. No other modifications were made to the simulator, although it should be noted that, while working in the simulator, the operators were not exposed to the distractions that are common in their daily work environment.

Experimental Platform: Process Tracing

A commercial process historian software package was employed to record regular samples of the variables associated with the simulated process. These include the controlled variables (e.g., regulatory and advanced controls) and process variables (e.g., pressures, temperatures, and analyzer readings). Samples were taken of the full suite of parameters approximately every 5 seconds, the fastest rate afforded by the application.

Experimental Team

The author and a cohort comprised the experimental team. The author administered the questionnaires, provided training, controlled the audio equipment, and operated the data collection software.

A skilled former operator of the process served as a cohort in the experiment. The cohort was a qualified operator of both the E1 and E2 units and was known to all of the participants. His role was threefold. First, he acted as the simulated field operator, interacting with field equipment and instruments (as they are represented in the simulator) and communicating with the participants in a way that closely resembled the dynamic of an operations team. Second, he acted as the panel operator for the units immediately upstream of the target process. Third, the cohort controlled the process simulation and introduced faults as required by the protocol.

Because the cohort was an integral part of the data collection activity, his role merits detailed explanation. The cohort was the regular administrator of the simulator and had 5 years of experience in this role. No other person in the plant was qualified to run the simulator and his participation was therefore not optional. Prior to the start of the experiment, the cohort was briefed on the purpose of the experiment to a similar extent as the participants. That is, he knew that the purpose of the study was to compare the new interfaces to the current interface across normal and abnormal scenarios. He was not, however, aware of the hypothesized results and was not privy to the data collection beyond what he could observe while running the simulator. Moreover, the cohort had been provided with no training on the new displays, although he was one of the operators who reviewed prototype designs and provided feedback on them based on his prior experience.

During the data collection, the cohort and experimenter were seated in a room adjacent too the simulator room, out of the direct line of sight of the participants. Communication between the participants and the cohort was conducted verbally through an open door. The cohort was aware of which scenario was presented to the operator because he initiated each one in the simulator. He was advised that his interaction with the participants could bias their actions and admonished not to assist with fault diagnosis, despite the fact that such interaction is normal in other simulator training. The participants were similarly instructed to treat the cohort as a novice operator who would provide accurate reports of plant conditions and correctly execute control actions in the field, but would not assist with detection and diagnosis of any disturbances.

Generally speaking, the cohort performed his role in a highly professional manner with few notable errors. On one occasion he identified a fault to a participant before the operator could report his diagnosis (the data point was dropped). On several other occasions the experimenter cautioned him to respond to operator questions more carefully. However, he tended to be very cautious about what he said to the operators and frequently “cleared” his answers to their questions with the experimenter. Considering the challenges of performing this type of research, the cohort performed exceptionally well and made an invaluable and irreplaceable contribution to the effort.

Training and Supplementary Materials

Participants assigned to the P+T and P+F+T interface conditions were given approximately 2 hours of training on their respective interface. The training addressed the purpose, appearance, and behavior of the interface views. Each participant in this condition was given 45 minutes of self-directed practice with the novel interface. The practice session included observing an example process disturbance during which cooling water flow through a heat exchanger was lost. The participants were encouraged to follow the effects of the disturbance. At the end of the practice session, the participants were asked to demonstrate proficiency in completing several monitoring and control tasks, as well as to predict the expected appearance of the graphics under generic process anomalies (see Appendix D for specific queries). Refresher training was provided until each operator was able to complete all of these tasks without aid and provide correct predictions of appearance of all views for the anomalies. Neither the training disturbance nor the proficiency questions related directly to any of the experimental scenarios.

The cohort advised the current and former E2 operators of several small differences between their process and the E1 process. They were also provided with an ‘equipment mapping’ sheet to help them to relate equipment names from their unit to equipment names in the E1 unit. The operators reviewed this list of 8 items prior to the start of the scenarios and were allowed to refer to it during the test.

All of the written procedures that pertain to the reactor section were made available to the operators in the study. Two of those procedures are relevant to the Anticipated scenarios.

Experimental Task

At the beginning of the experimental session, the operators were told that the test scenarios were all localized to the reactor section of the simulated process. They were advised that no disturbances would be introduced in other sub-units that would typically be part of their purview. Otherwise, they were instructed to complete each of the three scenarios as if they had encountered the event in their normal operating capacity. They were told at the beginning of each scenario whether there was a planned activity for that scenario (the Anticipated Normal scenario) or not (the two Abnormal scenarios). In the latter cases, they were to monitor for and respond to any process disturbances. They were also instructed that disturbances could occur at any time during any of the scenarios or not at all.

To enhance the representativeness of the experimental environment, the participants were instructed that the cohort would act as both the field operator on duty and the operator of the upstream process. They were advised that the field operator would provide accurate observation about the state of the process equipment, but would not be as helpful with fault diagnosis as would normally be expected. Finally, they were asked to "think aloud" about their monitoring, diagnosis, and control behavior.

Description and Justification of Measures

Three types of measures were taken; demographic measures, objective performance measures, and qualitative preference measures. The demographic measures were collected to ensure that the participants have the anticipated level of understanding of the process and the scenarios. These measures address operating experience and familiarity with the scenarios presented among others. Dependent measures include hydrogen and acetylene slippage (important system functions), detection time, diagnosis time, diagnosis accuracy, performance score, and number and type of control actions. Finally, qualitative evaluations were conducted using Likert scales to assess the participants' attitude towards the interface that they used.

Most of the performance measures have been used in prior studies of EID (see Vicente, 2002). These include detection time, diagnosis time, diagnosis accuracy and control action counts. We assume that superior performance is indicated by prompt detection of events, quick and accurate deduction of the root cause of events, and efficient control behavior. The two 'slip' measures were added to assess possible differences in the economic performance of the interfaces. Acetylene passing through the reactor can contaminate the process stream and it is desirable to reduce excess hydrogen to reduce waste and prevent unwanted reactions. Thus, better performance would be associated with low values for these measures. The performance score was a new measure suggested by the availability of procedures to identify a canonical set of steps to be completed[5]. The scoring method was conceived a priori with the assistance of two highly experienced operators. A perfect score could be achieved by completing all of the expected steps.

[5] See the performance score results for an explanation of how this measure was used for the Unanticipated Abnormal event.

RESULTS

This chapter presents the quantitative results of the study described in the previous chapter. First, several preliminary issues such as missing data and treatment of the data are discussed. Then, results are presented by performance measure, with measures having significant effects appearing first.

Missing Data

There are several missing data points. Table 4 provides a summary of these points and an explanation of their source. The quantity of missing data was small compared to the amount of data that was successfully collected. The impact of the missing data was not substantial except in the assessment of fault diagnosis time (see below).

Table 4: Account of missing data points.

Subject #	Scenario/Questionnaire	Source	Impact
15	Anticipated Normal	Simulator failure	No process data No verbal protocol
4	Anticipated Normal Anticipated Abnormal Unanticipated Abnormal	Data collection device failure	No process data
43	Anticipated Abnormal	Participant did not complete trial, did not understand instructions	No process data No verbal protocol
3, 7, 9, 16, & 34	Anticipated Normal Anticipated Abnormal Unanticipated Abnormal	Audio recording operation failure	No verbal protocol
11	Post-Test questionnaire	Experimenter error	Preference questions do not include T views.
36	Post-Test questionnaire	Participant declined to respond	No preference responses

Overview of Data Analysis Issues

In many cases involving time-based measures, counts of control actions, and slip, Shapiro-Wilks tests revealed that the data were distributed in a non-normal fashion. The spread of values increased as the mean of the values increased (i.e., there was a right skew to the data). This is a common observation in measurements with a zero baseline and monotonic trend (as are each of the above measures). In these cases, log transformations of the data values were employed to improve the normality of the distributions. The log transformation pulls in extreme high values and spreads out low values. In cases where the log transformation is used, the vertical axis is labeled as, “ln (parameter(unit))”. The axis itself is linear in all cases.

Also, two items related to data presentation deserve mention. First, the results are reported with 95% confidence intervals where possible, as suggested by Vicente and Torenvliet (2000). These intervals are expressed in the error bars. Confidence intervals were used in place of null hypothesis significance testing because, in addition to communicating statistical significance, they provide three types of information: 1) they graphically identify the range over which the actual population mean is likely to lie, 2) they give an indication of the precision of measurement, and (3) they express relationships across a set of group means (Vicente & Torenvliet, 2000). Second, for each outcome measure, least squares means (lsmeans) were used to account for the effects of significant covariates that were unbalanced in the design. The plots reflect these adjusted means. For measures where covariates were determined to be significant through an ANCOVA, their impacts are described.

Trial Completion Time

A trial completion time (TCT) for each trial was calculated based on the process tracing data. Start and stop criteria for each of the scenarios are given in Table 5. The Anticipated Normal trial started and ended when the operator started and completed the procedure. The two Abnormal trials started with the event onset and ended when the participant had either secured or recovered the process.

Table 5: Start/stop criteria for trial completion time.

Scenario	Start Criteria	Stop Criteria
Anticipated Normal	1st move command to MOV on Reactor A	MOV on Reactor B closed AND TC410.PV<58° C
Anticipated Abnormal	Hydrogen Flow (FC413.PV) goes to 284.54 kg/hr	Flare open AND EITHER HV41001 closed OR MS410 closed
Unanticipated Abnormal	TC410.PV<TC410.SP-0.5° C	Flare open AND HV41001 closed OR Temp turnaround AND two cycles where ACR412.PV<2.5 ppm

Figure 20 shows the results of the adjusted log-mean trial completion times by Interface[6]. The graph shows a steady decrease in TCT between interfaces. The difference between the Current and P+F+T levels of Interface is significant, while the difference between the Current and P+F levels is not significant, although nearly so. The difference between the P+F and P+F+T interfaces is not significant.

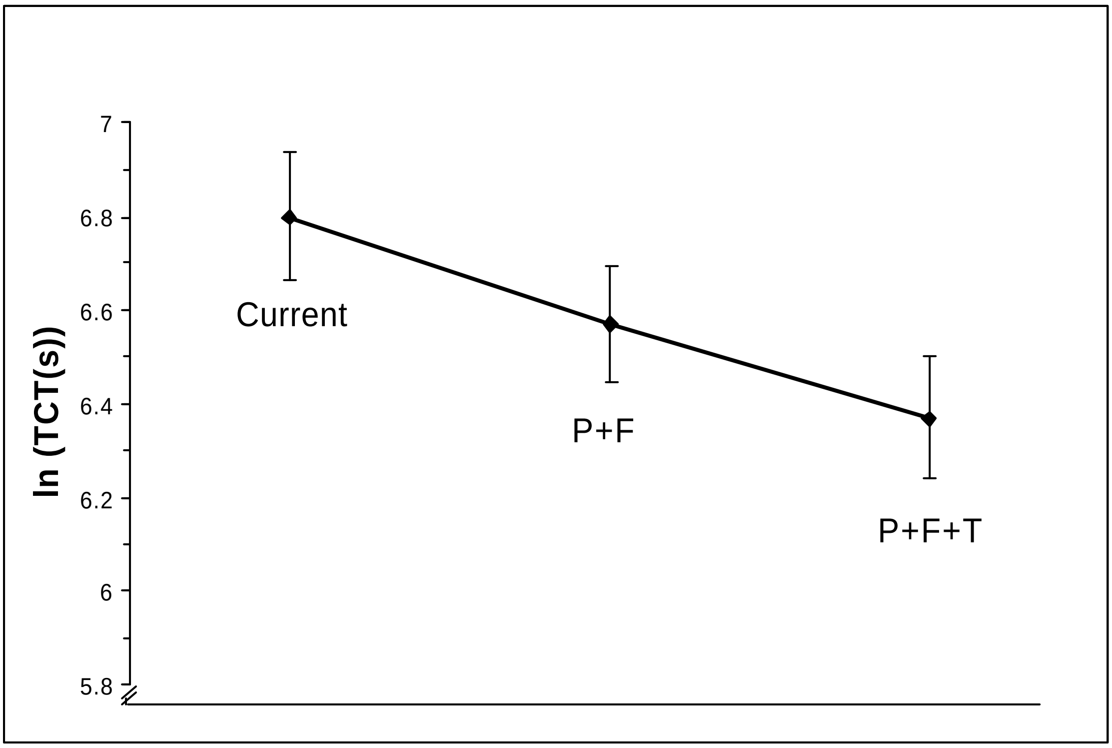

Figure 20: Log transformed trial completion time by interface.

There are also some significant differences in the Interface x Scenario interaction (see Figure 21). In both the Anticipated Normal and Anticipated Abnormal conditions, the P+F+T yielded significantly faster trial completion times than the Current interface condition. The P+F interface did not differ significantly from the Current interface in these two Scenarios, although the difference is very nearly significant in the Anticipated Abnormal case. All three displays showed similar TCTs in the Unanticipated Abnormal condition.

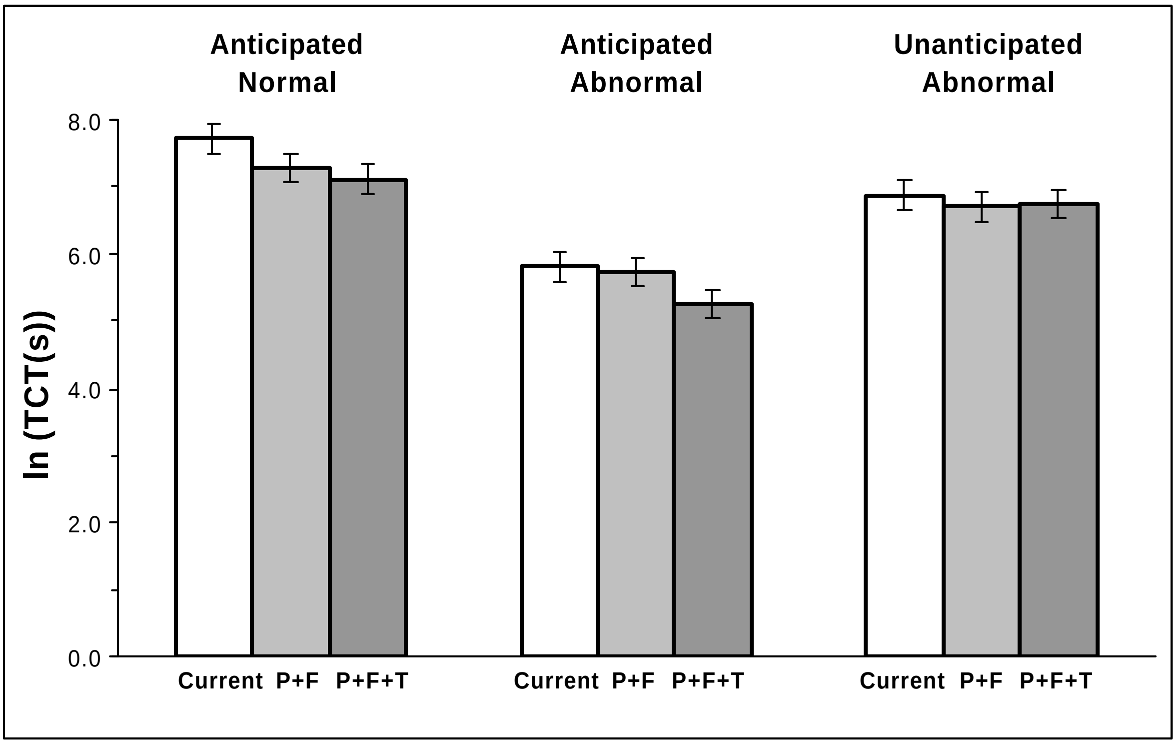

Figure 21: Log transformed trial completion time.

Table 6 shows the best estimates and confidence intervals for each of the differences between the P+F+T, P+F, and Current interfaces. These intervals are returned to the original scale (i.e., seconds as opposed to log seconds) to facilitate comparison using a procedure outlined by Ramsey and Shafer (1997). Thus, the best estimate indicates that participants in the

---

[6] Note that on this figure, and in several other figures in this section, the graph baseline is offset. Using a zero baseline made it difficult to compare the confidence intervals.

Current condition took 1.70 times as long (in seconds) to complete the trials as those in the P+F condition, and 2.70 times as long as those in the P+F+T condition. Participants in the P+F condition took 1.60 times as long to complete trials compared to those using the P+F+T interface. Despite these differences, care should be taken to note that the lower bounds on the confidence intervals for the Current–PF and PF-PFT comparisons are very close to including zero (on the log scale).

Table 6: Confidence intervals for differences in trial completion time.

	Best Estimate (log scale)	C.I. Range (log scale)	Best Estimate (geometric mean)	C.I. Range
Current-P+F	0.24	(0.05, 0.41)	1.70	(1.11, 2.59)
Current-P+F+T	0.44	(0.24, 0.62)	2.70	(1.75, 4.14)
P+F-P+F+T	0.20	(0.02, 0.04)	1.60	(1.05, 2.39)

An ANCOVA revealed that trial completion times also showed several significant covariate affects (in fact, the corrected lsmeans were used in Figure 20 and Figure 21[7]). Participant Unit was a significant covariate for TCT (see Figure 22), with E1 operators demonstrating faster completion times than either E1 or E2 operators. The difference between E2 and E3 operators was not significant. Interestingly, both participant age, $F(1,23)=8.25$, $p<0.01$, and industry experience, $F(1,23)=7.69$, $p=0.01$ were also significant covariates. However, contrary to expectations, the effects are in opposite directions from each other. Despite the log transformed age and log transformed years of industry experience having a strong positive correlation ($r^2=0.68$), the correlation between age and TCT was positive ($r^2=0.06$) while the correlation between years of industry experience and TCT was negative ($r^2=-0.03$). Thus, older participants were slower while more experienced participants were faster. This apparent discrepancy (we would expect both covariates to have the same sign) is very difficult to interpret, although the relatively small effect size implies that the relationship may not be a critical one.

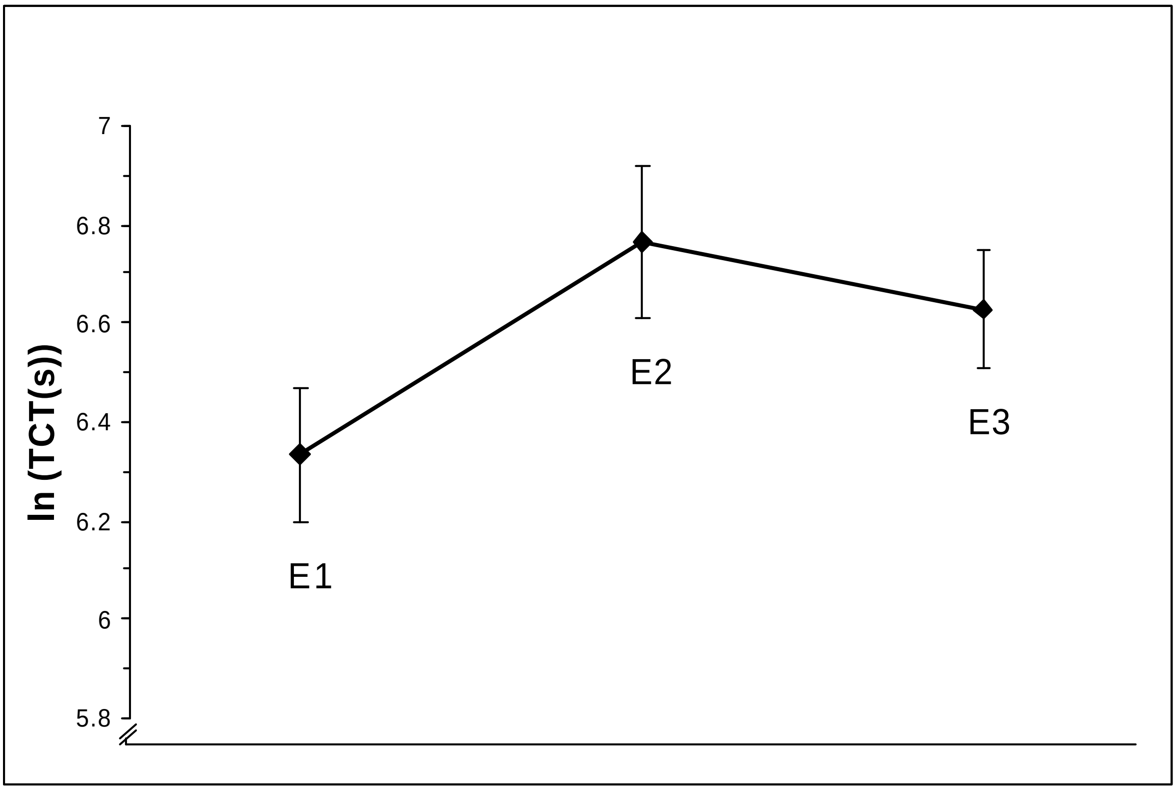

Figure 22: Log transformed trial completion time by unit.

Finally, Figure 23 shows that the differences between the log transformed mean TCTs for each Scenario were all significant. However, these means are not particularly interesting because they primarily reflect a difference in the time constants of the events as opposed to any differences in the difficulties of dealing with the particular types of Scenario.

Anticipated Abnormal and Unanticipated Abnormal Trial Completion

One of the difficulties in establishing trial completion times came from the differences in outcomes of the two fault scenarios. In the Anticipated Abnormal case, for example, some

---

[7] It is worth noting that the qualitative results for TCT were the same regardless of whether corrected or uncorrected means were used.

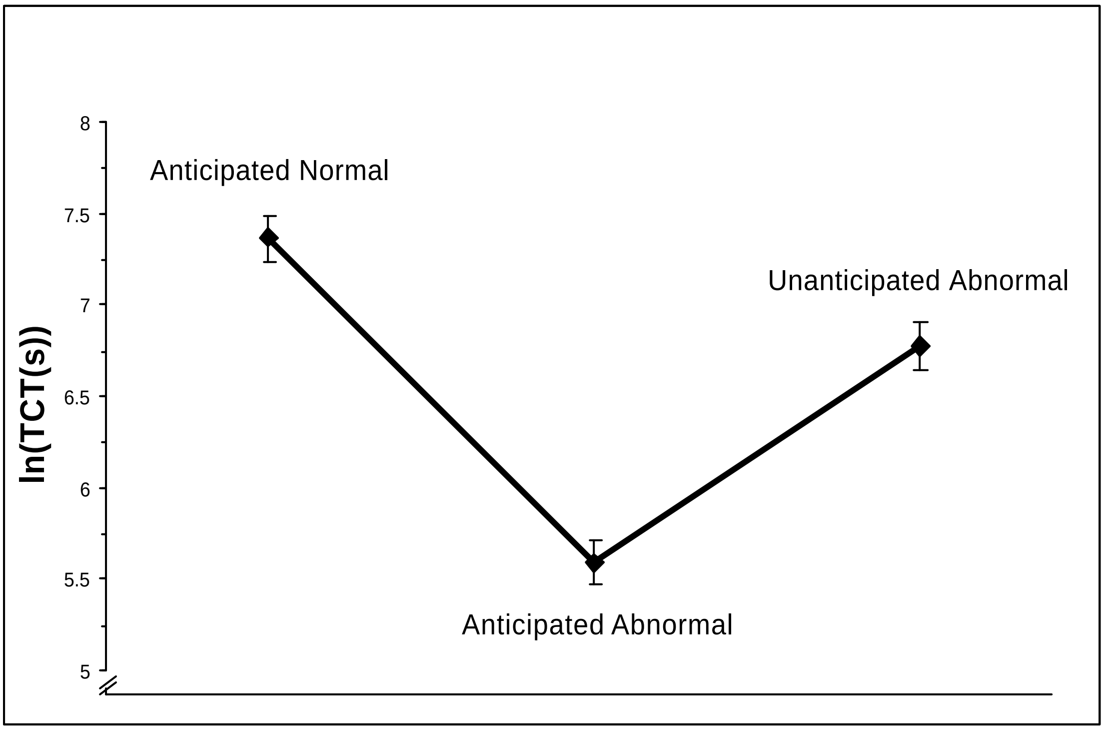

Figure 23: Log transformed trial completion time by scenario.

of the operators never identified the large hydrogen source into the reactor. Of those who did, only a subset managed to isolate that source. This motivated an alternative consideration of trial completion performance. The numbers of participants who failed or passed the hydrogen isolation criterion in the Anticipated/Abnormal Scenario is listed in Table 7. Performance against the criterion was assessed by reviewing the process data for the variable associated with hydrogen flow rate into the reactor. If that value was reduced to 0 following the onset of the fault, isolation was confirmed. A Fisher's Exact Test on the pass/fail proportions indicates that the success or failure in isolating the hydrogen source is affected by the Interface ($p<0.05$). More participants in the P+F and P+F+T groups managed to isolate the fault.

Table 7: Hydrogen isolation pass/fail in Anticipated Abnormal scenario.

Interface	Fail to Isolate	Isolate	Total
Current	5	5	10
P+F	2	8	10
P+F+T	0	9	9
Total	7	22	29

A similar criterion-based evaluation was attempted for the Unanticipated Abnormal Scenario. The proportion of operators recovering the process or failing to recover was recorded and tested using the Fisher's Exact Test. The result for the table shown in Table 8 is not statistically significant ($p<0.30$).

Table 8: Process recovery pass/fail in Unanticipated Abnormal scenario.

Interface	Fail to recover	Recover	Total
Current	4	6	10
P+F	7	3	10
P+F+T	3	7	10
Total	14	16	30

Control Actions

More significant differences were observed in terms of the number of control actions made by the operators. Operator control actions on the five primary controlled variables in the AHR section were identified from the process data. Those variables were:

TC410 – Reactor inlet temperature control setpoint or valve output

WRC413 – Hydrogen weight ratio control setpoint

FC413 – Hydrogen flow control setpoint or valve output

PC412 – Flare valve control setpoint or valve output

FC135 – Indirect control of carbon monoxide in hydrogen stream

<u>Total control actions</u>. Figure 24 shows the log transformed total number of control actions by Interface. There is a decrease in the number of control actions as the amount of information provided in the interface increases. Participants in the P+F and P+F+T conditions

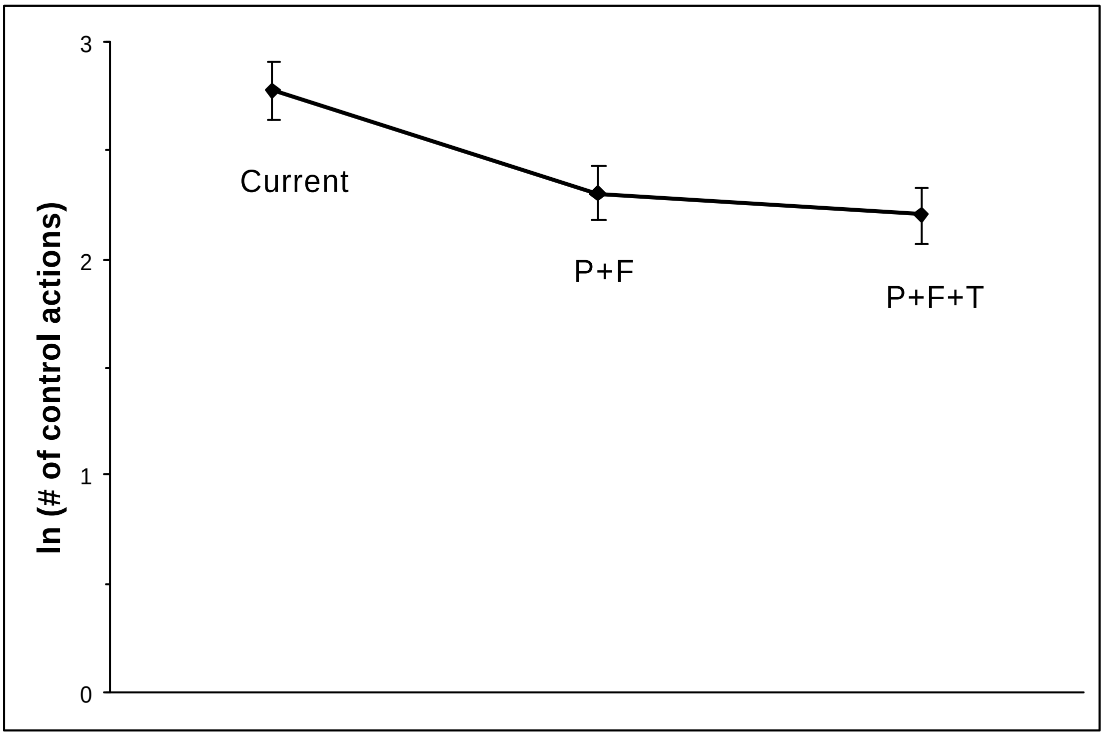

Figure 24: Log transformed total control actions by Interface.

took significantly fewer control actions than participants in the Current condition. The difference between the P+F and P+F+T groups was not significant.

Table 9 shows the best estimates and confidence intervals for each of the differences between the P+F+T or P+F interface and Current interface. As with the trial completion times, these intervals are returned to the original scale (i.e., moves as opposed to log moves) to facilitate comparison. Thus, the best estimate indicates that participants in the Current condition made three times as many control actions as those in the P+F condition, and almost four times as many control actions as those in the P+F+T condition.

Table 9: Confidence intervals for differences in number of control actions.

	Best Estimate (log scale)	C.I. Range (log scale)	Best Estimate (geometric mean)	Range
Current-P+F	0.477	(0.285, 0.8753)	3.00	(1.52, 5.92)
Current-P+F+T	0.580	(0.1823, 0.7723)	3.80	(1.93, 7.50)

MOV control actions for the Anticipated Normal scenario. The reactor swing procedure involves opening and closing two motor-operated valves in steps. Direct support of this task was provided in the P+F+T interface MOV View. In addition, both the P+F and P+F+T displays provide functional information about the two reactor temperature profiles, which are directly affected by these valve movements. Therefore, the number of control actions taken during this procedure may reveal some performance differences between operators in each of the Interface conditions. In particular, it may reveal a difference between the P+F and P+F+T interfaces, since the former did not contain the detailed information to support of the MOV move task.

Figure 25 presents the number of MOV control actions for by Interface for the Anticipated Normal scenario. Similar to the count of total number of control actions, the plot shows a significant difference between the Current and P+F+T groups, and a nearly significant difference between the Current and P+F groups. However, although the P+F+T group made fewer control actions than the P+F group, this difference is not statistically significant.

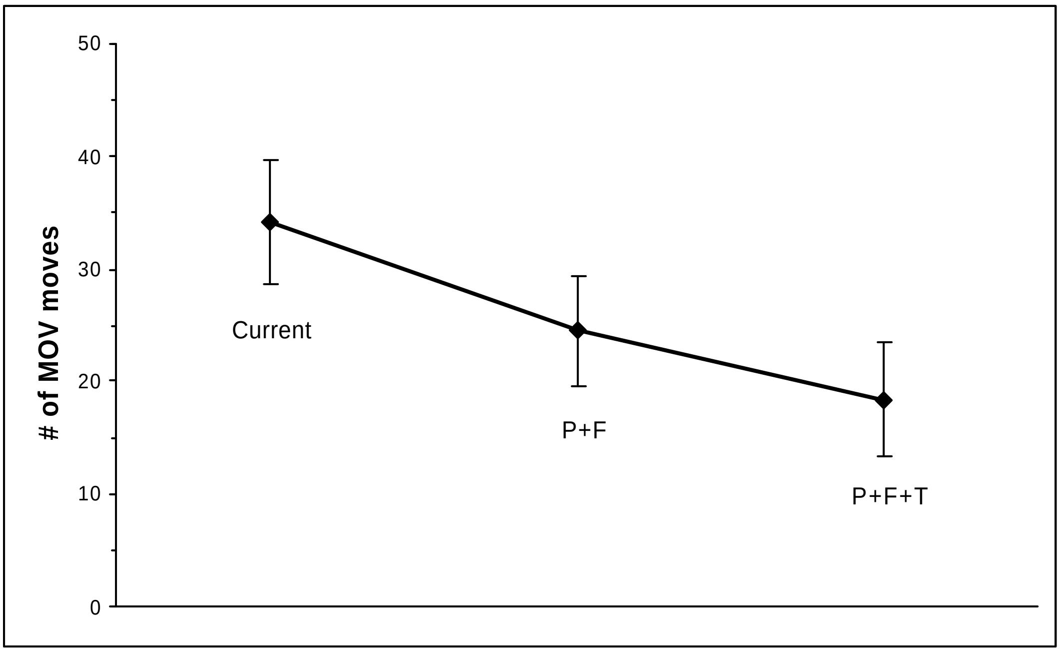

Figure 25: MOV control actions in the A/F scenario.

This particular comparison was of interest because control of the MOV is directly supported by the MOV View, which the P+F+T interface included and the P+F interface did not. Thus, a View that was drawn exclusively from the task-based work analyses is associated with a large drop in the number of control actions that make up the task. However, we are also seeing a nearly significant effect for the P+F group, and no difference between P+F and P+F+T. A possible explanation for this is that both displays contain the Reactor Temperature Profile View. The primary criterion for selecting the timing and magnitude of inputs to the MOV is the temperature profile in the reactor. Thus, both ecological interfaces support the assessment of the function whose state is the primary driver of MOV control actions. And operators using either of these displays show at least marginally significant reductions in the number of control actions to the component that affects that function. Participants in the P+F+T group show a further reduction in control actions when using an interface that explicitly supports execution of that action, although the additional reduction is not statistically significant.

Control actions by scenario. Similar to the TCT metric above, the main effect of control actions by Scenario was also significant, but not highly informative (see Figure 26). More control actions were taken in the Anticipated Normal Scenario than in either of the other two Scenarios. The difference between the number of control actions in the Anticipated Abnormal and Unanticipated Abnormal is very nearly significant. However, these differences reflect differences in the nature of the tasks. The Anticipated Normal scenario involves a stepwise opening and closing of two valves accompanied by several temperature setpoint changes over a long Scenario. The Anticipated Abnormal Scenario demands a few, well-timed control actions to isolate the AHR. The Unanticipated Abnormal Scenario offers opportunities for several moves to protect the stability of the process while the operator either diagnoses and corrects the fault, or enters into the isolation steps. Thus, the three Scenarios call for control actions under three very different task conditions and comparing between them to assess categorical differences between the Scenario types would have low face and construct validity.

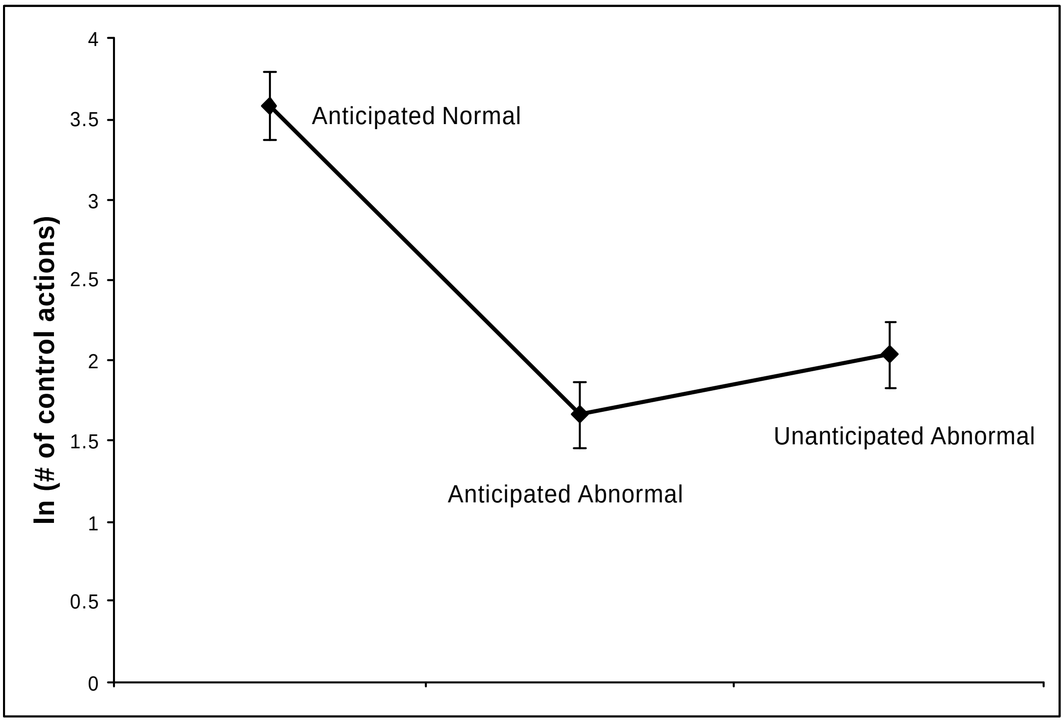

Figure 26: Log transformed number of control actions by scenario.

Diagnosis Accuracy

Diagnosis accuracy was scored according to a scaling technique described by Pawlak and Vicente (1996) and adapted to this application. The scoring criteria are:

0- The operator says nothing relevant to the fault.

1- The operator provides a vague, but correct description of the effects of the fault. Example: "The inlet temperatures to the reactor is falling."

2- The operator provides a correct statement of the specific functional impact of the fault. Example: "I seem to be losing heat exchange in E411."

3- The operator provides a correct localization of the faulty component. Example: "There must be a non-condensable fluid on the shell side of E411."

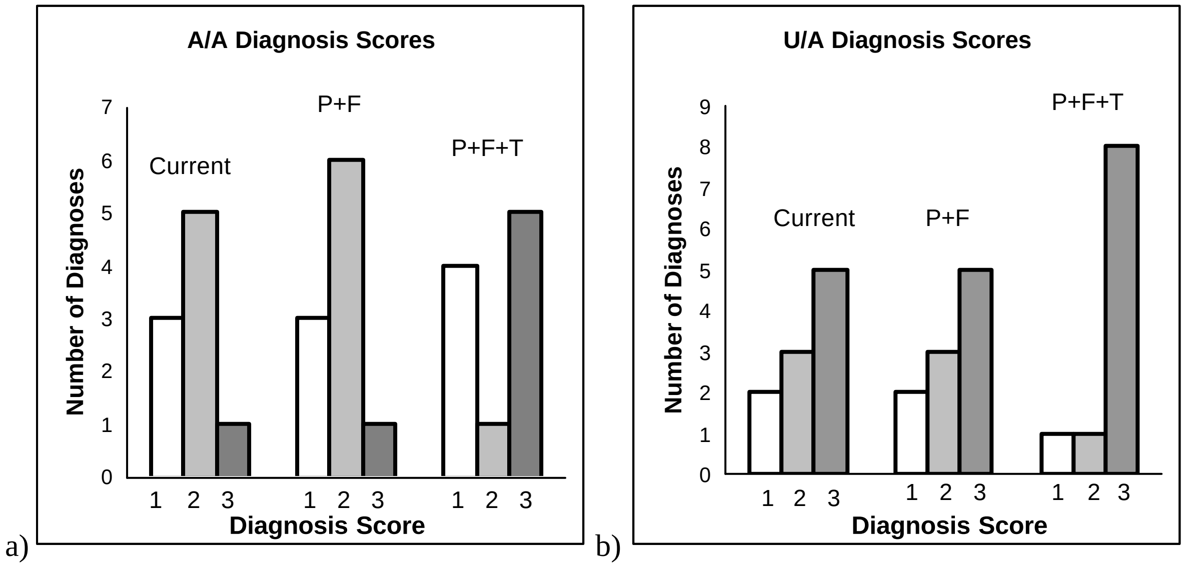

Figure 27: Fault diagnosis accuracy scores for the a) Anticipated Abnormal and b) Unanticipated Abnormal scenarios.

Scores for the two fault scenarios are presented in Figure 27. In previous studies of EID (e.g., Burns, 2000), non-parametric tests have been performed on diagnosis data of this type by combining the scores from the scenarios and applying a Chi-Square test. However, the present case presents a challenge to one of the criteria for employing the Chi-Square test. Siegel and Castellan (1988) point out that the expected frequency table that is calculated in the preparation of the Chi-Square test statistic should not have cell values that are too small (i.e., no more than 20% of the cells should have an expected frequency of less than 5). The distribution of diagnosis scores obtained in this study yields an expected frequency table that fails to meet this criterion. However, Siegel and Castellan (1988) offer a solution to this problem in a process of combining categories to raise the expected cell values. When applied to the present data, a combination of the 1 and 2 scores leads to an expected frequency table with no cell values less than 5 (with the exception of the cells for the zero scores, which were excluded from the following test). The contingency table with the combined categories is shown in Table 10.

Table 10: Frequency of diagnosis score across interface (scores of 1 and 2 combined).

	Display			
Score	Current	P+F	P+F+T	Row Total
0	0	0	0	0
1 or 2	13	14	7	34
3	6	6	13	25
Column Total	19	20	20	59

The contingency table is statistically significant (χ^2=6.35, df=2, $p<0.05$)[8] and we conclude that the diagnosis score depends on the Interface treatment. In both scenarios, there are more 3 scores and fewer 1 or 2 scores in the P+F+T condition. In contrast, the Current and P+F conditions look virtually indistinguishable from one another.

Performance Score

The performance of each operator on all three trials was scored on a scale from zero to four according to the criteria listed in Table 11. The steps for the Anticipated Normal and Anticipated Abnormal scenarios were drawn from the written procedures. Although these procedures call for a specific order of operations, only step completion was used in determining the scores because, in both cases, the task constraints on order are very weak. The scoring criteria for the Unanticipated Abnormal scenario could not be based on an existing procedure and were therefore developed in collaboration with the simulator operator and the senior operator in charge of procedures. The scale for this scenario was based on a combination of their expert judgment and a review of the range of responses observed.

Given the more subjective process of identifying the criteria for the Unanticipated Abnormal scenario, it is possible that the scale reflects what these two experts expected or observed the participants to do as opposed to what might be included in a procedure explicitly

[8] It is worth noting that the cell consolidation procedure noted above does tend to increase the power of the Chi-square test. However, the decision to follow the procedure was based on cell size and not power.

developed for the novel event. In addition, there was also some concern that the score could be influenced by the presence of control action sequences. For example, when securing the reactor (regardless of the root cause of a disturbance) it is normal to both open the flare valve and close HV41001. Thus, scoring each of these as one actionmay tend to return higher scores because the operator may be executing a sort of control script.

Table 11: Performance scoring criteria.

Scenario	Scoring Criteria
Anticipated Normal	1 point given for each of the following steps: Open MS410 Close MS411 Reduce inlet temperature < 58° C Increase hydrogen flow ratio
Anticipated Abnormal	1 point given for each of the following steps: Close MS410 Open Flare Valve Reduce inlet temperature Close HV41001
Unanticipated Abnormal	0. No action relevant to the fault 1. Response to surface characteristics of the fault (e.g., manipulate TC410.SP) and <u>fail to execute</u> flare procedure when TI417.PV>130.0° C 2. Response to surface characteristics of the fault (e.g., manipulate TC410.SP) and <u>execute</u> flare procedure correctly when TI417.PV>130.0° C 3. Vent non-condensables from E411 shell and execute flare procedure correctly when TI417.PV>130.0° C 4. Vent non-condensables from E411 shell and recover process without going to flare

A bar chart of the performance score distribution for the Anticipated/Normal scenario is shown in Figure 28. A visual inspection of the scores indicates that there is almost no difference between the Interfaces for the Anticipated/Normal Scenario. Given the high rate of complete task performances (i.e., scores of 4) across the groups, a possible explanation for not observing a benefit for the P+F+T Interface is that there is a ceiling effect for the Normal

Scenario. This is corroborated by the observation that nearly all of the participants read the procedure before they started the trial.

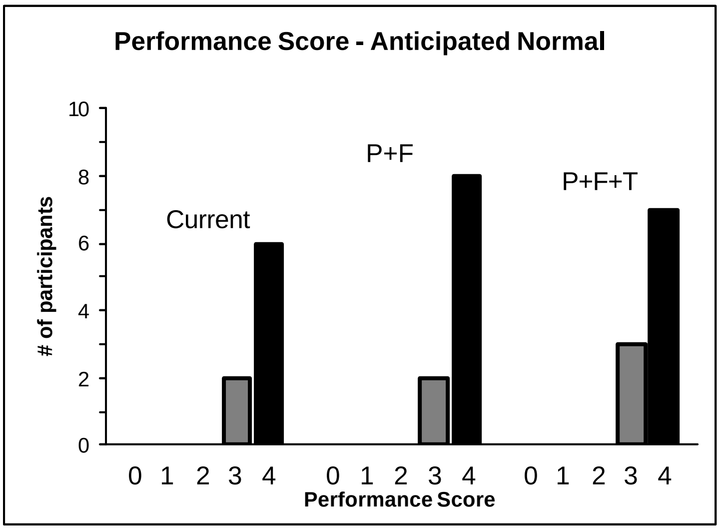

Figure 28: Performance score distribution for Anticipated/Normal scenario.

Bar charts of the performance score distributions for the Anticipated/Abnormal and Unanticipated/Abnormal scenarios are shown in Figure 29. A combined chart is shown in Figure 30. The scores for these two scenarios were entered into a contingency table to conduct a Chi-Square test (see Table 12). Once again, issues with minimum cell values in the expected frequency table were observed (which is why the two scenarios could not be examined independently). To compensate, performance scores of 0, 1 and 2 were combined and compared against scores in the 3 and 4 categories. The contingency table is statistically significant (χ^2=10.62, df=4, p<0.05). We conclude that performance scores in the abnormal scenarios depends on Interface Group. However, the cause of this difference is less clear than in the case of diagnosis score. It appears that the P+F+T display is generating the greatest number of complete performances. However, the Current Interface has a fair number of complete scores in the Unanticipated/Abnormal case as well. Notably, the distribution of scores in the Unanticipated/Abnormal condition is wider and possibly multi-modal for the Current and P+F+T interfaces.

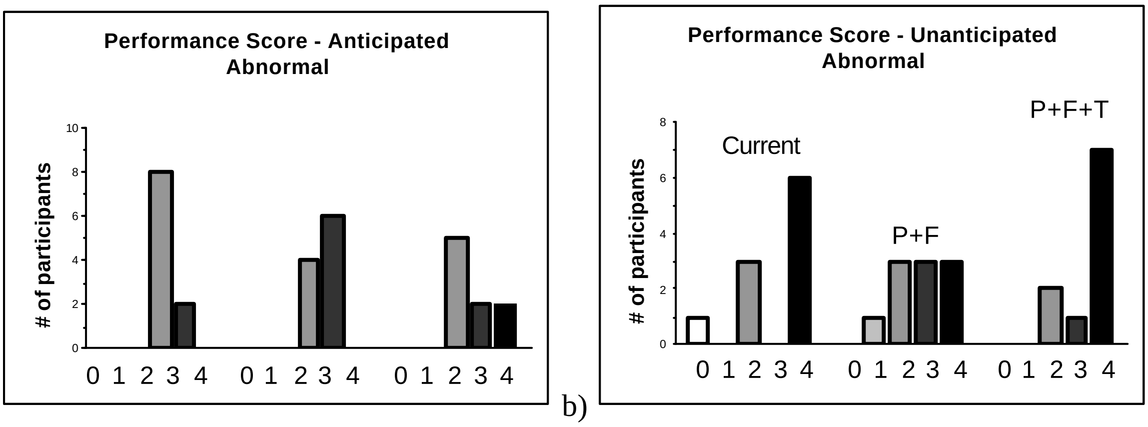

Figure 29: Performance score distributions for a) Anticipated Abnormal, and b) Unanticipated Abnormal scenarios.

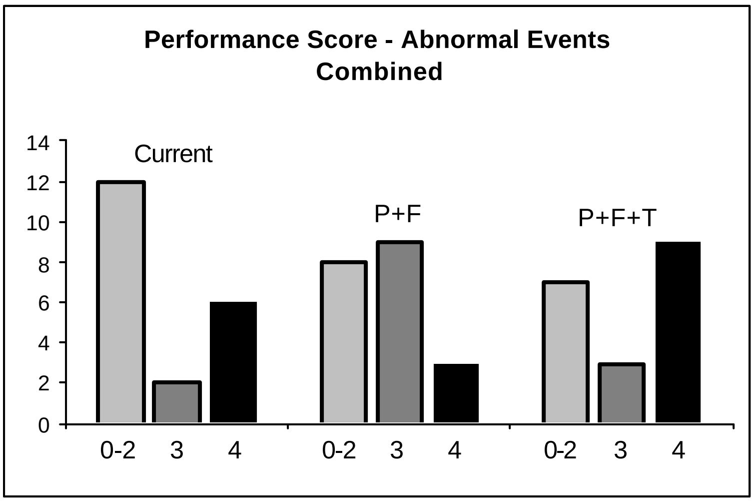

Figure 30: Performance score distribution for combined Anticipated Abnormal and Unanticipated Abnormal scenarios.

Table 12: Frequency of performance scores across Interface for abnormal scenarios.

	Display			
Score	Current	P+F	P+F+T	Row Total
0,1, or 2	12	8	7	27
3	2	9	3	14
4	6	3	9	18
Column Total	20	20	19	59

Finally, it is also interesting to note that performance scores for all three Interface groups appear to be higher in the Unanticipated/Abnormal Scenario compared to the Anticipated Abnormal Scenario. Hypothesis H1 predicted that the scores for the Current and P+F groups would decline more than the performance for the P+F+T group across these two events. Had this trend towards higher scores appeared in both the Anticipated Abnormal and Unanticipated Abnormal cases, this might have signaled the influence of a scoring inaccuracy as mentioned above. However, given that these high scores appear only in the Unanticipated Abnormal condition and not in the Anticipated Abnormal condition (where the same control script would have been equally likely), there is no consistent evidence that such a scoring inaccuracy was present.

Fault Detection Times

The time elapsed between event start and the first indication that the operator noticed an anomaly is captured as fault detection time (DT). The detection indications were sometimes overt but were often judged by the experimenter. An indication could include any oral or physical behavior that conveyed surprise (e.g., "Uh Oh!"), focused attention (e.g., sitting up sharply)[9], or confusion (e.g., "What the…"). The detection times were recorded in real time

[9] This indication was used in only one case where the operator's behavior was highly exaggerated and observed by both the experimenter and the simulator operator. The participant in question did not make any verbal indications until well after he started making control actions that were consistent with detection of the fault.

during the data collection sessions and, when audible, on the verbal protocol. The cohort provided start times (upon initiating the event in the simulator) and the experimenter judged detection times via the audio system. Because faults were only presented in the two Abnormal Scenarios, only times for these two Scenarios are included in the Group means and confidence intervals presented in Figure 31. While the DTs for the two ecological interfaces are faster than the Current interface, the mean differences are not significant.

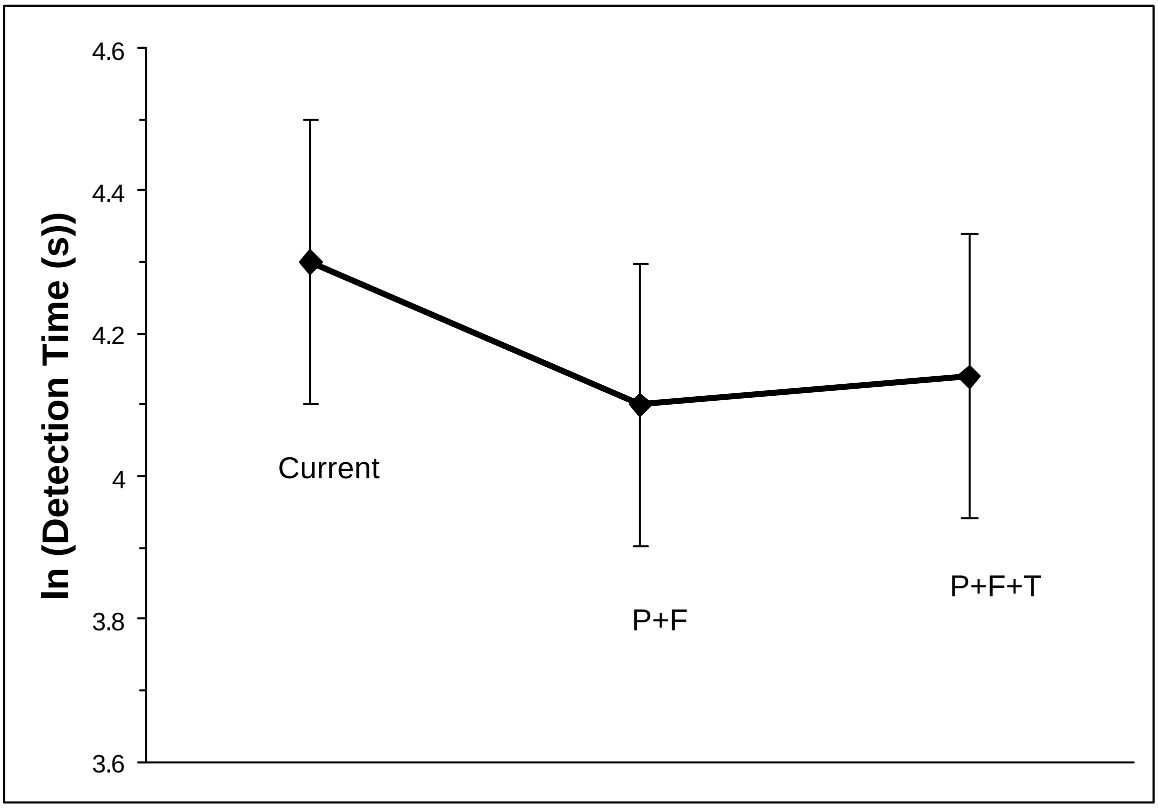

Figure 31: Log transformed fault detection time.

As with the TCTs above, the means in Figure 31 are least squares means, corrected for significant covariate effects of Finishing Experience and Order. Finishing Experience (r^2=-0.34) and Order (r^2=-0.21) were both negatively correlated with Detection Time. Thus, operators with more experience controlling an AHR (regardless of the unit in which it was acquired) detected events more quickly. Operators who saw a fault trial later in the sequence of three trials were also faster at detecting faults. This suggests that some learning was taking place during the study. Notably, no specific scenario order effects were observed. That is, detection times were faster for later trials regardless of the specific order of the earlier trials.

Hydrogen and Acetylene Slip

Given that the functional purpose of the AHR is to keep acetylene concentration below 5 ppm, a measure of the amount of acetylene passed during a scenario is a measure that offers a high degree of face validity. In addition, an operational constraint on the process is to keep the concentration of hydrogen in the reactor effluent below 200 ppm. Thus, the total amount of acetylene and hydrogen output from the reactor, or 'slip', was measured for each scenario. However, no group differences were found for either of these measures (see Figure 32).

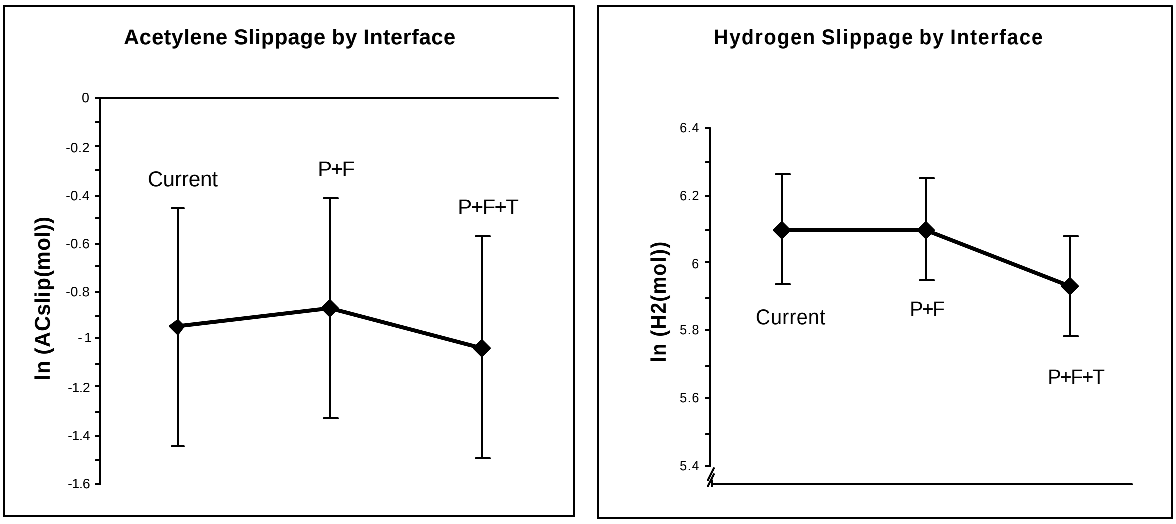

Figure 32: Log transformed mean acetylene and hydrogen slippage.

Time to Diagnosis

In several previous studies, the time to reach a complete diagnosis (Score of 3) with an ecological interface was observed to be significantly faster than the time to reach the same level of diagnosis with a contemporary interface. Thus, time to diagnosis (TTD) was included as an outcome measure in this study. Unfortunately, however, we were unable to perform this comparison for the A/U scenario for two reasons. First, only 3 operators in the Current and

P+F interface groups achieved a diagnosis score of 3. Second, of those who did, one of them was a participant for whom the audio recording had been lost and the other two did not provide complete diagnoses until the written questionnaire at the end of the trial. Therefore, we are only able to compare the diagnosis times for the Unanticipated Abnormal scenario. In this case, there were still several missing data points, but enough remaining to draw a comparison. The mean times by group for that scenario are shown in Figure 33. No mean differences are suggested in the data.

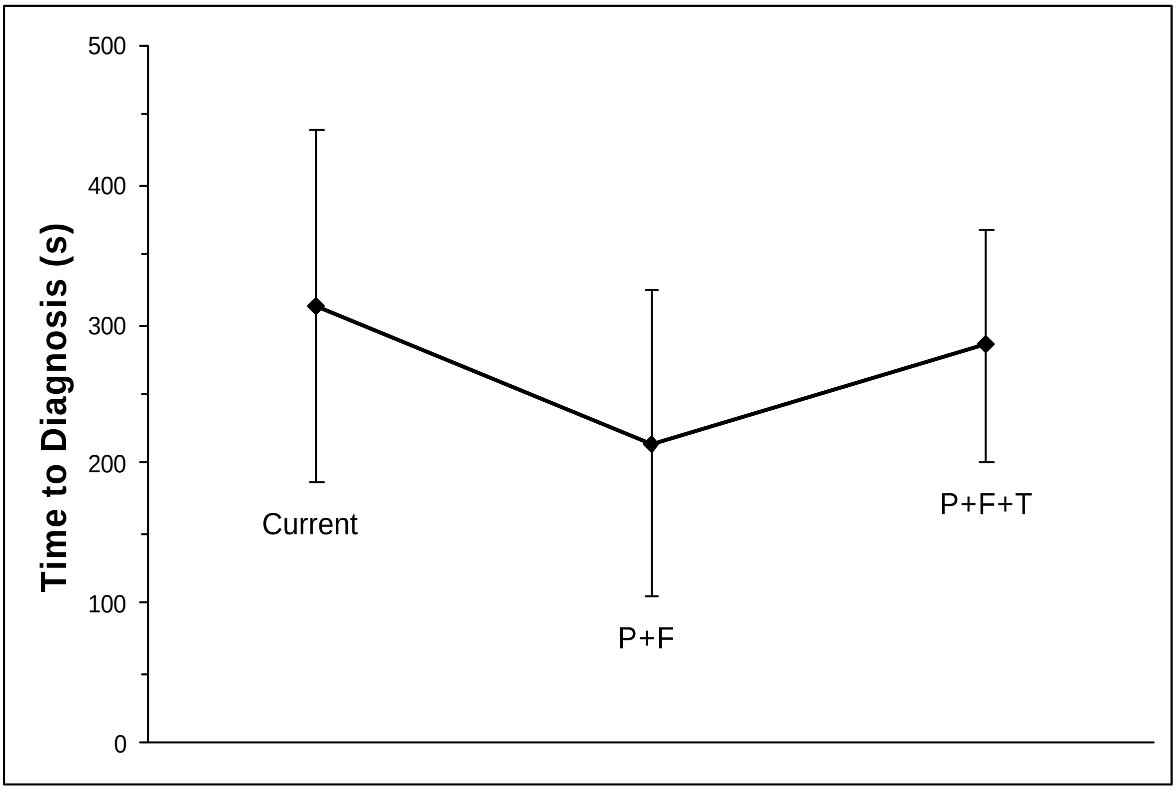

Figure 33: Time to complete diagnosis for Unanticipated Abnormal.

QUESTIONNAIRE RESPONSES

The previous chapter provided the results of various times, counts, and scores. While such objective measures of performance are most useful in assessing our hypotheses, we were also interested in gaining some insight into operator preferences about the interfaces, their perceived effectiveness, and usefulness. In this chapter, we discuss subjective feedback collected from the operators.

Interface Use

Throughout the data collection activity, operators provided a wealth of feedback about the usefulness and usability of the interface views. This feedback took two forms, questionnaire responses and spontaneous comments. First, on the post-test form, operators in the P+F and P+F+T conditions ranked the graphics according to their usefulness. Second, most of the operators offered their insights into the usefulness and usability of specific graphics, views, and whole interfaces. These free form comments came during training, both during and following scenarios, and at the end of the experimental session. Usually these comments were directed to the experimenter, although they were occasionally addressed to the cohort. An account of both of these types of feedback is important, both for understanding the performance data and for continually improving the impact of ecological interfaces in the process industries. Results from the questionnaires are presented below. A summary of the free-form comments is included as Appendix E.

Questionnaire responses

Interface effectiveness. At the conclusion of the final scenario, the participants were asked to respond to a series of questions and to perform a ranking task (see Appendix D for detail). The Post-Test Questionnaire containing these queries was issued to each participant after he completed the Post-Scenario Questionnaire for the third scenario. The experimenter was generally either in the vicinity of the simulator room or the adjacent control room while the Post-Test Questionnaire was completed. The cohort was typically moving between the two rooms during this time to reset the simulator. He had been instructed not to speak to the participants until the experiment was completed and the participant debriefed. For questions 3-6, a four-point Likert type scale was used. The question and response distributions for each of these questions are given in Figure 34.

The four charts in Figure 34 convey three points. First, at least 70% of the participants in each of the Interface conditions indicated that the novel and/or current interfaces were either somewhat or very effective at supporting both procedure execution and non-procedure based activity. For the P+F and P+F+T groups, this number was at least 80%. Second, there are no apparent differences in effectiveness rank between the novel interfaces and the current ones (compare charts top to bottom). Third, operators perceived both the novel and current interfaces to be slightly less effective at supporting non-procedure based activity than they were at supporting procedure execution (compare charts left to right). However, a series of Fisher's Exact Tests showed no significant effect of Interface, unit, total prior operating experience, or prior operating experience in E1 or E2 on effectiveness rating for any of these questions.

Given the overwhelmingly positive participant responses to these questions, it is likely that some degree of confirmation bias is present. The operators were aware that the study involved the evaluation of new displays and it is reasonable to assume that they were prone to respond positively to the queries. It is, however, worth noting that many of the operators showed little hesitancy in criticizing the content and form of the novel displays either directly to, or in the presence of, the experimenter. The extent to which these opinions are also reflected in their survey responses cannot be determined.

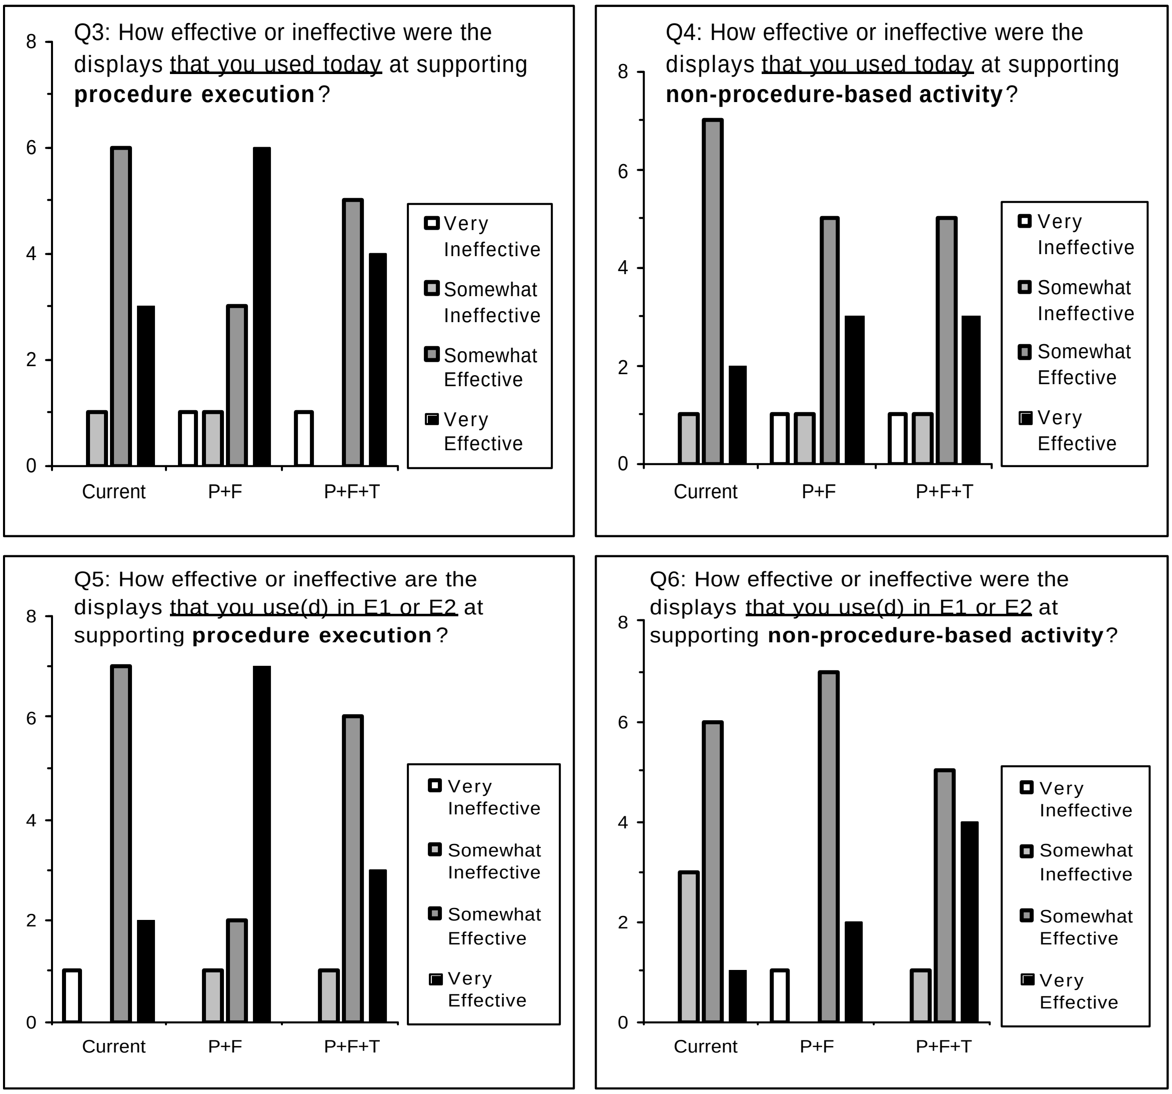

Figure 34: Response distributions for post-test questionnaire items 3, 4, 5 & 6.

Attitude towards change. It is a mantra in the petrochemical industry that operators have a strong aversion to change. If this is true, it could have an impact on the partcipants' acceptance of the novel displays employed in this study. Questions 7 and 8 in the Post-test Questionnaire addressed this concern. The question and response distributions for these two queries are given in Figure 35.

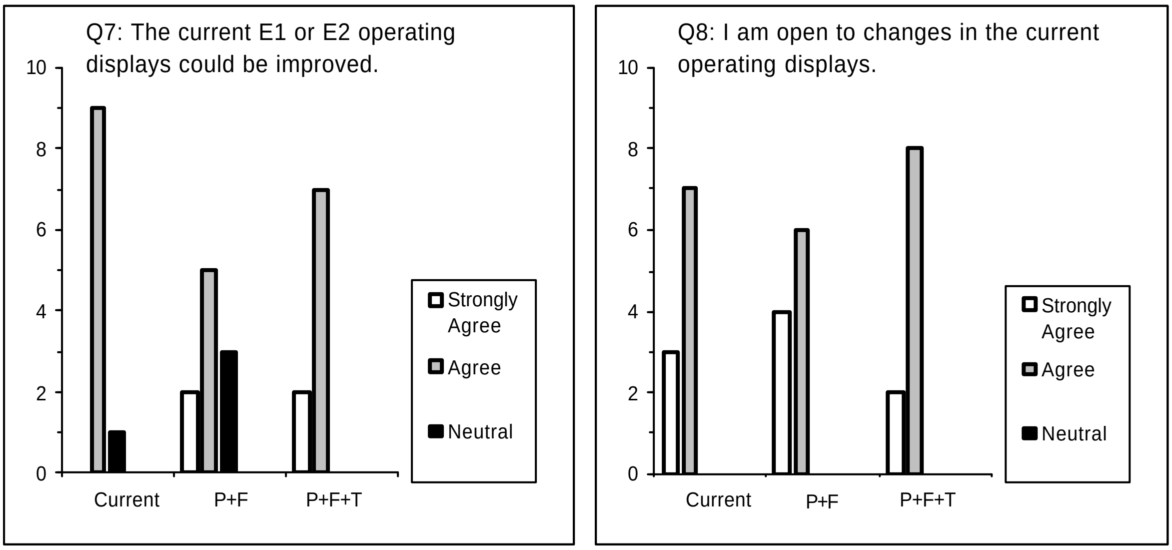

Figure 35: Response distributions for post-test questionnaire items 7 and 8.

The responses to items 7 and 8 mark a strong contrast between the reported attitudes of the participants in this study to the perceived attitudes of the industry. At least 70% of the operators in each of the Interface groups report agreement to some extent that the current operating displays could be improved, and 100% report openness to changes in the displays. Notably, not a single operator expressed disagreement with either of these statements. As with the effectiveness assessments, a series of Fisher's Exact Tests showed no significant effect of Group, unit, total prior operating experience, or prior operating experience in E1 or E2 on self-reported attitudes toward change.

We must again exercise caution in interpreting the responses to these two questions. It is again reasonable to assume that an operator would desire to be seen as someone who is open to change. Whether this attitude would persist if the participants were actually faced with a change in interface cannot be determined with any certainty from this questionnaire. However, that the trend in responses completely counters the industry stereotype of process operators makes the result interesting enough to report.

An interesting distinction not addressed in either of the acceptance questions, but revealed in the operator's ad hoc comments, is in the possible difference in attitude towards novel content and novel form. It was clear from interacting with the operators that their expectation of "new graphics" was new forms. They were expecting to see the same content that they were used to working with, recast in new forms to emphasize important relationships, and they generally responded positively to this idea. However, some of the content of the ecological interfaces was also different, most notably the abstract function information in the Mass and Energy and Heat Exchanger views. The operators were far more critical of this novel content than they were of the novel forms. Presumably, this could derive from a lack of knowledge and experience in putting this new information to use. In the next section, the operator attitudes towards the specific views corroborate this explanation.

View preferences for P+F and P+F+T. Participants in the P+F and P+F+T conditions provided rankings of the process views in order of usefulness. Figure 36 shows a side-by-side comparison of the rankings for each view, with rankings by P+F participants on the left and rankings by P+F+T participants on the right. It should be pointed out that the P+F+T interface contained 3 views that were not included in the P+F interface. Thus, the P+F+T horizontal axes have 12 ranking levels while the P+F axes have only 9. This makes direct comparisons of actual rankings between these two conditions difficult. However, there are some consistent patterns within and between the groups that warrant mention. The views in Figure 36 are ordered from most useful to least useful according to P+F respondents.

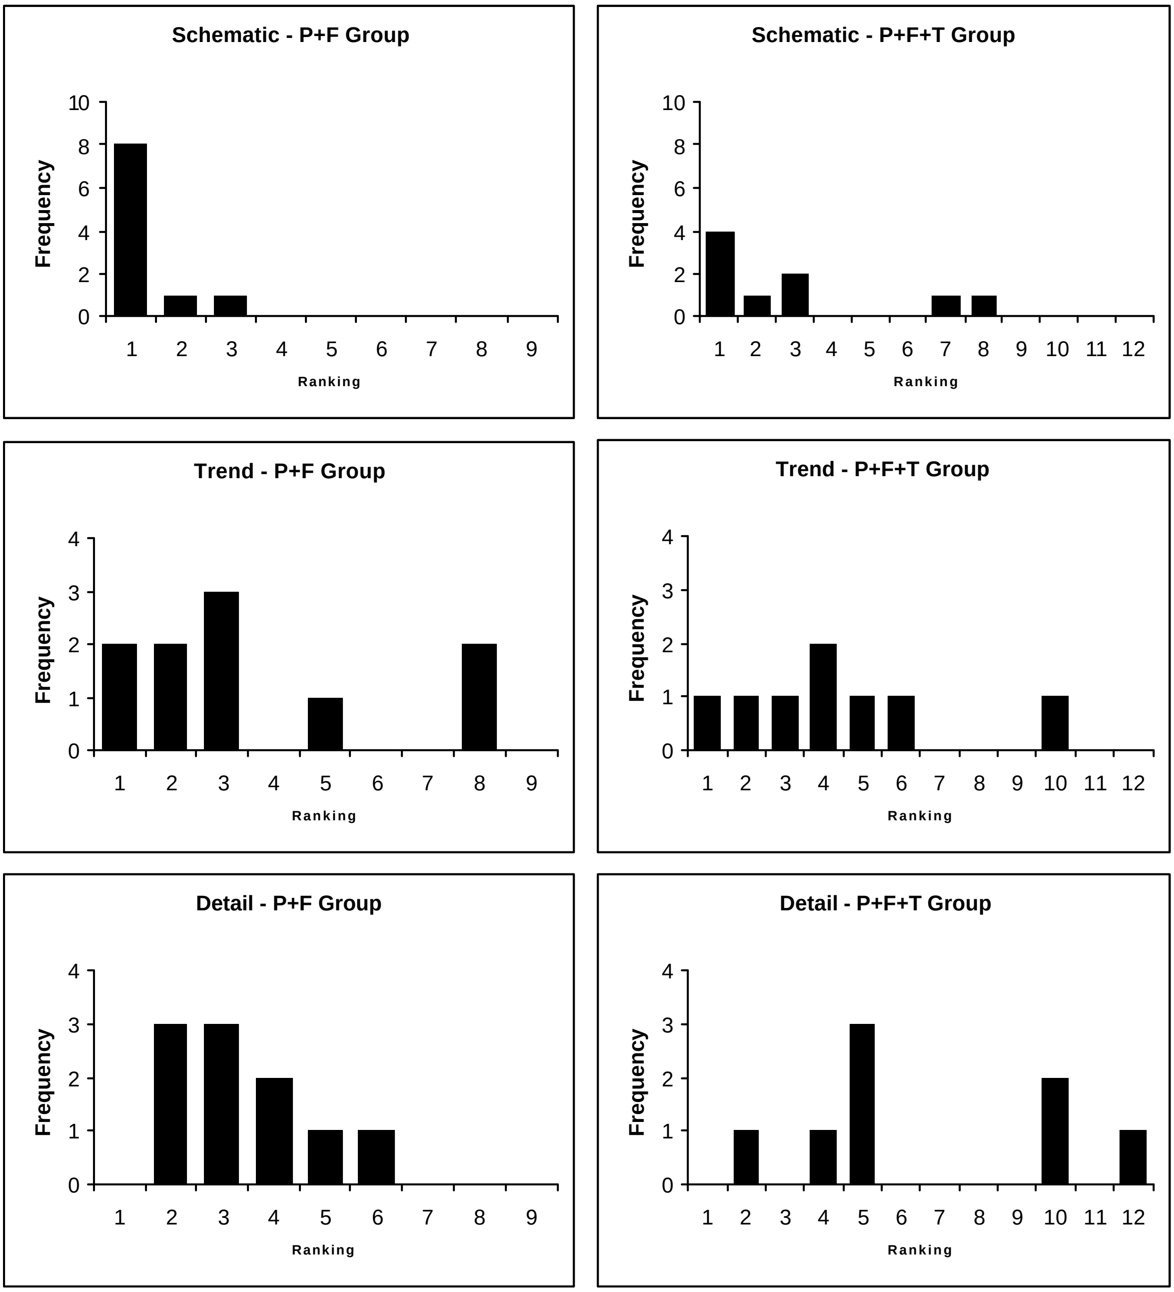

Figure 36: Usefulness rankings for P+F and P+F+T participants.

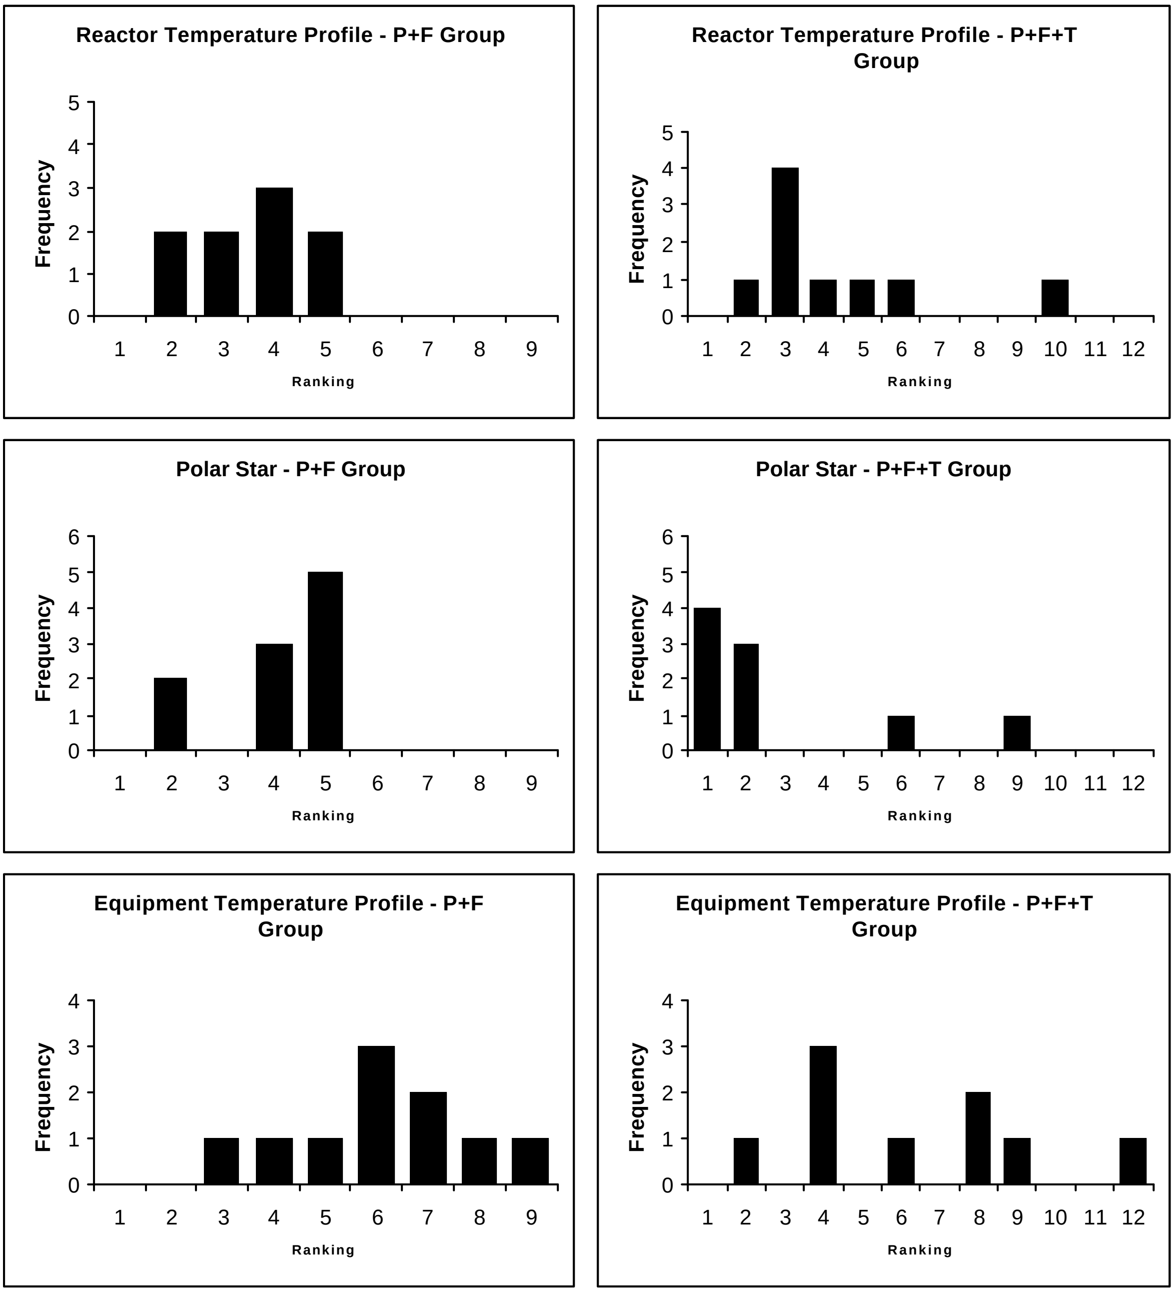

Figure 36 (cont.): Usefulness rankings for P+F and P+F+T participants.

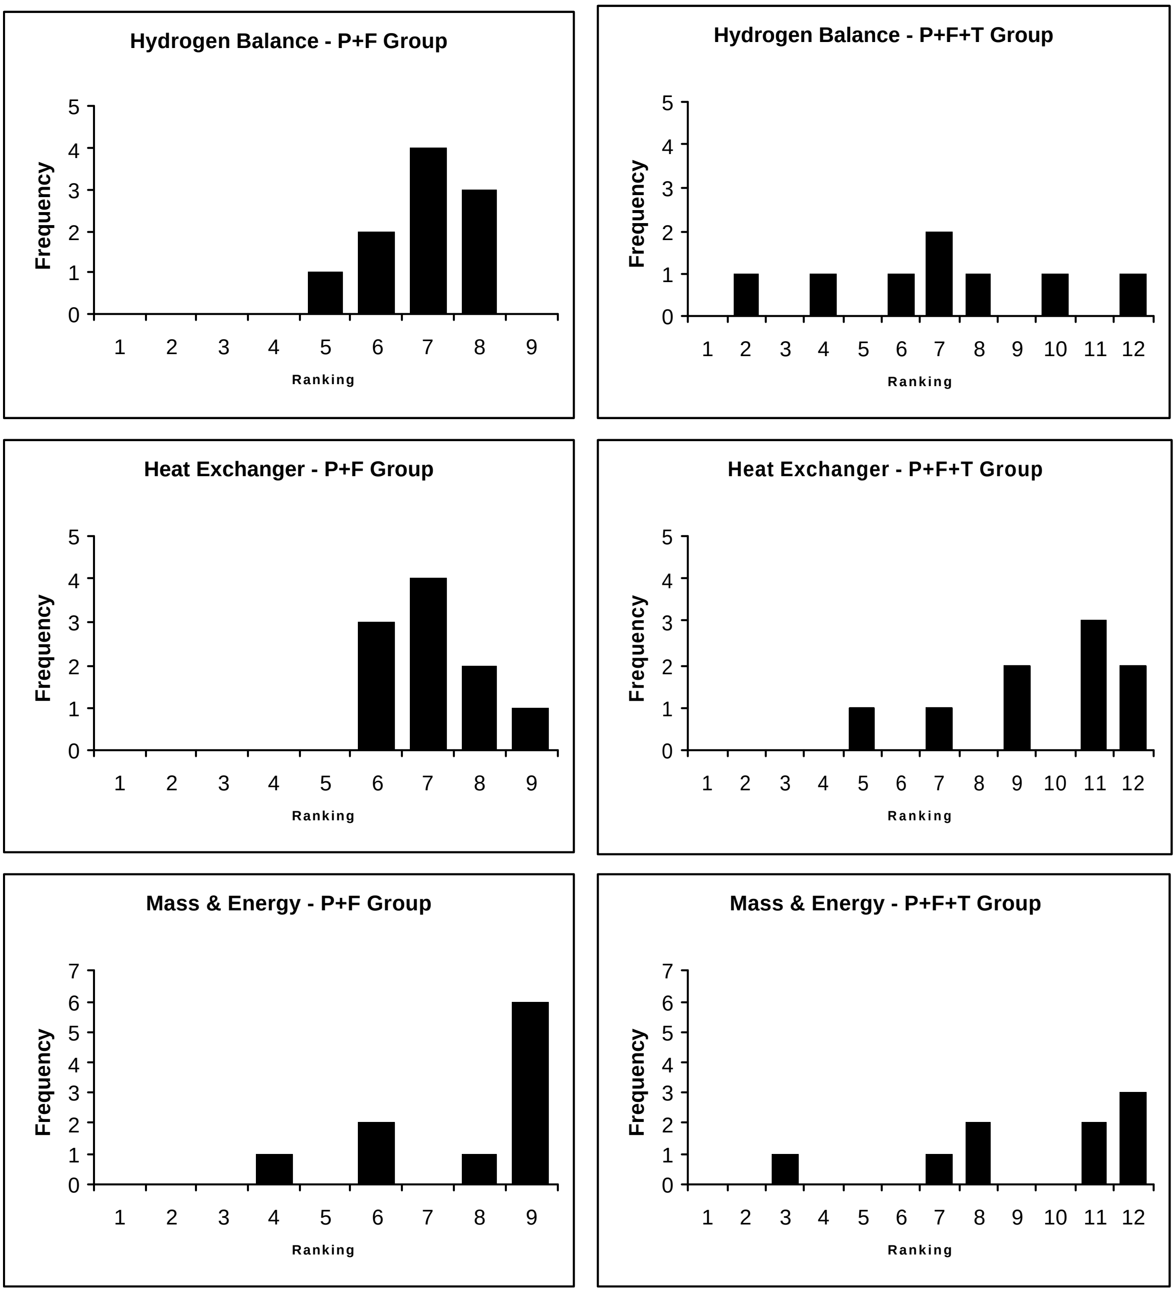

Figure 36 (cont.): Usefulness rankings for P+F and P+F+T participants.

There are several notable patterns. First, the three interface views ranked as being the most useful are those that are most similar to the Current interface. The Schematic, Trend, and Detail views are similar in content to display screens used in the E1 and E2 units. Moreover, they are essentially identical (in both content and form) to those employed in E3. Second, the Heat Exchanger View and Mass and Energy View were consistently reported to be the least useful to the operators. This is particularly interesting because the Heat Exchanger View provides information to assist in the diagnosis of the tube leak while the mass flow surge that causes the temperature runaway is visible in the Mass and Energy View. Third, the views received fairly consistent rankings between the groups, especially for the most and least popular displays. One interesting exception to this is the Polar Star View, which received mediocre rankings from the P+F participants, but fairly high rankings from the P+F+T participants. Fourth, there appears to be more variability in responses in the P+F+T case. Finally, it is worth noting that the Hydrogen Balance View in the P+F+T interface contained a flow control graphic and analyzer update timers that were not available in the P+F version. The rankings for this display are slightly better in the P+F+T condition, suggesting that either the increased information content or the more effective integration of information (cf., Burns, 2000) may have affected the perceived usefulness of the view.

Figure 37 shows the rankings for the three views only available to the P+F+T participants. Two of these views are part of the Procedure View tool. Given that only one participant used this view in any way (and then only for reviewing the reactor swing procedure prior to execution), it is likely that the responses regarding these displays reflect two factors. First, the operators may have been ranking the paper procedures that were made available to them, rather than the Procedure View tool itself. Second, because the operators could simply ignore these views, they did not develop a distaste for them, as with some of the other views. Most of the operators started by ranking the several most useful and one or two least useful views. They then placed the remaining views in the lower to middle range of ranks. This is predominantly where we see all three of the Task-specific views.

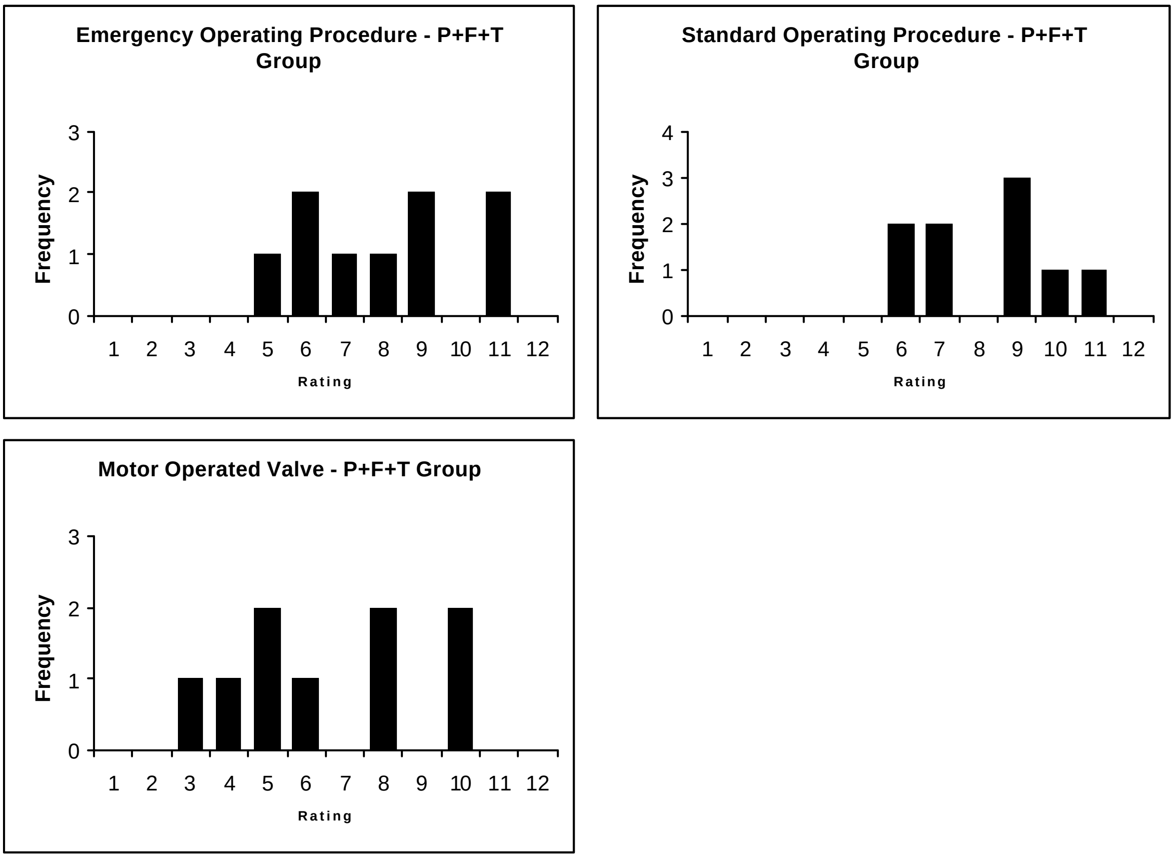

Figure 37: Usefulness rankings for views used only by P+F+T participants.

The MOV view was the most popular of the three views exclusive to the P+F+T display. It was also the only one of those that was constantly visible to the operators. Most of the operators in the P+F+T condition used the MOV view effectively in completing the reactor swing scenario. Both of these factors likely contributed to its higher rankings than either of the Procedure Views.

DISCUSSION

Hypotheses

The discussion of the above results progresses is reverse order of the hypotheses laid out in the Experimental Method Chapter.

H3: An effect of Interface was predicted. The results of this study show several significant performance advantages for the two ecological interfaces (i.e., P+F and P+F+T) over the Current interface in an industrial process simulation with professional operators. These advantages include significant (P+F+T) or nearly significant (P+F) reductions in trial completion times; more successful trial outcomes for the Anticipated Abnormal scenario (P+F+T only); fewer control actions taken across all scenarios (P+F and P+F+T); significant (P+F+T) or nearly significant (P+F) reductions in MOV control actions in the Anticipated Normal scenario; more accurate fault diagnoses (P+F+T only); and somewhat better control performance scores for abnormal scenarios (P+F+T only). The advantages were more pronounced in the P+F+T condition, although there was no statistic on which a direct P+F to P+F+T comparison showed statistical differences. This suggests that the additional task-based information was, at best, moderately beneficial to effective control.

Despite the trend toward improved performance in the ecological conditions for these measures, performance on several other outcome measures was not differentiated by Interface. For example, no significant fault detection time differences were observed, although the trend in the data suggested faster detections for the ecological interfaces. Also, no time to diagnosis benefit was observed for the ecological interfaces in the Unanticipated Abnormal scenario. Most notably, however, the two process measures of acetylene and hydrogen slip were not sensitive to group differences. In retrospect, however, the relatively short exposure times of the scenarios (compared to actual process operations) and the over representation of fault scenarios may not have been sufficient to expose any differences between the Interface groups for these two measures.

As a final comment on H3, it is notable that the Current interface was not significantly better than either of the two ecological interfaces on any of the outcome measures. This is surprising as the participants had, on average, 3.6 years of operational experience with this display. Most of them had more years of experience (an average of 6.6) with displays of similar design. In contrast, participants in the P+F and P+F+T groups were given two hours of training and one hour of practice with the novel interface prior to the experimental trials. It would be reasonable to expect that one or more of the outcome measured would have shown some benefit for the Current interfaces. We have seen no evidence to confirm this expectation.

H2: An effect of Scenario was predicted. With regard to the hypothesis of a main effect of Scenario, the results obtained provide no valid evidence. While several of the measures (e.g., trial completion time, number of control actions, and hydrogen and acetylene slip) were sensitive to scenario type, these differences appear to be best explained by the characteristics of the particular scenarios as opposed to characteristics of the scenario types or the classes of process events that they were intended to exemplify[10]. There is no consistent trend among these results, despite their consistent statistical significance.

It is unlikely, however, that a more comparable set of scenarios could have been developed while meeting the constraints of the scenario taxonomy. This is an inescapable challenge of having used a simulator study, particularly one that so closely aligns itself with a

[10] To draw conclusive statements about this relationship would require a study with several scenarios of each type that could be compared.

combination of known events. Our experience here suggests that this question would be better addressed in a more controlled experimental environment.

H1: An interaction between Interface and Scenario was predicted. The hypothesized interaction between the interface content and scenario type is only weakly suggested by the data. The mean trial completion time for P+F+T in the Anticipated Abnormal scenario is faster than that for the other two displays, there are more complete diagnoses for P+F+T in the Anticipated Abnormal scenario, and all participants in the P+F+T group successfully isolate the hydrogen source in Anticipated Abnormal. All three of these observations are consistent with the expected added benefit of task-based information for anticipated events. However, there is no Interface x Scenario interaction for number of control actions. In addition, the evidence suggests some benefit for task-based information where we would not expect to see it. For example, the slight P+F+T advantage in performance score stems primarily from differences in the Unanticipated Abnormal scenario, where we would not have expected to see a difference compared to P+F.

The relative absence of this expected interaction is surprising considering that two of the procedures used to identify the task-based information requirements were directly pertinent to the Anticipated Normal and Anticipated Abnormal scenarios. Thus, while the interface designs themselves demonstrate that task-based information can be effectively integrated into an ecological interface, the advantages appear to be, at best, only loosely coupled to anticipated events. Rather, it appears that the advantage is more generally derived either from the presence of the additional information, regardless of the event, or the better integration of that data (cf., Burns, 2000).

The hypothesized interaction is also not clear with respect to the benefit of functional information. The expected benefit for the P+F group compared to the Current group in managing Abnormal events is not visible in the trial completion times for the Anticipated Abnormal event, and diagnosis accuracy scores for both fault events are indistinguishable from the Current interface. Moreover, neither ecological interface led to faster trial completion times compared to the Current interface in the Unanticipated Abnormal scenario. These mixed interaction results suggest that added functional information is not always conducive to improved performance on abnormal events.

An unexpected result. One unusual observation made here was that of a performance advantage of ecological displays over the current display in a Normal event condition. Most of the previous studies of EID in laboratory settings have failed to show mean differences between Interfaces under normal operations. In a notable counter-example to this trend, Janzen and Vicente (1998) showed that, under normal operating conditions, attention to abstract function information was positively correlated with faster trial completion times. However, a direct comparison between that study and the one presented here is not possible because neither of the present ecological interfaces were parsed according to the levels of the abstraction hierarchy, as was the interface in the Janzen and Vicente (1998) study.

There are several possible explanations for this observation. For example, it is possible that the increased level of domain complexity in this study, coupled with the use of professional operators, revealed an effect that does not manifest itself in less-complicated laboratory simulations. Alternatively, the particular Anticipated Normal scenario that was used in this study may have been complex enough to require problem solving; a characteristic of abnormal events for which ecological interfaces are designed to provide support. A third possible explanation is a combination of the first two. Perhaps there are few events in highly complex systems that do not require some degree of problem solving (although Ham and Yoon's (2001a, 2001b) results suggest that the benefit of functional information content increases as task difficulty increases). This explanation raises the possibility that operators faced by any event in a complex system may benefit from an ecological interface; and that the distinction between Normal and Abnormal events may be less important than the criteria for what qualifies as an Events as opposed to Routine. Regardless, it is interesting to note that, in the present case, the one view that contains exclusively Abstract Function information (i.e., the Mass and Energy View), received very poor usefulness ratings from the participants in the two ecological interface conditions. This is the very content that previous studies have suggested is a key to the benefit of ecological interfaces (see Ham & Yoon, 2001a, 2001b).

This last observation brings us back to the question of whether the openness to changes in process displays expressed by the participants reflects openness to novel content as well as to novel form. We presently have no way of assessing whether the operators were using the abstract function information in the ecological interfaces. Given that the existing empirical evidence supports the contention that such information is one of the key advantages to ecological interfaces, whether or not they were using the abstract function information could have a substantial impact on the potential value of introducing ecological interfaces to industry settings. If operators are not receptive to the novel content that, in part, distinguishes this design framework, then ecological interfaces may face an uphill battle with respect to operator acceptance of the technology. On the other hand, if performance advantages were to persist despite operators not using the abstract function information, we would have evidence that the observed benefits of EID are not solely the product of this content.

Questionnaire Responses

While the Post-test Questionnaire cannot be considered rigorous, it does help to place the results of this study in some context. The operators showed no bias in ranking the Current interface suite over the novel interface suite in terms of effectiveness. This, combined with the high effectiveness ratings and the performance data, confirms that the EID framework can lead to a demonstrably effective interface environment in which professional operators can act to complete their work activities. Moreover, in this limited sample of a single plant site, operators reported themselves to be open and receptive to learning and using complex configurable graphics to monitor, control and trouble-shoot a realistic industrial process. Taken together, the questionnaire responses confirm that professional process plant operators are reportedly willing and able to employ ecological interfaces.

Limitations

There are several limitations to the conclusions drawn from this study. First, novel interfaces for a complex system are both enabled by and limited by the skills of the designers. In this case, the lead designer (i.e., the author) had a fair amount of experience in designing graphics based on system-based requirements and little experience in designing based on task-based requirements. Although the other major contributor to the design had some experience in task-based design, he had no prior experience in petrochemical systems. The design process was also resource bounded. The resource limitation primarily affected the form of the P+F+T interface in that the designers were unable to fully exploit the task-based information requirements. Previous studies (e.g., Burns, 2000) have shown the importance of effective integration of related information in graphical displays; a standard that may not have been sufficiently met in the P+F+T interface (see page 44). With experienced designers, additional time and implementation funds, a more dynamic task-based support could have been provided.

A second limitation lies in the scenarios selected for the empirical evaluation. Taken together, the selection criteria (as noted above) and the need to match scenarios to the event taxonomy over-constrained the choice of event. No single one of the Scenarios was a perfect match. In the end, the three scenarios that best matched these constraints were chosen. Even then, only one scenario in each event category was possible. Another potentially important factor was that the Anticipated Abnormal event was caused by a dual fault while the Unanticipated Abnormal event was caused by a single fault. Previous studies of EID have shown that dual faults are more difficult to manage than single faults. This effect may have been washed out by the already substantial differences between the faults that prevented us from making comparisons between the Scenarios. Finally, it is also worth nothing that this study is similar to many other studies of fault management in quasi-realistic settings in that it vastly over-represents the frequency of fault occurrence.

A third limitation to the interpretation of the findings of this study is the paucity of training provided to the users of the novel interfaces. Whereas participants in the Current interface condition had years of experience with their information system, participants in the P+F and P+F+T conditions had a few hours. It is unlikely that this difference could be entirely remedied in any study with expert users of an information system.

Finally, although this study created an experimental environment that is more realistic than that afforded by prior studies of EID, the scope remains limited in comparison to the environments in which the participants regularly work. The role of auditory alarms (intentionally removed in this study), collaborative diagnosis and problem solving (actively discouraged in this study), and the influences of different work cultures (held constant here) are all factors that remain unexplored in EID research (Vicente, 2002).

Contributions

Despite these limitations, this study makes several novel and significant contributions to the EID literature. Most importantly, it appears to be the first empirical evaluation of the ecological interface design approach in a simulated industrial setting with professional operators. This is important because it begins to address the question of whether the benefits that have been observed with ecological interfaces in laboratory settings scale up to industry applications. One particular aspect of such applications is the use of operating procedures. The results obtained here confirm that performance benefits for ecological interfaces (as compared to a contemporary interface) persist even when detailed operating instructions are available to operators prior to and during the execution of the events addressed by those procedures. A second characteristic of industry settings is the presence of skilled operators. This study showed no consistent pattern of results to suggest that an operator's age or experience precludes him from, or preferences him towards, benefiting from an ecological interface.

The design component of this dissertation represents its own contribution. The vision for EID includes the expansion of the information content from the traditional system structure and operator competencies, to include task, strategy, and socio-organizational information requirements (see Vicente, 1999). The P+F+T interface represents an important step forward in that vision by incorporating both system- and task-based information. This is the first example of this sort of integration for an application with identified constraints in both of those categories (cf., Rasmussen, Pejtersen, & Goodstein, 1994).

Future Directions

Return to the laboratory. A key shortcoming of this study is that the anticipated interaction between event type and interface content was not observed. Although it is possible that no such effect exists, it is more likely that the additional sources of variability in the more realistic setting that prevent us from seeing small interaction effects (cf., Xu et al. (1999), who observed an interaction between semantic hypertext representations (part-whole and means-ends) and task difficulty in a problem-solving search task). In order to take a closer looks at this question, it may be necessary to return to the laboratory where the experimenter could exert greater control over participant assignment, train the operators to criteria, take advantage of greater flexibility in creating scenarios and procedures that meet the theoretical event categorizations, and increase contact time with the participants.

Procedures. Operating procedures are and will continue to be a key component of the information suite available to professional operators. Solution providers in the industry are developing tools to automate procedures or support them with electronic media in the control room and the field. The potential role for EID in supporting automated or electronic procedures should be investigated to further the relevance and impact of the design approach.

Operator strategies. In light of our success in integrating task-based information into EID, a next meaningful challenge would be to integrate support for various control strategies. This represents an important challenge for the approach because it may be extremely difficult to identify strategy constraints that can provide meaningful guidance to operators across a range of process events.

Conclusions

The objective of this dissertation was to determine whether, and under what conditions, there is an advantage to adding task-based information requirements to system-based ecological interface design for petrochemical process control. An attempt was made to answer this question empirically in the context of a high-fidelity process simulation with professional operators. The data from that investigation suggest that operators using an ecological interface containing task based information showed broader performance benefits compared to the current interface than a traditional ecological interface. However, the traditional ecological interface group also showed performance benefits compared to the current interface for some of the same outcome measures. The study was less successful in articulating whether the event conditions of anticipatedness and normalcy interacted with the information content of the two ecological interfaces

Aside from this main objective, this dissertation makes contributions to three corollary issues. First, the process of integrating task-based information into an ecological interface is an important step in the continued evolution of the EID approach. That the integration was accomplished through the use of information requirements to compare the products of the system-based and task-based analyses suggests that, as additional types of analyses are incorporated into EID, information requirements may continue to serve as a useful organizing tool. Second, implementing ecological interfaces for a representative petrochemical process furthers the state of the art of interface design in this domain. That performance benefits were obtained with these displays confirms that ecological interfaces have a role to play in improving the safety and productivity of industry operations. Third, the study raises the fidelity and scale of the empirical study of EID. A majority of prior EID research has been conducted in microworld settings with students. In this study, professional operators used ecological interfaces implemented in a commercial process instrumentation and control environment to monitor for and respond to realistic process events. These characteristics enhance the generalizability of the study to industry applications.

Our findings represent a conservative estimate of the potential benefits of EID in an industry setting. It is almost certain that these benefits would increase if the level of training and experience on the ecological interfaces were comparable to that on the Current interface. For experienced process operators (with thousands of hours of training and operational experience on the current interface) to have demonstrated better performance with an ecological interface with which they had only a few hours of training and practice is encouraging. It suggests that their control expertise is robust to changes in information presentation and, more importantly, that their expertise can be better supported with ecological interfaces.

REFERENCES

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

Bullemer, P. T., & Nimmo, I. (1994). Understanding and Supporting Abnormal Situation Management in Industrial Process Control Environments: A New Approach to Training.

Burns, C. M. (2000). Putting it all together: Improving Display Integration in Ecological Displays. Human Factors, 42, 226-241.

Buttigieg, M. A. & Sanderson, P. M. (1991). Emergent features in visual display design for two types of failure detection tasks. Human factors, 33, 631-651.

Carswell, C. M., & Wickens, C. D. (1987). Information integration and the object display. Ergonomics, 30, 511-527.

Cochran, T., & Bullemer, P. (1996). Abnormal situation management: Not by new technology alone... Paper presented at the AIChE Plant Safety Symposium, Houston, TX.

Geddes, N. (1989). Understanding human operators' intentions in complex systems. Unpublished Ph.D. dissertation, Georgia Institute of Technology.

Ham, D. H., & Yoon, W. C. (2001a). The effects of presenting functionally abstracted information in fault diagnosis tasks. Reliability Engineering and System Safety, 73, 103-119.

Ham, D. H., & Yoon, W. C. (2001b). Design for information content and layout for process control based on goal-means domain analysis. Cognition, Technology & Work, 3, 205-223.

Hutchins, E. L., Hollan, J. D., & Norman, D. A. (1986). Direct manipulation interfaces. In D. A. Norman & S. W. Draper (Eds.), User centred system design: New perspectives on human-computer interaction, (pp. 87-124). Hillsdale, NJ: Erlbaum.

Itoh, J., Monta, K., Sakuma, A., & Makino, M. (1993). An ecological interface design for intelligent man-machine systems for BWR plants. In Proceedings of the Topical Meeting on Nuclear Plant Instrumentation, Control and man-Machine Interface Technologies (pp.97-103), LaGrange Park, IL: American Nuclear Society.

Jamieson, G. A., & Ho, W. H. (2001). An integrated task- and system-based ecological interface for a petrochemical process (CEL-01-02). Cognitive Engineering Lab: University of Toronto.

Jamieson, G. A., & Miller, C. (2000). Exploring the culture of procedures. In Proceedings of the Fifth International Conference on Human Interaction With Complex Systems (pp.141-145). Los Alamitos, CA: IEEE Computer Society Press.

Jamieson, G. A., Reising, D. V. C., & Hajdukiewicz, J. R. (2001). EID design rationale project: Case study report (CEL_01_03). Toronto, Canada: Cognitive Engineering Lab.

Jamieson, G. A., & Vicente, K. J. (2001). Ecological interface design for petrochemical applications: supporting operator adaptation, continuous learning, and distributed, collaborative work. Computers and Chemical Engineering, 25, 1055-1074.

Janzen, M. E. & Vicente, K. J. (1998). Attention allocation within the abstraction hierarchy. International Journal of Human-Computer Studies, 48, 521-545.

Lind, M. (1994). Modeling goals and functions of complex industrial plants. Applied Artificial Intelligence, 8, 259-283.

Miller, C. A., & Vicente, K. J. (1998a). Integrated abstraction hierarchy and plan-goal graph model for the DURESS II system: A test case for unified system- and task-based modeling and interface design (CEL-98-07). Cognitive Engineering Lab: University of Toronto.

Miller, C. A., & Vicente, K. J. (1998b). Toward an integration of task- and work domain analysis techniques for human-computer interface design. In Proceedings of the Human Factors and Ergonomics Society 42nd Annual Meeting (pp. 336-340). Santa Monica, CA: Human Factors and Ergonomics Society.

Miller, C. A., & Vicente, K. J. (1998c). Comparative analysis of display requirements generated via task-based and work-domain based analyses: A test case using DURESS II (CEL-98-08). Cognitive Engineering Lab: University of Toronto.

Miller, C. A., & Vicente, K. J. (1998d). Abstraction decomposition space analysis for NOVA's E1 acetylene hydrogenation reactor (CEL-98-09). Cognitive Engineering Lab: University of Toronto.

Miller, C. A., & Vicente, K. J. (1999a). Task "versus" work domain analysis techniques: A comparative analysis. In Proceedings of the Human factors and Ergonomics Society 43rd Annual Meeting. Santa Monica, CA: Human Factors and Ergonomics Society.

Miller, C. A., & Vicente, K. J. (1999b). Comparative Analysis of display requirements generated via task-based and work domain-based analyses in a real world domain: NOVA's acetylene hydrogenation reactor (CEL-99-04). Cognitive Engineering Lab: University of Toronto.

Miller, C. A., & Vicente, K. J. (2001). Comparison of display requirements generated via hierarchical task and abstraction-decomposition space analysis techniques. International Journal of Cognitive Ergonomics, 5, 335-356.

NOVA(1999). User interface and subpicture library for Nova Chemicals: Functional requirements specification (rev. 4). Red Deer, Alberta, Canada: Nova Chemicals.

Paulsen, J. L. (1992). A pipe reactor design and the interface system. 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.

Ramsey, F. L., & Schafer, D. W. (1997). The statistical sleuth: A course in methods of data analysis. Belmont, CA: Duxbury Press

Rasmussen, J. (1976). Outlines of a hybrid model of the process plant operator. In T. B. Sheridan & G. Johannsen (Eds.) Monitoring behavior and supervisory control. New York: Plenum.

Rasmussen, J. (1983). Skills, rules, and knowledge: signals, signs, and symbols, and other distinctions in human performance models. IEEE Transactions on Systems, Man, and Cybernetics, SMC-13, 257-266.

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., Pejtersen, A. M., & Goodstein, L. P. (1994). Cognitive systems engineering. New York: Wiley.

Reising, D. V. C. (1999). The impact of instrumentation location and reliability on the performance of operators using an ecological interface for process control. Unpublished Ph. D. dissertation, University of Illinois at Urbana-Champaign.

Shneiderman, B. (1983). Direct manipulation: A step beyond programming languages. IEEE Computing, 16, 57-69.

Sharp, T. D., & Helmicki, A. J. (1998). The application of the ecological interface design approach to neonatal intensive care medicine. In Proceedings of the 42nd Annual

Meeting of the Human Factors and Ergonomics Society (pp. 350-354). Santa Monica, CA: Human Factors and Ergonomics Society.

Siegel, S., & Castellan, N. J. (1988). Nonparametric statistics, 2nd Ed. McGraw-Hill: New York.

Shepherd, A. (1989). Analysis and training in information technology tasks. In Diaper, D. (Ed.). Task analysis for human-computer interaction (pp. 15-55). Ellis Horwood, Ltd.: Chichester.

Vicente, K. J. (1990). Ecological interface design as an evaluation tool. In Proceedings of the Topical Meeting on Advances in Human Factors Research on Man/Computer Interactions: Nuclear and Beyond (pp. 259-265), LaGrange Park, IL: American Nuclear Society.

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 computer based work. Mahwah, NJ: Erlbaum.

Vicente, K. J., & Rasmussen, J. (1992). Ecological interface design: Theoretical foundations. IEEE Transactions on Systems, Man, and Cybernetics, SMC-22, 1-18.

Vicente. K. J. (2002). Ecological interface design: Progress and challenges. Human Factors, 44, 62-78.

Wickens, C. D., & Andre, A. D. (1990). Proximity compatibility and information display: Effects of color, space, and objectness of information integration. Human Factors, 32, 61-77.

Wickens, C. D., & Carswell, C. M. (1995). The proximity compatibility principle: Its psychological foundation and relevance to display design. Human Factors, 37, 473-494.

Wolff, H. S. (1967). Signal received but not understood. In Proceedings of the IEE Professional Group 6 Colloquium on the Interpretation of Biological Signals (pp. 14-16).

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

Xu, W., Dainoff, M. J., & Mark, L. S. (1999). Facilitate complex search tasks in hypertext by externalizing functional properties of a work domain. International Journal of Human-Computer Interaction, 11, 201-229.

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. Santa Monica, CA: Human Factors and Ergonomics Society.

APPENDIX A: CONTROL TASK ANALYSES

The Decision Ladders for each of the controllers (with the exception of the WRC) are shown in this section.

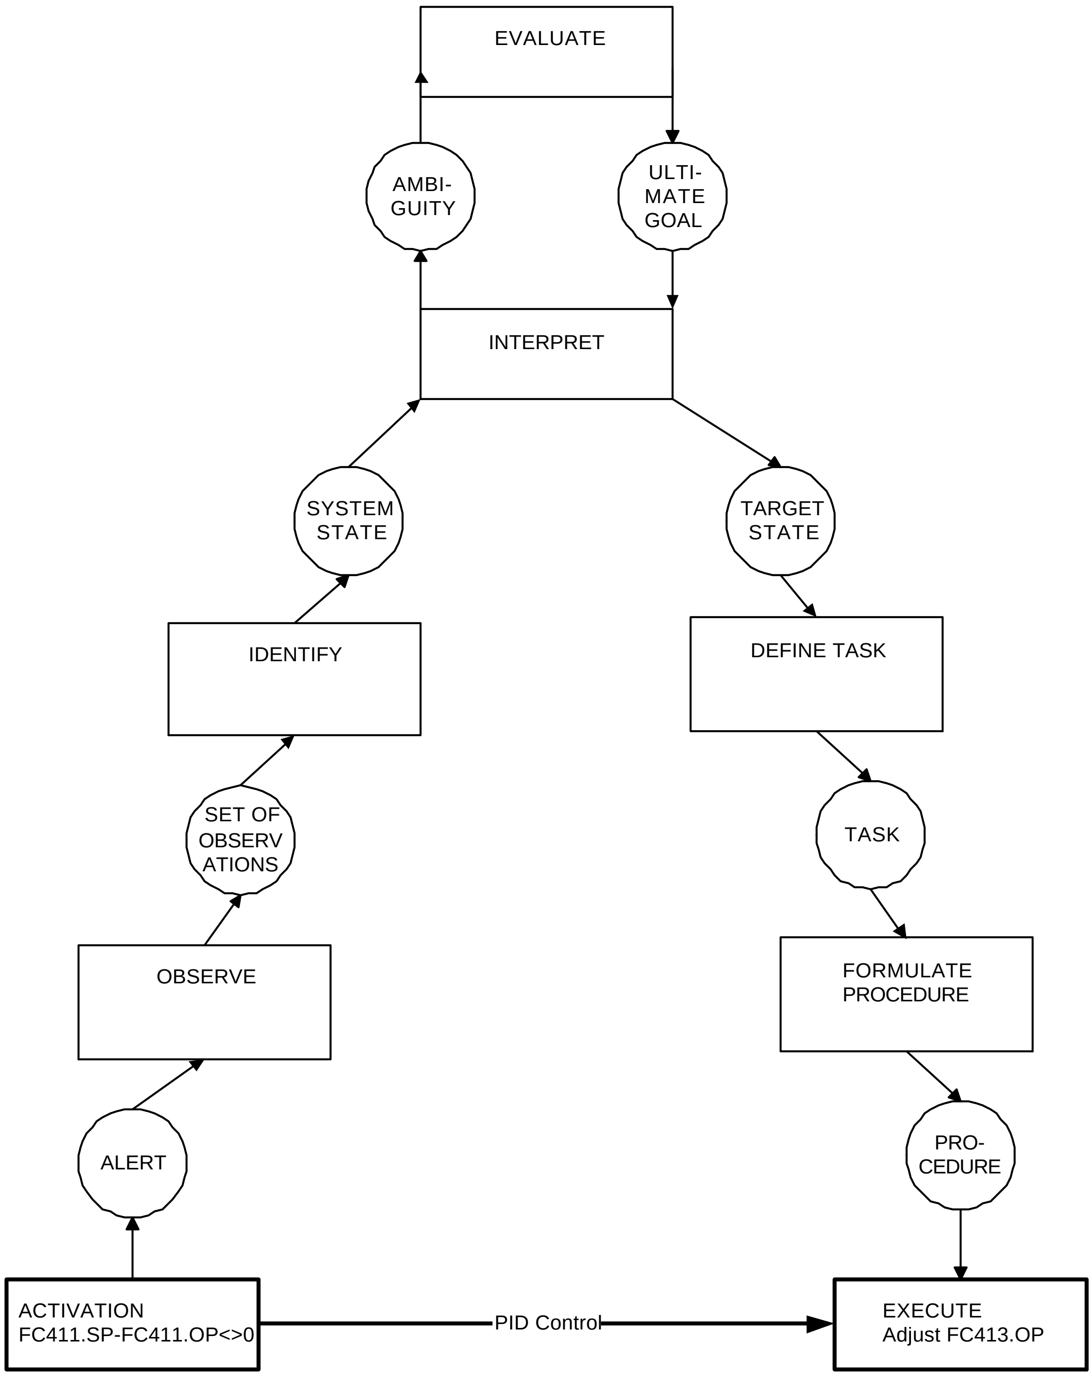

Figure 38: Decision ladder for reactor inlet temperature control.

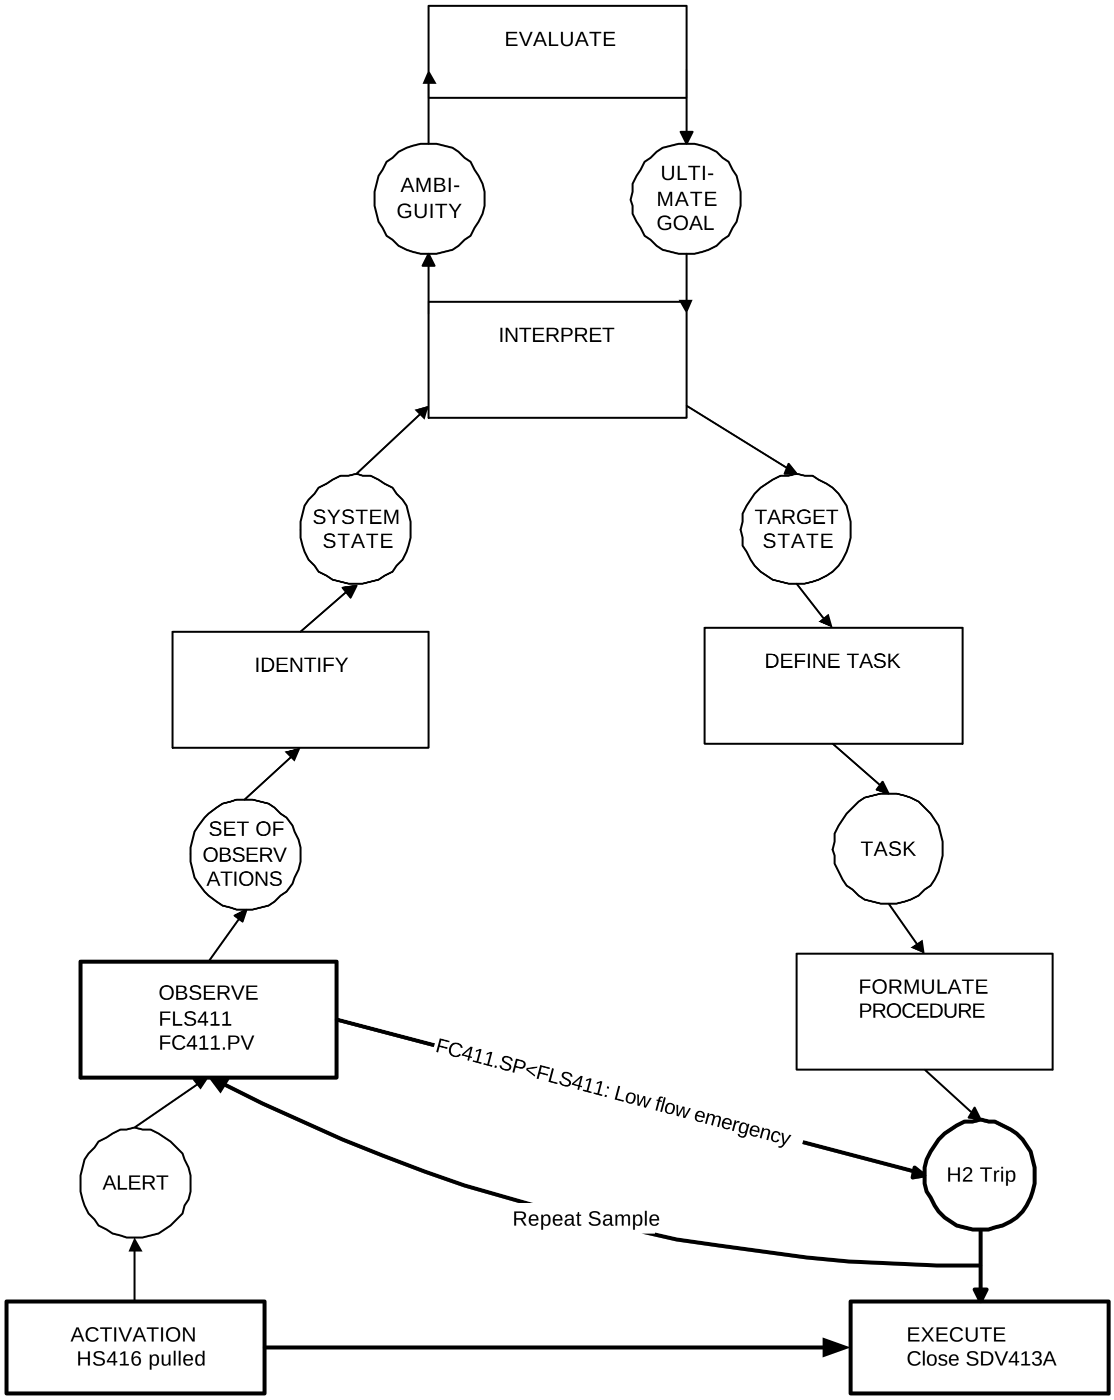

Figure 39: Decision ladder for hydrogen flow trip control.

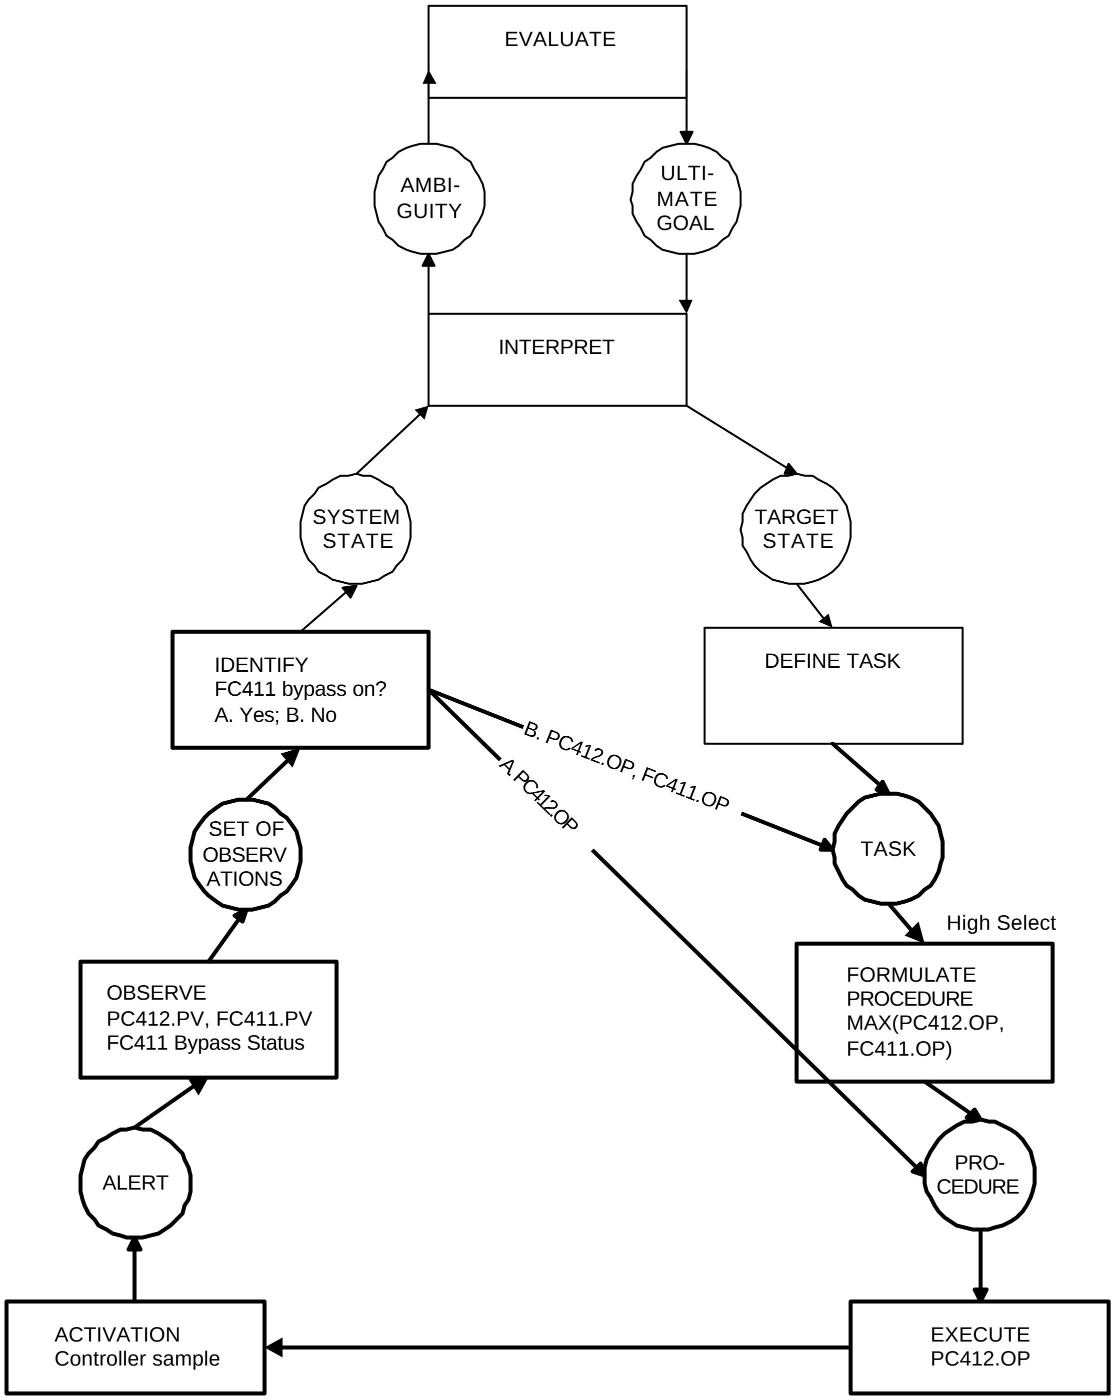

Figure 40: Decision ladder for process to flare control.

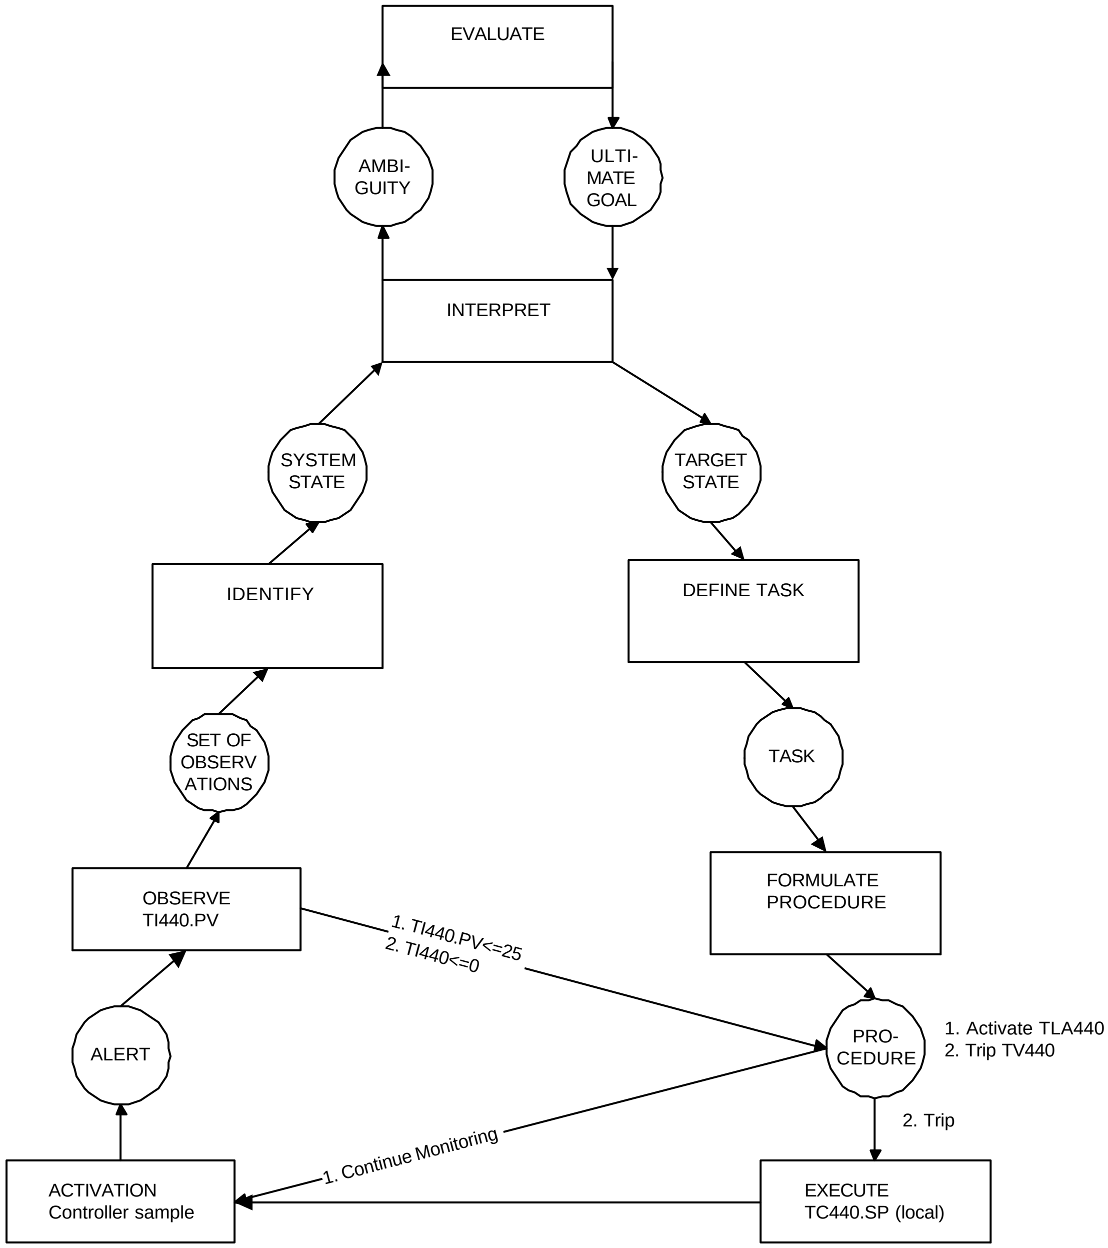

Figure 41: Decision ladder for hydrogen feed temperature control.

APPENDIX B: INFORMATION REQUIREMENTS MATRIX

THIS PAGE INTENTIONALLY LEFT BLANK

IR	ADS	HTA	CTA	E1	DCS?	Sim?	Schematic	Trend	Detail	RxProf.	M&E	Paulsen	H2Bal.	HX	Proc.	MOV
Provide identity, appearance, and location for Physical Form-Component functions																
E410	✔	✔		✔	N/A	N/A	✔				✔	✔				
E411	✔	✔		✔	N/A	N/A	✔				✔	✔			✔	
E412	✔	✔		✔	N/A	N/A	✔				✔	✔				
E413	✔	✔		✔	N/A	N/A	✔				✔	✔				
R410 A	✔	✔	✔	✔	N/A	N/A	✔			✔	✔	✔			✔	✔
R410 B	✔	✔	✔	✔	N/A	N/A	✔			✔	✔	✔			✔	✔
FV135	✔	✔	✔	✔	N/A	N/A	✔						✔			
PV441	✔	✔			N/A	N/A	✔						✔			
TV440	✔	✔	✔		N/A	N/A	✔						✔			
FV413	✔	✔	✔		N/A	N/A	✔						✔		✔	
TV410	✔	✔	✔	✔	N/A	N/A	✔						✔		✔	
SDV413 A	✔	✔	✔		N/A	N/A	✔						✔			
PV410 A	✔	✔			N/A	N/A	✔								✔	
PV410 B	✔	✔			N/A	N/A	✔								✔	
PV412	✔	✔	✔	✔	N/A	N/A	✔						✔		✔	
HV41001	✔	✔	✔	✔	N/A	N/A	✔					✔				
MV410	✔	✔	✔	✔	N/A	N/A	✔			✔		✔			✔	✔
MV411	✔	✔	✔	✔	N/A	N/A	✔			✔		✔			✔	✔
VS1 (rf35503)	✔	✔			N/A	N/A	✔									
VH3 (e2h2vlv)	✔	✔			N/A	N/A	✔									
VM1 (rf35504)	✔	✔			N/A	N/A	✔									
VM4	✔	✔			N/A	N/A	✔								✔	
VM5	✔	✔			N/A	N/A	✔								✔	
VM6 (rf35501)	✔	✔			N/A	N/A	✔								✔	
VM8 (rf35502)	✔	✔			N/A	N/A	✔								✔	

IR	ADS	HTA	CTA	E1	DCS?	Sim?	Schematic	Trend	Detail	RxProf.	M&E	Paulsen	H2Bal.	HX	Proc.	MOV
VM10 (rf35402)	✓	✓			N/A	N/A	✓									
VM13 (rf35404)	✓				N/A	N/A	✓									
VM14 (rf35403)	✓				N/A	N/A	✓									
VW2 (rf35401)	✓	✓			N/A	N/A	✓									
Identify state, and range of possible states, for Physical Function-Component functions																
E410 shell side INLET temperature	✓			✓	✓	✓	✓	✓	✓			✓		✓		
E410 shell side OUTLET temperature	✓			✓	✓	✓	✓	✓	✓			✓		✓		
E410 tube side INLET temperature	✓			✓	✓	✓	✓	✓	✓			✓		✓		
E410 tube side OUTLET temperature	✓			✓	✓	✓	✓	✓	✓			✓		✓		
E411 shell side INLET temperature	✓	✓		✓	✓	✓		✓	✓			✓		✓		
E411 shell side OUTLET temperature	✓	✓				✓		✓	✓					✓		
E411 tube side INLET temperature	✓	✓		✓	✓	✓	✓	✓	✓			✓		✓		
E411 tube side OUTLET temperature	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓	✓	✓		
E412 shell side INLET temperatures	✓					✓		✓	✓					✓		
E412 shell side OUTLET temperature	✓					✓		✓	✓					✓		
E412 tube side INLET temperature	✓			✓	✓	✓	✓	✓	✓	✓		✓		✓		
E412 tube side OUTLET temperature	✓			✓	✓	✓	✓	✓	✓			✓		✓		
E413s shell side INLET temperature	✓			✓	✓	✓		✓	✓					✓		
E413 shell side OUTLET temperature	✓					✓		✓	✓					✓		
E413 tube side INLET temperature	✓			✓	✓	✓		✓	✓					✓		
E413 tube side OUTLET temperature	✓	✓		✓	✓	✓	✓	✓	✓					✓		
R410 A temperature profile	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓				
R410 A days in service	✓	✓		✓	✓	✓	✓	✓	✓							
R410 B temperature profile	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓				
R410 B days in service	✓	✓		✓	✓	✓	✓	✓	✓							
FV135 setting	✓	✓	✓	✓	✓	✓	✓	✓	✓				✓			

IR	ADS	HTA	CTA	E1	DCS?	Sim?	Schematic	Trend	Detail	RxProf.	M&E	Paulsen	H2Bal.	HX	Proc.	MOV
PV441 state	✓	✓	✓			✓		✓	✓				✓			
TV440 setting	✓	✓	✓			✓		✓	✓				✓			
FV413 setting	✓	✓	✓	✓	✓	✓	✓	✓	✓				✓			
TV410 setting	✓	✓	✓	✓	✓	✓	✓	✓	✓				✓	✓		
SDV413A state	✓	✓	✓	✓	✓	✓	✓	✓	✓				✓			
PV410 A state	✓	✓														
PV410 B state	✓	✓														
PV412 setting	✓	✓	✓	✓	✓	✓	✓	✓	✓				✓			
HV41001 setting	✓	✓	✓	✓	✓	✓	✓	✓	✓							
MV410 setting	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓					✓	✓
MV411 setting	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓					✓	✓
VS1 setting	✓	✓				✓		✓	✓							
VH3 setting	✓	✓				✓		✓	✓							
VM1 setting	✓	✓				✓		✓	✓							
VM4 setting	✓	✓				✓		✓	✓							
VM5 setting	✓	✓				✓		✓	✓							
VM6 setting	✓	✓				✓		✓	✓							
VM8 setting	✓	✓				✓		✓	✓							
VM10 setting	✓					✓		✓	✓							
VM13 setting	✓					✓		✓	✓							
VM14 setting	✓					✓		✓	✓							
VW2 setting	✓					✓		✓	✓							
Identify topological connections between Physical Function-Component functions																
DMDS is added to the raw feed stream (at FV135) before upstream of the furnaces	✓				N/A	N/A							✓			

IR	ADS	HTA	CTA	E1	DCS?	Sim?	Schematic	Trend	Detail	RxProf.	M&E	Paulsen	H2Bal.	HX	Proc.	MOV
The C2 feed stream and the H2/CO stream are split after the pyrolysis furnaces, but before reaching the remaining AHR components. The streams enter the AHR separately	✓	✓		✓	N/A	N/A	✓						✓			
Two different sources of H2/CO stream are available to the AHR—the stream from the pyrolysis furnaces and the E2 stream	✓	✓			N/A	N/A	✓									
The E1 H2/CO stream enters E413 directly, with no limiting valve available	✓			✓	N/A	N/A	✓									
The E2 H2/CO stream can be routed into the H2/CO stream for the AHR after (that is, bypassing) E413 by appropriate settings on VH3.	✓				N/A	N/A	✓									
The H2/CO stream can be diverted to flare immediately after E413 by means of PV441.	✓	✓			N/A	N/A	✓						✓			
FV413 is located before the SDV valves and before the point where H2/CO and C2 feed mix.	✓			✓	N/A	N/A	✓						✓			
SDV A is positioned to interrupt the flow of H2/CO to the C2 stream.	✓				N/A	N/A	✓						✓			
The C2 feed stream enters the tube side of E410 directly, with no limiting valve available.	✓			✓	N/A	N/A	✓									
From E410, the C2 stream is connected to shell side of E411, with no limiting valve available.	✓			✓	N/A	N/A	✓					✓				
After E411, the C2 stream intersects the H2/CO.	✓			✓	N/A	N/A	✓					✓	✓		✓	
Utility steam supplies both E411 and E413, but the pathway splits before entering either heat exchanger.	✓			✓	N/A	N/A	✓									
VS1 is positioned before E413 and can affect the flow of steam to it.	✓				N/A	N/A	✓									
TV410 is located before E411and can affect the flow of steam into it.	✓			✓	N/A	N/A	✓								✓	
Steam and condensate flow out of E411.	✓			✓	N/A	N/A	✓							✓		

IR	ADS	HTA	CTA	E1	DCS?	Sim?	Schematic	Trend	Detail	RxProf.	M&E	Paulsen	H2Bal.	HX	Proc.	MOV
Mixed H2/CO and C2 feed can be connected to flows downstream (thereby bypassing) the reactors by means of appropriate settings on VM1 and MV410 or MV411.	✔	✔			N/A	N/A	✔									
The two reactors (R410 A and B) are located side by side, but there is no direct connection between them.	✔			✔	N/A	N/A	✔			✔					✔	
MV410 and VM4 are located on a pipeline entering R410A, while VM6 is located on a pipeline leaving that reactor. R410 B has a similar configuration using MV411, VM5 before and VM8 and after.	✔	✔		✔	N/A	N/A	✔								✔	
Pressure relief valves PV410A (R410A) and PV410B (R410B) provide a path to flare	✔	✔			N/A	N/A	✔								✔	
Reactor outflow lines connect after VM6 and VM8.	✔				N/A	N/A	✔								✔	
A connection exists from after the connection of the effluent flows to PV-412, which regulates flare flow	✔			✔	N/A	N/A	✔						✔		✔	
Connections exist after the reactor to the tube side of E412 (VM10 lies on this connection), or to the pipelines after E412 (VM13 lies on this connection), or to the flare (PV412 lies on this connection).	✔			✔	N/A	N/A	✔					✔				
The tube side of E412 is connected to a pipeline leading to the shell side of E410.	✔			✔	N/A	N/A	✔					✔				
Cooling water enters the shell side of E412 via a pipeline and leaves via a pipeline which contains VW2.	✔	✔			N/A	N/A	✔									
The connection from E412 to E410 intersects a connection bringing reacted feed to or from E2. VM14 lies on this connection.	✔	✔			N/A	N/A	✔									
The connection from E412 enters and exits the shell side of E410 where it leaves the AHR. HV41001 lies on the exiting connection.	✔			✔	N/A	N/A	✔					✔				
Identify values, and range of possible values for the Generalized Function-Component functions																
The flow of DMDS into the pyrolysis furnaces	✔	✔	✔	✔	✔	✔		✔	✔				✔			
The flow of H2/CO entering the AHR [H2 Flow 1]	✔	✔				✔	✔	✔	✔		✔		✔			

IR	ADS	HTA	CTA	E1	DCS?	Sim?	Schematic	Trend	Detail	RxProf.	M&E	Paulsen	H2Bal.	HX	Proc.	MOV
The flo w of steam into E413 [Steam flow 1]	✓	✓				✓		✓	✓					✓		
The flow of heat into the H2/CO stream in E413	✓					✓		✓	✓		✓			✓		
The flow H2/CO to flare [Flare flow 1 at PSV441]	✓	✓				✓		✓	✓							
The flow of heat and H2/CO to or from E2 [E2 H2 flow]	✓	✓				✓										
The flow of heat and H2/CO to the Turbo Expanders	✓					✓										
The flow of heat and H2/CO through SDV413a valve	✓	✓	✓		✓	✓							✓			
The flow of cracked gas entering the AHR at E410	✓	✓				✓		✓	✓				✓			
The flow of heat from the reacted feed stream into the C2 feed stream at E410 [Heat xfer 1]	✓	✓			✓	✓					✓			✓		
The flow of steam into E411 [Steam flow 2]	✓	✓	✓			✓								✓		
The flow of heat into the C2 feed stream at E411 [Heat Xfer 2]	✓	✓			✓	✓					✓			✓		
The flow of heat and mixed feed through the reactor bypass [feed bypass]	✓	✓				✓		✓	✓							
The flow of CO into the reactor [Heat and CO flow 3]	✓	✓		✓	✓	✓		✓	✓				✓			
The flow of H2 into the reactor [Heat and H2 flow 3]	✓	✓		✓	✓	✓		✓	✓				✓			
The flow of C2H2 into the reactor [Heat and C2H2 flow 2]	✓	✓		✓	✓	✓		✓	✓				✓			
The flow of C2H4 into the reactor [Heat and C2H4 flow 2]	✓	✓		✓	✓	✓		✓	✓				✓			
The flow of C2H6 into the reactor [Heat and C2H6 flow 2]	✓	✓		✓	✓	✓		✓	✓							
Flow of heat into the reactor	✓					✓		✓	✓		✓	✓				
Pressure in the reactor	✓	✓	✓	✓	✓		✓			✓						
The rate of C2H2 conversion	✓				✓	✓		✓	✓				✓			
The rate of C2H4 conversion	✓					✓		✓	✓				✓			
The rate (or presence) of C2H4 decomposition	✓					✓							✓			

IR	ADS	HTA	CTA	E1	DCS?	Sim?	Schematic	Trend	Detail	RxProf.	M&E	Paulsen	H2Bal.	HX	Proc.	MOV
The flow of CO out of the reactor	✔	✔			✔	✔		✔	✔				✔			
The flow of H2 out of the reactor.	✔	✔			✔	✔		✔	✔				✔			
The flow of C2H2 out of the reactor	✔	✔		✔	✔	✔		✔	✔				✔			
The flow of C2H4 out of the reactor	✔	✔				✔							✔			
The flow of C2H6 out of the reactor	✔	✔				✔		✔	✔							
Flow of heat out of the reactor	✔				✔	✔		✔	✔		✔					
The flow of reacted feed around E412 and toward E410	✔					✔										
The flow of reacted feed to flare (via PV-412) [Flare Flow 3]	✔	✔	✔	✔	✔	✔		✔	✔				✔			
The flow of reacted feed into E412	✔					✔	✔	✔	✔					✔		
The flow of reacted feed out of E412 toward E410	✔					✔	✔	✔	✔					✔		
The flow of cooling water into and out of E412	✔	✔				✔		✔	✔					✔		
The flow of heat from the reacted feed in E412	✔	✔				✔		✔	✔		✔			✔		
The flow of reacted, cooled feed to or from E2	✔					✔										
The flow of reacted, re-cooled feed (after E410) to the rest of the plant [Reacted cooled feed outflow 2]	✔	✔	✔			✔		✔	✔				✔	✔		
Identify connections between Physical and Generalized Function-Component functions																
The effect of valve setting on flow through the valve for all valves.	✔	✔	✔	✔	✔	✔	✔	✔	✔				✔			
The effect of the state of PV441 on flow of H2/CO through TV440.	✔					✔	✔						✔			
The effect of the setting of TV440 on flow of H2/CO through CV6 and downstream, and on flow through VH3.	✔	✔	✔			✔	✔						✔			
There is no physical function connection to H2/CO flow to E319	✔				N/A	N/A	✔									
The effect of the setting of SDV413a on the flow of H2 to the C2 stream.	✔		✔	✔	✔	✔	✔						✔			

IR	ADS	HTA	CTA	E1	DCS?	Sim?	Schematic	Trend	Detail	RxProf.	M&E	Paulsen	H2Bal.	HX	Proc.	MOV
The fact that the flow (and temperature) of the steam into E411 alone affects the flow of heat into the C2 feed stream in that heat exchanger.	✓			✓	✓	✓	✓					✓		✓		
The effect of the setting of VM1 on flow of mixed feed through MV410 and MV411.	✓	✓				✓	✓									
The effect of the state of pressure relief to flare valves (PV410A and PV410B), secondary reactor input valves (VM4 and VM5) and reactor outlet valves (VM6, VM8) have on flow through the primary reactor input valves: MV410 and MV411.	✓					✓	✓								✓	
The effect of the setting of the reactor flare valves (PV410A and PV410B) on flow through VM4 and VM5 into the reactors.	✓					✓	✓								✓	
The effect of the settings of the reactor outlet valves (VM6 and VM8) have on flow through the reactors.	✓	✓				✓	✓								✓	
The effect of the setting of PV412 on flow of reacted feed through VM10 and VM13.	✓				✓	✓	✓								✓	
The effect of the setting of VM13 on flow of reacted feed through VM10.	✓					✓	✓									
The effect of the setting of VM10 on flow through VM13 and through E412.	✓					✓	✓									
The effect of the setting of VW2 on flow of cooling water through E412.	✓	✓				✓	✓									
The effect of the setting of VM14 on flow of cooled, reacted feed through VM13 and into E410.	✓					✓	✓									
The effect of the setting of HV41001 on flow of reacted, cooled feed through E410 and, therefore, on the heat transfer from the reacted feed to the unreacted C2 feed stream.	✓			✓		✓	✓								✓	
Identify topological connections between Generalized Function-Component functions																
The effect of the rate of DMDS flow on CO concentration	✓	✓			✓	✓							✓			

IR	ADS	HTA	CTA	E1	DCS?	Sim?	Schematic	Trend	Detail	RxProf.	M&E	Paulsen	H2Bal.	HX	Proc.	MOV
Diverting H2/CO flow to or from E2 means affecting the heat input to either Heat Transfer 3 or downstream to feed mixing.	✔	✔				✔										
Diverting H2/CO flow to flare means affecting the heat and H2 and CO flows downstream.	✔					✔	✔				✔	✔				
Diverting H2/CO flow to the Turbo Expanders means affecting the heat and H2 and CO flows downstream.	✔					✔	✔					✔				
H2/CO Flow lockout means affecting the heat and H2 and CO flows downstream.	✔				✔	✔	✔					✔	✔			
The rate of heat transferred in Heat Transfer 1 can affect the rate transferred in Heat Transfer 2 (and vice versa).	✔			✔	✔	✔					✔	✔				
The effect of routing mixed feed to the Reactor Bypass on the heat, component materials and pressure in the reactors (it will reduce or eliminate them).	✔					✔	✔			✔	✔	✔	✔			
The effect of routing mixed feed to the flare on the heat, component materials and pressure in the reactors (it will reduce or eliminate them).	✔	✔		✔	✔	✔	✔			✔	✔	✔	✔			
The effect of acetylene conversion on acetylene flow out of the reactor. Same for CO and CO conversion.	✔			✔		✔							✔			
The effect of C2H4 decomposition on the CH4 flow (it is one of two ways to produce it), C2 production (it is the only way to produce it) and heat flow out of the reactor (it is extremely exothermic).	✔					✔							✔			
The effect of the availability of H2 on C2H2 conversion and C2H4 conversion—it is required for both and limits them if not available.	✔			✔		✔							✔			
CO presence affects the availability of H2, and therefore affects the incidence or rate of C2H2 and C2H4 conversion.	✔			✔		✔							✔			
Ethane does not participate in any reaction, thus any increase in ethane across the reactor must be the result of ethylene conversion, and any loss of ethane across the reactor is indicative of a leak or unexpected reaction.	✔					✔							✔			
The effect of all reactions on heat flow out of the reactor: all are exothermic.	✔			✔	✔	✔				✔	✔	✔	✔			

IR	ADS	HTA	CTA	E1	DCS?	Sim?	Schematic	Trend	Detail	RxProf.	M&E	Paulsen	H2Bal.	HX	Proc.	MOV
The effect of flare flow after the reactor (at Flare flow 3) on reactor pressure, heat and reactant availability. It also reduces the availability of heat and reacted feed downstream.	✔				✔	✔	✔				✔					
The effect of bypassing E412 on the heat flow to E2 or to Heat Transfer 1.	✔					✔	✔				✔	✔				
The effect of heat and feed flow to Heat Transfer 1 on the rate of heat transfer to the incoming C2 Feed.	✔	✔		✔	✔	✔					✔	✔				
Identify Physical Function-Component:Subsystem part-whole relationships																
H2/CO Heat and Supply Input Unit: FV135, VH3, E413, VS1, TV440, CV6, FV413, SDV413 A, PV441	✔				N/A	N/A	✔				✔	✔	✔			
C2 Heat and Supply Input Unit: VM1, MV410, MV411, VM4, VM5, E410, E411, TV410	✔				N/A	N/A	✔				✔	✔				
Reaction Unit: R410A, R410B, PV410A, PV410B, VM6, VM8	✔			✔	N/A	N/A	✔			✔	✔	✔				
Effluent Cooling Unit: E412, E410, VM10, VW2, VM14, VM13, HV41001, PV412	✔				N/A	N/A	✔				✔	✔				
Identify Generalized Function-Component: Subsystem part-whole relationships																
H2/CO Heat and Supply: DMDS Flow, Raw Feed Flow, CO Flow 1, E2 H2 Flow, E2 CO Flow, Steam Flow 1, Heat Transfer 3, Heat and H2 Flow 1, Heat and CO Flow 1, TE Flow, Heat and H2 Flow 2, Heat and CO Flow 2, Flow Lockout, Flare Flow 1	✔				N/A	N/A					✔	✔				
C2 Heat and Supply: C2H2 Flow 1, C2H4 Flow 1, C2H6 Flow 1, Heat Transfer 1, Steam Flow 2, Heat Transfer 2, Feed Bypass, Heat and C2H2 Flow 2, Heat and C2H4 Flow 2, Heat and C2H6 Flow 2, Heat and H2 Flow 3, Heat and CO Flow 3, Flare Flow 2	✔				N/A	N/A					✔	✔				

IR	ADS	HTA	CTA	E1	DCS?	Sim?	Schematic	Trend	Detail	RxProf.	M&E	Paulsen	H2Bal.	HX	Proc.	MOV
Reaction: Reactor Pressure, C2H2 Conversion, C2H4 Conversion, C2H4 Decomposition, Heat and C2H2 Flow 3, Heat and C2H4 Flow 3, Heat and C2H6 Flow 3, Heat and CH4 Flow 1, Heat and H2 Flow 4, Heat and CO Flow 4, Heat and C2 Flow 1, Heat and C2H2 Flow 4, Heat and C2H4 Flow 4, Heat and C2H6 Flow 4, Heat and CH4 Flow 2, Heat and H2 Flow 5, Heat and CO Flow 5, Heat and C2 Flow 2,	✓				N/A	N/A					✓	✓	✓			
Effluent Cooling: Cooling Water Flow, Heat Transfer 5, After Cooler Bypass, Reacted Feed Outflow, Reacted E2 Feed Flow, Heat Transfer 1, Flare Flow 3	✓				N/A	N/A					✓	✓				
Identify Generalized Function-Subsystem functions																
H2/CO Heat and Supply Input	✓				✓	✓					✓		✓			
C2 Heat and Supply Input	✓				✓	✓					✓		✓			
Reaction	✓				✓	✓				✓			✓			
Product Cooling	✓				✓	✓					✓					
Identify topological connections between Generalized Function-Subsystem functions																
H2/CO and C2 Supply Input effect Reaction	✓				N/A	N/A							✓			
Reaction effects Effluent Cooling	✓				N/A	N/A				✓						
Effluent Cooling effects C2 Heat and Supply Input	✓				N/A	N/A								✓		
Identify values, and range of possible values, for the Abstract Function-Component functions																
Mass Source I (DMDS)	✓				✓	✓					✓		✓			
Mass Transport II (H2/CO stream)	✓					✓					✓		✓			
Mass Barrier 413	✓					✓								✓		
Mass Source II (Steam)	✓					✓								✓		
Mass Sink I	✓					✓								✓		
Mass Source III (E2 H2/CO)	✓					✓	✓									
Mass Sink II (E2)	✓					✓	✓									
Mass Transport III (VH3)	✓					✓	✓									

IR	ADS	HTA	CTA	E1	DCS?	Sim?	Schematic	Trend	Detail	RxProf.	M&E	Paulsen	H2Bal.	HX	Proc.	MOV
Mass Transport V (TV440)	✔					✔					✔					
Mass Barrier II (PV441)	✔					✔	✔									
Mass Barrier III (CV6)	✔					✔	✔									
Mass Sink III (E319)	✔					✔	✔									
Mass Transport VII (FV4 13)	✔				✔	✔		✔	✔		✔		✔			
Mass Transport VIII (MOV)	✔					✔		✔	✔		✔					
Mass Sink IV (Flare)	✔				✔	✔		✔	✔		✔					
Mass Balance III (reaction)	✔	✔		✔	✔	✔							✔			
Mass Barrier V (PV410A, PV410B)	✔					✔	✔									
Mass Transport IX (VM6, VM8)	✔					✔	✔									
Mass Barrier 412	✔					✔								✔		
Mass Source IV (CW)	✔					✔								✔		
Mass Sink IV	✔					✔	✔									
Mass Transport X (VM13)	✔					✔	✔									
Mass Barrier VIII (PV412)	✔				✔	✔	✔	✔	✔		✔		✔			
Mass Barrier 410/411	✔					✔								✔		
Mass Transport XI (C2 stream)	✔				✔	✔		✔	✔		✔		✔			
Mass Source V (Steam)	✔					✔								✔		
Mass Transport XII (HV41001)	✔					✔		✔	✔		✔					
Mass Sink V	✔					✔								✔		
Mass Sink VI (Product out)	✔					✔		✔	✔		✔					
Mass Source VI (E2 Product)	✔					✔										
Mass Sink VII (E2 Product)	✔					✔										
Mass Transport XIII (VM1)	✔					✔										
Energy Source I (DMDS)	✔					✔										
Energy Transport II (H2/CO stream)	✔					✔					✔					

IR	ADS	HTA	CTA	E1	DCS?	Sim?	Schematic	Trend	Detail	RxProf.	M&E	Paulsen	H2Bal.	HX	Proc.	MOV
Energy Balance 413	✓					✓					✓			✓		
Energy Source II (Steam)	✓					✓								✓		
Energy Sink I	✓					✓								✓		
Energy Source III (E2 H2/CO)	✓					✓										
Energy Sink II (E2)	✓					✓										
Energy Transport III (VH3)	✓					✓										
Energy Transport V (TV440)	✓					✓		✓	✓		✓					
Energy Barrier II (PV441)	✓					✓		✓	✓		✓					
Energy Barrier III (CV6)	✓					✓										
Energy Sink III (E319)	✓					✓										
Energy Transport VII (FV413)	✓				✓	✓		✓	✓		✓					
Energy Transport VIII (MOV)	✓					✓		✓	✓		✓					
Energy Sink IV (Flare)	✓				✓	✓		✓	✓		✓					
Energy Balance III (reaction)	✓	✓		✓	✓	✓	✓	✓	✓		✓					
Energy Barrier V (PV410A, PV410 B)	✓					✓										
Energy Transport IX (VM6, VM8)	✓					✓										
Energy Balance 412	✓					✓		✓	✓		✓			✓		
Energy Source IV (CW)	✓					✓								✓		
Energy Sink IV	✓					✓								✓		
Energy Transport X (VM13)	✓					✓										
Energy Barrier VIII (PV412)	✓					✓										
Energy Balance 410/411	✓				✓	✓		✓	✓		✓			✓		
Energy Transport XI (C2 stream)	✓					✓		✓	✓		✓					
Energy Source V (steam)	✓					✓								✓		
Energy Transport XII (HV41001)	✓					✓		✓	✓		✓					
Energy Sink V	✓					✓								✓		

IR	ADS	HTA	CTA	E1	DCS?	Sim?	Schematic	Trend	Detail	RxProf.	M&E	Paulsen	H2Bal.	HX	Proc.	MOV
Energy Sink VI	✔					✔		✔	✔		✔					
Energy Source VI	✔					✔										
Energy Sink VII	✔					✔										
Identify Abstract Function-Subsystem functions																
Mass & Energy Source 1	✔					✔					✔					
Mass & Energy Source 2	✔					✔					✔					
Mass & Energy Balance	✔					✔					✔					
Mass & Energy Sink	✔					✔					✔					
Identify connections between General and Abstract Function-Subsystem functions																
Mass & Energy Source 1: H2/CO Heat & Supply	✔					✔					✔					
Mass & Energy Source 2: C2 Heat & Supply	✔					✔					✔					
Mass & Energy Balance: Reaction	✔					✔					✔					
Mass & Energy Sink: Cooling	✔					✔					✔					
Identify topological connections between Abstract Function-Subsystem functions																
Mass & Energy Sources 1 & 2 both affect Mass & Energy Balance 1	✔				N/A	N/A					✔					
Mass & Energy Balance 1 affects Mass & Energy Sink 1	✔				N/A	N/A					✔					
Mass & Energy Sink 1 affects Mass and Energy Source 2	✔				N/A	N/A					✔					
Identify value, and range of possible values, for Functional Purpose function																
Target and current value of C2H2 content upon leaving the AHR system	✔	✔		✔	✔	✔	✔	✔	✔				✔			
Identify Connections between Abstract Function and Functional Purpose functions																
That changes in mass into and out of the reaction balance can affect the functional purpose.	✔				✔	✔							✔			

IR	ADS	HTA	CTA	E1	DCS?	Sim?	Schematic	Trend	Detail	RxProf.	M&E	Paulsen	H2Bal.	HX	Proc.	MOV
The fact that use of a mass sink (specifically, the flare can prevent bad product from flowing out of the AHR).	✔				✔	✔	✔						✔			
Control task requirements not included in ADS																
FC135.SP		✔	✔	✔	✔	✔		✔	✔				✔			
FC135.OP		✔	✔	✔	✔	✔		✔	✔				✔			
Error (FC135.SP - FC135.PV)			✔		✔	✔			✔				✔			
FC-135 control mode (manual/auto/cascade/backup)		✔	✔	✔	✔	✔		✔	✔							
TC-410 SP		✔	✔	✔	✔	✔		✔	✔				✔			
TC-410 OP		✔	✔	✔	✔	✔		✔	✔				✔			
Error (TC-410 SP - TI-410 PV)			✔		✔	✔			✔				✔			
TC-410 control mode (manual/auto/cascade/backup)		✔	✔	✔	✔	✔		✔	✔							
TIC440.SP		✔	✔			✔						✔				
TIC440.OP			✔			✔							✔			
Status of TI-440 PV versus 0 C and 25 C			✔			✔						✔				
TC-440 control mode (manual/auto/cascade/backup)			✔			✔		✔	✔							
H2/C2 weight ratio SP		✔	✔	✔	✔	✔		✔	✔				✔			
H2/C2 weight ratio PV		✔	✔	✔	✔	✔		✔	✔				✔			
H2/C2 weight ratio error (SP-PV)			✔		✔	✔			✔				✔			
H2/C2 weight ratio OP (to FC413 SP)			✔	✔	✔	✔		✔	✔				✔			
WRC411 control mode (manual/auto/cascade/backup)		✔	✔	✔	✔	✔		✔	✔							
FC413.SP			✔	✔	✔	✔		✔	✔				✔			
FC413.OP			✔	✔	✔	✔		✔	✔				✔			

IR	ADS	HTA	CTA	E1	DCS?	Sim?	Schematic	Trend	Detail	RxProf.	M&E	Paulsen	H2Bal.	HX	Proc.	MOV
Error (FC413.SP-FC413.OP)			✔		✔	✔			✔				✔			
FC-413 control mode (manual/auto/cascade/backup)		✔	✔	✔	✔	✔		✔	✔							
HS-416 A status		✔	✔	✔	✔	✔										
FLS-411 SP			✔	✔	✔	✔			✔							
EV-413A SP			✔	✔	✔	✔			✔							
ZSC-413 A status			✔	✔	✔	✔							✔			
C2 low flow manual bypass status		✔	✔	✔	✔	✔			✔							
C2 low flow SP (FC-411)			✔	✔	✔	✔			✔							
PC-412 OP			✔	✔	✔	✔		✔	✔				✔			
PC-412 SP		✔	✔	✔	✔	✔		✔	✔				✔			
PC-412 PV		✔	✔	✔	✔	✔		✔	✔				✔			
FC-411 OP			✔	✔	✔	✔		✔	✔				✔			
FC-411 SP			✔	✔	✔	✔		✔	✔				✔			
MAX (PC-412.OP, FC-412.OP)			✔	✔	✔	✔		✔	✔				✔			
PC-412 control mode (manual/auto/cascade/backup)		✔	✔	✔	✔	✔		✔	✔							
HV-41001 OP		✔	✔			✔	✔	✔	✔							
HV-41001 control mode (manual/auto/cascade/backup)		✔	✔	✔	✔	✔		✔	✔							
MS410.OP		✔	✔			✔			✔	✔					✔	✔
MS411.OP		✔	✔			✔			✔	✔					✔	✔
MV-410 control mode (manual/auto/cascade/backup)			✔	✔	✔	✔		✔	✔							
MV-411 control mode (manual/auto/cascade/backup)			✔	✔	✔	✔		✔	✔							

IR	ADS	HTA	CTA	E1	DCS?	Sim?	Schematic	Trend	Detail	RxProf.	M&E	Paulsen	H2Bal.	HX	Proc.	MOV
Unique HTA items																
Procedure steps		✓			N/A	N/A									✓	
Sequence constraints on procedure steps		✓			N/A	N/A									✓	
Overview of procedure steps with notation of current step		✓			N/A	N/A									✓	
Panel/Outside Operator roles and responsibilities		✓			N/A	N/A									✓	
Acetylene outlet analyzer update schedule		✓			✓	✓							✓			
Hydrogen outlet analyzer update schedule		✓			✓	✓							✓			
Carbon dioxide analyzer update schedule		✓			✓	✓							✓			
Position of MOV valve during swing		✓				✓										✓
Timer on MOV movements		✓			✓	✓										✓
Status of E2 H2 production		✓			✓	✓										
Status of double-block & bleed and vent on reactors		✓				✓									✓	
Location, control, and status of 1" pressure-up lines on reactors		✓				✓									✓	
Column Totals: 313 IRs	260	132	73	177	106	233	121	115	125	20	66	46	92	36	47	35

APPENDIX C: P+F AND P+F+T INTERFACE DESCRIPTIONS

In this appendix, we provide a detailed description of the novel interfaces used in this study. This includes a discussion of terminology, an overview of the interface, and a preview of the graphics practices used repeatedly in the displays[11]. The subsequent sections each describe a major component of the interface in detail.

Terminology

Woods (in preparation) describes the organization of a user interface as follows:

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. Often referred to simply as 'view'.
3. Workspace: The set of viewports and classes of process views that can be seen in parallel or in sequence.

In addition the above definitions, the following definition is used within this document:

4. Graphic: A particular visual element within a Process View.

Workspace

The workspace is comprised of two standard 21" monitors (CRT) at a resolution of 1280x1024 each (see Figure 42). The two monitors sit one atop the other; Monitor 1 denotes the lower monitor and Monitor 2 denotes the upper monitor. Monitor 1 is divided into three viewports and Monitor 2 is divided into seven viewports. Each of the viewports that make up Monitor 1 will be described briefly in the following section. The viewports that make up

[11] Images in this section have been modified to protect NOVA Proprietary information.

Monitor 2 are described in much greater detail. The reason for this (and the obvious design differences between the Monitors) is that the authors designed the views in Monitor 2 while NOVA constructed those in Monitor 1 (see section on color below). The differences in level of detail reflect the original purpose of this description, which was to provide the implementer of the displays with the detail necessary to code the desired behaviors.

Commonalities

There are several graphical standards employed throughout the views. These are described here generally and are assumed to hold for each view unless otherwise specified. The term 'standard' is used here to indicate that a method is used consistently within the displays and not in referring to broader standards for graphical interface design.

Text and numerical formats. Two types of text are employed in all of the views. Labels and data are printed in 11-point Arial font. This includes all digital tag values and most of the labels used to identify data and graphical elements. Most scale values are printed in 8-point Arial font. It should be noted that font sizes are not precisely specified in the ActiveX implementation. They should approximate 11-point or 8-point type, but may differ slightly from this standard. All digital values are printed with a single decimal place, e.g., ###.#. Exceptions are noted in the specification for each view.

Lines. Most lines are drawn as width 1, the thinnest line allowed by low level WIN32 API "Create Pen" function used to draw them. (Within VB this is nearly equivalent to the BorderWidth property of a standard line object.) This includes the lines on all scales. Exceptions to this general rule for lines are noted where required.

Colors. Color selections were made based on qualitative appearance and were not based on any standard. This document should not be used to select a color scheme for graphical displays.

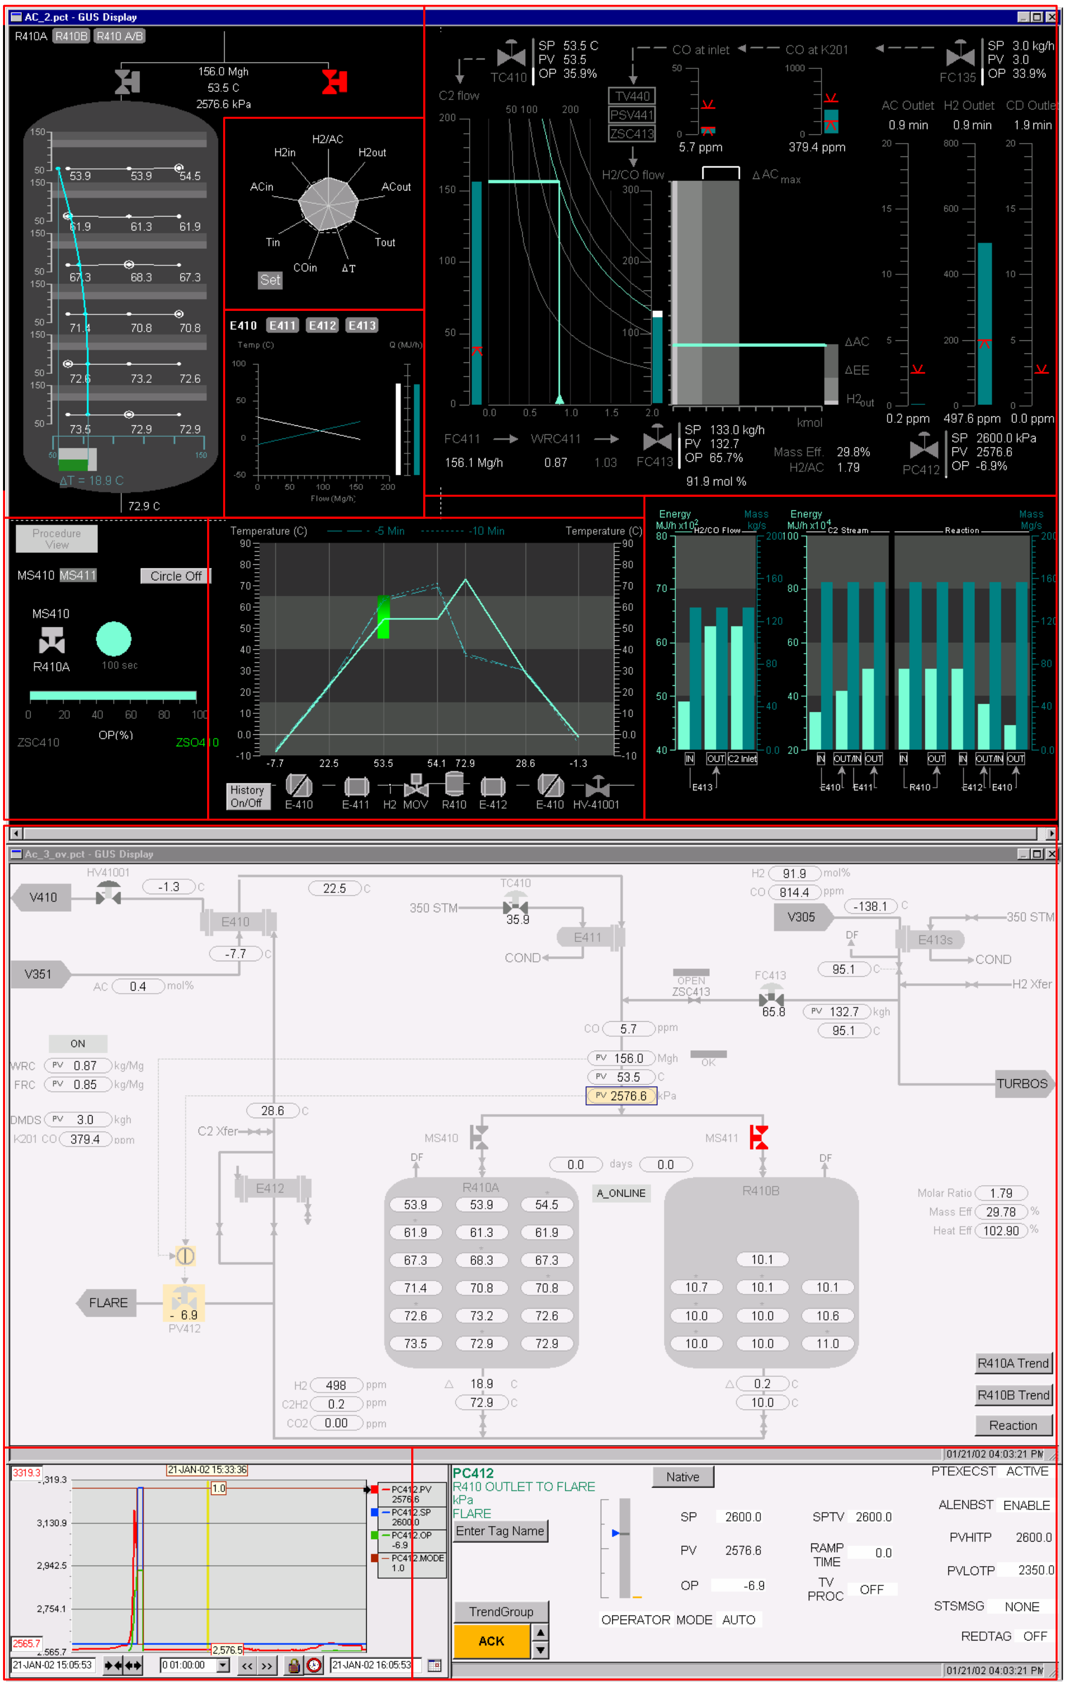

Figure 42: Overview of monitor 1 (lower) and monitor 2 (upper) with viewports delineated.

The two Monitors employ different background colors; black for Monitor 2 and light gray for Monitor 1. Black was used for Monitor 2 based on unequivocal comments from operators consulted during the design review process. When the decision was made to have NOVA develop the views in Monitor 1, the plan was to have them use existing graphics and controls developed for a new process. They were to modify the color scheme of those views from light gray to black backgrounds. However, it proved to be impossible to make these modifications. By the time this became evident, it was too late in the implementation process to change the background colors in Monitor 2.

Focus Regions. 'Tooltips' and 'on-click' behaviors are specified in many of the displays. The regions exhibiting these behaviors are active and are called 'focus regions'. Focus regions are shown graphically in the spec when they cannot be clearly described in the text. A 'tooltip' refers to a small text box that appears in the vicinity of a graphical element that identifies the element and, in some cases, provides the current value of the data tag associated with that element. A single 'on-click' behavior is employed. When the left mouse button is pressed while the pointer is located in a focus region, the data tag associated with that region is called up in the Detail View (see below). In the interface description below, the term 'standard focus region' refers to a region that yields a tooltip with tagname only, and sends the associated tag to the Detail View. Focus regions with tooltips containing both tagname and current value are called out wherever they occur.

Persistence. Each control contains some graphical elements that must be configured, e.g., scale ranges, default values, critical parameters, colors, etc. Once these values have been set, it is desirable that the controls retain that configuration when they are closed and later re-launched. This is accomplished in two ways in the controls. Some controls rely on the GPB frame to maintain persistence, others have an associated .INI file, and still others rely on a combination of these two strategies. A consistent strategy would be preferable for controls that are to be distributed to multiple users.

DETAIL VIEW

Purpose

The Detail View allows the operator to access the parameters of a data tag and make changes to controller modes, setpoints and outputs.

Description

The contents of the Detail View are the parameters of a data tag. The parameters visible at any time in the view depend on the tag selected and its type. A control valve contains the largest number of parameters and will serve as an example for the view description (see Figure 43). Generally speaking, any field with a white background is one that allows either text input or pull-down menu selection.

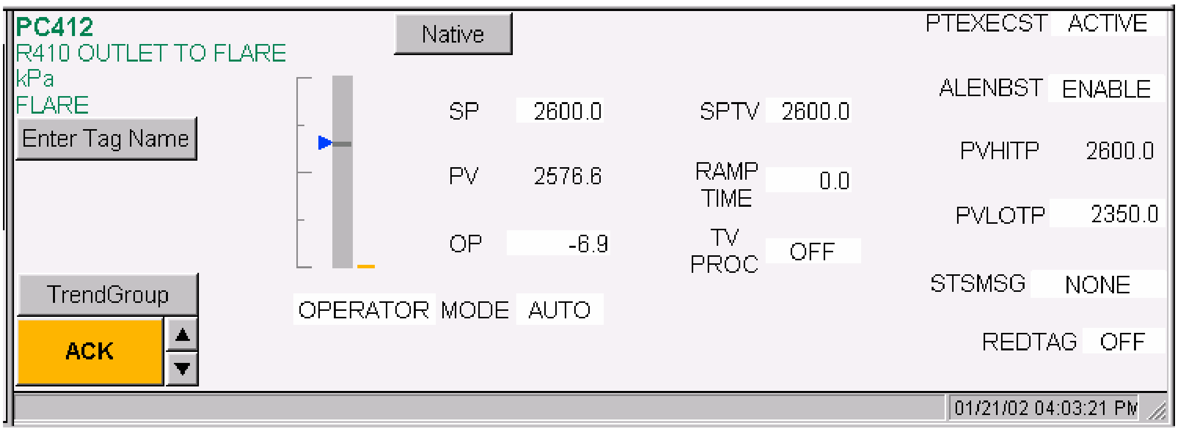

Figure 43: Detail view.

Tags are brought to the Detail View in one of two ways. First, clicking on most of the objects in the various views will send the tag associated with that object to the Detail. Second, a button labeled "Enter Tag Name" allows the operator to specify a tag to be shown in the view.

The tag name appears in the upper left-hand corner of the view, followed by a description of the tag and the units. To the right of this information is a scale showing the setpoint (SP), output (OP), and current value (PV) of the controller. The range of the scale is defined by the high and low level alarms and corresponds to the gray column. The SP is shown by a blue triangle to the left of the column. A dark gray line in the column shows the PV for the tag. To the right of the column, a thinner gold column reflects the OP of the valve (0-100%). Digital values for each of these parameters are provided to the right. Further to the right, values are provided for setting target values (SPTV) and ramp times, and commands for executing those ramps (TV PROC).

The Controller Mode display appears at center bottom of the view. This field is primarily used to set the control mode for the various controllers in the reactor sectio n. The Alarm Enable Status (ALENBST) allows the operator to disable the alarms that are associate with the individual tags. The high and low level alarm (PVHITP & PVLOTP) setting fields were disabled in this implementation.

Several other items that appear in the Detail View were not functional in the application described here. These include the "Native Window", "Trend Group", and alarm acknowledge (ACK) buttons", as well as the "STSMSG" and "REDTAG" fields.

SCHEMATIC VIEW

Purpose

The Schematic View provides an overview of the equipment organization and status in the reactor section.

Description

The Schematic View is based on an annotated mimic display concept. The graphics reflect the reactor section process flow and process tag values are displayed in the locations at which they are measured in the field.

Process Flow. The main content of the view is the equipment and lines that make up the reactor section process. Inflows and outflows are noted with arrows with upstream sources and downstream sinks labeled. Tag identifiers for valves are provided along with PV values, if available. Heat exchangers are labeled and readings from nearby temperature indicators are placed upstream and downstream. Inline valves controlled from the field are included on the process line and analyzer readings are placed as close to their field positions as possible. Two gray bars in the upper right corner of the view serve as hydrogen flow and low-flow trip warnings. These bars turn red when the respective trip occurs.

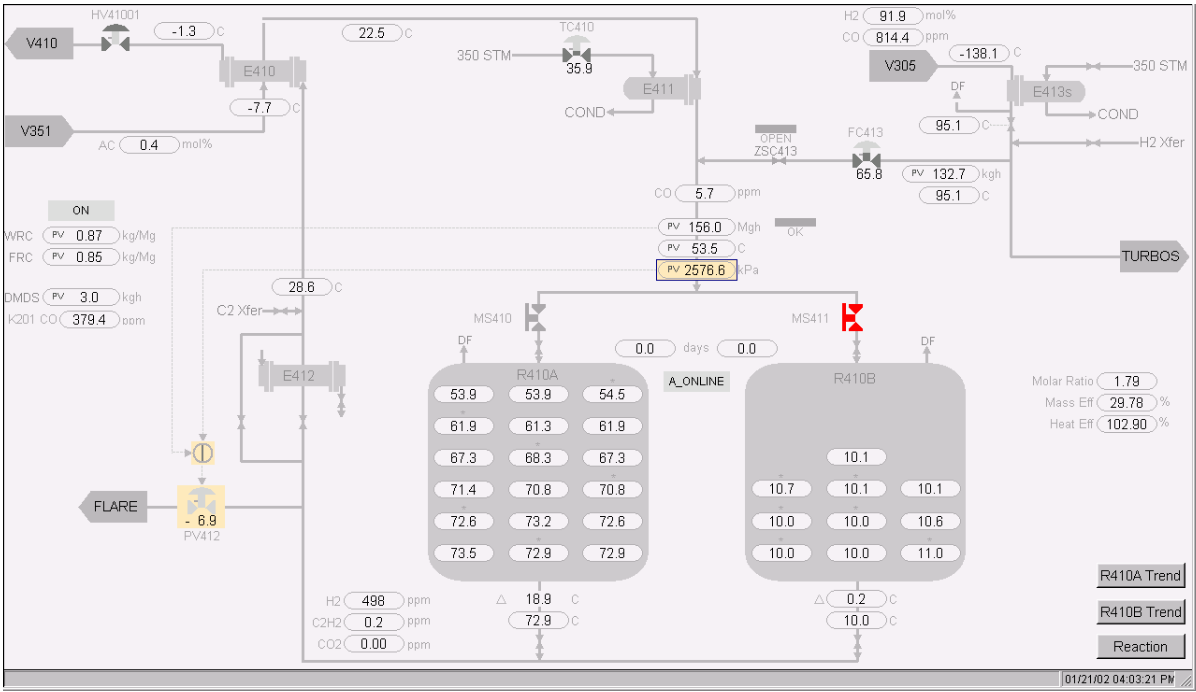

Figure 44: Schematic view.

Annotations. Most of the graphics have several types of annotation to reflect focus, control modes, alarms states, and off-normal settings. These annotations are discussed in great detail in Nova's design document (NOVA, 1999) and will be covered only briefly here.

The tag that is "in focus" at any given time is indicated by a manila background. This tag is yoked to the Detail View. For some tags, several objects will have this background coloring. For example, a control valve in focus appears with a manila background behind the valve as well as the process variable value (see Figure 44).

Magenta is used to denote an off-normal condition. A temperature, pressure, or flow tag with a magenta oval around the value denotes a point for which the alarms have been disabled. A control valve with a magenta box around it is in an off-normal control state (e.g., the valve is in manual (MAN) whereas it is typically in advanced control (CAS)).

Alarms are shown in three colors, blue for low priority alarms, orange for high priority alarms, and red for emergency alarms. A colored box appears around the tag parameter value and valve graphic. This box flashes as long as the alarm is in. Auditory alarm indications were not provided in the interface.

Valves. There are six valves shown on the schematic. Four are control valves and two are motor-operated valves. The control valves are shown in two-tone gray and have curved caps. The proportion of the valve shown in the darker gray represents the output setting for the valve. Thus, if the valve is 50% open, the icon's lower 50% is dark gray and upper 50% is light gray. The two MOV valves, which have rectangular caps, are color coded according to OP as well. They appear in light gray when the valve is open, red when it is closed, and yellow when it is in between the open and close states. Additional detail on how these states are determined is included on page 162.

Supplemental information. There are several locations at which supplemental text or icons are provided in locations not tied to the schematic. On the left-hand side of the View are two controller settings for the hydrogen flow ratio controller. Below that is a flow value for the DMDS, which is controlled from a valve that is not in the same geographic location as the rest of the reactor section equipment. An upstream carbon monoxide concentration is also provided.

In the lower right of the view are three buttons for pre-set trend groups (R410A Trend, R410B Trend, Reaction). Pressing any of these buttons signals the Group Trend (see below) to display a trend with a preset group of tags related to the label on the button. For the two reactor buttons, the group trend is of the average row temperatures for the reactor. For the reaction trend, the mole ration, acetylene and hydrogen outlet analyzers, inlet temperature, delta T, flow, temperature and pressure are all traced. Finally, the molar ratio, heat efficiency, and mass efficiency appear at center right.

TREND VIEW

Purpose

The trend view is the primary means for providing historical information in the Interface. Two trend views are available to be displayed in the trend window, a Detail Trend and Group Trend. The first is yoked to the Detail View so that the value(s) associated with the tag appearing in the Detail View are shown on the trend. The second reflects either the selection of a predetermined set of tags (chosen from buttons on the Schematic View) or a group of tags assembled by the operator.

Description

Tag parameters. As many as seven tag parameters can be plotted simultaneously. Boxes to the right of the plot identify the parameters in the current trend. The tag name and parameter is noted first, followed by the current value of that tag. A short line segment and a colored box indicate the line color associated with that tag. For most tags (e.g., a temperature indicator or analyzer reading), only a single parameter is plotted. For a control valve, the PV, SP, OP, and controller mode are all plotted together (see Figure 45). Clicking on the tag boxes identifies the currently selected tag and a black arrow pointing at the color box shows which of the tags is selected.

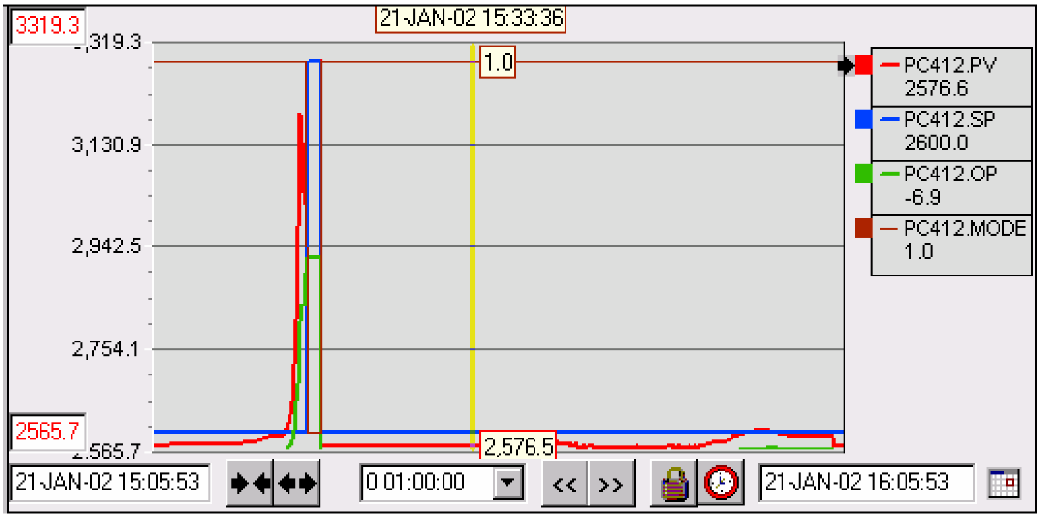

Figure 45: Trend view.

Scope. The bottom row of indicators in the view shows the time range included in the current plot. Fields on the left and right ends of the view give the date and time for the oldest and most recent points of the plot, respectively. To the right of the time on the left are two double-arrow icons. Clicking on the icon with the arrows pointing together reduces the time frame of the plot while clicking on the icon with the arrows pointing away from each other expands the time frame. A pull-down menu to the right of the arrows allows the user to select a time frame from a pre-determined set of ranges (e.g., 10 min., 20 min., 1 hour, 2 hours, etc.).

Further to the right are two more double-arrow icons with the heads pointed in the same direction. These two icons advance or rewind the plot in steps sized according to the current selected time frame. Finally, two additional icons allow the user to lock a time frame in place or to force the plot to advance to the current time.

Plot area. The time frame selected and the historical values of the tag parameter(s) that appear on the plot determine the range of the vertical plot. The value range of the plot depends on the currently selected parameter. This range adjusts automatically to show all of the historical values for the tag over the given time frame. The user can reset this range by entering new values in the scale limit boxes to the left of the plot. The color of the text in these boxes corresponds with the trace color for the selected tag. When a new parameter is selected, the range of the vertical automatically updates and the high and low values are shown in the boxes.

Group tends. The Group Trends graphic is essentially identical in appearance and function to the Trend Detail. The parameters appearing in the plot are determined either by selecting a pre-defined group on the Schematic View or building a new group manually. Any tag in the Schematic can be added to a group by right-clicking on the tag and selecting the parameter to add to the trend. Up to seven traces can be added to a group.

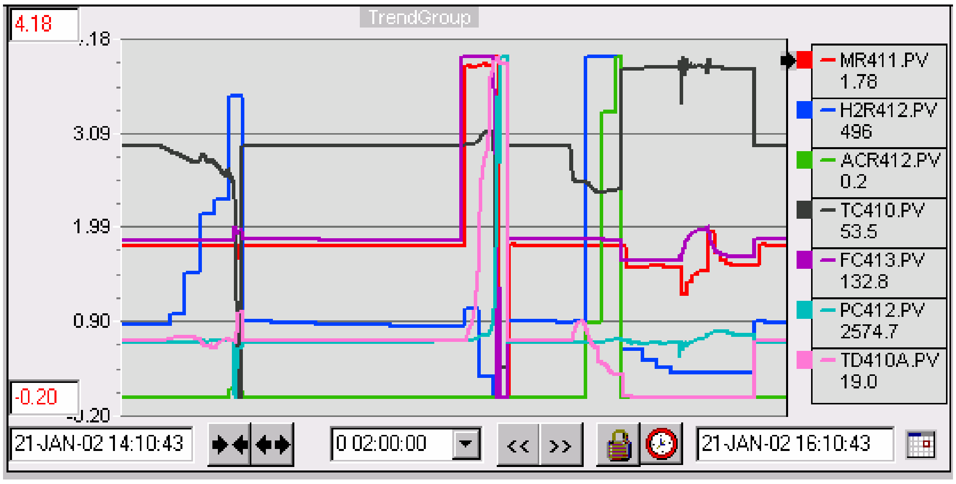

Figure 46: Trend view: Reaction trend group.

HYDROGEN BALANCE VIEW

Purpose

The Hydrogen Balance View (Figure 47) provides information about the types and rates of reactions taking place in the reactor. It also provides information about the input and output streams. The view can be described in terms of four separate graphics: the Ratio Controller, Hydrogen Balance, Analyzers, and Control Valves. The many standard tooltips in this view are shown graphically at the end of the description.

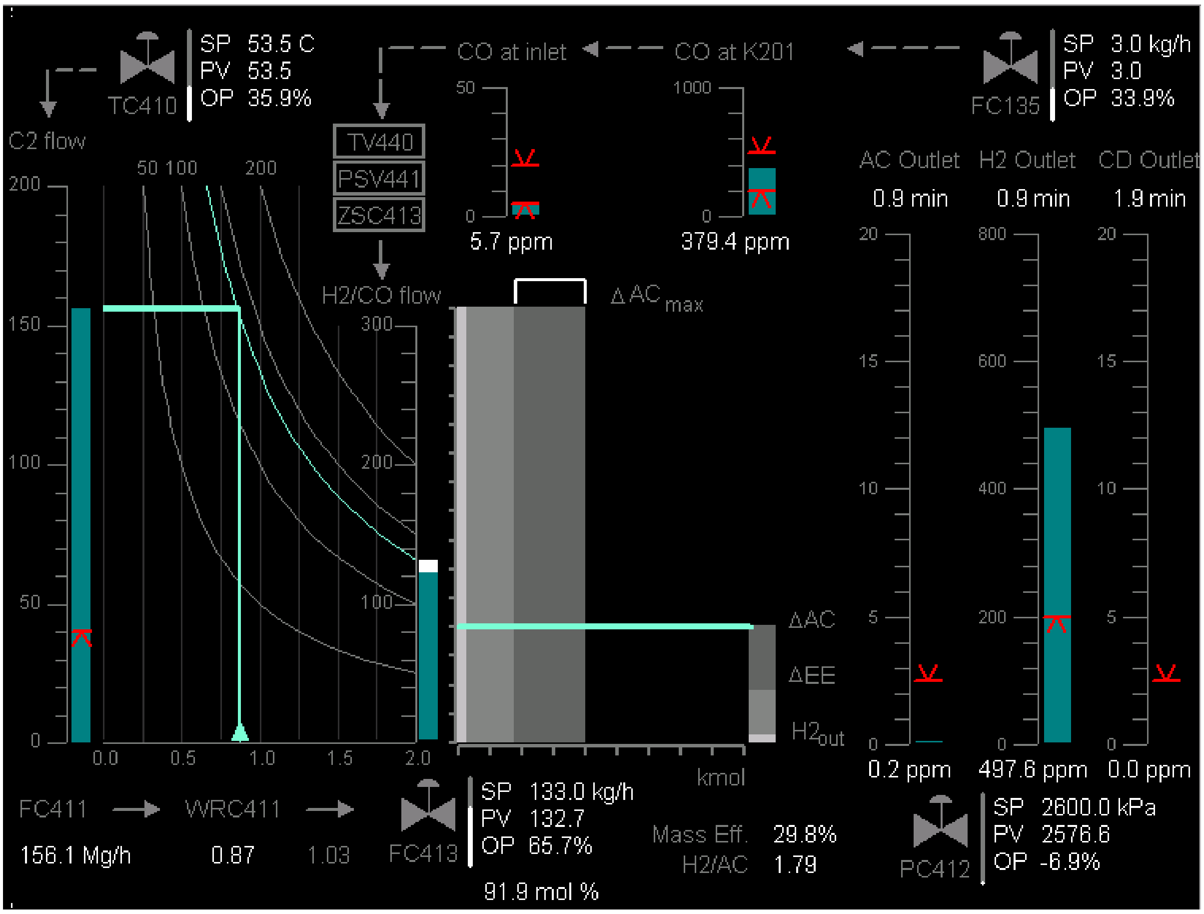

Figure 47: Hydrogen balance view.

Description

Ratio controller. The Ratio Controller Graphic allows the operator to monitor the control of the H2/CO stream (see Figure 48). Two flow-rate graphics are used to display the flow-rate of the C2 stream (left side of Figure 48) and the H2/CO stream (right side of Figure 48). Each graphic has a gray fixed scale, ranging from 0 to 200 in Mg/hr for the C2 flow (FC411) and 0 to 300 in kg/hr for the H2/CO flow (FC413). Above each of the scales, the flow is denoted in dark gray. A minimum limit for the FC411.PV is set at 40 Mg/hr. The limit is indicated by red limit arrow formed by a chevron above a horizontal line segment, both of width 2. This graphic is used several times in the View.

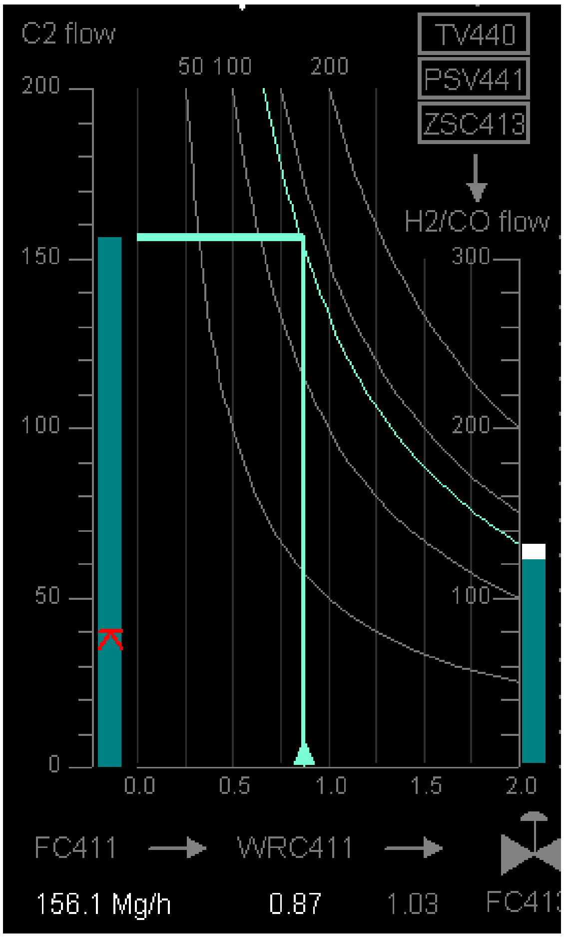

Figure 48: Hydrogen balance view: Ratio controller.

For the C2 flow, a cyan bar indicates the flow rate. A divided cyan and white bar is used to describe the flow rate of the H2/CO stream. The total height of this bar is given by FC413.PV. The cyan portion of the bar represents the amount of H2, the white portion represents the amount of CO and impurities. The proportion of the column drawn in cyan is given by H2M304.PV. It is assumed that H2M304.PV will take a value between 0 and 100. For values greater than 100, the graphic behaves as if H2M304.PV=100. For values less than 0, the graphic behaves as if H2M304.PV=0.

The behavior of the ratio controller is represented by a graphic that lies between the two flow rate bars (Figure 48). The baseline of the graphic reflects the value of the weight ratio controller (WRC) setpoint. It is drawn in a medium gray50 and has a range of 0.0 to 2.0. Darker vertical gray lines extend from the baseline at 0.25 increments, extending to a point equal to the maximum values of the C2 flow scale. Four constant value lines are drawn in medium gray on the plot at values of 50, 100, 150, and 200. At every point along each of the lines, the product of the value along the baseline and the value along the C2 flow scale is a constant. The lines are static. A fifth constant value line is drawn in green. The value of this line is determined by the current value of FC413.PV. A small green triangle rests on the baseline at the present value of the setpoint of the weight ratio controller. A green line of width 2 extends vertically from the triangle, up to the present constant value line. If the WRC entry is less than 0, this line will behave as if WRC=0; if it is greater than 2.0, no green line will be plotted on the graph. A second green line is drawn horizontally from the top of the C2 flowrate column across to the current constant value line. The two lines should meet near the current constant value line defined by FC413.PV. Immediately above three of the four static lines (50, 100, and 200) is a text indicator of its value written in 8 point Arial font in medium gray.

Below the Ratio Controller graphic is a combination of text indications and connecting arrows (see bottom of Figure 48). These include static text for the C2 flow and weight ratio controller in medium gray font. Actual values for the C2 flowrate and WRC setpoint are printed in white font. A weight correction factor (WF413COMP) appears in medium gray text.

Hydrogen balance. The middle portion of the view shows the Hydrogen Balance Graphic (see Figure 49). The graphic displays the amount of H2 entering the reactor (H2F410.PV), what reactions occur inside the reactor (ACD411.PV, EAD411.PV), and the H2 accounted for after the reactor (H2F411.PV). As a whole, the graphic depicts the mass efficiency calculation (EFM410.PV). The actual value for the mass efficiency tag is given in white 11-point Arial font at the bottom of the graphic. The graphic is based on a kmol-kmol scale with a range of 0-150 kmol/h on the vertical axis. No scale values are shown on this axis. However, the axis itself forms a focus region for H2F410.PV, with the tooltip showing both tagname and current value. This current value forms a point on this axis. Although this point is not shown graphically, it serves in the construction of a later graphic.

There are three anticipated dispositions for the hydrogen, two of which are reactions:

1. Ethylene Conversion (ΔEE): C2H4 + H2 -> C2H6 + heat
2. Acetylene Conversion (ΔAC): C2H2 + H2 -> C2H4 + heat
3. Some hydrogen may pass through the reactor without reacting ($H2_{out}$)

A baseline scale below the rectangles denotes the amount of H2 accounted for in kmol/h. The increments on this scale are the same as those on the vertical axis. The width of the three rectangles represents the rates of H2out, ΔEE, and ΔAC, respectively. The rectangles are colored in increasingly darker shades of gray. The height of the rectangles is fixed at the height of the H2/CO scale. Each rectangle also serves as a focus region, with the tooltip

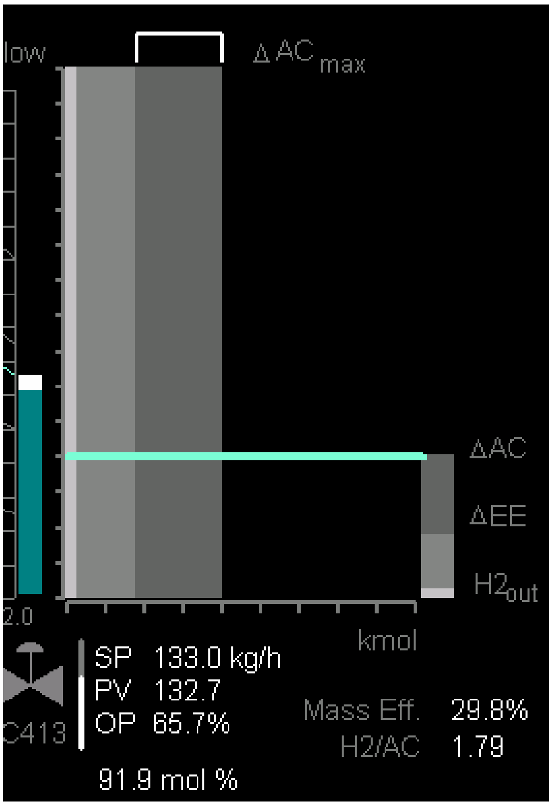

Figure 49: Hydrogen balance view: Hydrogen balance.

showing both tagname and current value. To the right of the scale, a stacked bar graph reflects the width of each of the three disposition rectangles, using the same colors. This vertical bar represents the total H2 consumption and can be compared to the H2 input on the left side of the reactor graphic. A cyan line of width 2 is drawn from the point on the vertical kmol scale corresponding to H2F410 to the top of the stacked H2 output bar. If all of the H2 is accounted for, the line will be horizontal. If the line is sloped negatively, the graphic indicates that some of the incoming H2 is not being accounted for and there may be a problem. If the line is sloped positively, the graphic indicates that more outgoing H2 is being measured than incoming H2. A light gray horizontal reference line is drawn from the point corresponding to H2F410.PV to help operators detect deviations in the balance line.

The Hydrogen Balance graphic also includes an indication of the maximum amount of acetylene conversion possible (i.e., the total number of acetylene moles entering the reactor).

This is shown by a white bracket located just above the ΔAC rectangle. The left edge of the bracket lines up with the left edge of the rectangle and its width is equal to the number of moles of inlet acetylene. Any discrepancy between the end of the bracket and the rectangle indicates that some acetylene is not reacting.

Analyzers. Analyzer readings for acetylene (AC), hydrogen (H2), and carbon monoxide (CO) in the outlet stream are located on the far right of the H2 Balance View (see Figure 50). White bars indicate the value of the concentration of each in PPM. The scales are fixed from 0 to 20 for acetylene and CO, and 0-800for H2. A limit graphic (described above) on each scale indicates a limit for the concentration. For AC and CO, these are maximum limits of 2.5ppm, for H2 a minimum limit is prescribed at 200ppm. Digital values for the instrument readings are provided below the scales. Above each of these three scales is a digital update timer for the measuring instrument. The digital value is presented in white font in the form #.# minutes and counts down to the next update for the analyzer.

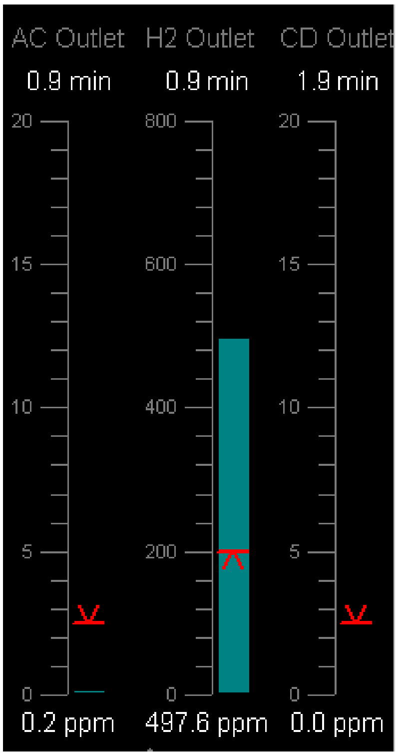

Figure 50: Hydrogen balance view: Analyzers (outlet).

Additional Analyzer graphics appear above the Hydrogen Balance graphic to represent the incoming CO concentration and H2/CO flow lock-out (see Figure 51). The two CO concentration graphics use the same color and style as those described above (less the analyzer update timer). However, their vertical height is reduced. The CO at inlet scale ranges from 0-50 with minimum and maximum limits at 5ppm and 20ppm, respectively. The CO at K201 scales ranges from 0-1000 with a maximum limit at 500ppm and a minimum limit at 200ppm. Dashed gray lines connect the text indications together and to the H2/CO flow rate scale.

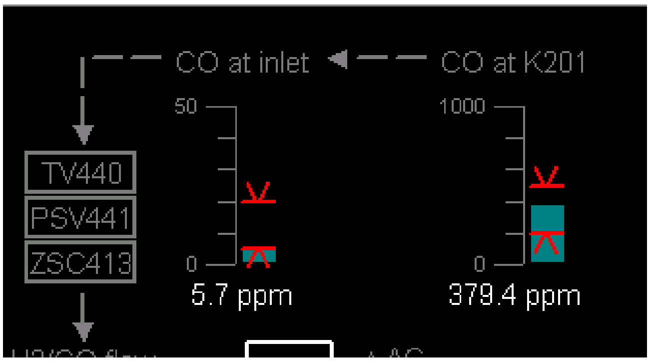

Figure 51: Hydrogen balance view: Analyzers (inlet).

Three medium gray outline rectangles lie on the flow path line leading to the H2/CO flow rate column (see Figure 51, left). Each represents one possible location for blocking or venting H2/CO flow. The box colors for TV440 and ZSC413 can change depending on process values. The TV440 box changes to yellow when 0<TV440.PV<=25.0C; the box turns red when TV440.PV<=0. The ZSC413 box turns red when ZSC413.PV=ON. The box color for PSV441 cannot change color because the relief valve that it represents is not instrumented.

<u>Control valves.</u> Four standard control valve representations appear above the three graphics discussed above (see Figure 52). These are for the reactor temperature inlet control (TC-410), the DMDS control (FC-135), H2/CO flow control (FC-413), and reactor inlet pressure control (PC-412). The standard control valve representations consist of a valve icon and a simple graphic for the valve position. This consists of a gray column to represent the

range of possible positions (0-100% open) and a white column whose length is proportional to the valve OP position. Text values for the SP, OP, and PV tags are shown in white 11-point Arial font.

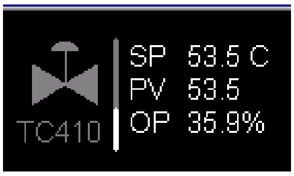

Figure 52: Hydrogen balance view: Standard control valve representation.

Mapping

Figure 53 shows the mapping between graphics in the Hydrogen Balance View and DCS tags.

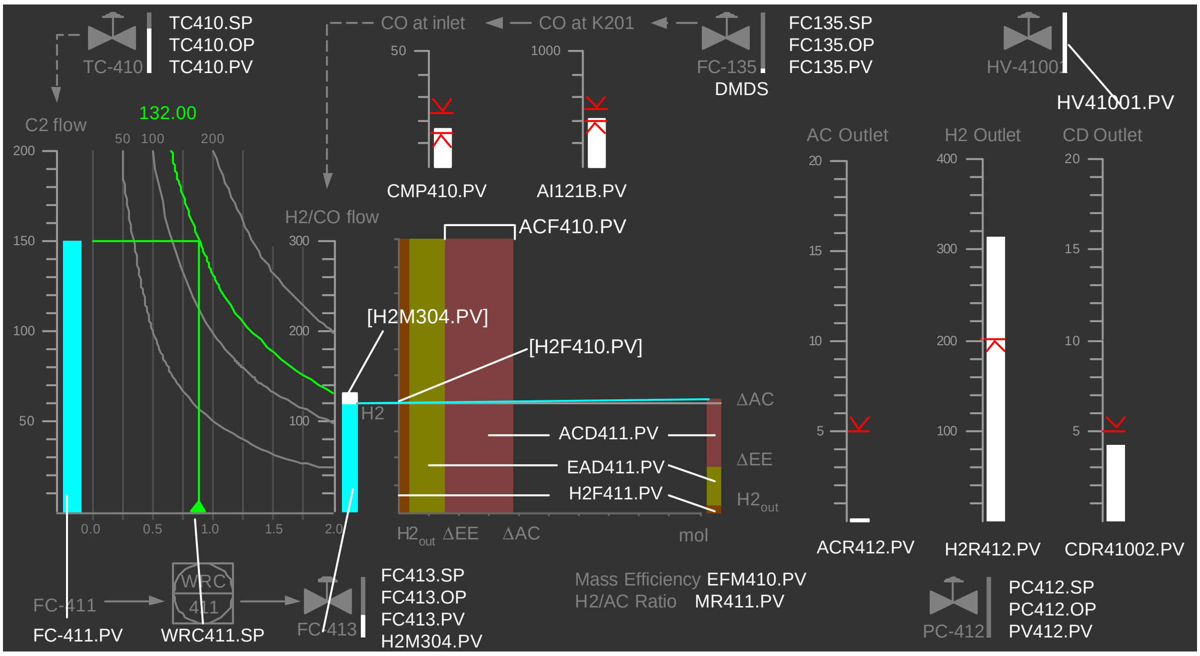

Figure 53: Hydrogen balance view: Tag mapping.

Tool-tip and 'on-click' actions are defined for the objects associated with the following tags:

TC-410	CMP410	AI121B	FC-135	HV-41001	ACF410
FC-411	WRC-411	FC-413	EFM-410	MR411	ACR412
H2R412	CDR41002	PC-412	ACD411	EAD411	H2F411

Approximate focus areas for each of these tags are shown in Figure 54. The tooltips for most of these focus areas are of the standard variety. However, for H2F410.PV, ACD411.PV, EAD411.PV, and H2F411.PV, the tooltip shows both tagname and current value.

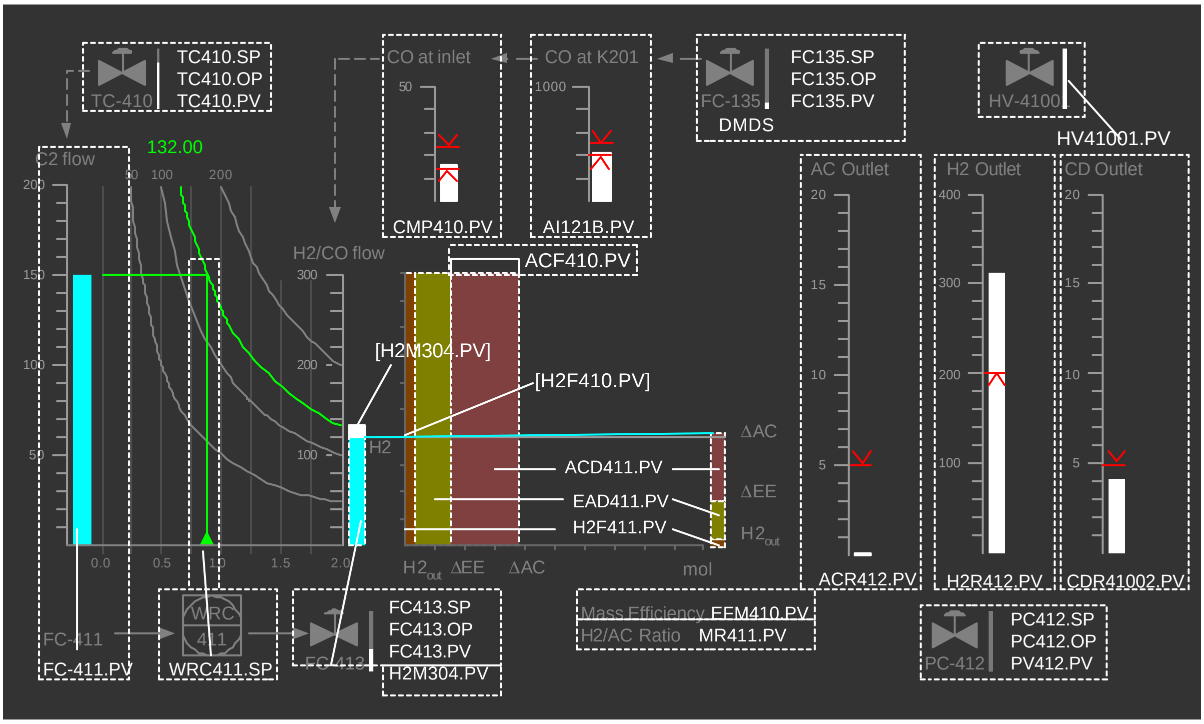

Figure 54: H2 balance view: Approximate focus areas.

MASS AND ENERGY VIEW

Purpose

The Mass and Energy View (Figure 55) provides an overview of mass and energy flows through the AHR. This view allows the operator to monitor mass and energy flows across the reactor subsystems and detect anomalies.

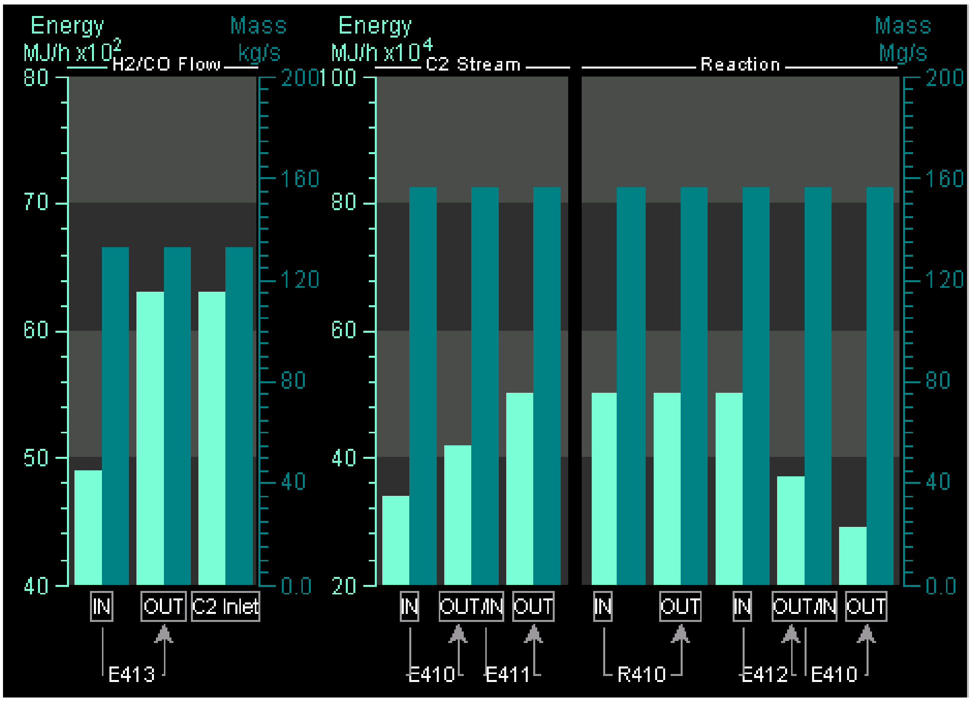

Figure 55: Mass and energy view.

Description

The view consists of a series of paired bars that represent the mass and energy flow plotted at several locations in the process steam. The energy bars are colored cyan and the mass bars dark cyan. A focus area is defined for each bar, with tooltips showing both tagname and value. The view is divided into two graphs, one for the H2/CO (Offgas) flow and one for the C2 Stream and Reaction subsystems. This is because the mass flow rate in the C2 stream is several orders of magnitude larger than the Offgas stream. For each of the two graphs, an cyan energy scale is located on the left and a dark cyan mass scale is located to the right. The text on these scales is printed in the same color as the scale. The Offgas flow graph plots energy from 4,000 to 8,000 MJ/h and plots mass from 0 to 200 kg/s. The C2 Stream and Reaction graph plots energy from 20,000 to 100,000 MJ/h and plots mass from 0 to 200 Mg/s. Horizontal blocks of alternating dark and darker gray are placed at several intervals to improve the readability of the graph.

The labels at the bottom of the graph help the operator locate the mass and energy values in the process flow. The equipment labels and “IN” and “OUT” pointers are printed in white text. The lines connecting the equipment labels and the “IN” and “OUT” test are white. The box around the “IN” and “OUT” labels is made up of solid straight lines in medium gray. The line pointing to the “OUT” box in each case has a small arrow at the tip.

Mapping

Figure 56 shows the mapping of mass and energy graphics to their respective tags in the DCS. There are 11 sets of mass and energy flow tags, which are printed below the columns on the graph to which they are connected. For each pair, the energy flow tag is listed on top, and the mass flow tag below. Many of the tags are repeated because the sensor set in the existing system is under-specified in comparison to this design. This means that assumptions of conservation of mass and energy have been made.

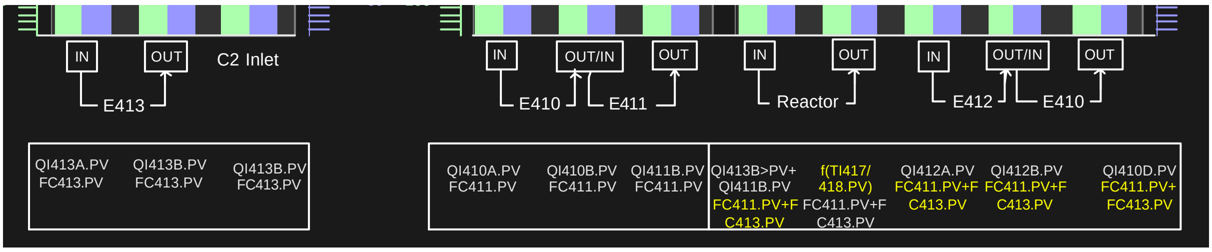

Figure 56: Mass and energy view: Tag mapping.

PAULSEN TEMPERATURE VIEW

Purpose

The Paulsen Temperature View (Figure 57; adapted from Paulsen, 1992) shows the temperature profile of the C2 stream across the AHR equipment. The view provides an overview of temperatures in the system and places them in the context of established temperature limits.

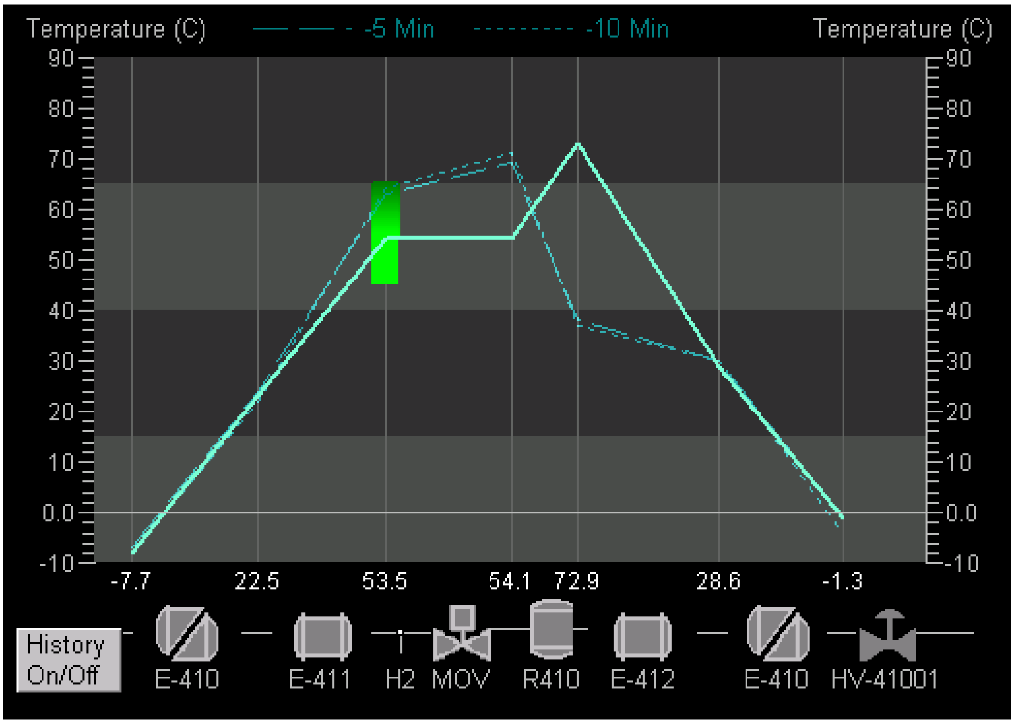

Figure 57: Paulsen temperature view.

Description

Graph. The Paulsen Display view plots seven temperature readings for the C2 stream. A simplified schematic of the AHR components is laid out at the bottom of the view to provide a visual context for the temperature readings. Hard-coded text identifiers appear in gray text below the components and also act as standard focus areas. The temperature readings on the plot are lined up with the appropriate location on the schematic, allowing operators to quickly determine where the temperature reading originates. Reference lines in dark gray are drawn from the appropriate location on the schematic to the top of the graph to clearly indicate the location of the temperature readings. A horizontal gray baseline is drawn at 0 degrees. Digital values of the temperatures are placed below the baseline of the graph at the point where the reference lines meet. The reference lines and digital values also serve as standard focus areas.

The C2 stream temperature profile is plotted as a continuous cyan line of width 2 on the graph with inflection points at the seven temperature indications.

Scale. The temperatures are plotted on a fixed scale from –25 to 175°C. A light gray scale is placed on each side of the graphic. Horizontal blocks of alternating dark and darker gray are placed at 25-degree intervals to improve the readability of the graph.

Limits. The view also provides a graphical indication of specified temperature limits for the system. Two long red lines are placed at 150°C to indicate the high temperature limit for the heat exchangers. Another limit graphic exists for the reactor temperature. The yellow line, placed at 110°C, indicates the alarm temperature. The red line, placed at 130°C indicates the reactor high temperature trip. See the Reactor Temperature Profile View section for a detailed description of these limits.

Operators must also monitor the C2 feed temperature to control reactor activity. The normal feed temperature is from 45°C to 55°C with a maximum of 60°C and a minimum of 35°C. A rectangular graphic to display the normal and limit temperatures is provided in the display. The graphic is placed at the temperature reading for the C2 stream just after E-411, straddling the reference line. The solid green portion indicates the target temperature zone of 45°C to 55°C. The green gradient portion of the graphic indicates the acceptable temperature range of 55°C to 65°C.

History threads. A pair of temperature history threads is provided to support detection of slow temperature profile changes. Each thread shows the temperature profile at a fixed lag behind current time (defined as lag[1] and lag[2]) where lag[1] must be more recent than lag[2]. Each thread is composed of a dashed cyan line. The more recent line (lag[1]) has longer dashes and shorter spaces, while the older data line (lag[2]) has larger spaces and shorter dashes.

Mapping

Figure 58 provides the mapping of each of the digital tag values in the graphic to its tagname. The reactor inlet temperature can come from one of two tags, depending on which reactor is online. When the value of tag ACX410.S1=A_ONLINE, the value of TAV4061 is displayed; when the value of ACX410.S1=B_ONLINE, the value of TAV4101 is displayed.

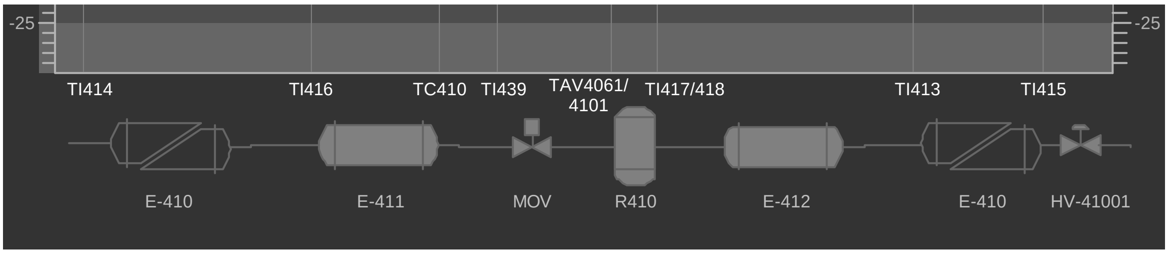

Figure 58: Paulsen temperature view: Tag mapping.

REACTOR TEMPERATURE PROFILE VIEW

Purpose

The Reactor Temperature Profile View (Figure 59) provides the operator with the vertical and horizontal temperature distributions of the reactor bed. It also includes the common feed input temperature, flow, and pressure, and the outlet temperature reading. The view allows the operators to monitor the status of the reactors and determine if temperatures are within desired operating conditions. The design also assists operators in understanding the state of the reactions, planning future actions, responding to potential problems, and troubleshooting.

Description

The design consists primarily of a dark gray schematic-like display of the reactors with supporting equipment icons and digital values. A detailed temperature display graphic is integrated into the schematics of each of the two reactors.

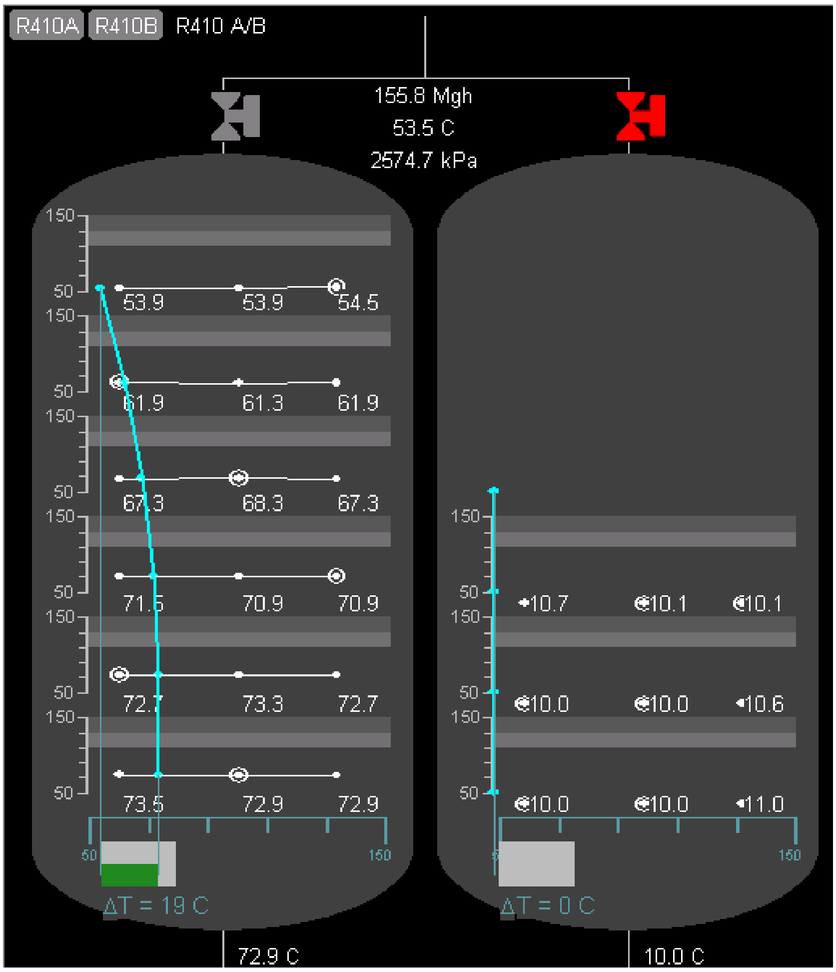

Figure 59: Reactor temperature profile view (A&B tab).

There are three tab-based views, the temperature profile of reactor R410A with associated polar star and Heat Exchanger panel (see Figure 60, left), R410B with associated polar star and Heat Exchanger panel (see Figure 60, right), or both R410A and R410B simultaneously with no polar stars or Heat Exchanger panel (Figure 59). The operators can switch between the view options by clicking on the tabs. The default operating condition is for presentation of data pertaining to the on-line reactor [tag ACX410.S1]. Thus, the display will appear as the left-hand side of Figure 60 if R410A is on-line [ACX410.S1 value=1] or as the right-hand side of Figure 60 if R410B is on-line [ACX410.S1=2].

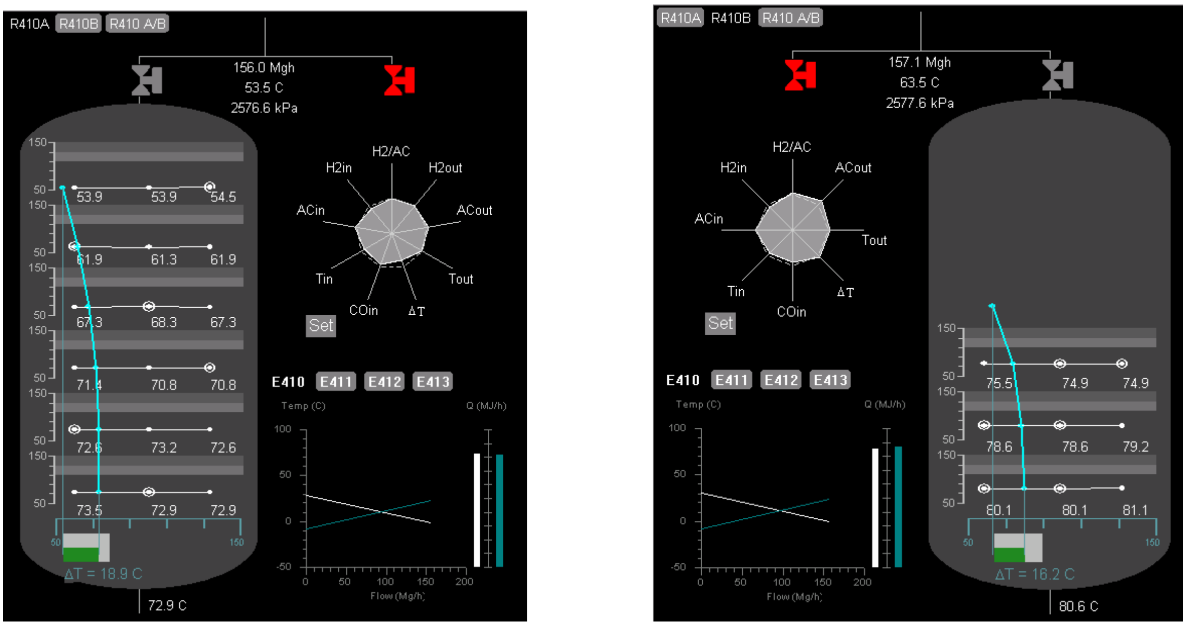

Figure 60: Reactor temperature profile view: R410A (left) and R410B (right) with associated polar stars.

Reactor temperature display graphic. The reactor temperature display graphic provides a visual indication of the temperature profiles in the reactors (see Figure 61). It is composed of the horizontal and vertical temperature profile graphical elements (described below). Each reactor has 3 rows of temperature sensors, each row having 3 sensors corresponding to a temperature tag.

Horizontal temperature profile graphical element. Each row of three temperature tags is displayed against a temperature scale (see Figure 62). Both scale lines and text are light gray. Values for scale limits for each row are drawn from the INI file with a default range of 50°C to 150°C. Three individual temperature readings are displayed as white dots placed at their appropriate relative position on the vertical scale. These dots act as focus regions with tagnames only appearing in the tooltip. One (R410A) or two (R410B) dots in each row has a white circle around them to indicate that the associated tags are connected to the PLC410 Reactor Temperature Trip Function. The temperature readings for a given row are linked together with straight white lines of width 2. The digital value of each temperature indication

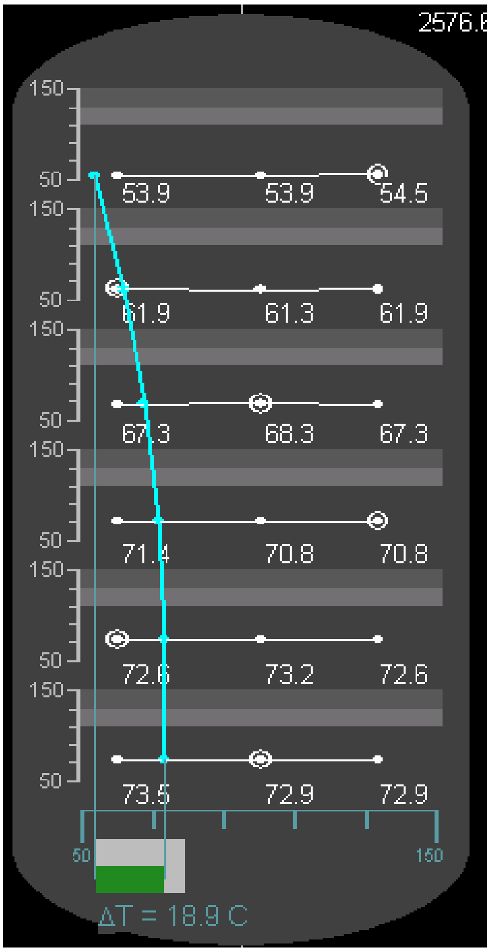

Figure 61: Reactor temperature profile view: Temperature display graphic.

is displayed in white below the dots. The digital values are positioned below the scale and are fixed (i.e., the position of the digital values will not change). The white dots are the predominating visual element in the view, followed by the connecting lines.

The two gray bands located near the top of the scale are indicators of alarm and trip temperatures (see Figure 62). The lighter gray band begins at 110.0°C and ends at 129.9°C and indicates the alarm/warning temperature range. The darker gray band begins at 130.0°C and continues to the top of the scale (in this case, 150.0°C). These bands provide an indication of how close the temperatures are to a possible alarm or trip.

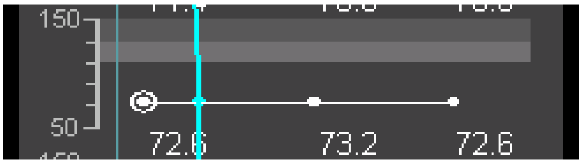

Figure 62: Reactor temperature profile view: Horizontal temperature profile graphical element.

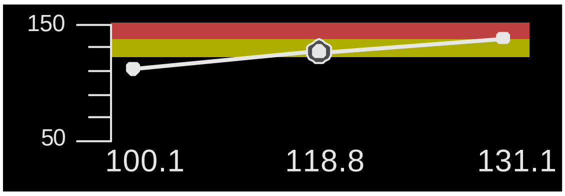

Figure 63: Reactor temperature profile view: Alarm/warning temperature bands.

Vertical temperature profile graphical element. The Reactor Temperature Display graphic also contains a vertical temperature profile (see Figure 64). The Vertical Temperature Profile graphic displays the average (mean) temperature tag for each row, but on a separate scale placed at the bottom of the graphic. The Vertical Temperature Profile Display element is superimposed on top of the Horizontal Temperature Profile Display element. The mean temperature for each row is depicted by a cyan dot placed in each horizontal row. The cyan dot acts as a focus region with the text, “Mean=” followed by the tag value. The dots are connected by straight cyan lines of width 2 to help the operator visualize the temperature profile across the height of the reactor.

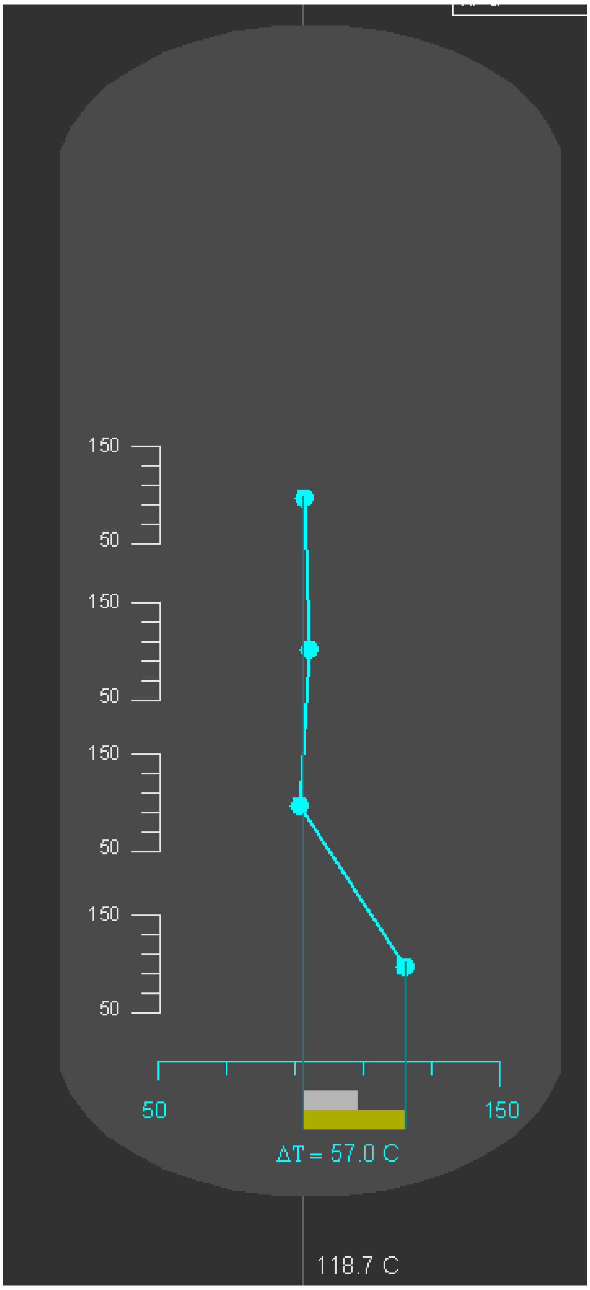

Figure 64: Reactor temperature profile view: Vertical temperature profile graphic.

The ΔT across the reactor height is indicated at the bottom of the display (see Figure 65). This value is the difference between the average temperature of the bottom row and the average temperature of the top row. Dark cyan lines are drawn from the average temperature indicator of the top row and the average temperature indicator of the bottom row to indicate the ΔT value on the scale.

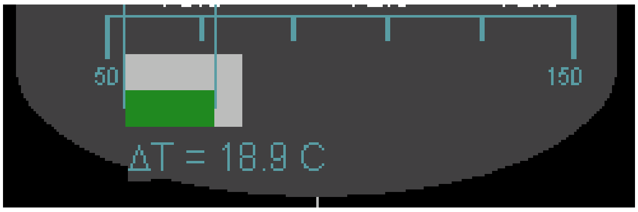

Figure 65: Reactor temperature profile view: DT graphic.

The ΔT should not exceed 25°C. A light gray box placed below the scale is used to indicate the maximum-desired ΔT. The left edge of the box begins at the dark cyan line extended from the average temperature indicator of the top row and extends to the right for 25°C. A dark green bar is drawn between the two dark blue lines to indicate the ΔT.

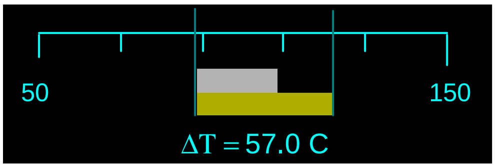

Figure 66: Reactor temperature profile view: DT graphic with high DT.

Supplementary data and graphics. In addition to the graphical elements discussed above, light gray flow path lines are included above and below the reactor as shown in Figure 67. Five digital tag values are presented in the proximity of these lines. The common reactor inflow rate (FC411.PV), temperature (TC410.PV), and pressure (PV412.PV) indications are positioned at the top of the display. Below each reactor is a temperature indicator, TI417 for Reactor A, and TI418 for Reactor B. All five of these digital values functions as a focus region with a tooltip showing the associated tag name.

Mapping

Figure 67 provides an indication of the TDC tag associated with each digital and graphical information element in the view.

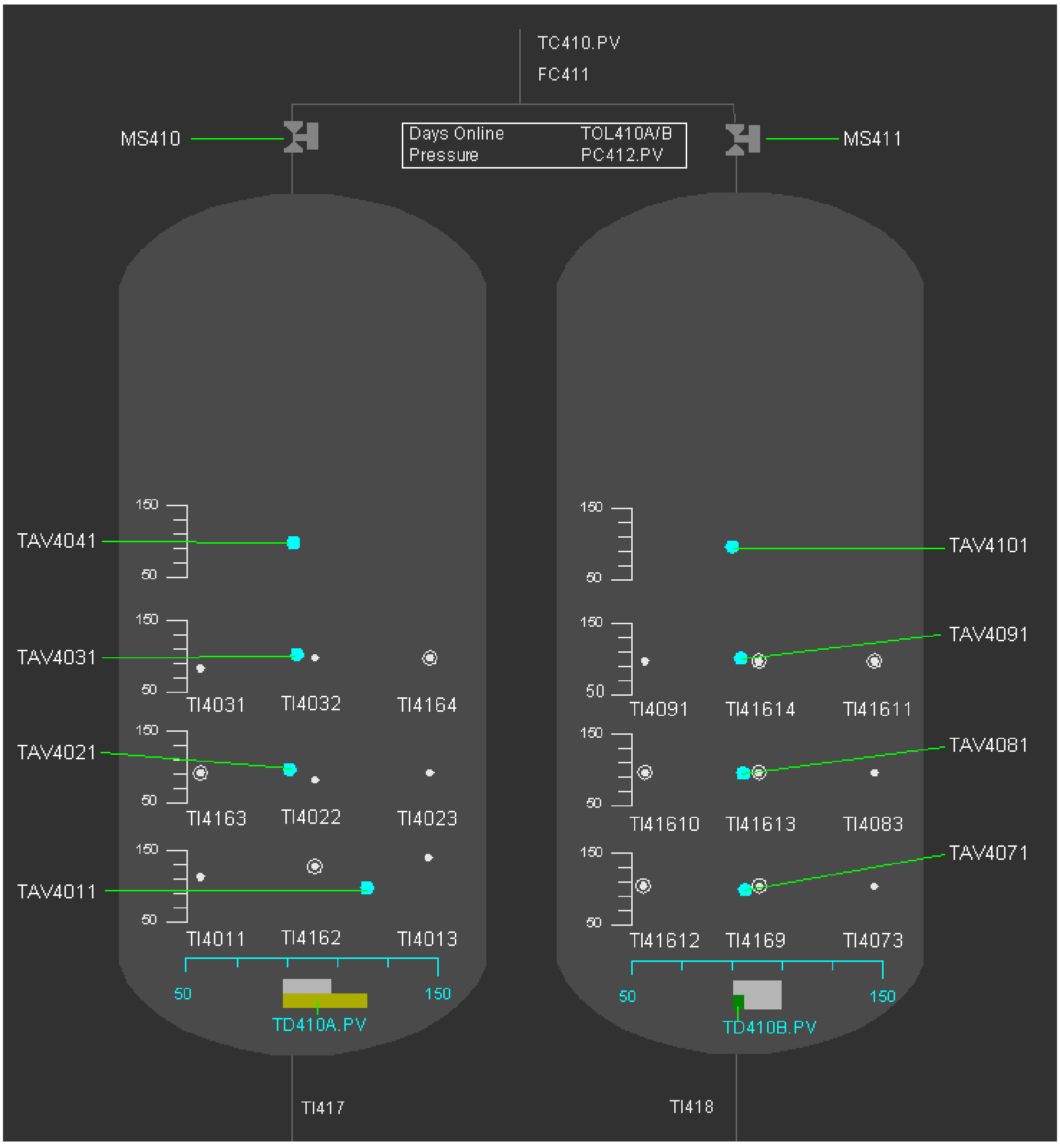

Figure 67: Reactor temperature profile view: Tag mapping.

<u>MOV valve icons.</u> An MOV valve icon appears above each reactor schematic. Each acts as a focus region with only tag name appearing in the tooltip. The current status of the valve is taken from the tags MS410.PV (for R410A) and MS411.PV (for R410B). These PVs have three states and an important BADPV condition. The states are 0-STOP, 1-OPEN, and 2-CLOSED. The 0 state is effectively non-functional and would only appear in the case of a sensor failure. A BADPV state will occur when the MOV does not meet the conditions that define the states 0, 1, and 2. This will occur whenever the valve is between its fully open and

fully closed position. Three MOV icons are defined to convey four possible states of the PV parameter. The color mapping is: 0-yellow, 1-gray, 2-red, BADPV- yellow.

POLAR STAR VIEW

Purpose

The Polar Star View (Figure 68) provides an "at a glance" overview of the reactor performance (Wolff, 1967). It is intended to call the operator's attention to anomalies in the reactor performance. The Polar Star View appears in the Reactor Temperature Profile View. Its left/right position is specified by the tag selected (R410A, R410B, R410 A&B)

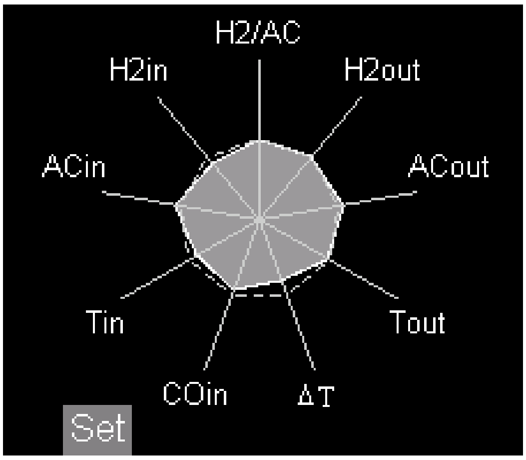

Figure 68: Polar star view.

Description

The polar star is composed of four to twelve evenly spaced radii, or "spokes", with a coincident vertex. Each radius is drawn as a solid light gray line. Each of the radii is associated with a process variable tag. A descriptive tag designation appears in a caption at the end of the spoke in white text. This description text acts as a focus area with the tooltip containing both tag name and value. A dashed light gray "reference" polygon is superimposed on the radii with its center point at the common vertex of the radii and its edges at the midpoint of the radii. All of these elements are static in the display (i.e., they do not disappear, change color, flash, etc.).

The operator configures the polar star using a configuration dialog. In the dialog, the operator can select the number of spokes, assign tags to those spokes, and change the vertex, midpoint, and tip. In addition, the configuration dialog allows the user to set a “fixed” end at either the vertex or tip of each spoke. Identifying the fixed end is important to determining how the individual spokes are scaled (see Use Case: Configure axis properties).

The current value of the parameter is established on the axis in a position that is proportional to its value relative to the limits. No indication for that position is placed on the display. Rather, the current value positions of each axis are connected to the current value positions on each of its neighboring axes using a white line. This process, repeated around all of the axes, forms the outline of a solid “data” polygon. A gray fill is added to the interior of the data polygon.

A "Set " button is included in the display. Pressing this button automatically rescales the spokes so that the midpoint value is equal to the current value. This rescaling results in a symmetric data polygon at the reference polygon (See Use Case: Establish midpoints).

Mapping

The default configuration for the polar star contains the following tags:

Spoke 1:	MR411.PV	H2/AC
Spoke 2:	H2R412.PV	H2out
Spoke 3:	ACR412.PV	ACout
Spoke 4:	TI417.PV	Tout
Spoke 5:	TD410A.PV	ΔT
Spoke 6:	AI121B.PV	Coin

Spoke 7:	TC410.PV	Tin
Spoke 8:	ACM351.PV	ACin
Spoke 9:	WF413.PV	H2in

HEAT EXCHANGER VIEW

Purpose

The Heat Exchanger View (Figure 69) provides overview information about the performance of the four heat exchangers in the reactor section. The view supports evaluation of the operating performance of the heat exchangers. A tab-based design is employed to allow serial access to information about each exchanger. The view is located on the R410A and R410B tabs below the polar star (Reactor Temperature Profile View). It is not visible when the R410A/B tab is selected.

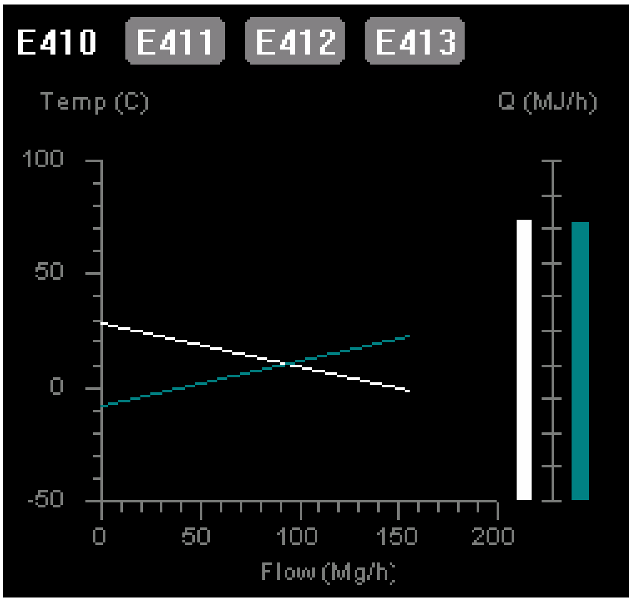

Figure 69: Heat exchanger view.

Description

The view is comprised of two graphics, an inlet-outlet temperature plot and a heat exchange comparison plot. The inlet-outlet temperature plot is at left in the display. It represents flow rate on the horizontal axis and temperature (in degrees C) on the vertical axis. Labels are printed in gray above the vertical axis and below the horizontal axis (see Figure 69). Flow units for E410, E411, and E412 are Mg/hr; units for E413 are kg/hr. Temperature units are degrees C. Scales for each axis are given in Table 13 for each exchanger. The axes and scale markings are medium gray.

Two lines of width 2 are plotted on the co-ordinate field. Each represents the temperature change on one side of a heat exchanger. The hot side line is drawn in white and the cold side in cyan. The temperature lines are essentially point-to-point and are not descriptive of a function. There are two styles of lines; the one used for each side of the exchanger depends on whether condensing is taking place on that side of the exchanger. The simple lines are for sensible temperature changes (i.e., no phase change). These lines are drawn from the inlet temperature to the outlet temperature for an exchanger side (see Figure 69).

Exchanger sides (i.e., hot sides of E411 and E413) with phase changes are more complicated. Figure 70 shows temperature change lines for a condensing heat exchanger. The cold side line is as described above. The hot side line, however, is composed of three line segments. The first (left-most) segment reflects the temperature drop of superheated steam to the saturation temperature (138°C). The second line segment (center) shows the constant temperature latent heat exchange due to condensation. The final segment (right) shows the subcooling of the condensed vapor. If the length of the line is divided into five equal portions, the first and third segments are each 1/5 of the length of the line, while the second segment is

3/5 of the length of the line. Condensing exchangers have a gray dotted line drawn across the plot at the saturation temperature of the pressurized steam (138°C).

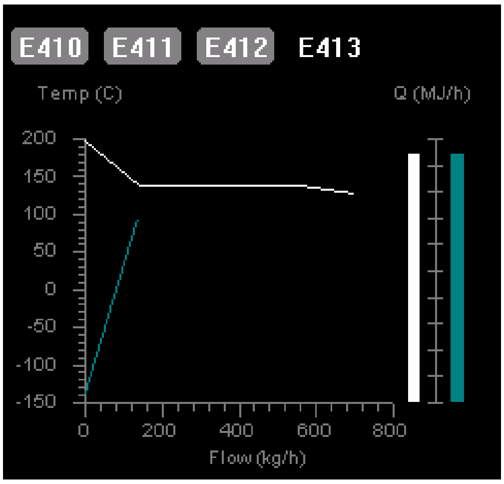

Figure 70: Heat exchanger view: condensing heat exchanger.

The second graphical element in the Heat Exchanger View is the heat exchange comparison plot (see Figure 70). This plot is composed of two columns, each representing the heat exchange on one side of the exchanger. Each column acts as a focus region with the tooltip showing both tag name and value. The columns are color-coded to match the temperature change lines; white for the hot side and cyan for the cold side. The label "Q(MJ/hr)" appears above the columns in gray 11-point Arial text. A gray scale is drawn between the columns. The scale for E410 and E411 ranges from 0-10,000 MJ/h. The scale for E412 ranges from 0-15,000 MJ/h. The scale for E413 ranges from 0-1,500 MJ/h.

The plots for the four exchangers are accessed serially through a tabbed interface (see Figure 70). Each tab displays the plot for one exchanger.

Mapping

Inlet and outlet temperature tags and heat exchange values for these lines are given in Table 13 for each exchanger. The on-line reactor determines the temperature indicator used for the E412 hot side inlet temperature. When the value of tag ACX410.S1=A_ONLINE, the value of TI417 is displayed; when the value of ACX410.S1=B_ONLINE, the value of TI418 is displayed.

Table 13: Axis ranges and data tags for each heat exchanger (all tags are PVs).

	x-axis range	y-axis range	Hot in TI-	Hot out TI-	Cold in TI-	Cold out	Hot Flow	Cold Flow	Q hot QD-	Q cold QD-
E410	(0,200) Mg/hr	(-50,100)	413	415	414	TI416	FC411	FC411	412B	410B
E411	(0,200) Mg/hr	(0,200)	911A	911B	416	TC410	FI911S	FC411	911A	411B
E412	(0,600) Mg/hr	(0,200)	417/ 418*	413	912A	TI912B	FC411	FI912S	412A	912B
E413	(0,800) kg/hr	(-150,200)	913A	913B	340	TC440	FI913S	FC413	913A	413B

MOV VIEW

Purpose

The MOV View (see Figure 71) provides feedback for real-time control of the MOV valves (MS410 or MS411). It uses a predictive display to create a representation of the MOV state that is not afforded by the existing instrumentation.

Description and Mapping

The view is comprised of two tabs, one for each of the MOVs. An icon shows the MOV state at the bottom left (MS410) or right (MS411) of the view. The set of status icons is identical to those employed in the Reactor Temperature View and the appropriate icon for the current valve state is selected in the same way. Above the icon is a white label denoting the valve name; below the icon is a white label denoting the reactor to which the valve controls flow.

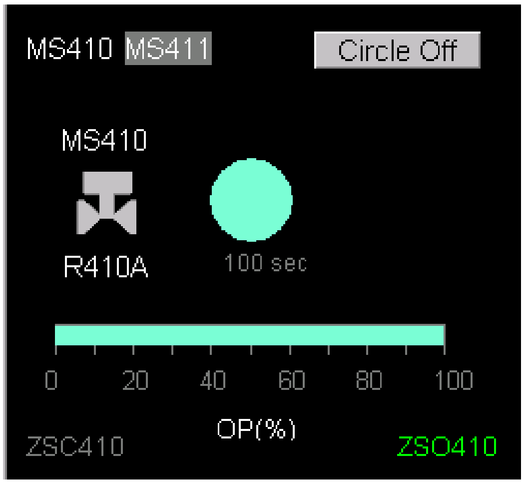

Figure 71: MOV view.

The MOV valve itself is controlled in the DetailView window. The activation of either tab sends the associated valve to the Detail View. The commands are OPEN, CLOSE, STOP, and ENTER. The OPEN, CLOSE, and STOP states map onto the MS410.OP or MS411.OP values. The ENTER button sends those commands to the control system. The MOV window is sensitive to changes in these tags.

Extending across the bottom of the view is a gray horizontal scale indicating the percent output (0-100) of the valve controller. The label "OP %" appears in white text below this scale. A cyan bar drawn against the scale shows the current degree to which the valve is open. This value is not given by a tag. Rather its position is proportional to the total amount of time that the valve has been moving in one direction (OPEN or CLOSE) divided by 100 seconds (the time required to fully open or close the valve). This calculation must be performed in the control. The animated valve position representation moves to the right when the current MS410.OP or MS411.OP tag value is OPEN, and to the left when the value is CLOSE. It stops moving when the value is STOP.

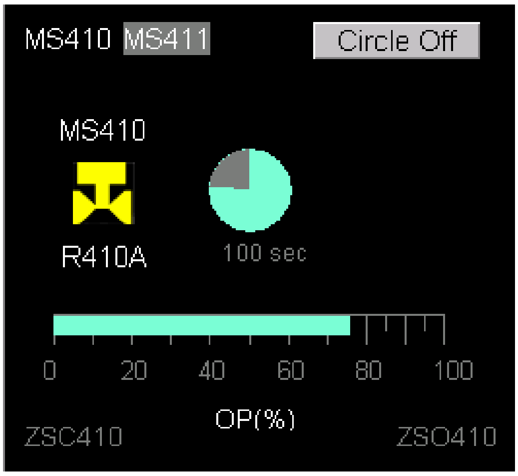

Figure 72: MOV view: Valve in travel.

Indicators for the close (ZSC410 and ZSC411) and open (ZSO410 and ZSO411) switches for either valve are placed below the scale at the left and right, respectively. The color of the indicators depends on the value of their associated tag. The mapping between the indicators and their tags is given in Table 14.

Table 14: Open and close indicator mappings.

MOV	Open Indicator	Closed Indicator	ZSO ON	ZSC ON	ZSO/ZSC OFF
MV411	ZSO411.PV	ZSC411.PV	Green	Red	Gray
MV410	ZSO410.PV	ZSC410.PV	Green	Red	Gray

In addition to the valve animation, there is a timer representation for the valve movement. A gray circle is used to indicate the total amount of time required to completely open or close the valve (i.e., 100 seconds). Once the operator has ENTERED an OPEN command, a cyan sector will expand in a clockwise direction from the 12:00 position to show the execution period relative to the total amount of time required to fully open or close the valve. A similar animation is shown in a counter-clockwise direction when a CLOSE command is entered. A numerical reading for the elapsed time in the execution period is placed above the circle in cyan font.

PROCEDURE VIEW

Purpose

The Procedure View (Figure 73) supports the operator in executing the reactor swing and reactor temperature runaway procedures. The procedure steps are included in the view in a color-coded tree structure. Each step is associated with a static image of the reactor equipment state.

Description

The view is divided into two areas, the Steps window on the left and the Images window on the right (see Figure 73).

Steps window. The Steps window presents the steps of the procedure printed as a tree structure in medium gray 11-point Arial font. The steps are hierarchically organized in the tree structure to reflect their organization in the procedure. Each step is numbered and its text content and associated images (see below) are provided in a text file. The step text is color-coded based on the progress of the operator through the procedure. The color mappings are as follows (with values in parentheses indicating the maximum number of steps that can have that text color):

Active Step (1)	magenta
Acknowledged Step (1)	chartreuse
Previous Steps (all)	dark green
Future Steps (all)	medium gray

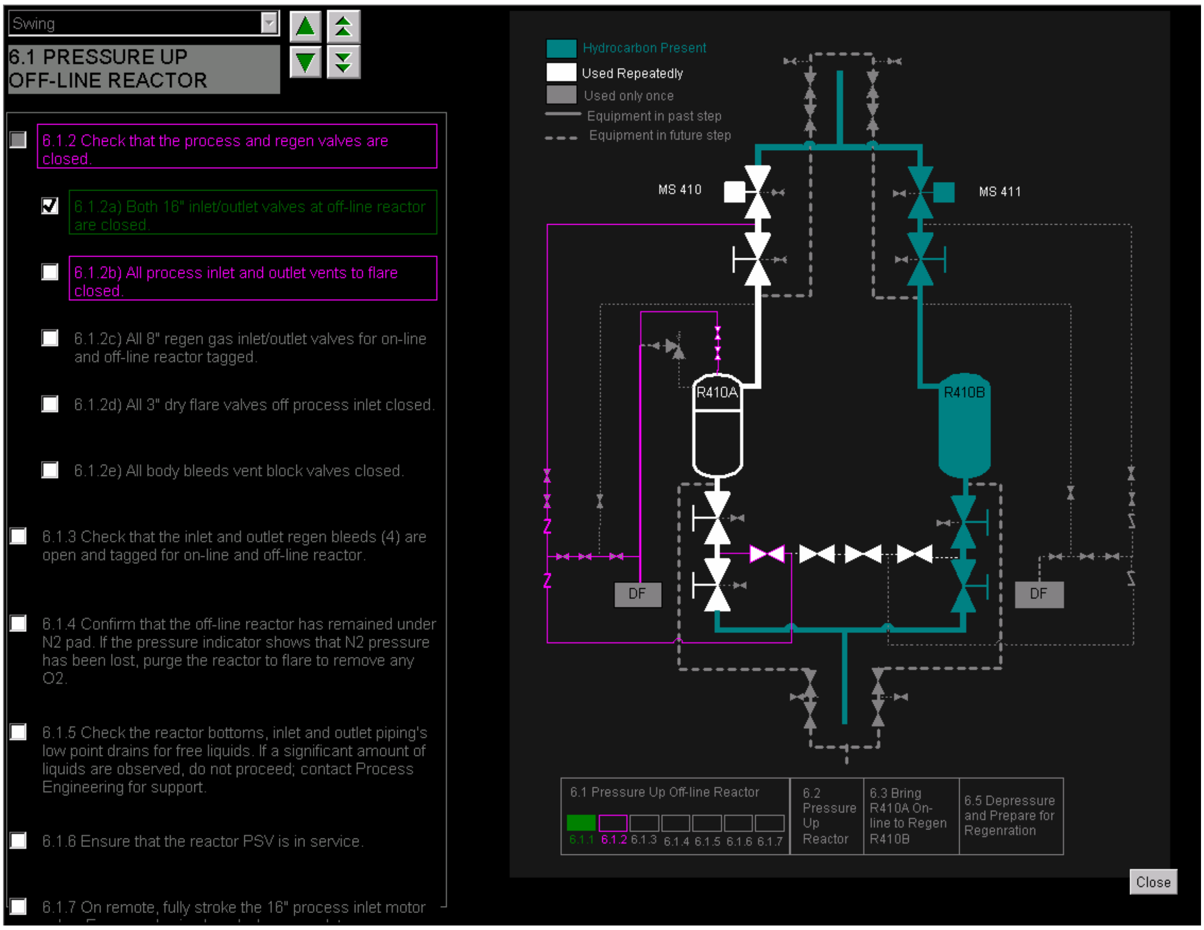

Figure 73: Procedure view.

There is a check box to the left of each step. Interaction with this check box determines the color of the step text. The 'active' step is the step that the operator is currently working on, but has not been confirmed to be complete. When a step is confirmed to be complete, clicking in the check box changes the status of that step to 'acknowledged'. Following a 2-second delay, the next step in the sequence is placed in 'active' status. Once the succeeding step is made 'active', the formerly 'acknowledged' step transitions to 'previous' step status. A 'future' step is one that has not yet been reached in the procedure. A use case for checking off steps and stepping between them is provided below.

At the top of the scrolling window, the high level (e.g., 6.X) step is shown in black text printed in 14-point bold Arial font on a gray background. This text is not part of the scrolling window and no images are associated with it. It is determined by the active step (6.XX or 6.XXx). The text updates to the next high-level step when the first sub-step of that level is activated.

A box surrounds the text of the 'active' and 'previous' steps. The color of the box matches the color of the step text. No box is drawn around 'future' steps.

Several steps act as 'parent' steps with several 'children'. Parent steps have grayed out check boxes that a user cannot check off. An alternative set of step states are defined for parent steps. These are:

Active Parent (1)	magenta
Acknowledged Parent (1)	chartreuse
Previous Steps (all)	dark green
Future Steps (all)	gray50

A parent step becomes an 'active parent' in the same manner that a regular step does (i.e., when the previous step has been completed). However, the first child step under the parent also becomes 'active' at the same time. The parent remains in the 'active parent' state while any of its children are 'active'. The standard step progression semantics follow for the child steps until the last child step is completed. When the final child step is checked off, it transitions to 'acknowledged' and the parent transitions to 'acknowledged parent' for 2 seconds. Then both steps transition to 'previous'. Once the parent step is 'previous', a check mark is placed in its check box. Throughout the child steps, the child steps determine the images. No images are associated with the parent step itself.

Check boxes for parent steps are grayed out. A check mark will automatically be placed in the check box of the parent step once the last sub-step of the parent step has been checked off. This should require no direct interaction by the operator.

Images window. The images window shows a series of static bitmap images that are mapped to each step of the procedure. The files are identified by name in the formatted procedure text file (see above). Two images are identified for each non-parent level 1 and 2 step. One image maps to the 'active' step and one maps to the 'acknowledged' step. No interaction with the bitmap images is required.

For each step, the 'active' image highlights the pieces of equipment that are associated with that step. To do so, the relevant lines and outlines are colored magenta while the step is active. The 'acknowledge' image highlights the same pieces of equipment in chartreuse for several seconds before the next active image is shown. Other codes employed in the jpegs will be described later.

The steps overview provides the operators with a brief idea of how far they have gone with respect to the entire swing process. It is located at the bottom of the Images Window. A small rectangle represents each substep. The rectangles are color-coded to match the status of the procedure steps to which they refer.

A legend is also included in the upper left hand corner of the Images window. It helps the operator to map the color graphic codes in the images to their meaning.

Message window. A hidden message window can appear below the MOV window when error messages are required (e.g., to call attention to a missing procedure step or image file). The Window appears below the MOV window, pushing the window up and forcing a re-sizing of the Steps window. When the user clicks on the close window button, the MOV returns to its normal position and the Steps window re-sizes accordingly.

Access and procedure selection. The Procedure View is launched from a "Procedure View" button in the lower-left corner of the Reactor Temperature Profile View. Once launched, the view appears in the upper-right corner of the right-hand monitor. The Steps and

Images windows are initially empty. Above the Steps window, a combo box pull-down menu can be accessed to select from the available procedures. “Select Procedure” should be written on the combo box. Once a procedure is selected, the Steps and Images windows are populated with the appropriate text and jpeg. An alternative procedure can be selected from the pull-down menu before the first step of the procedure is checked off. Once any step of the populated procedure has been ‘acknowledged’, the pull-down menu is disabled.

The control is closed by left-clicking on a “Close” button in the MOV window.

Mapping

No mapping information is required for the Procedure View because it contains no real-time data.

P+F INTERFACE

The display suite described in this Appendix contains information from both task-based and work-domain based analyses. We refer to this interface as the P+F+T version. One of the design objectives was to create a second interface that would allow us to assess the differential value of adding the task-based information. For that reason, we created a P+F version by removing graphics associated with information requirements that were identified only by the task-based analyses. Figure 74 shows the appearance of Monitor 2 in this version. The rationale behind this approach is given by Jamieson, Reising, and Hajdukiewicz (2001).

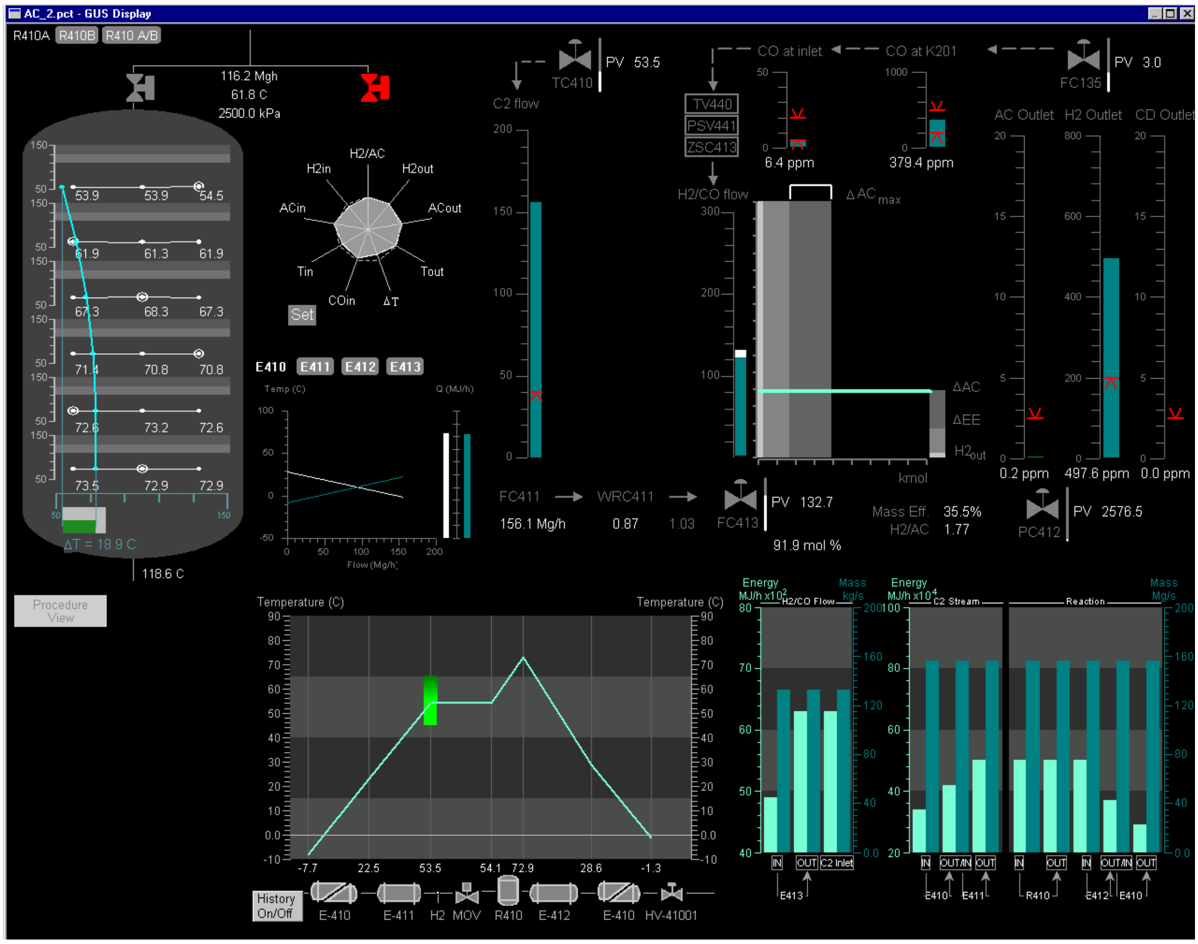

Figure 74: The P+F version of the AHR ecological interface.

The graphics removed from the displays are the following:

- Procedure View
- MOV View
- H2 Balance View: Weight Ratio Control graphic
- H2 Balance View: Analyzer update timers
- H2 Balance View: Setpoint and Output data on control valves

APPENDIX D: EXPERIMENTAL MATERIALS

Demographic Questionnaire

1. Age: _____
2. Do you have 20/20 vision (corrected or uncorrected)?: Y N
3. Do you have a color vision deficiency?: Y N
4. Unit (circle one): E1 E2 E3
5. Team (circle one): 1 2 3 4 5 Other:______
6. How many years of work experience do you have in the petrochemical industry?____
7. How long have you been working at NOVA? _____ years, _____ months
8. Please rate your level of expertise in using Microsoft Windows applications: 1=novice, 5=expert: _______
9. Please check the panel positions for which you have been qualified. For each, check whether you are current and for how long you worked in that job.

Unit	Qualified	Current	Time spent on panel (years, months)
E1 Cracking			
E1 Finishing			
E2 Cracking			
E2 Finishing			
E3 Cracking			
E3 Finishing			
Other: _______________			
Other: _______________			

Interface Training Assessment

1. Send a tag to the detail view.
2. Input a step change to any tag.
3. Trend a single variable.
4. Trend multiple variables.
5. Change the time scale on a trend plot.
6. Open/Close an MOV.
7. Explain what the polar star would look like if a CO spike hit the reactor.
8. How would you identify a temperature excursion in the reactor?
9. Explain why the mass and energy profiles are shaped as they are.
10. How long until the next acetylene analyzer update?
11. How would the Paulsen view appear during a loss of cooling water incident?
12. How would you detect an unexpected hydrogen source?
13. Name three places in the interface where you can get information about the acetylene output.
14. Locate and explain three tags that are not available in the current displays.
15. Cycle through the heat exchangers. What would fouling in a heat exchanger look like?
16. Call up the Procedure View and check off the first three steps of the Swing.

Post-Scenario Questionnaire

1. Please describe what happened in the scenario that you have just completed.

2. Was there a fault in the scenario? YES (go to 3) NO (go to 4)

3. If there was a fault, please indicate the root cause of that fault. Be as specific as you can.

4. Please describe what action(s) you took during the scenario.

For the following 6 questions, please refer to the display view key and refer to each view by its designation.

5. Which view was most useful to you in *monitoring* the plant? ________________
6. Which view was most useful to you in *diagnosing* the plant state? ________________
7. Which view was most useful to you for *planning* your control actions? ________________

8. Which view was least useful to you in *monitoring* the plant? ________________
9. Which view was least useful to you in *diagnosing* the plant state? ________________
10. Which view was least useful to you for *planning* your control actions? ________________

Post-Test Questionnaire

1. In your operating experience, have you encountered any of the three scenarios that you experienced today? (Check one):

   ___Yes *Please proceed to 2.a)*
   ___No *Please proceed to 3)*

   2.a) Identify the scenario(s) that you have seen before. ________________
   ____

2. In your training, have you practiced or planned for any of the three scenarios that you encountered today? If so, which one(s)?______________________________________
   ____

3. How effective or ineffective were the displays <u>that you used today</u> at supporting **procedure execution**? (Check one):

   ___Very Ineffective
   ___Somewhat Ineffective
   ___Somewhat Effective
   ___Very Effective

4. How effective or ineffective were the displays <u>that you used today</u> at supporting non-**procedure based activity**? (Check one):

   ___Very Ineffective
   ___Somewhat Ineffective
   ___Somewhat Effective
   ___Very Effective

5. How effective or ineffective are the displays <u>that you use(d) in E1 or E2</u> at supporting **procedure execution**? (Check one):

   ___Very Ineffective
   ___Somewhat Ineffective
   ___Somewhat Effective
   ___Very Effective

6. How effective or ineffective were the displays <u>that you use(d) in E1 or E2</u> at supporting **non-procedure based activity**? (Check one):

   ___Very Ineffective
   ___Somewhat Ineffective
   ___Somewhat Effective
   ___Very Effective

7. The current E1 or E2 operating displays could be improved. (Check one):

   ___Strongly Agree
   ___Agree
   ___Neither agree nor disagree
   ___Disagree
   ___Strongly disagree

8. I am open to changes in the current operating displays. (Check one):
    ___Strongly Agree
    ___Agree
    ___Neither agree nor disagree
    ___Disagree
    ___Strongly disagree

Questions 9 and 10 should only be answered by users of the new displays:

9. Please put the views in order with respect to their usefulness. Start from 1=most useful, and go to up to the least useful:
    ___ Detail
    ___ Emergency Procedure
    ___ Equipment Temperature Profile
    ___ Heat Exchanger
    ___ Hydrogen Balance
    ___ Mass & Energy
    ___ MOV
    ___ Polar Star
    ___ Reactor Temperature Profile
    ___ Equipment Schematic
    ___ Swing Procedure
    ___ Trend

10. What difficulties did you find in using the new displays? ______________________
    ____
    ____
    ____
    ____
    ____

APPENDIX E: INTERFACE COMMENTARY

Schematic View. Operators used this view more than any other; either because it was most similar to the schematic displays that they use in the current control system or because it was most useful for control Requested improvements to this view focused on better labeling and organization of the supplementary information. The molar ratio and efficiency data should be collocated with the ratio control and inlet analyzer data on the left hand side of the view. The dispersion of this data imposed an unnecessary visual search and memory task on the operators.

Detail View. The tendency of the control input fields to time out caused some confusion. When an operator is planning to make a move, he often opens a change window or activates a pull-down menu well in advance of his intended action execution. These fields in the Detail View time out after several seconds of inaction. The delay should be increased substantially (or removed) based on a task-analysis of control changes.

Trend View. Displaying only a single scale when several traces are on a trend forced the operators to interact with the display unnecessarily. Current trend applications display up to four separate scales, color-coded to match the traces to which they apply. The absence of this information induces a display manipulation task.

Overlaying the Detail Trend and Group Trend limited the benefits of the view yoking between the Detail View and the Trend Detail. When the Group Trend was in front, the operators were sometimes confused by not seeing traces for the parameters of the tag in the Detail View. The desirable solution would be to allocate each of the trends to a dedicated viewport (which is, in fact, how they are organized in E3).

The Unanticipated Normal event revealed a design problem with this view. The auto-scaling function does not work well when the controller PV is not tracking the SP. With the scaling based on the PV, the SP trace was sometimes not fully visible on the plot.

<u>Reactor Temperature Profile View.</u> This view was popular with the operators in both normal and abnormal situations. It appears to have successfully met the goal of providing the operator with a perceptual indication of the temperature profile function. The primary issue with this display is the limited range of the temperature scales. Wider ranges or auto-scaling are required to make this view into a viable application.

<u>Polar Star View.</u> The operators reported using the polar star for normal monitoring, diagnosis, and response assessment. They grasped the display concept quickly and were able to make use of it almost immediately. The biggest concern expressed about the view was the scaling on the axes. The general consensus was that the ranges should be tighter to make slight deviations during normal operations more salient.

It was unfortunate that the polar stars were obscured when the R410A&B tab was selected. This was usually the case during the reactor swing. Several operators suggested that having both polar stars visible would be very beneficial to the performance of this task. They wanted to have a polar star for the old reactor to monitor its performance during the swing. As well, they wanted a 'target' configuration for a polar star for the new reactor. This star would have its midpoints set at the target values for the reactor being fully on-line. In effect, they wanted to use the graphic to monitor their progress towards a multi-dimensional target state.

<u>Heat Exchanger View.</u> Despite the addition of information not available in the current process, this view was not put to much use, largely because of the absence of any time-based information in the view. Since both the form and some of the content were new, the operators

had no expectation for what a normal appearance of the graphic would be. Many of them commented that they might have been able to use the display if there were some sort of indication of what the display looked like in the recent past.

MOV View. The MOV view was unique in that many of the operators indicated during training that they did not see much use for it, and then proceeded to use it heavily during the reactor swing scenario. This was obvious from the sharp drop in communication events between the panel operator and the cohort acting as the field operator.

Several modifications to the display would improve its usefulness, however. First, the view can be hidden during most operations. The MOVs are only moved during reactor swings and during emergency situations (and even in those events, the MOV movement is of secondary importance). Second, during a reactor swing, it would be beneficial to have indications for both valves visible simultaneously. This would reduce the likelihood of confusion about which valve was being represented. It would also facilitate the strategy of moving both valves simultaneously, a strategy employed by a few of the operators.

Equipment Temperature Profile View. Feedback on the Equipment Temperature Profile varied considerably. Several operators liked the profiles and the ability to compare current and past conditions. These individuals tended to agree that the display would be even more useful if it included a longer train of process equipment. Even this suggestion did not improve the outlook of many of the other operators. Oddly, most of those who did not find the display useful could not explain why that was the case. Thus, other than changing the appearance of the target inlet temperature region, no insights for improving the display were generated.

<u>Mass & Energy View.</u> Operators did not find this view useful. This may be surprising to people familiar with the EID literature. Abstract Function level information (of which the View is comprised) has often been shown to be particularly useful new information for operators. The difference in this case is that it is not new information. Because the reactor section is part of a gas plant, process flows are already measured in mass units. In prior applications of EID, flows have been measured in volumetric flows (e.g., steam cycle systems). Thus, the mass information was both new and useful because it revealed information at a level where conservation laws could be applied. With the energy information being based on flow rate and temperature, this content essentially reflected little more than temperature change when the mass flow was held constant.

<u>H2 Balance.</u> The operators found little use for this view for three reasons. First, none of the scenarios exploited the hydrogen balance graphic, the part that they expected to be most useful. Second, the display is populated with data that appears in the schematic. Third, it is located in the upper right corner of Monitor 2, a location where operators say they are unaccustomed to looking for detailed information.

Several operators commented that they anticipated that the hydrogen balance graphic could be used for process tuning.

<u>Procedure View.</u> The Procedure View was rarely viewed outside of training. A few operators looked at it to review the swing procedure, but almost none of them opened it during a scenario. To a certain extent, this observation validated our decision to remove as many navigation requirements from the interface as possible. When faced with an array of new graphics, an operator is not eager to find something else to look at. One operator did suggest that the display would be very useful for training and preparing for a planned procedural execution.

Cognitive
Engineering
Laboratory

CEL TECHNICAL REPORT SERIES

CEL 93-01 "Egg-sucking, Mousetraps, and the Tower of Babel: Making Human Factors Guidance More Accessible to Designers"
• Kim J. Vicente, Catherine M. Burns, & William S. Pawlak

CEL 93-02 "Effects of Expertise on Reasoning Trajectories in an Abstraction Hierarchy: Fault Diagnosis in a Process Control System"
• Klaus Christoffersen, Alex Pereklita, & Kim J. Vicente

CEL 94-01 "Cognitive 'Dipsticks': Knowledge Elicitation Techniques for Cognitive Engineering Research"
• Klaus Christoffersen, Christopher N. Hunter, & Kim J. Vicente

CEL 94-02 "Muddling Through Wicked Problems: Exploring the Role of Human Factors Information in Design"
• Catherine M. Burns

CEL 94-03 "Cognitive Work Analysis for the DURESS II System"
• Kim J. Vicente & William S. Pawlak

CEL 94-04 "Inducing Effective Control Strategies Through Ecological Interface Design"
• William S. Pawlak

CEL 94-05 "Research on Factors Influencing Human Cognitive Behaviour (I)"
• Klaus Christoffersen, Christopher N. Hunter, & Kim J. Vicente

CEL 94-06 "Ecological Interfaces for Complex Industrial Plants"
• Nick Dinadis & Kim J. Vicente

CEL 94-07 "Evaluation of a Display Design Space: Transparent Layered User Interfaces"
• Beverly L. Harrison, Hiroshi Ishii, Kim J. Vicente, & Bill Buxton

CEL 94-08 "Designing and Evaluating Semi-Transparent 'Silk' User Interface Objects: Supporting Focused and Divided Attention"
• Beverly L. Harrison, Shumin Zhai, Kim J. Vicente, & Bill Buxton

CEL 95-01 "An Ecological Theory of Expertise Effects in Memory Recall"
• Kim J. Vicente & JoAnne H. Wang

CEL-SP "Strategic Plan"
• Cognitive Engineering Laboratory

CEL-LP "Cognitive Engineering Laboratory Profile"
• Cognitive Engineering Laboratory

CEL 95-04 "A Field Study of Operator Cognitive Monitoring at Pickering Nuclear Generating Station-B"
• Kim J. Vicente & Catherine M. Burns

CEL 95-05 "An Empirical Investigation of the Effects of Training and Interface Design on the Control of Complex Systems"
• Christopher N. Hunter

CEL 95-06 "Applying Human Factors to the Design of Medical Equipment: Patient-Controlled Analgesia"
• Laura Lin, Racquel Isla, Karine Doniz, Heather Harkness, Kim J. Vicente, & D. John Doyle

CEL 95-07 "An Experimental Evaluation of Transparent Menu Usage"
• Beverly L. Harrison & Kim J. Vicente

CEL 95-08 "Research on Factors Influencing Human Cognitive Behaviour (II)"
• Christopher N. Hunter, Michael E. Janzen, & Kim J. Vicente

CEL 95-09 "To the Beat of a Different Drummer: The Role of Individual Differences in Ecological Interface Design"
• Dianne Howie

CEL 95-10 "Emergent Features and Temporal Information: Shall the Twain Ever Meet?"
• JoAnne H. Wang

CEL 95-11 "Physical and Functional Displays in Process Supervision and Control"
• Catherine M. Burns & Kim J. Vicente

CEL 96-01 "Shaping Expertise Through Ecological Interface Design: Strategies, Metacognition, and Individual Differences"
• Dianne E. Howie

CEL 96-02 "Skill, Participation, and Competence: Implications of Ecological Interface Design for Working Life"
• Peter Benda, Giuseppe Cioffi, & Kim J. Vicente

CEL 96-03 "Practical Problem Solving in a Design Microworld: An Exploratory Study'
• Klaus Christoffersen

CEL 96-04 "Review of Alarm Systems for Nuclear Power Plants"
• Kim J. Vicente

CEL 96-05 "DURESS II User's Manual: A Thermal-hydraulic Process Simulator for Research and Teaching"
• Lisa C. Orchanian, Thomas P. Smahel, Dianne E. Howie, & Kim J. Vicente

CEL 96-06 "Research on Factors Influencing Human Cognitive Behaviour (III)"
• Dianne E. Howie, Michael E. Janzen, & Kim J. Vicente

CEL 96-07 "Application of Ecological Interface Design to Aviation"
• Nick Dinadis & Kim J. Vicente

CEL 96-08 "Distributed Cognition Demands a Second Metaphor for Cognitive Science"
• Kim J. Vicente

CEL 96-09 "An Experimental Evaluation of Functional Displays in Process Supervision and Control"
• Catherine M. Burns and Kim J. Vicente

CEL 96-10 "The Design and Evaluation of Transparent User Interfaces: From Theory to Practice"
• Beverly L. Harrison

CEL 97-01 "Cognitive Functioning of Control Room Operators: Final Phase"
• Kim J. Vicente, Randall J. Mumaw, & Emilie M. Roth

CEL 97-02 "Applying Human Factors Engineering to Medical Device Design: An Empirical Evaluation of Two Patient-Controlled Analgesia Machine Interfaces"
• Laura Lin

CEL 97-03 "ADAPT User's Manual: A Data Analysis Tool for Human Performance Evaluation in Dynamic Systems"
• Xinyao Yu, Farzad S. Khan, Elfreda Lau, Kim J. Vicente, & Michael W. Carter

CEL 97-04 "Research on the Characteristics of Long-Term Adaptation"
• Xinyao Yu, Renée Chow, Greg A. Jamieson, Rasha Khayat, Elfreda Lau, Gerard Torenvliet, Kim J. Vicente, & Michael W. Carter

CEL 98-01 "Applying Human Factors Engineering to Medical Device Design: An Empirical Evaluation of Patient-Controlled Analgesia Machine Interfaces"
• Laura Lin

CEL 98-02 "Building an Ecological Foundation for Experimental Psychology: Beyond the Lens Model and Direct Perception"
• Kim J. Vicente

CEL 98-03 "Cognitive Work Analysis: Toward Safe, Productive, and Healthy Computer-based Work"
• Kim J. Vicente

CEL 98-04 "Ecological Interface Design for Petrochemical Processing Applications"
• Greg. A. Jamieson & Kim J. Vicente

CEL 98-05 "The Effects of Spatial and Temporal Proximity of Means-end Information in Ecological Display Design for an Industrial Simulation"
• Catherine M. Burns

CEL 98-06 "Research on Characteristics of Long-term Adaptation (II)"
• Xinyao Yu, Gerard L. Torenvliet, & Kim J. Vicente

CEL 98-07 "Integrated Abstraction Hierarchy and Plan-Goal Graph Model for the DURESS II System: A Test Case for Unified System- and Task-based Modeling and Interface Design"
• Christopher A. Miller & Kim J. Vicente

CEL 98-08 "Comparative Analysis of Display Requirements Generated via Task-Based and Work Domain-based Analyses: A Test Case Using DURESS II"
• Christopher A. Miller & Kim J. Vicente

CEL 98-09 "Abstraction Decomposition Space Analysis for NOVA's E1 Acetylene Hydrogenation Reactor"
• Christopher A. Miller & Kim J. Vicente

CEL 99-01 "Development of an Analytical Framework and Measurement Tools to Assess Adaptive Performance of Individual and Teams in Military Work Domains"
• John R. Hajdukiewicz, Kim J. Vicente, & Robert G. Eggleston

CEL 99-02 "Applying Perceptual Control Theory and Ecological Interface Design to the Control Display Unit"
• Sandra Chéry

CEL 99-03 "Research on the Characteristics of Long-Term Adaptation (III)"
• Gerard L. Torenvliet & Kim J. Vicente

CEL 99-04 "Comparative Analysis of Display Requirements Generated via Task-Based and Work Domain-based Analyses in a Real World Domain: NOVA's Acetylene Hydrogenation Reactor"
• Christopher A. Miller & Kim J. Vicente

CEL 99-05 "A Cognitive Engineering Approach for Measuring Adaptive Behavior"
• John R. Hajdukiewicz & Kim J. Vicente

CEL 00-01 "Differences Between the Eye-fixation Patterns of Novice and Expert Operators of the DURESS II Physical Interface"
• Madhava Enros & Kim J. Vicente

CEL 00-02 "If Technology Doesn't Work for People, then It Doesn't Work"
• Kim J. Vicente