Integrating Task- and Work Domain-Based Work Analyses in Ecological Interface Design: A Process Control Case Study

Greg A. Jamieson, *Member, IEEE*, Christopher A. Miller, Wayne H. Ho, and Kim J. Vicente, *Senior Member, IEEE*

***Abstract*—In this paper, we present a case study wherein several work analysis methods were incorporated in the design of a graphical interface for a petrochemical production process. We follow this case from the application of the work analysis methods, through the consolidation of information requirements, to the design of a novel interface that integrates the requirements. The findings confirm earlier assertions that task-based and work domain-based analysis frameworks identify unique and complementary requirements for effective information systems that are intended to support supervisory control of complex systems. It further provides the first industrial demonstration of ecological interface forms based on integrated task- and work domain-based work requirements.**

***Index Terms*—Graphical user interfaces, man–machine systems, petroleum industry, process control.**

I. Introduction

JAMES GIBSON defined a display as "... a surface that has been shaped or processed so as to exhibit information for more than just the surface itself" [1, p. 42]. This definition prompts two questions for the display designer. First, what information is to be displayed? Second, how should the surface be shaped to exhibit that information [2], [3]? Thus, interface design (i.e., the design of both displays and controls) can be generically described as a two-stage process: 1) the specification of information content and structure and 2) the subsequent semantic mapping of those requirements onto visual forms. Much of the human–computer interaction literature can be viewed as attempts to refine and answer both of these questions.

The basis for human-centered approaches to the specification of interface content is work analysis. Although methods for conducting work analysis have multiplied and evolved, direct comparisons of the products of these methods are only a very recent addition to the literature [4], [5]. Moreover, the challenge of *integrating* the products of different methods in the design of interfaces is almost entirely unexplored. This latter challenge is of particular concern to Ecological Interface Design (EID) [2], which has been almost invariably coupled with a single work analysis method. The continued development of this design framework depends largely on its ability to incorporate the results of other work analysis methods.

This paper builds on the recent comparative evaluations of work analysis methods by providing a case study of the process of conducting several alternative work analyses for a complex system and integrating the results for the purpose of interface design. Contributions are made in two key areas: First, we continue to explore the academic problem of comparing work analysis methods by focusing on the challenges of integrating the results of several alternative analyses. Second, we provide new insight into the practical challenges and benefits of designing graphical interfaces that communicate the results of the analysis efforts.

A. Two Approaches to Content Specification

Work analysis methods for complex systems have typically fallen into two general categories: task-based and work domain-based[1] approaches [4]. Task-based approaches reflect the prevalent perspective in human factors engineering. They include methods, such as Hierarchical Task Analysis (HTA) [6] and timeline analysis techniques [7]. Work domain-based analysis approaches are less common, although they have enjoyed increasing attention in the past decade [8]–[10]. These include the Abstraction–Decomposition Space (ADS) [11] and the viable systems model [12].

Choosing between task- and work domain-based methods involves a tradeoff between efficiency and robustness [9]. The goal of both forms of analysis is to provide an understanding of the ways in which known goals can be achieved in various contexts of use, so that interfaces or training materials can be designed to facilitate efficient, safe, and productive work in the domain. However, the two methods focus on different aspects of the domain. Task analyses focus on the tasks, activities, or actions that must be performed to meet higher level goals in specific contexts [13]. By focusing on the actions known to be critical to successful task completion, given existing or projected equipment and procedures, the process of analysis

Manuscript received September 24, 2005; revised May 26, 2006. This work was supported by cash and in-kind contributions from the Natural Sciences and Engineering Research Council, NOVA Chemicals, Ltd., and Honeywell, Inc. This paper was recommended by Associate Editor S. Guerlain.

G. A. Jamieson and K. J. Vicente are with the Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, ON M5S 3G8, Canada (e-mail: jamieson@mie.utoronto.ca).

C. A. Miller is with Smart Information Flow Technologies, Minneapolis, MN 55401 USA.

W. H. Ho is with IBM Canada Ltd., East Markham, ON L3R 9Z7, Canada.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TSMCA.2007.904736

[1]Elsewhere, we have used the term "system-based" approaches. We opt to use "work domain-based" here to be consistent with the more broadly employed use of the terms. They should be considered to be equivalent.

TABLE I
RELATIVE ADVANTAGES AND DISADVANTAGES OF TASK-BASED AND WORK-BASED DOMAIN FORMS OF WORK ANALYSIS (AFTER [4])

Work Domain Analysis	Task Analysis
Provides explicit knowledge about domain constraints and capabilities, but may fail to exploit procedural knowledge	Provides compiled procedural knowledge, but may hide the rationale behind procedures
Supports adaptation to unforeseen contingencies and recovery from errors	Supports efficient execution of procedures developed for foreseen contingencies
Identifies information requirements more readily and directly	Identifies priority, procedural, and temporal constrains more readily
Is independent of context of use	Is less independent of context of use

and design captures acquired knowledge about how to operate a facility. This approach is efficient, because it capitalizes on existing knowledge about specific behaviors known to produce good results in at least some circumstances. By embedding this acquired knowledge in information artifacts, the analyst conveys that knowledge to operators and facilitates their use of it in response to those events.

In contrast, the work domain-based methods focus on the overall environment or "ecology" in which the work is performed and the goals are achieved [9]. Knowledge is captured, not about specific methods known to produce good results in some circumstances but rather about the innate capabilities of the plant or process *in all circumstances*. The analysis is conducted independently of specific events, tasks, activities, and actions that might be relevant. The advantage of this approach is that the results of the analysis are more robust, i.e., they are applicable to all events, even those that the analyst is unable to anticipate [2]. Artifacts based on the results convey this knowledge to operators—providing greater support for the full range of circumstances that the operator might encounter, but not necessarily highlighting or providing optimal support for known or common methods of using the system. Comparatively speaking, task-based approaches trade off robustness in favor of efficiency, whereas work domain-based methods trade off efficiency in favor of robustness [14].

B. Explicit Comparisons of Task- and Work Domain-Based Techniques

Despite the importance of work analysis in human factors engineering and despite the many comparisons of the merits and deficiencies of various approaches, the literature offers very few direct comparisons of those approaches [4]. In a notable exception, Burns and Vicente compared the abstraction hierarchy, multilevel flow modeling, and the decision ladder by applying each to a generic power plant boiler [5]. They classified abstraction hierarchy and multilevel flow modeling as work domain models but emphasized that each focuses on a different object of analysis and each is constructed using different relations between model elements. The abstraction hierarchy is a model of *work domain structure* and is based on *structural means—ends relations*, whereas multilevel flow modeling models *work domain goals* based on *goal-achievement relations*. Similarly, the decision ladder can be characterized as a *task model* employing *action means–ends relations*.

Burns and Vicente applied each of these modeling tools to the example process and summarized the distinctions along the dimensions of: 1) the type of model structure; 2) the difference between goals and purposes; and 3) the types of information produced (in addition to those distinctions previously noted) [5]. This sort of direct comparison is particularly valuable, as it serves to make sense of distinctions between the models that have proven to be difficult for some. In understanding the different types of information produced by each model, analysts can choose the tool that most appropriately fits their problem [5].

Miller and Vicente set out to determine whether, how, and to what advantage, task- and work domain-based modeling techniques could be integrated [4]. They compared an HTA of the DURESS II thermal hydraulic process simulator with an ADS[2] model that was previously described in the literature [15]. Miller and Vicente compared the *information requirements* (IRs) identified by each technique [4]. An IR[3] is a statement of a specific piece of information or relationship that should be incorporated into the suite of information artifacts intended to support effective control of a target process. In the case of DURESS II, the general conclusion was that the two analysis techniques each produced unique IRs that had complementary strengths and weaknesses. A summary comparison of the two frameworks is presented in Table I.

C. Outline for Paper

Two key limitations of the preceding comparisons motivated the work described here. First, both comparisons were conducted using simple process models. Second, in neither case were the results of modeling efforts used to develop operator interfaces that integrated the results of the two analyses. Thus, it remains to be demonstrated that the results of task and work-domain models of a representative industrial process can be effectively translated into an integrated interface that is both understandable to process operators and compatible with their sensory and cognitive limitations [16], [17]. In the balance of this paper, we address the *understandability* criterion by discussing the design of a set of representational forms based

[2]In this paper, we employ the term "Abstraction–Decomposition Space" over the less comprehensive (but more common) term "Abstraction Hierarchy." Abstraction Hierarchy refers to one of the two dimensions of hierarchy that make up the Abstraction–Decomposition Space, with the other dimension being the Aggregation–Decomposition Hierarchy.

[3]The phrase "information requirements" may provoke some confusion. After all, as we will see here, it is rare for all IRs to be included in an interface design—even one that has been used to operate a plant safely and efficiently for years. The sense in which all IRs identified by these analytic processes are *requirements* is dependent on what is considered *effective* control. As these expectations increase, so must the information coverage.

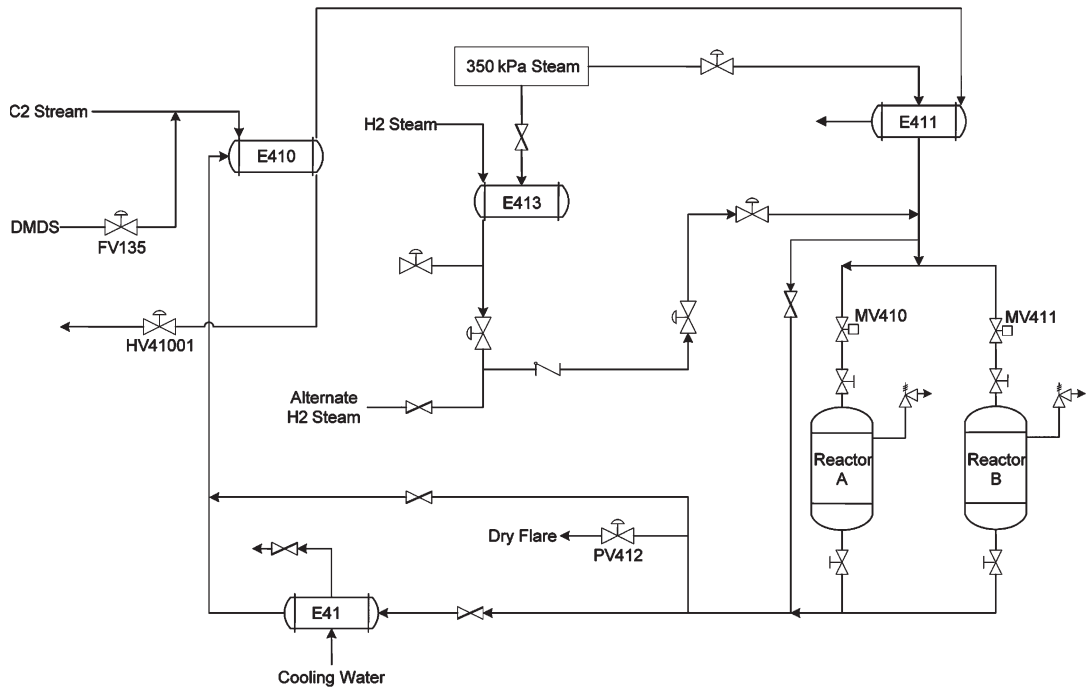

Fig. 1. Acetylene hydrogenation reactor.

on the IRs generated by ADS, HTA, and decision ladder models of an industrial petrochemical process. In an accompanying paper, we explore the *compatibility* and *effectiveness* criteria by reporting on an empirical evaluation of those forms in the hands of professional operators [18].

We begin with a discussion of three work analyses that were conducted on a representative industrial manufacturing process. Next, we consider the differences in the IRs that were derived from each of those analyses and show how they can be consolidated to inform design. This paper culminates in the presentation of an ecological interface view based on the requirements from a combination of work domain-based and task-based approaches.

II. Method

A. Test Bed

The target system for the work analysis is an acetylene hydrogenation reactor (commonly referred to as the "reactor"), which is a system within an ethylene refining process. The reactor is a moderately complex domain by contemporary chemical engineering standards. It is, however, arguably the most critical part of this particular ethylene refining process. Upsets in the reactor are the most frequent cause of costly upsets in the refinery.

The ethylene refinery converts a natural gas stream, consisting primarily of methane (CH_4), 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 subsystems treat the process flow and extract impurities. The reactor is located in the latter half (or "finishing end") of the refining process. It receives a stream consisting of ethane, ethylene, and various trace elements including acetylene. This is referred to as the "C2 stream" since single carbon molecules (e.g., CH_4) have been removed. Subsystems downstream of the reactor separate the ethane and ethylene from the trace elements, but these processes are particularly sensitive to acetylene in the process flow. The purpose of the reactor is to remove this acetylene.

To accomplish this, the C2 stream is reacted with hydrogen to convert the acetylene into ethylene or ethane (C_2H_6) in a process called hydrogenation. As illustrated in Fig. 1, the C2 stream enters the reactor section at E410, which is a cross-flow heat exchanger that recovers heat from the reactor outlet flow to preheat incoming feed. The C2 stream then proceeds to a second heat exchanger (E411) to raise its temperature to the level required for the reaction to occur. Immediately after E411, the C2 stream meets the hydrogen stream, which has been preheated in heat exchanger E413. The combined flow passes through either the MV410 or MV411 valve into a reactor vessel where it contacts a solid catalyst bed and reacts. Reactor effluent is cooled in heat exchanger E412 before passing through E410 to preheat the incoming feed. Further details on the control objectives of the reactor are introduced below.

B. How Analyses Were Conducted

The work domain analysis and the first of two task analyses were performed by just one of the authors. The industry sponsor provided him with extensive documentation of the reactor process, including piping and instrumentation diagrams, equipment specifications, operating manuals, and the plant procedures. He also visited the work site on several occasions, conducting unstructured interviews with process engineers and procedure writers, and task observations of operators. It is worth noting that this author is a psychologist, untrained in the supporting knowledge of process engineering, chemistry, and control theory required to understand the domain.

In the course of the research project, this author completed the ADS and HTA analyses and passed on detailed technical reports to two of the other authors. The initial plan was for this pair of investigators to translate the IRs identified in the technical reports directly into the interfaces. Both of them had prior experience with the analysis frameworks, and one had prior work analysis experience in petrochemical processing. In addition, both of these authors were familiar with the design framework that had been selected for use in the project. However, for the reasons discussed below, these authors were not able to initiate the design process immediately (see Section III-D1). Instead, they repeated much of the previous analysis work and introduced the second task analysis before embarking on the design.

While we cannot be sure of the organizational processes that were employed by other researchers and practitioners (such details are rarely reported in the literature), our experience in this case differed markedly from our other work analysis efforts, where a single analyst (or analysts) conducted both the analysis and design phases of the project. Several implications that emerged from this experience are discussed below (see Section III-D1).

III. Work Analyses

In this section, we describe the process for identifying IRs for control of the acetylene hydrogenation reactor using one work domain-based and two task-based work analysis tools. For each of the analyses, we provide an overview of the modeling effort and highlight key insights gained from completing it. We do not provide detailed descriptions of the analysis frameworks themselves. Such descriptions are readily accessible in the literature cited for each framework. We also do not provide a complete recounting of the results of the analyses. Complete details of the models are contained in technical reports [19]–[21].

A. Work Domain-Based Analysis: ADS

A work domain-based analysis of the reactor was conducted using the ADS framework [9]. Fig. 2 provides an overview of the five levels of abstraction (shown on the vertical axis) and three levels of decomposition (shown on the horizontal axis) that we used in this analysis. Consistent with many previous applications of the ADS framework,[4] models were constructed for only a subset of the cells (indicated by crosses in the appropriate cells) of the matrix, and these cells generally fall along the diagonal from top left to bottom right. The dashed cross in the Physical Form–Component cell indicates that, although a model was completed for this cell, it was ultimately not factored into the design of the information system because the object of analysis shifted from the actual physical process to a simulation of it (as discussed in [18]).

In the following paragraphs, we briefly describe the models included in the ADS.

1) Decomposition Analysis: Fig. 3 shows the part-whole decomposition relationships in graphical form. Fifty-one components have been aggregated into nine subsystems, which make up the overall reactor system. It is worth noting that, in the scope of the entire ethylene refining unit, the reactor itself is considered a subsystem. Thus, the terminology depends on the granularity of analysis.

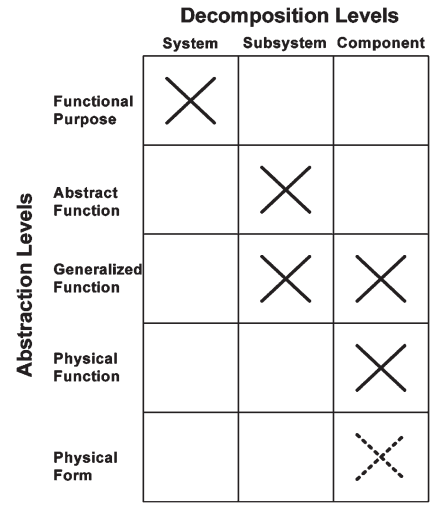

Fig. 2. Overview of ADS.

2) Physical Function—Component: The Physical Function model for the components of the reactor is shown in Fig. 4. As in most prior applications of the ADS to physical systems, the Physical Function model resembles a piping and instrumentation diagram with the instrumentation and automated controllers removed. The elements represent physical components in the reactor, and the topological links denote physical connections between those elements without reflecting actual dimensions, distances, or orientation.

The model in this cell highlights an unusual system boundary characteristic of the reactor. We chose to include the pyrolysis furnaces and the addition of DiMethyl DiSulfide (DMDS) as a part of the reactor system, even though they are temporally and spatially distant from the other reactor components. We made this choice because the proportion of CO in the H_2 stream is critically influenced by the performance of the furnaces. The reactor operator can influence CO concentration by controlling the amount of DMDS that is added in the furnaces from a single common header. While this makes drawing boundaries around our "system" in geographical terms conceptually problematic, it is meaningful in functional terms. This modeling decision reflects the fact that, in highly coupled systems, connections between components and functions are not always exclusively determined by physical proximity.

3) Generalized Function—Component: Fig. 5 shows the Generalized Function model for the reactor components. The boxes represent the process functions of the equipment components with the topological connections between them denoting causal relationships [9], [10].

This representation repeats two innovations introduced in the first application of the ADS framework to a petrochemical system [22]. The first is that flows of material are differentiated according to type. While other ADS analyses have generally not been concerned with the composition of different elements within a flow, those differences are critical in many

[4]See [9] or [15] for discussions.

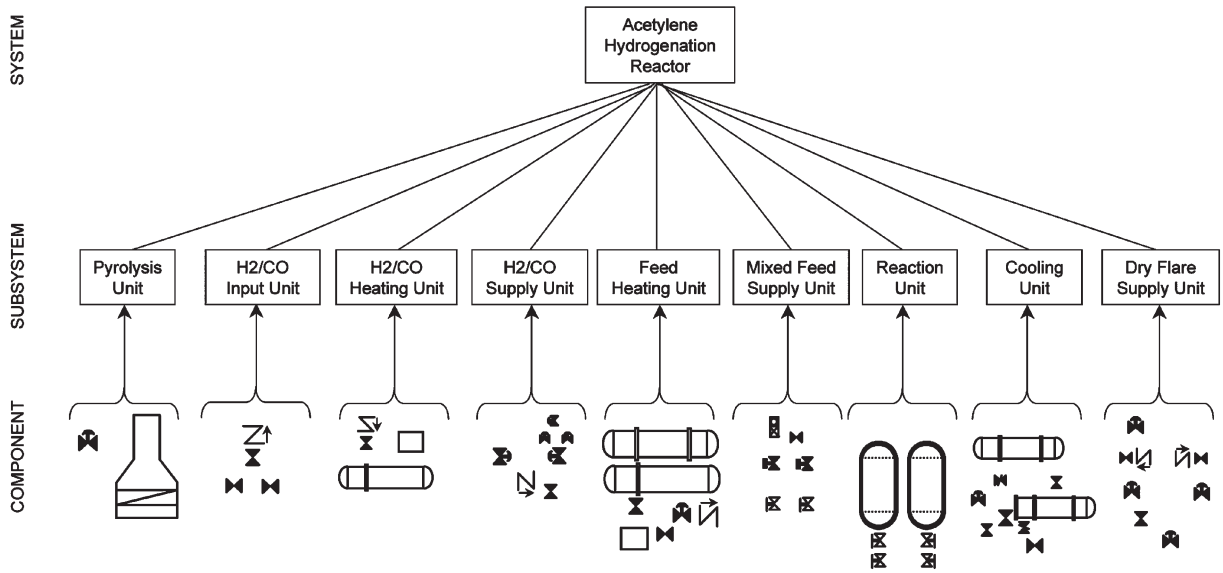

Fig. 3. Part-whole decomposition of the reactor.

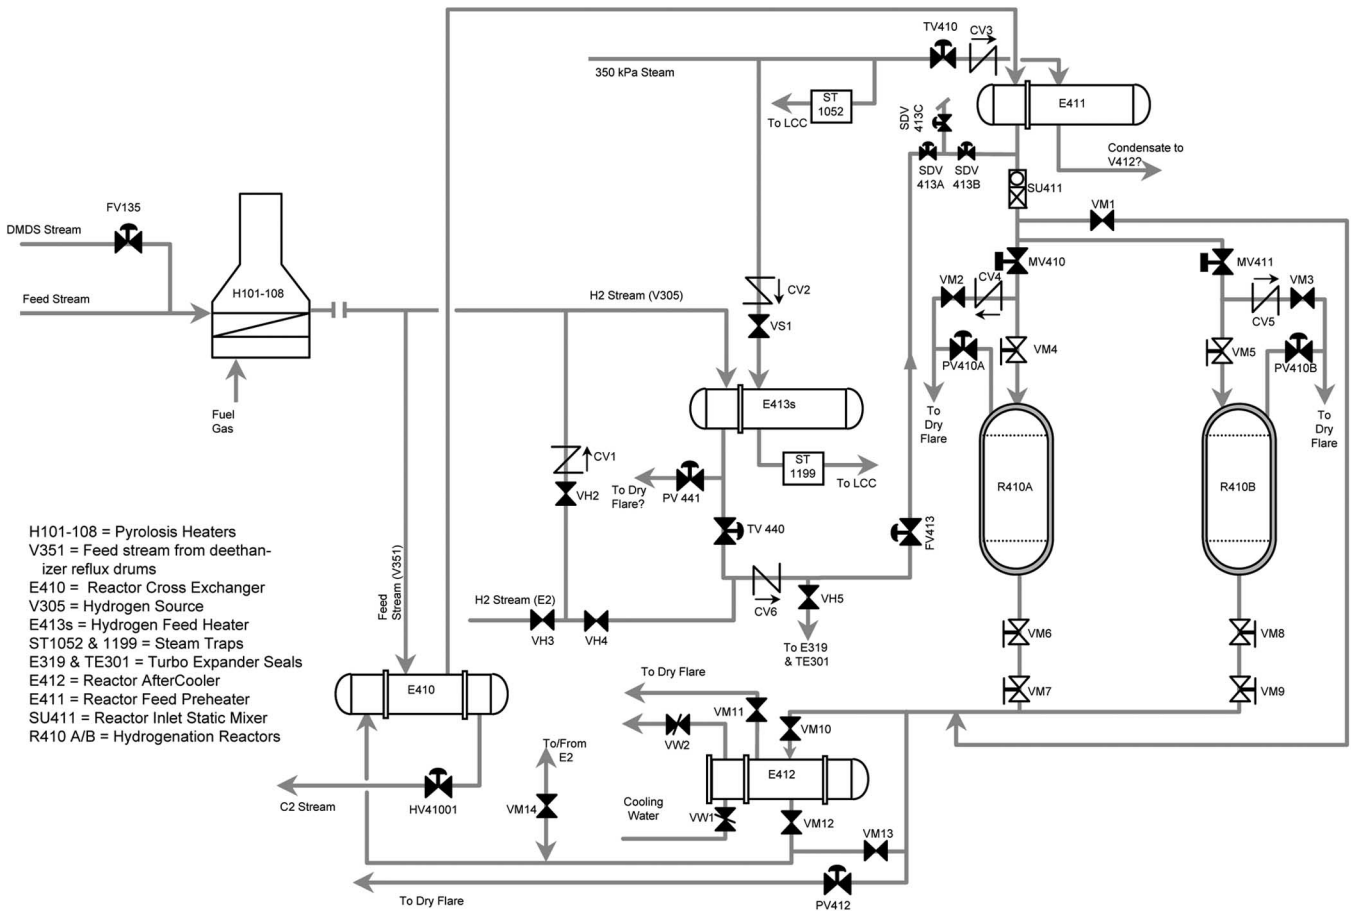

Fig. 4. Physical Function–Component model of the ADS.

petrochemical processing applications including the reactor. Thus, pyrolysis produces five different types of flows (C_2H_2, C_2H_4, C_2H_6, CO, and H_2), which are of interest to the reactor operator. These are separated into two streams and enter the reactor by two different routes. The C2 products are intermingled in the same pipeline flow; thus, they are shown here as going through the same heat transfer processes (Heat xfer 1 and 2). The H_2 and CO are mixed in a second flow (i.e., the H_2/CO stream) that enters the reactor via a different route. It is therefore shown undergoing a different heat transfer function (Heat xfer 3) before the two streams are mixed and sent together to the reactor functions.

The second innovation repeated here is the modeling of chemical reactions themselves. The most natural representation

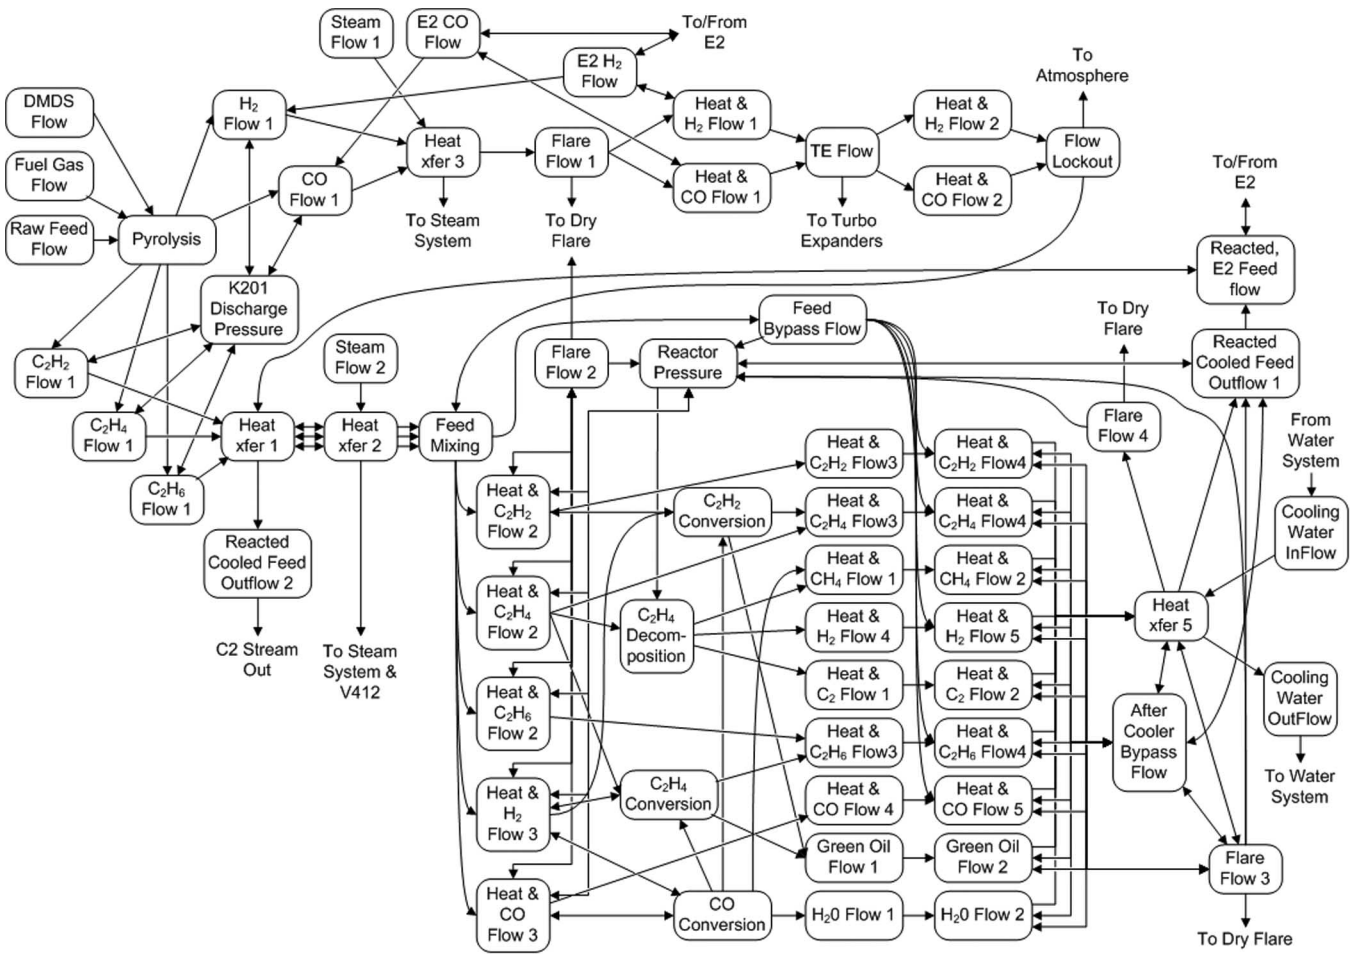

Fig. 5. Generalized Function–Component model of the ADS.

for chemical reactions is to place them at the Generalized Function level of decomposition. They are, after all, a "general function" of the plant, and although they are tied to specific components (such as the reactors themselves), they are not uniquely identified with physical equipment—thus, it seems difficult to class them as Physical Functions. They are subject to the mass and energy first principles represented at the Abstract Function level, but to collapse them into generalized representations of mass and energy exchanges would be to eliminate most of the critical information about which components react with which and how—information that is critical to understanding the reactor. Thus, we have represented the four reactions possible in the reactor vessels. While they are difficult to include in the figure, the actual chemical equations and other knowledge about the molecular ratios of reactants and the pressure, heat, and catalyst conditions required for each reaction are explicitly included in the analysis at this level. Among this additional information is the critical fact that all four reactions are exothermic, although at different rates. This fact is partially implied in Fig. 5 by the fact that heat, in addition to the reaction products, flows out of the reactors.

4) Generalized Function—Subsystem: Fig. 6 represents the first step along the decomposition axis of Fig. 3. In this model, we have moved up a level of aggregation and are now showing relationships between subsystem entities. However, because this is still a model at the Generalized Function level of abstraction, the same types of functions described in Fig. 5 are described here.

5) Abstract Function—Subsystem: Abstract Functions in process control domains are generally concerned with mass and energy relationships in the plant, and the reactor is no exception. Six key functions serve to describe fundamental processes of mass and energy: Source, Sink, Balance, Barrier, Storage, and Transport. These functions are denoted by boxes, with the connections between the boxes reflecting flows of mass or energy [9], [10]. Again, following precedent, separate models for the mass and energy relationships were constructed. The model for mass relationships is shown in Fig. 7. The model for energy relationships is nearly identical, with the primary exception being that the heat exchangers (all shell-and-tube type) allow movement of energy between process flows (i.e., balances) while simultaneously acting as barriers to the interaction of the two mass flows.

6) Functional Purpose—System: The final cell of the ADS to be modeled in our analysis is the Functional Purpose abstraction of the full reactor system. This "model" is a very simple statement of the purpose of the reactor as a whole: Reduce the concentration of acetylene in the process flow to less than 5 ppm. This purpose serves both safety and production objectives that would be distinguished in an analysis of the ethylene

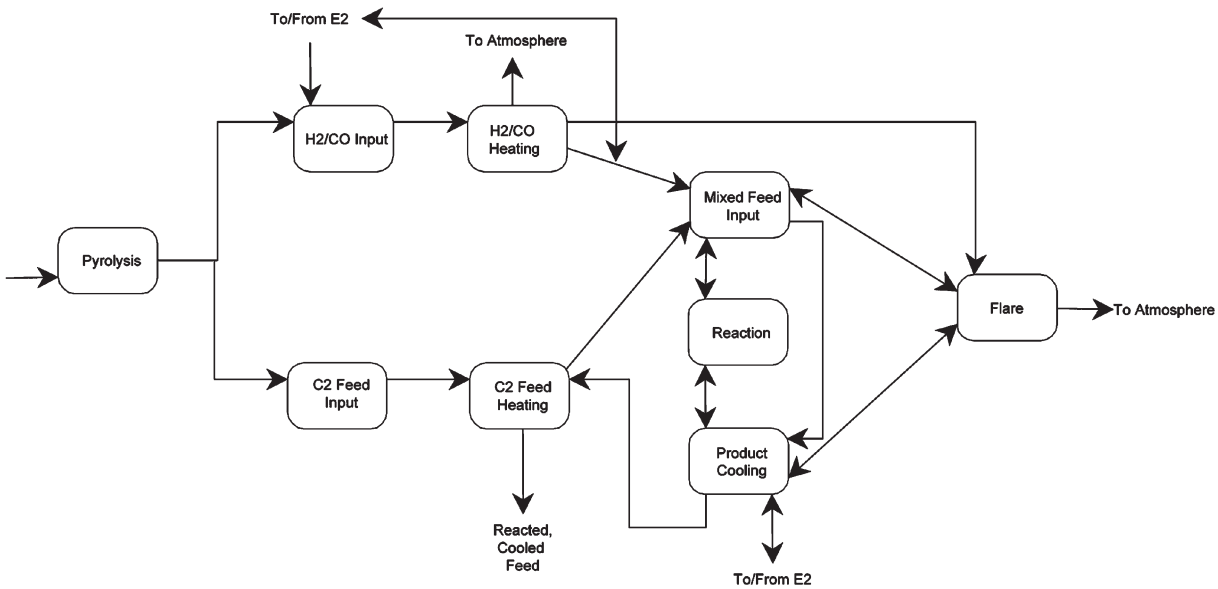

Fig. 6. Generalized Function–Subsystem model of the ADS.

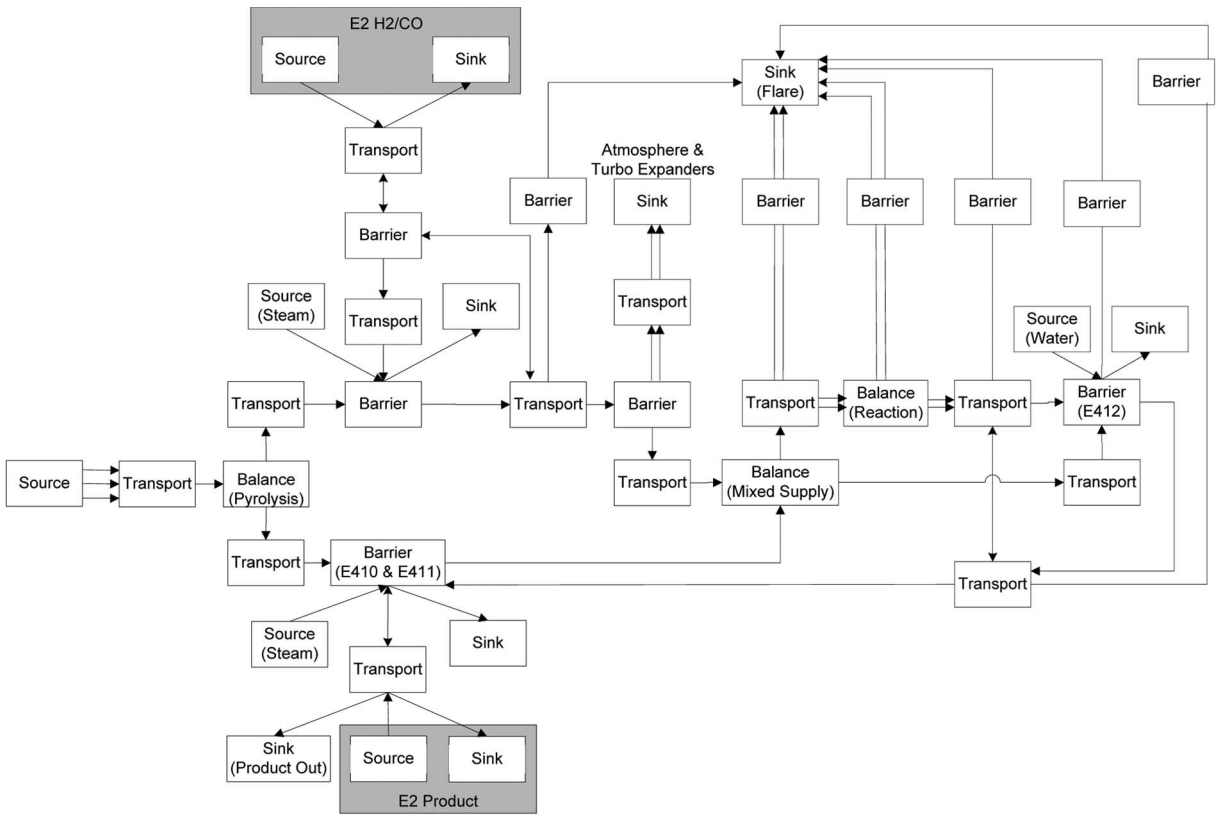

Fig. 7. Abstract Function–Subsystem ADS model of mass relationships.

refinery as a whole. Acetylene in concentrations greater than 5 ppm may destabilize downstream reactions (i.e., it is a safety concern), and it will degrade the quality of the ethylene produced by the plant (i.e., it is a production concern).

B. Task-Based Analysis: HTA of Procedures

Having completed the work domain analysis of the reactor, we next applied HTA to develop a task model for the purposes of comparison. HTA is a simple informal task analysis method, yet one that can be readily extended to capture and organize IRs. Despite its simplicity, HTA contains most of the characteristics of even the most complex task modeling approaches. HTA also has the advantage of being widely known and used in the task analysis community.

We analyzed four primary control tasks: start-up, normal operations, shut down, and fault management [20]. These are very common task distinctions in industrial process domains

Task	Plan	Timing	Actors	IRs
	5. Adjust DMDS with FC-135	seq	CRO	DMDS flow rate and control. Valve FV135 position and flow. Knowledge of previous or typical DMDS flow setting '3-5' is typical. CO ppm vs. threshold of 500 ppm would be nice to know.
	Plan 1.5.1.5 Adjust DMDS with FC-135			
	Plan 1.5.1.5: If CO> ~500ppm, do 1. Else, do 2.			
	1. Adjust FV135 open		CRO	position and flow through FV135, CO ppm.
	2. Adjust FV135 closed		CRO	position and flow through FV135, CO ppm.

Fig. 8. Example of the tabular format of the HTA.

and comprise a majority of the control activity performed on the reactor. In an effort to control the scope of our study, we excluded the catalyst regeneration task, which is a maintenance task performed periodically on a reactor vessel that has been removed from service.

1) Purpose and Sources: The purpose for which one performs a task analysis can have a profound impact on the types of information collected [13]. Our primary purpose in this exercise was to provide knowledge about how an interface to support the tasks should be designed. A task analysis for producing design requirements places more emphasis on identifying the information needs of users performing the tasks in the analysis but less emphasis on ensuring that the tasks are decomposed to a fine enough level to, for example, train a novice.

The primary source of data for the HTA was the written procedures for operation of the reactor. Secondary information about reactor operation was also provided by procedure writers, plant engineers, and operations personnel.

2) Scope of Control: In performing the HTA, we posited a hypothetical reactor operator and performed the analysis from his/her perspective with the goal of identifying requirements for displays that he/she might use. In practice, a single control room operator has responsibility for the whole "finishing end" of the refining process. This emphasis on the "reactor operator" required us to select (sometimes quite small) portions of procedures that were pertinent to the operation of the reactor. In practice, a control room operator who is responsible for the finishing end would be monitoring several subsystems, the status of any of which might inform him or her about the need to perform an action on the reactor. For our purposes, this simply took the form of a required communication about the status of another unit or about the timing for an action on the reactor.

Similarly, the control room operator works with several field operators who perform jobs that cannot be done from the control room, such as adjusting hand valves, inspecting for leaks, reading field instruments, etc. We generally avoided detailed expansion of the information needs of field operators and concentrated on the control room operator's needs. These include communications from the field operators about the status of equipment or the progress of field actions.

3) Analysis Results: A small portion of the results of the HTA for the reactor are presented in two different formats (after Shepherd [6]). Fig. 8 presents an example of the tabular form. While it is harder to visualize task relationships in this format, it is easier to link additional information to tasks. We have included three additional columns of information beyond the task relationships themselves. The first, labeled "Timing," contains information about the tasks' sequencing. The second column, labeled "Actors," contains information about the personnel, by role, who will be performing this task. The most common roles in these procedures are "CRO" (for control room operator) and "FO" (for field operator). The final additional column, labeled "IRs," contains IRs that were identified for each task. We discuss these requirements in a later section.

Fig. 9 presents the graphical form of the same tasks that were represented in Fig. 8. The graphical form emphasizes the "layout" of the tasks—their hierarchical and aggregate relationships. In this format, each layer of the hierarchy represents a series of tasks or actions that accomplish the higher level ("parent") task in some fashion. A "Plan" is always placed along the vertical line connecting the child tasks to their parent to show how/when/in what order they must be performed to accomplish their parent task. The plan is where information about the parallel or sequential relationships among the tasks and their initiation and completion conditions is located.

Since the complete HTA is far too large to include in this paper, only these two examples are provided here. In contrast to the ADS application, which required some adaptation of the modeling framework to the characteristics of the domain, the standard HTA representation was sufficient to describe the reactor operator tasks.

C. Task-Based Analysis: Control Task Analysis (CtrlTA) for Automation

As a final step in our modeling effort, we extended the task-based analysis to address the behavior of the automated controllers in the process. Understanding the intended behavior of these automated controllers is critical to both understanding normal plant operations under automatic control and to

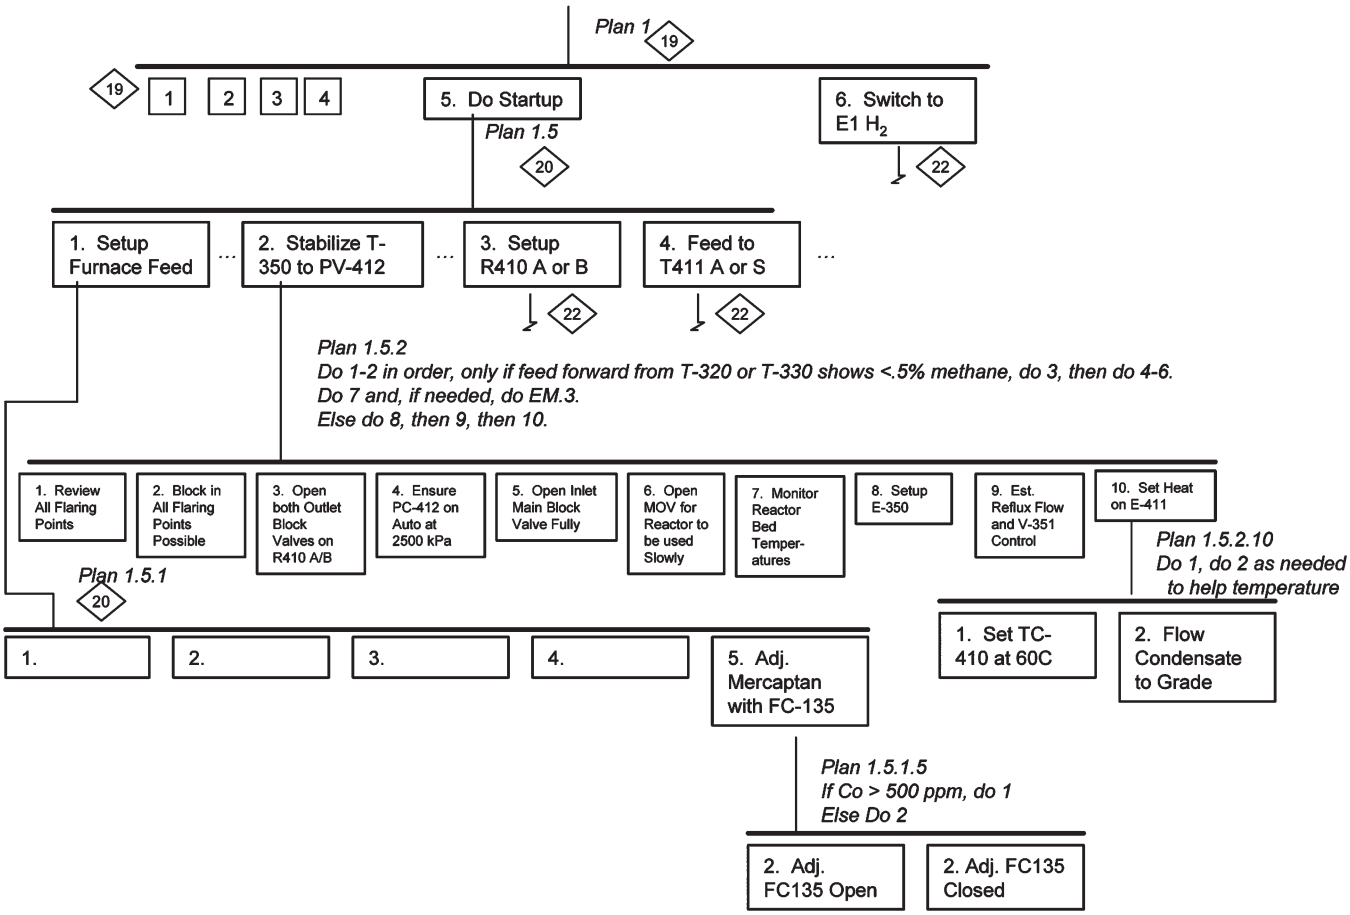

Fig. 9. Example of the graphical format for the HTA (empty boxes were expanded on other pages).

detecting and diagnosing failures. The HTA had not included the controllers because the procedures assume that the behavior of the controllers is a) known to the operator and b) completely reliable. The ADS also had not captured the activity of these controllers because a work domain analysis examines the object of control (i.e., the system) and not the control itself [14] (however, see [23] and [24]). Thus, the critical role of the automation had not been translated into requirements for the novel interface that was to be designed.

1) Decision Ladder Modeling: To fill in this gap, a CtrlTA was conducted to supplement the HTA. Also a task-based approach, CtrlTA identifies the requirements associated with known modes of operation (in this case, the control actions). The Decision Ladder framework was used to perform the analysis [25]. The decision ladder is a nonlinear sequence of information processing steps and intermediate states of knowledge. The framework is non-linear because the steps do not have to be completed in a canonical sequence. Rather, an actor (a generic term that incorporates the operator or the automation) can take advantage of shortcuts in association between nonsequential steps. 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 steps associated with interpreting and evaluating the signal can be bypassed.

We developed decision ladders for seven automatic controllers in the reactor, including one temperature controller, two ratio controllers, and four safety systems (low H_2 flow, low/high H_2 temperature, high reactor temperature, and high pressure). No automatic controllers were excluded from the analysis. The Weight Ratio Controller (WRC) was the most complex controller that was modeled.

2) Weight Ratio Control: The WRC serves as an example of the decision ladder concept. The WRC receives flow rate and chemical composition inputs from the C2 stream, calculates a hydrogen flow rate to match a setpoint ratio (H2[5] : C2), and outputs a control signal to a valve on the hydrogen stream (FC413) to achieve this ratio. Fig. 10 summarizes the Decision Ladder for the WRC, using the convention that boxes are information processing stages and circles are states of knowledge that serve as input and output. The task is cyclical, with changes in Target State acting as an input to the cycle. The Target State defines the desired ratio of the weight of flow of the hydrogen stream 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 the H2/CO flow. These data form a Set of Observations. The Identify process calculates a target flow rate for the mixed H2/CO stream, compares that to the actual flow rate, and uses the error to determine the necessary change to the output of the

[5] H_2 is referred to as "H2" in the reactor's legacy plant information system. For consistency, the following description maintains that practice.

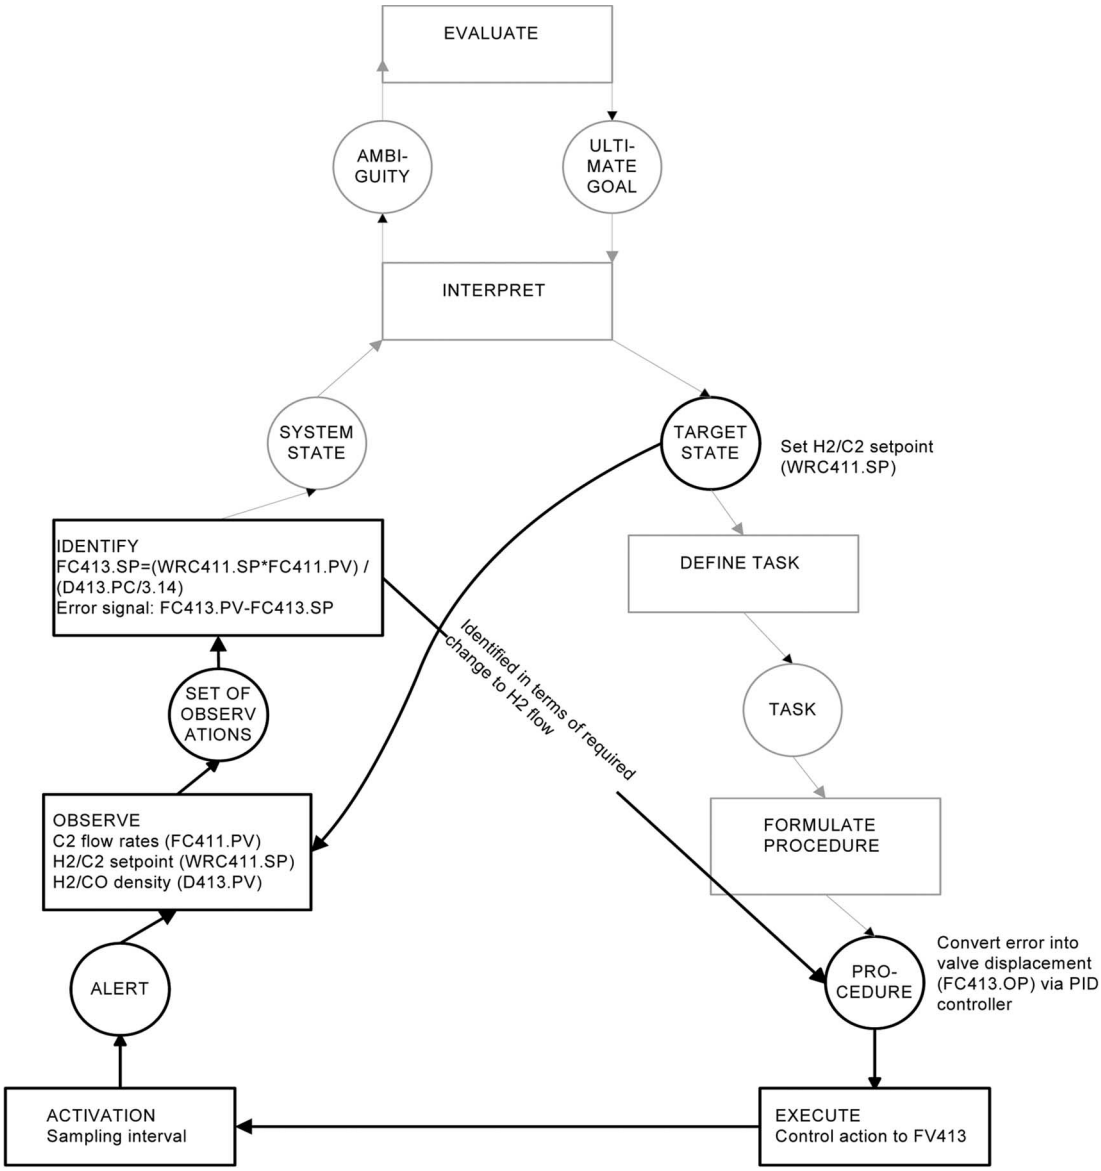

Fig. 10. Decision ladder for the Control Task Analysis of the weight ratio controller (adapted from [25]).

control valve on the H2/CO stream. Once the target state has been identified, the control proceeds directly to a predefined Procedure because the Define Task and Formulate Procedure steps are built into the automatic controller. 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), which is a value calculated by a proportional–integral–derivative control algorithm. The next step in the task is to return to the Activation process, signaling the controller to continuously repeat sampling cycles until a new target is entered.

D. Discussion

1) Lessons Learned From the Work Domain-Based Modeling: The chief lesson learned from the work domain-based modeling effort was that the ADS can be effectively applied to an operational industrial petrochemical process. Since this work is one of only three projects to date that apply the ADS to petrochemical processing operations [22], [26], this finding is important and encouraging for users in similar domains. Such domains pose challenges rarely encountered in past applications of the ADS (such as chemical reactions, geographically distributed plant operations, and complex temporal interactions). For example, the ability to represent multiple possible reactions as Generalized Functions and to integrate them into the overall system framework via the ADS—showing their means at the Physical Function level and their ends at the Abstract Function and Functional Purpose levels, and aggregating them in a more holistic "reaction" subsystem along the part-whole dimension—is an important, but natural, addition to the ADS framework. The reactor case study presented catalyzed reactions that are substantially more complex than those treated in a previous application of the ADS framework in this domain [22]. This is a valuable addition to the repertoire of ADS techniques, extending the range of chemical, petrochemical, and even biological realms where the ADS can be applied.

This paper also serves as an illustration of the utility of performing the ADS as a knowledge acquisition process. We

found that the structure of the ADS framework itself acted as a guide to acquiring the knowledge necessary to understand the operation of the reactor. The framework helped to direct the search for deep knowledge about the work domain, particularly for the domain novice who performed a majority of the knowledge acquisition work.

In a surprising contrast to this, however, we found that the completed ADS analysis was quite difficult to use as a means of learning about the process from scratch. Despite the fact that the two authors receiving the ADS analysis had prior experience with the framework, they found it difficult to familiarize themselves with the reactor process using the analysis results. While they were able to make sense of the representations, they were unable to reconstruct a comprehensive understanding of how the process worked. It was necessary to return to both the source material and domain experts to reconstruct the knowledge that had been acquired by the original analysts. It was not the case that the difficulty was due to inaccuracies in the original ADS effort, as few changes were made to the results.

We realized through these two experiences that much of the benefit to work analysis lies in the structured discovery process that it fosters. Sanderson has asserted that the value of performing work domain analysis is that it "prepares your mind to recognize the profitable design insight" [27]. Our experience (both in teaching and practicing the framework) is consistent with this characterization. In this sense, the ADS framework operates something like a "scaffold" for learning in the Vygotskian sense [28]. In addition to the value of the product of the analysis, the process whereby that product is produced appears to deliver value that is not obtainable from the product itself. In the present case, the list of IRs (discussed below) proved to be too far removed from the domain complexity to yield sufficient design insight. By reworking the ADS, the second group of analysts experienced the same discovery process that the initial analysts had experienced, albeit somewhat more rapidly.

2) Lessons Learned From Task-Based Modeling: The key insight generated in the transition from work-domain- to task-based analysis was the importance of matching the scopes of each. Although the scope that we set for each of these phases was consistent with their accepted practice, we inadvertently left out the automated controllers.

To compensate for this gap in our analysis, we introduced CtrlTA to incorporate the behavior of the automation. However, it is important to note that CtrlTA does not differentiate between automated and manual executions of control actions. The IRs for control are identical whether a human operator or a machine carries out the task—although, of course, the IRs attended to by various agents might differ by task and context.

Manual flow control of the hydrogen stream does occasionally occur. In such cases, the operators use heuristic and closed-loop control behaviors to manage the process. For example, operators engaging in manual control of the hydrogen flow tend not to factor in the analyzer readings and assume that the flows are pure hydrogen and C2 streams. This is a fair approximation, and it vastly simplifies the ratio calculation (which is typically performed in the head). To further ease the mental arithmetic task, operators also round off the desired setpoint ratio. Thus, instead of calculating a ratio of 2.1, they will use the value of 2.0 to simplify the arithmetic. At times, even these heuristic approaches become excessively demanding and operators shed the feedforward control task altogether, reverting to feedback control of the acetylene concentration at the reactor output. Although this can be an effective control strategy, it fails to exploit the engineering insight into the control task that has been built into the information acquisition and analysis automation.

Operators fall back on suboptimal control strategies because the flow calculation task is not supported by the current interface, which merely presents the controller input and output variables. Had we not conducted the CtrlTA, we might have concluded from the ADS and HTA that this was sufficient for supporting monitoring of the automation (i.e., the task defined in the procedures). However, supporting both manual and supervisory control requires that all of these IRs be made available in the display [29]. We demonstrate in Section IV-B how a graphical form was designed to support the task of achieving hydrogen flow control, regardless of whether the control is allocated to the automation or to the human operator.

3) Recommendations: As work analysis tools and perspectives move out of the laboratory into consulting and industry environments, managers may wish to consider the aforementioned insights in organizing work analysis and design efforts. Our experience suggests three recommendations. First, novices to a domain should participate in the knowledge acquisition and domain modeling efforts, even if (or perhaps, particularly if) they lack the engineering training normally expected of designers in that domain. The ADS framework itself appears to help direct and structure their learning. Second, there should be some continuity between the analysis and design teams. Received wisdom contained in completed models and IRs appears to be an insufficient basis for developing novel design insights. Given that our Australian colleagues have demonstrated that a work domain-based analysis is relevant across many phases of the systems life cycle [30]–[33], the implications of these recommendations may be very broad.

A third recommendation speaks to the problem of selecting the suite of work analysis tools. Analysts should take care to choose tools from each category to ensure that all relevant aspects of the work domain are considered. To this end, multiphase analysis frameworks such as Cognitive Work Analysis [9] suggest specific tools to minimize the gaps in a comprehensive work analysis.

E. Summary

The products of the work analyses include three separate sets of IRs for interfaces to support effective monitoring and control of the reactor. In the following section, we describe the process of transforming those IRs into a list of interface content requirements. Subsequently, we introduce an example interface form that combines elements from all three analyses.

IV. Information Requirements Analysis

A comprehensive set of IRs forms a necessary body of information for effectively supporting the task of supervisory control

TABLE II
INFORMATION REQUIREMENTS FOR THE WEIGHT RATIO CONTROLLER

# on Fig. 13	Information requirement to be shown	ADS	HTA	CtrlTA
1.	FV413 (appearance and location)	✔	✔	✔
2.	FV413 position	✔	✔	✔
3.	The flow of H2 into the reactor	✔	✔	
4.	The flow of CO into the reactor	✔	✔	
5.	Aggregate H2/CO Supply Input	✔		
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.	Aggregate C2 Supply Input	✔		
10.	Low flow limit on FC411.PV of 40 Mg/hr		✔	
11.	H2/C2 weight ratio SP		✔	✔
12.	H2/C2 weight ratio OP (also FC413 SP)			✔
13.	FC413.OP			✔
14.	Error (FC413.SP-FC413.PV)			✔
15.	H2 density in H2/CO stream			✔
16.	H2 density correction			✔

of the reactor. The aim of integrating task- and work domain-based analysis methods is an effort to move closer to identifying a sufficient set of IRs that, if satisfied in an interface, would allow for effective adaptation under the broadest range of events [9]. In this section, we discuss the consolidation and resolution of the IRs generated by each of the three work analyses.

A. Content Specification

1) Information Requirements Matrix: A matrix of IRs was formed by listing the requirements from each of the three analyses (i.e., ADS, HTA, and CtrlTA). Each requirement was cross checked against the requirements drawn from the other two analyses, and duplicate IRs were removed. Frequently, requirements from one analysis had to be assimilated or parsed to reconcile them with requirements from another analysis (see Section IV-A2). The complete IR Matrix is again too large to present here, but Table II presents an example for the subset of the reactor around the weight ratio controller described in the control task analysis.

For each IR, the analysis or analyses that identified that requirement are noted in the matrix. Thus, the matrix contains a column for each analysis (ADS, HTA, and CtrlTA) and checkmarks in the appropriate cell when the analysis identified the IR. 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 might 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 include a representation of the state of the valve in question. In this case, both the ADS and HTA would be credited with having identified the requirement for the valve and its current state. In general, we 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.

2) Analysis Overlap: A Venn diagram (see Fig. 11) 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 CtrlTA) identified the same 10% of the total IRs. However, there was no overlap between ADS and CtrlTA that was not also shared with the HTA. Generally speaking, there is a large overlap in the IRs that each analysis generated. However, there is also substantial uniqueness in each of them. Both of these observations conform to those of Miller and Vicente [4], who observed this combination of overlap and uniqueness in a simpler system.

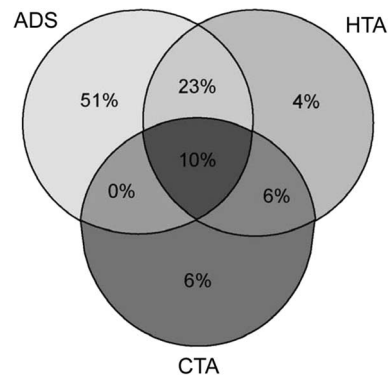

Fig. 11. Proportion of IRs captured by each analysis (circles not to scale).

B. Information Availability

A consistent observation that has been made by designers of advanced interfaces is that the analyses often identify a set of IRs that exceed the resources provided by contemporary instrumentation systems [2], [10], [22], [34]. This poses an important question when we consider integrating work domain-based approaches with task-based approaches. If the IRs that are produced by task-based approaches are indeed complementary to the IRs that are produced by work domain-based approaches, then it is necessary to examine whether these additional requirements are supported by the existing instrumentation systems. The results of that examination may have a substantial impact on system design. We begin by reviewing the current reactor interface and the information available in that interface. We continue by considering the information that was supported by the existing instrumentation system.

1) Description of the Current Reactor Interface: The existing reactor interface is based on a mimic display format supplemented with arrows representing the designed direction of flow, textual information pertaining to control parameters, and sensed (or analyzed) process values. Equipment is laid out according to the process sequence, and process variables are placed near the lines or equipment to provide information about measured variables at those locations. Information about the reactor is provided on two such views, which are supported by several others that are rarely accessed and provide static details about controller logic and alarm configurations. A second viewport in the workspace alternates between two views, with one providing a table of control parameters (e.g., the WRC setpoint) and the other providing a table of analyzer values.

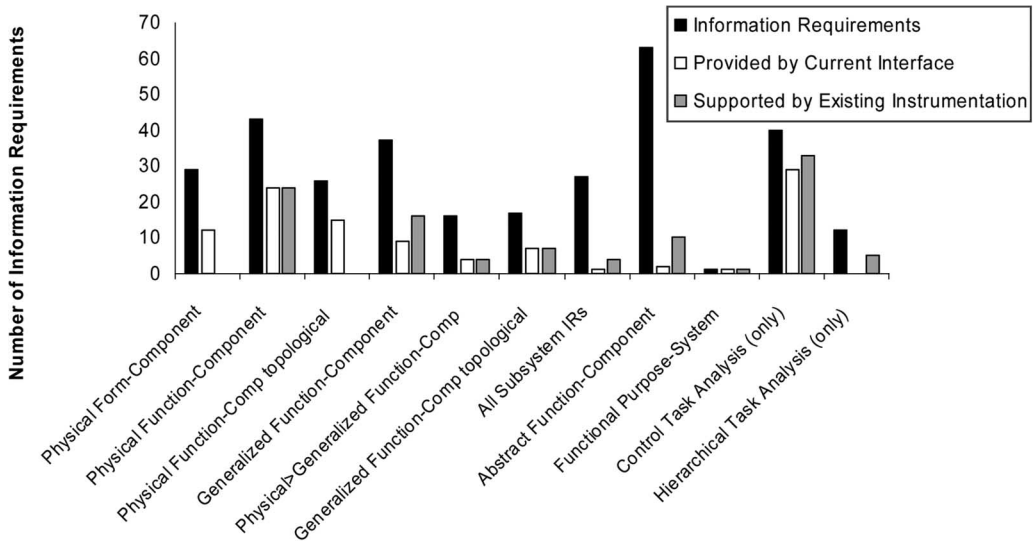

Fig. 12. Comparison of IRs identified by work analyses against (a) requirements included in the current reactor interface and (b) requirements that could be met with existing instrumentation.[6]

Finally, a "change" zone at the bottom of the workspace provides detailed information about a particular data point and allows for changes to controlled variables. All of these representations adhere to the single sensor single indicator or elemental display of data philosophy [35].

In addition to the mimic diagrams, two other information displays support monitoring and control by the operator. First, the alarm summary page provides an interface to the alarms that are implemented in the instrumentation and control software. Examples of these "soft" alarms include limits on equipment temperature, process pressure, and control configurations. A summary of the variables 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 distinct from the hard-wired panel alarms also in use for the most critical process parameters. Second, a trending package allows the operator to plot historical information for up to four process 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.

The design philosophy of the existing information system is consistent with that employed in most contemporary process control environments. Newer generation distributed control systems provide the opportunity for more advanced graphical techniques, but our observations indicate that few plants have taken advantage of the opportunity to provide higher level process information. Rather, what we tend to observe in the industry is the same mimic-display philosophy with a greater range of graphical user interface widgets (e.g., tabs, pull-down menus and pop-up windows) and graphics that contain little additional information pertaining to control.

2) Information Provided by the Existing Reactor Operator Interface: Fig. 12 compares the combined number of IRs that were generated from all sources to the number of IRs met by the existing interface. The first nine pairs of columns on the left are taken from the ADS, whereas the rightmost two are from the CtrlTA and HTA (as labeled). There is a clear shortfall in the information content of the current interface compared to the requirements identified. The reasons for this shortfall are not consistent across the categories. At low levels of abstraction, such as at the Physical Form level, the interface excludes many IRs to reduce display clutter and optimize control during normal operations. Many Physical Function and Generalized Function IRs at the component level are not displayed, because they are not measured, although they could be if additional sensors were provided. This shortage of information propagates, limiting the availability of aggregated and Abstract Function information. In contrast, a majority of the control task requirements are met in the existing interface, largely because they are required for automatic control and thus are available in the instrumentation and control system. Notably, none of the HTA-only IRs are met in the interface. This is because the written procedures meet these requirements.

3) Information Requirement Coverage Possible With Existing Instrumentation: Fig. 12 also assesses what proportion of the IRs *could be met* with the existing instrumentation if the necessary derivations of higher order properties were performed. It is clear that there is not a great deal of information that could be supported by the existing instrumentation system that is not already available. Thus, the current interface appears to be exploiting much of the available information. Providing more information to the operator would require an expanded sensor set. This observation is quite important, as the availability of sensors has been cited as a possible limitation in designing ecological interfaces [10].

C. Discussion

1) Lessons Learned: In retrospect, the process of integrating the IRs benefited from having the foundation established by the ADS. The 2-D hierarchical structure of this analysis

[6]An assessment of instrumentation availability is not applicable to the Physical Form–Component and Physical Form–Topological categories.

provided a framework within which the requirements that were 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 previously. 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 independently of operational tasks, it should be broader and more inclusive than any task-based analysis (see Table I). Thus, from an analytical perspective, work domain-based analyses may always hold an edge over task analysis in forming a foundation of IRs.

2) Implications: The IRs analysis reveals that, in order to extend the content of operator information systems for the reactor, additional sensors would have to be introduced to the system. This demonstrates that human factors contributions to system design can (and should) extend beyond the role of determining the interface form to determining the information content of the interface [36]. The analysis itself indicates where those sensors would be useful. This observation is not meant to suggest that the engineers who designed the current instrumentation and control system were negligent. Rather, their decisions were constrained by cost considerations—installing and maintaining sensors can be prohibitively expensive. Although there is substantial empirical evidence that well-designed interfaces containing more extensive information content can support improved operator performance [37], the literature is largely silent on the question of economy (although see [38]). Thus, a key area for future research is establishing the cost-benefit tradeoff of extending the information content of supervisory control interfaces [39].

V. Design Example

A framework for interface design must provide for both specification of information content and structure, as well as determination of interface form. Having identified the information that should be included in an interface for the reactor, we now turn to the challenge of designing a visual form. In particular, we are interested in the design of forms that integrate the work domain- and task-based content. This section offers a detailed description of one graphical form that appears in the reactor interface. For a more comprehensive description of the entire interface, see [21].

A. Ecological Interface Design Extended

Ecological Interface Design is a framework for designing user interfaces for complex systems [2]. EID uses the ADS to identify information content and structure for the interface. The design of the visual form is guided by the Skills, Rules, Knowledge framework [40], which describes three qualitatively different ways in which people process information. These ways of processing information lead to qualitatively distinct behaviors: skill-based behavior, rule-based behavior, and knowledge-based behavior. Ecological interface forms aim to exploit operators' powerful pattern recognition and psychomotor abilities by presenting information in such a way as to promote skill- and rule-based behavior. This allows operators to deal with task demands in a relatively efficient and reliable manner rather than forcing effortful cognitive processing. Knowledge-based behavior is also supported by embedding an ADS representation of the work domain in the interface. This provides operators with an external visualization of plant structure and dynamics that offers support during novel situations requiring adaptive problem solving.

While their paper on the EID framework emphasizes using the ADS for work analysis, Vicente and Rasmussen anticipated the eventual expansion of the work analysis to include the products of other techniques to determine interface content [2]. However, EID to date has predominantly been attempted using the ADS exclusively (although see [8]). 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.

B. Display of Integrated System- and Task-Based Constraints: Weight Ratio Control View

In this section, we provide an example of an ecological interface view for the WRC that was developed through an integration of both work domain- and task-based requirements. In doing so, we demonstrate conclusively that the EID framework accommodates the inclusion of task-based information and, by so doing, offers improvements over current practice. The example further demonstrates that ecological interfaces can be developed to support both supervisory and manual control of a complex system.

In describing the WRC view, we refer back to the contents of Table II and use the numbers in the left-hand column to refer to display elements in the graphic shown in Fig. 13. The central column of Table II then describes the IR that that display element is meant to satisfy. The following three columns indicate which of the analyses identified that requirement. The description of the appearance and behavior of the graphic in Fig. 13 exemplifies how the requirements [denoted in brackets] were instantiated.

The ratio control graphic supports manual or supervisory control of the incoming mixed stream of hydrogen (H2) and carbon monoxide (CO), as shown in Fig. 13. Two vertical bar graphs display the flow rate of the feed (C2) stream [IR9] (tag name FC-411) and the H2/CO stream [IR5] (tag name FC-413). The bar denoting the C2 stream (at left) aggregates the acetylene [IR6], ethylene [IR7], and ethane [IR8] components in the C2 flow (the individual components become visible as a tooltip when the mouse is held over the column). A minimum limit for the aggregate flow is denoted by a red limit icon [IR10]. The bar denoting the H2/CO stream is divided into a

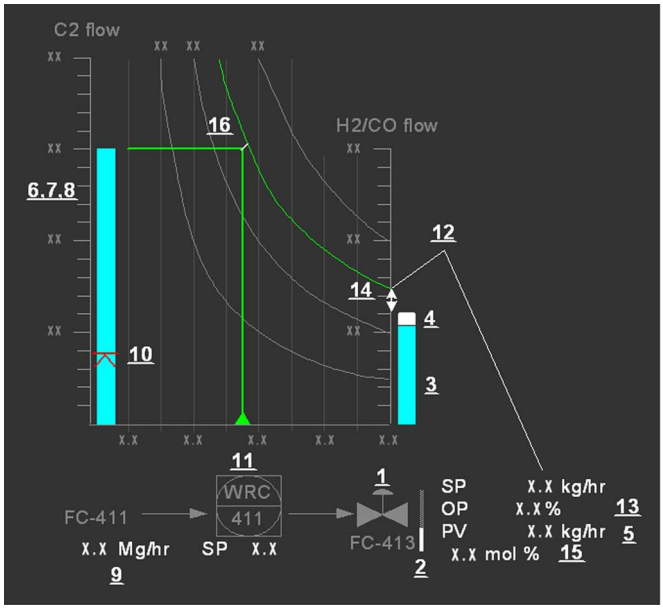

Fig. 13. Ratio controller graphic. The numbers correspond to the IRs in Table II.

cyan portion that represents the amount of H2 in the stream [IR3], and a white portion represents the amount of CO and impurities in the stream [IR4]. Distinguishing the proportion of this flow that is hydrogen is important because the controller attempts to regulate the hydrogen flow rate as opposed to the total flow rate.

A ratio controller determines the desired set point of a controlled flow (in this case, the H2/CO stream) by taking the product of an uncontrolled flow (in this case, the C2 flow) and the setpoint ratio chosen by the operator. This behavior is represented in the interface by a graphic lying between the two flow rate bars. The gray baseline of the graphic reflects the WRC setpoint range of 0.0–2.0, with gray vertical lines denoting 0.25 unit increments along this scale. A small green triangle rests on the baseline at the current setpoint value of the weight ratio controller [IR11]. A green line extends vertically from the triangle. A second green line extends horizontally from the top of the C2 flow rate bar to the right. Curves are constructed in the graphic such that, at every point along their length, the multiple of the values on the horizontal and vertical axes is a constant. Four static gray constant value curves traverse the space at values of 50, 100, 150, and 200. The placement of a fifth constant value line, which is drawn in green, is based on the calculations performed by the controller.

In the typical case of a ratio controller, the fifth constant value line would intersect with the vertical and horizontal green lines. However, the WRC factors in an adjustment [IR16] for the density of the mixed hydrogen stream [IR15] relative to a pure stream. The correction shifts the position of the curve up and to the right if the density is lower than anticipated (as is the case in Fig. 13), resulting in a higher flow setpoint for the H2/CO flow (the opposite applies for the high density condition). The offset is highlighted by a white line extending from the intersection of the horizontal and vertical green lines to the curved line. The curve thus terminates at the H2/CO flow scale at the value of the setpoint for the controlled flow [IR12].

Below the H2/CO flow scale, a valve icon refers to the hydrogen flow control valve [IR1]. The valve controller adjusts the output [IR13] to the valve stem [IR2] to minimize the steady-state error [IR14], which is reflected in the distance between the top of the H2/CO bar and the terminus of the curved green constant value line (as emphasized by the vertical double-headed arrow to the right of IR label #14).

C. Implications

1) Integration of IRs: The point of providing a detailed description of the WRC graphical forms is to demonstrate that the information content requirements identified by all three of the work analyses are indeed contained in the interface view. Thus, there is a direct correspondence between analysis and design. Often, this connection is not explicitly shown in accounts of ecological interfaces or configural displays [2]. Elsewhere, we have provided an account of the user-centered design process that led to this research product [41].

Beyond providing an accounting for IRs between analysis and design, the WRC example demonstrates the *effective* use of EID to convey content specified by both work domain- and task-based work analyses [16]. That is to say that the demonstration confirms that the design objective has been met. As far as we know, this is the first example of an ecological interface that satisfies these objectives for IRs from multiple work analysis methods.

2) Support for Automation Monitoring or Manual Control: The WRC graphic also serves as an example of how a single interface form can be used to support both supervisory and manual control. When the automation is engaged, the operator can monitor the graphic to confirm that the visual constraints previously described are maintained. If the terminus of the green curved line does not meet the H2/CO flow rate scale at the top of the stacked bar graph, this would suggest that the process variable is not tracking the controller setpoint (as is the case in Fig. 13). Alternatively, if the curved green line does not meet the endpoint of the white density correction line, this would suggest a problem between the flow analyzer and the WRC calculation (which is not the case in Fig. 13). Note that these two cases form "rules," allowing for rule-based monitoring of the automation by the operator. These rules can replace the knowledge-based reasoning (described in Section III-D2) that operators must currently rely on, thereby satisfying one of the design objectives of EID (see Section V-A).

When the WRC automation is not engaged, the graphic supports manual control of the H2/CO flow rate by graphically conveying the multiplication operation and density correction normally completed by the controller. In this case, the vertical green bar becomes a marker for the desired weight ratio. The operator's task is to adjust the setpoint of the hydrogen flow controller (FC413.SP), so that the curved green line intersects with the terminus of the white density correction line. In effect, this is a direct replication of the behavior of the automation. The computational effort, instead of being allocated to the operator, is allocated to the interface. The reduction in cognitive work afforded by this design should reduce operator reliance on heuristic and feedback control strategies under manual control.

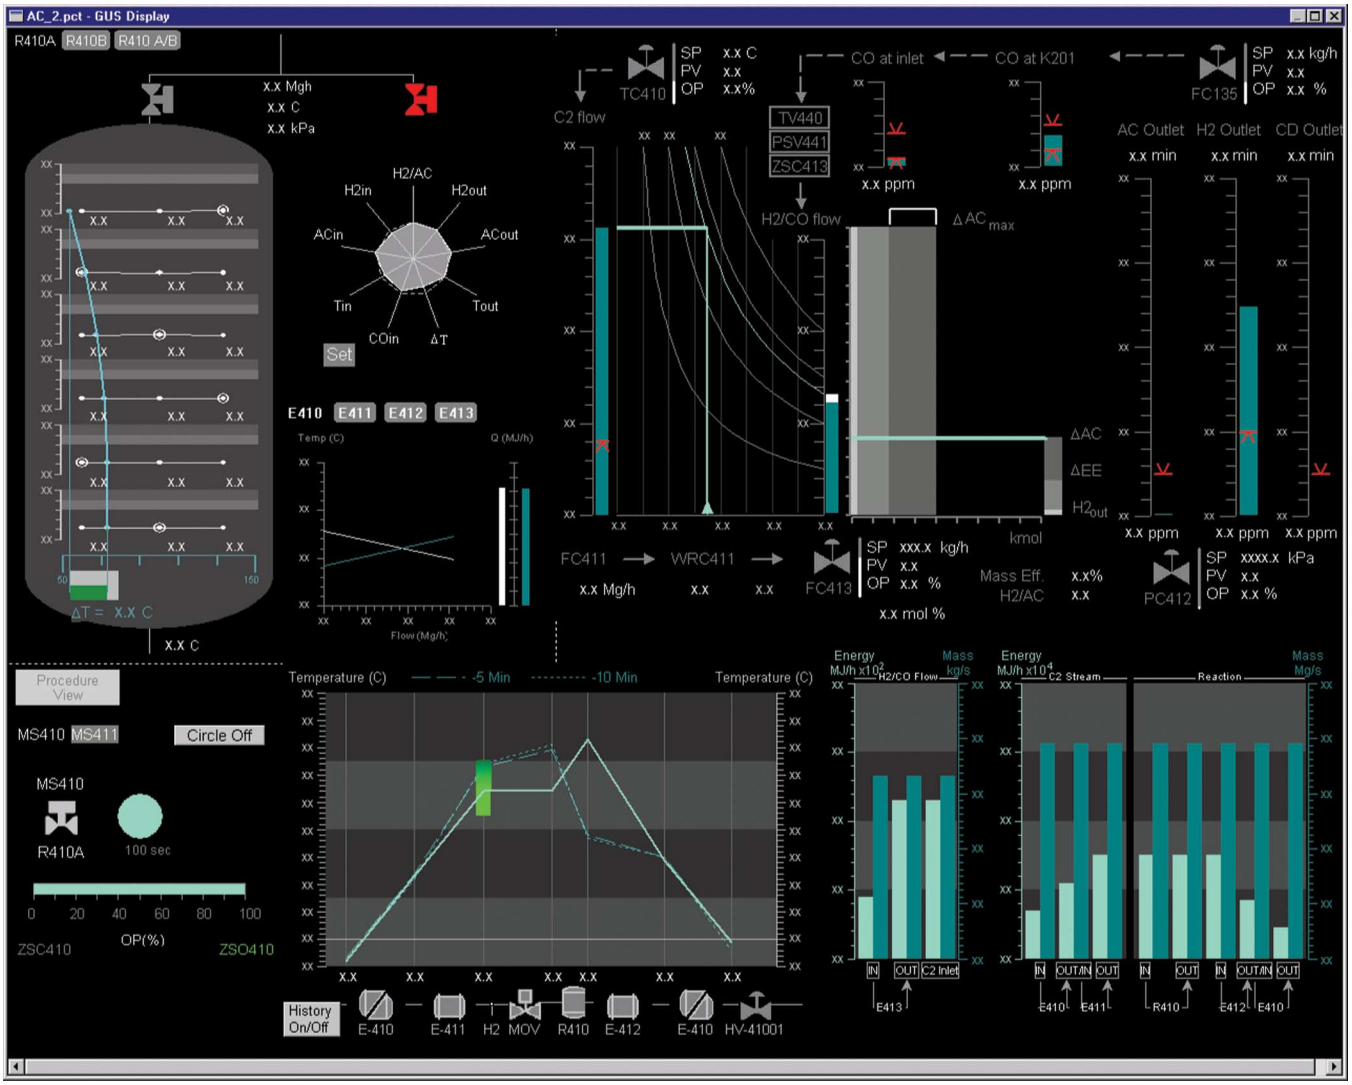

Fig. 14. $P + F + T$ interface for the reactor (upper monitor).

In allowing for this knowledge-based reasoning about the automation, the graphic satisfies the second objective of EID.

D. Ecological Interface for the Acetylene Hydrogenation Reactor

The WRC graphical form was combined with other forms to construct an ecological interface for the reactor (see Figs. 14 and 15). The interface communicates the entire content of the matrix of IRs in a single workspace. Ten process views are distributed across a workspace that spans two stacked 21-in monitors, each at a resolution of 1280 × 1024 (for a total screen size of 1280 × 2048).[7] The ten views are identified below. They are described briefly moving from left to right across the monitor in two rows.

[7]The background colors of the two monitors differ because one was implemented in-house using commercial object libraries, while the other was implemented by a contractor using custom graphics. We discovered late in the implementation phases that the in-house libraries did not allow for the range of color control required. The cost of extending the custom implementation was deemed excessive and unlikely to affect the results of the study.

On the upper monitor (Fig. 14), the *Reactor Temperature Profile View* in the upper left shows the horizontal and vertical temperature distributions across the reactor. A combination of configural graphics and digital values is used, and indications for high temperature alarms are also included. The *Polar Star View* just to the right supports overview monitoring of critical reactor variables though a concise configural display. Parameter values on the graphic can be "set" by the operator to increase its context specificity. Below the polar star, the *Heat Exchanger View* shows the first law of thermodynamics for heat exchangers through a combination of configural and non-configural graphics. This view allows for troubleshooting of heat exchanger problems. The WRC View appears next (see Section V-B), followed by the *Hydrogen Balance View* (upper right corner), which provides detailed graphical representations of the reactor functions and purpose. This view can be used to monitor the reactor and diagnose anomalies in its performance.

On the second row, the process of opening and closing motor-operated valves (MOV) is depicted through a predictive display in the *MOV View* located in the lower left corner. This view reduces the need for communication between control room and

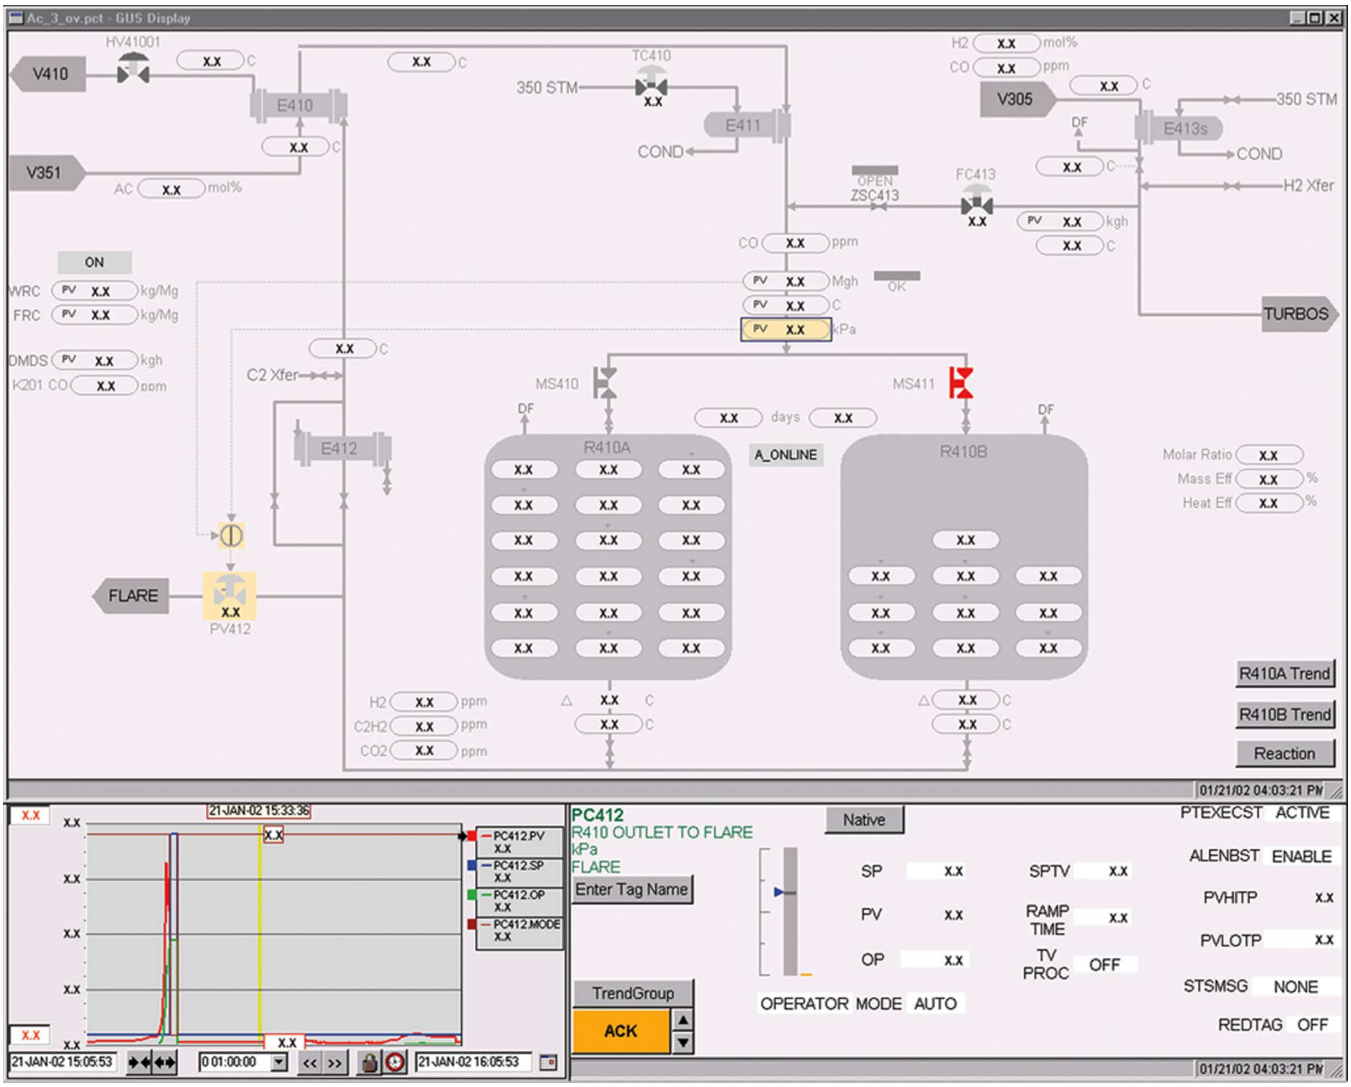

Fig. 15. $P + F + T$ interface for the reactor (lower monitor).

field operators. The *Paulsen View* [42] provides a mountainlike graphical profile of the equipment temperatures across the process along with digital values. Process constraints are shown where they are available. The *Mass and Energy View* (lower right) provides a graphical profile for mass and energy flow through the process.

On the lower monitor (Fig. 15), the *Schematic View* in the top two thirds of the image provides primarily physical and low-level functional information about the process. It is based on a mimic display that is annotated with digital values for critical variables and iconic representations of valve states. It is similar to the mimic display that is available in the current interface. In the second row, the *Trend View* on the left allows the operator to view a history of several process values in parallel. All changes to controlled process variables are made in the *Detail View* on the right. It provides the detailed process parameters associated with any data point.

Finally, the *Procedure View* (not shown) uses text, configural, and nonconfigural graphics to provide step-by-step procedure support. The Procedure View is the only view that is not always visible but rather appears in a window in the lower monitor.

VI. Conclusion

The literature on graphical interface design is rich with case studies of the application of specific analysis and design frameworks to representative problems. These cases serve as existence proofs for the applicability of those frameworks and serve as waypoints to other cognitive engineers in solving problems. Although these cases are exceedingly valuable, they tend to have one key drawback: They are almost exclusively based on a single work analysis method. To a large extent, this is a reflection of the time and financial constraints under which all analysis and design activities take place. Since work analysis methods are rarely the focus of scientific investigation, academic investigators tend to press on toward other research objectives. Similarly, industry practitioners press on toward solving a particular problem. A side effect of this focus is that analysts and designers become comfortable with, and associated with, a particular framework. Moreover, as they find value in repeating that framework, they become bound by their experience.

The case study presented here makes four distinct contributions. First, it confirms the key finding from a similar exercise

conducted on a less complex system [4]. Work domain- and task-based work analysis methods produce largely complementary IRs. Second, it leads to several recommendations for the conduct of work analysis and interface design in complex systems. Work domain modeling serves as a scaffold to help novices to structure their knowledge of a new domain: Project managers should endeavor to establish continuity across the analysis and design teams, and care should be taken to ensure that the mix of analysis tools provides the coverage of the work system required. Third, it introduces the Information Requirements matrix as an effective tool for integrating these IRs and resolving them into a comprehensive set of interface design requirements. Fourth, it demonstrates that these requirements can be effectively represented in an ecological interface. The case study also highlights several challenges to completing each of these steps and suggests measures to overcome them. In future work, we seek to determine whether having met these challenges results in performance advantages for the joint human–machine system [18].

ACKNOWLEDGMENT

This study was conducted in close association with engineers, operators, and managers at the Nova Chemicals, Ltd. Ethylene 1 plant in Joffre, AB, Canada. The authors would like to thank J. Cole, J. Cundict, K. Doe, J. Errington, and D. Lutz for the many hours of technical, engineering, and organizational assistance that they provided. M. Chignell, J. Shaw, A. Kirlik, P. Milgram, and three anonymous reviewers provided valuable feedback on previous drafts.

REFERENCES

[1] J. J. Gibson, *The Ecological Approach to Visual Perception*. Boston, MA: Houghton Mifflin, 1979.

[2] K. J. Vicente and J. Rasmussen, "Ecological interface design: Theoretical foundations," *IEEE Trans. Syst., Man, Cybern.*, vol. 22, no. 4, pp. 1–18, 1992.

[3] D. V. C. Reising and P. M. Sanderson, "Ecological interface design for Pasteurizer II: A process description of semantic mapping," *Hum. Factors*, vol. 44, no. 2, pp. 222–247, 2002.

[4] C. A. Miller and K. J. Vicente, "Comparison of display requirements generated via hierarchical task and abstraction–decomposition space analysis techniques," *Int. J. Cogn. Ergonom.*, vol. 5, no. 3, pp. 335–356, 2001.

[5] C. M. Burns and K. J. Vicente, "Model-based approaches for analyzing cognitive work: A comparison of abstraction hierarchy, multilevel flow modeling, and decision ladder modeling," *Int. J. Cogn. Ergonom.*, vol. 5, no. 3, pp. 357–366, 2001.

[6] A. Shepherd, "Analysis and training in information technology tasks," in *Task Analysis for Human–Computer Interaction*, D. Diaper, Ed. Chichester, U.K.: Ellis Horwood, Ltd., 1989, pp. 15–55.

[7] K. R. Laughery and K. R. Laughery, "Analytic techniques for function analysis," in *Handbook of Human Factors*, G. Salvendy, Ed. New York: Wiley, 1987, pp. 329–354.

[8] J. Rasmussen, A. M. Pejtersen, and L. P. Goodstein, *Cognitive Systems Engineering*. New York: Wiley, 1994.

[9] K. J. Vicente, *Cognitive Work Analysis: Toward Safe, Productive, and Healthy Computer Based Work*. Mahwah, NJ: Erlbaum, 1999.

[10] C. M. Burns and J. R. Hajdukiewicz, *Ecological Interface Design*. Boca Raton, FL: CRC Press, 2004.

[11] J. Rasmussen, "The role of hierarchical knowledge representation in decision making and system management," *IEEE Trans. Syst., Man, Cybern.*, vol. SMC-15, no. 2, pp. 234–243, Mar./Apr. 1985.

[12] A. Miller, "A work domain analysis framework for modeling intensive care unit patients," *Cogn. Technol. Work*, vol. 6, no. 4, pp. 207–222, Nov. 2004.

[13] B. Kirwan and L. Ainsworth, *A Guide to Task Analysis*. Bristol, PA: Taylor & Francis, 1992.

[14] J. R. Hajdukiewicz and K. J. Vicente, "A theoretical note on the relationship between work domain analysis and task analysis," *Theor. Issues Ergonom. Sci.*, vol. 5, no. 6, pp. 527–538, 2004.

[15] A. M. Bisantz and K. J. Vicente, "Making the abstraction hierarchy concrete," *Int. J. Hum.-Comput. Stud.*, vol. 40, no. 1, pp. 83–117, Jan. 1994.

[16] K. J. Vicente, "Ecological interface design as an evaluation tool," in *Proc. Top. Meeting Adv. Human Factors Res. Man/Comput. Interactions: Nuclear Beyond*, 1990, pp. 259–265.

[17] W. B. Rouse, "Computer-generated display system guidelines, Volume 2: Developing an evaluation plan," Elect. Power Res. Inst., Palo Alto, CA, NP-3701, 1984.

[18] G. A. Jamieson, *Ecological Unterface Design for Petrochemical Process Control: An Empirical Assessment*. to be published.

[19] C. A. Miller and K. J. Vicente, "Abstraction decomposition space analysis for NOVA's E1 acetylene hydrogenation reactor," Cogn. Eng. Lab.: Univ. Toronto, Toronto, ON, Canada, CEL-98-09, 1998.

[20] C. A. Miller and K. J. Vicente, "Comparative analysis of display requirements generated via task-based and work domain-based analyses in a real world domain: NOVA's acetylene hydrogenation reactor," Cogn. Eng. Lab.: Univ. Toronto, Toronto, ON, Canada, CEL-99-04, 1999.

[21] G. A. Jamieson and W. H. Ho, "A prototype ecological interface for a simulated petrochemical process," Cogn. Eng. Lab.: Univ. Toronto, Toronto, ON, Canada, CEL 01-02, 2001.

[22] G. A. Jamieson and K. J. Vicente, "Ecological interface design for petrochemical applications: Supporting operator adaptation, continuous learning, and distributed, collaborative work," *Comput. Chem. Eng.*, vol. 25, no. 7, pp. 1055–1074, 2001.

[23] M. Lind, "Modeling goals and functions of complex industrial plants," *Appl. Artif. Intell.*, vol. 8, no. 2, pp. 259–283, 1994.

[24] N. Mazaeva and A. M. Bisantz, "Modeling automation within an abstraction hierarchy," in *Proc. Human Factors Ergonom. Soc. 48th Annu. Meeting*, 2003, pp. 649–652.

[25] J. Rasmussen, "Outlines of a hybrid model of the process plant operator," in *Monitoring Behavior and Supervisory Control*, T. B. Sheridan and G. Johannsen, Eds. New York: Plenum, 1976, pp. 371–383.

[26] C. M. Burns, L. Garrison, and N. Dinadis, "From analysis to design: WDA for the petrochemical industry," in *Proc. Human Factors Ergonom. Soc. 47th Annu. Meeting*, 2003, pp. 258–262.

[27] P. M. Sanderson, "Tactics for conveying CWA concepts to undergraduate engineering versus psychology students," presented at the Workshop Cognitive Work Analysis, Seattle, WA, 2004.

[28] L. S. Vygotsky, *Thought and Language*. Cambridge, MA: MIT Press, 1962.

[29] G. A. Jamieson and K. J. Vicente, "Designing effective human-automation-plant interfaces: A control-theoretic perspective," *Hum. Factors*, vol. 47, no. 1, pp. 12–34, 2005.

[30] N. Naikar, B. Pearce, D. Drumm, and P. M. Sanderson, "Designing teams for first-of-a-kind, complex systems using the initial phases of cognitive work analysis: Case study," *Hum. Factors*, vol. 45, no. 2, pp. 202–217, 2003.

[31] N. Naikar and P. M. Sanderson, "Evaluating design proposals for complex systems with work domain analysis," *Hum. Factors*, vol. 43, no. 4, pp. 529–542, 2003.

[32] N. Naikar and P. M. Sanderson, "Work domain analysis for training-system definition and acquisition," *Int. J. Aviat. Psychol.*, vol. 9, no. 3, pp. 271–290, 1999.

[33] P. M. Sanderson, N. Naikar, G. Lintern, and S. Goss, "Use of cognitive work analysis across the systems life cycle: Requirements to decommissioning," in *Proc. 43rd Annu. Meeting Human Factors Ergonom. Soc.*, 1999, pp. 318–322.

[34] N. Dinadis and K. J. Vicente, "Designing functional visualizations for aircraft systems status displays," *Int. J. Aviat. Psychol.*, vol. 9, no. 3, pp. 241–269, 1999.

[35] L. P. Goodstein, "Discriminative display support for process operators," in *Human Detection and Diagnosis of System Failures*, J. Rasmussen and W. B. Rouse, Eds. New York: Plenum, 1981, pp. 433–449.

[36] K. J. Vicente, K. Christoffersen, and C. N. Hunter, "Response to Maddox critique," *Hum. Factors*, vol. 38, no. 3, pp. 546–549, 1996.

[37] K. J. Vicente, "Ecological interface design: Progress and challenges," *Hum. Factors*, vol. 44, no. 1, pp. 62–78, 2002.

[38] J. Errington, D. V. C. Reising, P. Bullemer, T. DeMaere, D. Coppard, K. Doe, and C. Bloom, "Establishing human performance improvements and economic benefit for a human–centred operator interface: An

industrial evaluation," in *Proc. 49th Annu. Meeting Human Factors Ergonom. Soc.*, 2005, pp. 2036–2040.

[39] G. A. Jamieson and D. V. C. Reising, "Cost-justifying investments in advanced human–machine interface technologies—Part I: Cost-benefit models for the process industries," in *Proc. 48th Annu. Meeting Human Factors Ergonom. Soc.*, 2004, pp. 1179–1183.

[40] J. Rasmussen, *Information Processing and Human–Machine Interaction: An Approach to Cognitive Engineering*. Amsterdam, The Netherlands: North Holland, 1986.

[41] G. A. Jamieson, W. H. Ho, and D. V. C. Reising, "Ecological interface design in practice: A design for petrochemical processing operations," in *Human–Computer Interaction: Theory Practice—Part I*, J. Jacko and C. Stephanidis, Eds. Mahwah, NJ: Erlbaum, 2003, pp. 133–137.

[42] J. L. Paulsen, *A Pipe Reactor Design and the Interface System*. Roskilde, Denmark: Syst. Anal. Dept., Risø Nat. Lab., 1992. RISØ-M-2931.

Greg A. Jamieson (M'03) received a B.S. degree in mechanical engineering and in psychology (with distinction) from the University of Illinois, Urbana, in 1996 and M.A.Sc. and Ph.D. degrees in mechanical and industrial engineering from the University of Toronto, Toronto, ON, Canada, in 1998 and 2002, respectively.

From 1998 to 1999, he was a Research Intern with Honeywell Laboratories, Minneapolis, MN, and subsequently held the position of Research Scientist from 1999 to 2002. He is currently an Associate Professor in the Department of Mechanical and Industrial Engineering, University of Toronto.

Dr. Jamieson is a member of the Human Factors and Ergonomics Society and the IEEE Systems, Man, and Cybernetics Society, where he serves as a member of the Human–Machine Systems SMC Technical Committee.

Christopher A. Miller received the B.A. (with honors) degree in psychology from Pomona College, Claremont, CA, in 1985 and the M.A. and Ph.D. degree from the University of Chicago, Chicago, IL, in 1991.

From 1989 to 2001, he held various positions culminating in being a Research Fellow at the Honeywell Laboratories, Minneapolis, MN. In 2001, he took coownership and became the Chief Scientist of Smart Information Flow Technologies, Minneapolis, MN, which is a small business specializing in research and development for human interaction with complex automation.

Dr. Miller is a member of the Human Factors and Ergonomics Society and the Association for Computing Machinery.

Wayne H. Ho received the B.A.Sc. degree in systems design engineering from the University of Waterloo, Waterloo, ON, Canada, in 1996 and the M.Eng. degree in mechanical and industrial engineering from the University of Toronto, Toronto, ON, Canada, in 2000.

From 1993 to 1995, he held various human factors and usability internship positions at Nortel, Microsoft, and Computing Devices Canada (General Dynamics). Since 1996, he has been a User Experience Specialist with IBM Canada Ltd., East Markham, ON, Canada. He is currently the User Experience Lead for IBM Rational software quality tools.

Mr. Ho is a member of the Human Factors and Ergonomics Society.

Kim J. Vicente (S'86–M'87–SM'01) received the B.A.Sc. degree in industrial engineering from the University of Toronto, Toronto, ON, Canada, in 1985, the M.S. degree in industrial engineering and operations research from the Virginia Polytechnic Institute and State University, Blacksburg, VA, in 1987, and the Ph.D. degree in mechanical engineering from the University of Illinois, Urbana-Champaign, in 1991.

From 1987 to 1988, he was with the Section for Informatics and Cognitive Science, Risø National Laboratory, Roskilde, Denmark, as a Visiting Scientist. From 1991 to 1992, he was with the faculty of the School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA. He is currently a Professor with the Department of Mechanical and Industrial Engineering, University of Toronto. He was an Associate Editor of the IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICS and was on the Editorial Boards of *Human Factors* and the *International Journal of Cognitive Ergonomics*. He is currently on the Editorial Board of *Theoretical Issues in Ergonomics Science*.