Cognitive Engineering Laboratory

Abstraction Decomposition Space Analysis for NOVA's E1 Acetylene Hydrogenation Reactor

Christopher A. Miller
And
Kim J. Vicente

CEL 98-09

Cognitive Engineering Laboratory Department of Mechanical & Industrial Engineering University of Toronto
5 King's College Rd. 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/

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

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 began 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, Natural Sciences and Engineering Research Council of Canada, Nova Chemicals, Rotoflex International, and Westinghouse Science & Technology Center. CEL also has collaborations and close contacts with the Mitsubishi Heavy Industries and Toshiba Nuclear Energy Laboratory. Recent CEL projects include:

- Studying the interaction between interface design and adaptation in process control systems.
- 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.
- Designing novel computer interfaces to display the status of aircraft engineering systems.
- Developing and evaluating advanced user interfaces (in particular, transparent UI tools) for 3-D modelling, animation and painting systems.

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.

Unified Modeling Project

UT/NCL/HTC/NSERC

Abstraction Decomposition Space Analysis for NOVA's E1 Acetylene Hydrogenation Reactor

A Report of work under Tasks 2 ("Select Work/Problem Domain") and Task 3 ("Abstraction Hierarchy Analysis") of the NOVA/UofT Task Breakdown (Jan. 13, 1998)

Release Date: 25 August, 1998
Document Version: 1.00
Filename: AcetyleneADS.doc

Submitted to:
Jamie Errington, NOVA Chemicals, Ltd.

Prepared by: Christopher A. Miller and Kim J. Vicente
Honeywell Technology Center &
University of Toronto Cognitive Engineering Laboratory
Department of Mechanical and Industrial Engineering

1. Document History

.01 First version. Primarily a full draft of section 6 and an outline for the rest of the document. Released to Kim Vicente for comment on 8/17/98.
.02 8/18/98. Added sections 4 and 5.
.03 8/21/98. Added drafts of sections 7-9.
1.00 8/25/98. Incorporated revisions based on Kim Vicente's comments to above sections, added sections 2 and 10. Released to Jamie Errington at NOVA and Peter Bullemer and Ian Nimmo at Honeywell.

2. Summary

The purpose of the work reported in this document was to identify a suitable real-world work domain within NOVA Chemicals refining operations and then use Rasmussen's (1985) Abstraction Decomposition Space (ADS) analysis technique to understand the structural and functional relationships which are important to achieving functional purposes in that work domain.

This work was one element in an overall effort to investigate whether the techniques used in the ADS, which are inherently focused on the physical characteristics of the work domain or 'system' can be profitably augmented by integration with analytic techniques which focus on the tasks that the user engages in. Other projects in this broader investigation, however, have been performed on a small, simulated laboratory system, a 'microworld' known as DURESS II. Thus, the reasons for applying the ADS to a real-world problem for NOVA Chemicals were two-fold: First, so doing would further demonstrate the utility of ADS for complex, real-world problems. Second, and more important for our research into unifying task- and work-domain based analysis techniques, was the need for an ADS model of a real system to work with. Developing such a model for a realistic work domain would put us at the same starting point used in more theoretical laboratory research, which began with an ADS model of DURESS II. Thus, we could begin to explore the application of our newly-developed laboratory techniques to the real-world domain.

In section 4, we describe the nature and utility of the ADS as the most common of a class of Work Domain Analysis (WDA) techniques. WDA approaches involve a thorough analysis of the constraints and capabilities which the physical plant imposes on work that can be done—with the belief that presenting this information to users via an interface provides them with a better understanding of the range of behaviors the system can exhibit, as well as more robust information about how they can achieve or avoid specific behaviors. An ADS is a two-dimensional modeling tool that can be used to conduct a WDA in complex sociotechnical systems. One dimension of the ADS is abstraction, represented as 'means-ends' links that show how lower level entities such as physical attributes of components, can accomplish higher level, more 'abstract' entities such as the flows of heat or material, or the overall purpose of the system. The other dimension is aggregation or, inversely, decomposition. In this dimension, part-whole links show how physical components can be aggregated into larger, but still physical, wholes (subsystems, systems, etc.). An ADS analysis consists of constructing a number of models of the overall work domain at various intersecting levels of abstraction and decomposition. Each model is 'complete' in the sense that it represents the entire work domain, but each provides a different view of the domain and highlights different sorts of information about it.

Next, in section 5, we describe our reasons for selecting the AHR as a work domain for analysis from among a small set of alternatives. These reasons essentially centered on the availability of knowledgeable support personnel to help us learn about the system, and on the size of the system—large enough to demonstrate interesting complexity yet small enough to make progress in the limited time available. Then we provide a general description of the functioning of the AHR: it's purpose of removing acetylene from C2 feed as a part of the overall ethylene refining process and the specific chemical reactions used to accomplish this end.

Section 6 provides the ADS analysis itself. We constructed six different models of the AHR. These were:

1. A partial Physical Form x Component model showing the location and appearance of individual components which comprise the AHR.

2. A Physical Function x Component model showing the role and name of the 51 individual components which make up the AHR and the physical links between them.
3. A Generalized Function x Component model which shows the general functions (such as material and heat flow, chemical reactions, etc.) which the physical components accomplish, as well as the causal links between these functions—that is, which functions can affect changes in the behavior of other functions and how.
4. A Generalized Function x Subsystem model which shows the same type of information as the Generalized Function x Component model, but does so at a higher level of aggregation. The 51 components reduce to nine subsystems, greatly reducing the complexity of the same types of functions described at the component level in the earlier model.
5. An Abstract Function x Subsystem model which shows the mass and energy relationships inherent in the system's subsystems—for example, this model identifies all sources and sinks for mass and energy in the system, but does not discriminate among the types of mass or energy involved.
6. A Functional Purpose x System model which simply states the overall purpose of the AHR system as a whole—to remove acetylene from the C2 Feed stream.

In addition to these six models, two transition diagrams are included. One showing the means-ends links between Physical Function components and Generalized Function components, and the other showing the part-whole links between Generalized Functions at the component and the subsystem levels.

The final sections of this report examine the utility of this AHR ADS analysis and models. In section 7, we review a specific incident that occurred at the E1 facility in 1985—a reactor runaway involving ethylene decomposition caused by an unanticipated hydrogen source. We found that the ADS models provide a good basis for understanding and explaining what occurred in this incident, understanding the rationale behind operators' actions in attempting to control the incident and predicting their results, and even of suggesting some alternate courses of action that might have been used instead—all this in spite of the fact that our models had not been built to include the unanticipated cause of the incident. This suggests both that the AHR ADS might provide a good tool for training, and that an interface designed from it would have provided rich and robust information for operators in this incident.

In section 8, we review the ADS models and re-formulate them as requirements for a display. The next obvious step (described briefly in section 9) would be to use the list constructed in section 8 to either review NOVA's current displays for the AHR to identify whether the ADS analysis requires information above and beyond that currently included in the displays, or to construct novel displays from the ADS requirements. We have not done either of those steps in this report due to limited time and lack of familiarity with NOVA's current displays.

Finally, in section 10, we summarize our conclusions and report lessons learned from this research.

3. Objectives, Rationale and Caveats

The purpose of the work reported in this document was to identify a suitable real-world work domain within NOVA Chemicals refining operations and then use Rasmussen's (1985) Abstraction Decomposition Space (ADS) analysis technique to understand the physical and functional relationships which are important to achieving functional purposes in that domain. Prior work (cf. Bisantz and Vicente, 1994 and Vicente, 1996 for a review) using this technique has shown that the relationships identified in this way can be readily translated into the display requirements for powerful interfaces for users who must interact with the physical work domain.

This work is one element in an overall effort to investigate whether the techniques used in the ADS, which are inherently focused on the physical characteristics of the work domain or 'system' (i.e., the 'plant' upon which workers' actions must take effect) can be profitably augmented by integration with an additional focus on the tasks that the user engages in with the system. In pursuit of those ends, we have previously investigated whether task-based and work domain-based analysis techniques can be integrated in a common framework (Miller and Vicente, 1998a), and whether a task-based analysis provides different information about display requirements than a work domain-based one does (Miller and Vicente, 1998b). Each of these investigations, however, were performed on a small, simulated laboratory system, a 'microworld' known as DURESS II. The utility of the ADS technique has, to date, only been occasionally demonstrated for complex, real-world work domains, and the utility of an integrated task- and work domain-based approach has not been completely demonstrated even in a laboratory setting.

Thus, the reasons for applying the ADS to a real-world problem for NOVA Chemicals was two-fold: First, so doing would further demonstrate the utility of ADS for complex, real-world problems. Second, and more important for our research into unifying task- and work-domain based analysis techniques, was the need for an ADS model of a real system to work with. As mentioned above, all of our work to date toward developing an integrated analysis technique has been done on a small laboratory simulation system. While this simplifies the research and is, therefore, important for our initial investigations, our goal is ultimately to apply our results to real-world problems. Our intention has been to do initial theory development in the lab and then move it to more realistic settings. We began with an ADS model of the DURESS II laboratory simulation, and have done pioneering work to show both that task analysis provides additional display requirement information and to suggest that task-based display elements can be added to a work-domain based interface. If we are to transition these finding to a real-world problem, then it is important to start at the same point there that we started with the laboratory system—with an ADS model. The work reported here was aimed at providing such a model.

4. The Abstraction Decomposition Space Analysis Technique

Rasmussen's (1985) abstraction-decomposition space (ADS) is the primary example of a class of analytic techniques intended to provide information about user needs and capabilities for the purpose of designing new systems, system automation or system interfaces. These techniques, which we have called Work Domain Analysis (WDA) technique or, sometimes, system-based techniques, emphasize the structure of the work domain—that is, the plant or equipment on and with which the user must achieve some set of functional goals. WDA techniques are in contrast to the other broad class of analytic techniques used in human-centered design: task-based techniques. These approaches place their emphasis on analyzing the set of tasks, activities and procedures which users either do or should engage in when operating the system.

The WDA approach involves a thorough analysis of the constraints and capabilities which the physical plant (aka 'system', aka 'work domain') imposes on work that can be done in the domain—with the belief that presenting this information to users via an interface provides them with a better understanding of the range of behaviors the system can exhibit, as well as more robust information about how they can achieve and avoid specific behaviors. An ADS is a two-dimensional modeling tool that can be used to conduct a WDA in complex sociotechnical systems. Rasmussen's approach, shares the Gibsonian (Gibson & Crooks, 1938) emphasis on the importance of the "field" or ecology in which an actor behaves for determining or

"constraining" the set of actions which are necessary or appropriate. There is a growing amount of empirical support showing that interfaces based on such work domain analyses can lead to better performance than traditional interface approaches, particularly in abnormal situations (Vicente, 1996).

The ADS is commonly, if somewhat incorrectly, referred to as the 'Abstraction Hierarchy' (AH). This name is incorrect because it refers to only one of the two hierarchical dimensions which are used simultaneously in an ADS analysis: abstraction and decomposition (or, inversely, aggregation). These two dimensions together form a matrix, as illustrated in Figure 4.1. This matrix is the 'Abstraction Decomposition Space'.

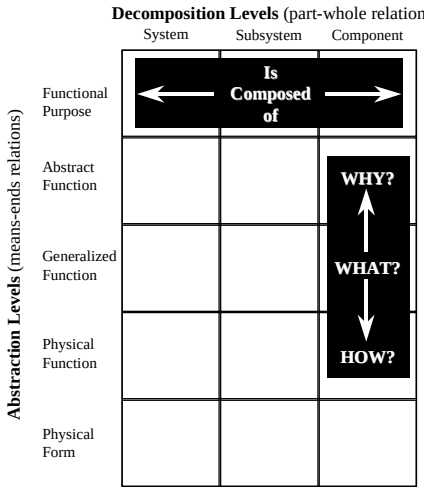

Figure 4.1. The Abstraction Decomposition Space (after Rasmussen, 1985, and Jamieson, 1998).

The part-whole, or 'decomposition' dimension of the ADS is straightforward. Here, we simply aggregate the physical entities in the plant at various levels moving up the axis (or, alternatively, decompose them moving down the axis). The relationship between an entity at an upper level and one at a lower level is 'is composed of'—the system as a whole is composed of subsystems which are composed of components. This analysis can be carried on to as many levels as is useful for the domain. Although this varies widely, users of the ADS space have frequently found three levels to be effective. By the same token, what constitutes a component or a system depends on the focus of the analysis as a whole. A pump will generally be a component if the analyst is interested in the 'system' which a board operator must deal with, but it might represent the whole system if the focus is on a maintenance worker whose role is to diagnose and repair pumps.

The means-ends or 'abstraction' dimension is somewhat more complicated. Here, moving up the dimension means moving from a more concrete to a more abstract description of the system, but the dimension of abstraction is one of functionality. This means that the lowest level descriptions are highly concrete descriptions of the form and appearance of plant components, but as one moves up the levels, one 'abstracts away' from these concrete details and adds more general information not present at the lower levels. For example, there may be no physical component responsible for producing a chemical reaction—thus, the reaction would not show up at the lower, Physical Function level. It would, however, appear at a higher General Function level and its effects (in terms of mass and energy) would show up at the still higher Abstract Function level. Movement upward along the abstraction dimension is toward progressively more general descriptions of the functions performed by specific, concrete entities.

One useful way of thinking about the abstraction dimension (after Rasmussen, 1985 and Jamieson, 1998) is as a hierarchy of means-ends relationships. This means, as illustrated in Figure 4.1, that the relationships between any three layers of the ADS can be characterized by the How-What-Why triad of questions. Attending to any given level means that that is 'What' the observer is currently focused on. The level above the current focus answers the question 'Why'—'why is that component or function present in the plant?' The answer can be obtained by looking at the higher level, more abstract functions that it accomplishes. Moving down a level from 'what' is being focused on answers the question 'How'—'how is the function I am interested in accomplished?' (in structural and functional terms, not in terms of user actions). Again, the answer can be obtained by looking at the lower level, more concrete functions which accomplish it. Note that this 'How-What-Why' window applies at any level of the ADS—that is, in a sense, it can be overlaid over any three vertical cells in the ADS space to answer the same set of questions about the relationships between entities in those cells.

The ADS provides a comprehensive analysis of the means-ends and part-whole relationships in the functional structure of the process being controlled. Each cell in the ADS represents a complete model of

the plant and could, conceivably, stand alone. However, much of the power of the ADS analysis comes from understanding the relationships between the cells. Thus, a typical ADS analysis will construct multiple models to populate several of the cells in the ADS matrix.

Constructing an ADS analysis requires a detailed knowledge of the plant and its interactions with the environment—and on the rules, equations or models governing these interactions. When these sources are inadequate, the analysis will be correspondingly inadequate—but this situation is less common than might be expected. The greatest threat to the safety of process control systems is events that are not familiar to operators and that have not been anticipated by designers (Vicente & Rasmussen, 1992). Under these challenging circumstances, the operator's role is one of adaptive problem solver. Because the event has not been anticipated by system designers, the available procedures, experience, and automated aids are not directly applicable. The one thing that does remain unchanged, however, is the functional structure of the plant and the principles that govern its interactions with the environment. Further, it is precisely within these constraints that the operator must improvise a solution.

5. Problem Domain: NOVA's E1 Acetylene Hydrogenation Reactor

5.1 Criteria for Problem Selection

In seeking a work domain to model for this effort, we had a number of different, and occasionally competing, criteria. Since the project was part of an attempt to integrate multiple modeling techniques—as well as an attempt to integrate university and business interests—criteria came from a number of perspectives. The criteria we considered were as follows:

A good problem from the work-domain modeling perspective should:

- have at least three natural layers of part-whole decomposition. That is, be viewable at (at least) system, subsystem and component levels.
- have easily recognizable/recognized system boundaries.
- have dynamics that are well understood—the better the understanding; the better the model.
- have substantial documentation and expert support available to aid us in understanding it.
- ideally, involve a chemical reaction and/or separation process. This would both be somewhat unique and would extend the ADS in ways that are critical for ascertaining its applicability to NOVA's processes and, therefore, result in pushing the development of the ADS itself.

A good problem from the task modeling perspective should:

- have multiple, well-defined procedures associated with it that may be appropriate in some conditions, but not in others.
- have a reasonably restrictive sequentiality—that is, you must do step 3 before 4, etc. Sequentiality needn't be absolute, but if anything can be done at any time, there is less advantage to many task modeling techniques.
- have strong ties between task or subtask steps and resource (including information) requirements—to do the task in this way means you'll need to know reactor temperatures, but doing it that way you don't have to.
- ideally, have procedures that are tied to different sets of operational equipment, personnel, etc. in a well-understood way—that is, one procedure uses one set of resources, while a second procedure uses a different set.
- have substantial documentation and expert support available to aid us in understanding it.

A good problem for both approaches should:

- be able to distinguish causal/functional practices (tasks) from those that are 'generally a good idea' (perhaps for safety reasons) or 'just the way we do things here'

- be small enough to enable us to build both a system and a task model in the time available, yet big enough to show some interesting complexity and variability.
- perhaps have frequent breakdown conditions. This makes the analysis of the system and the tasks associated with it more useful. It will also help the task modeling work if at least some of the appropriate responses to these breakdowns are known in advance.

A good academic research problem should:

- be publishable and openly discussable
- be realistic; have a grounding in the real world application

A good industrial research problem should:

- have obvious (though not necessarily near term) payoff for NOVA.

With these criteria in mind, three primary candidates for problem domains were identified and reviewed be representatives from NOVA and the University of Toronto: the Pelletizer unit of the NOVA's new Sclair process and either the full Finishing end of the Ethylene process or the Acetylene Hydrogenation Reactor (AHR) portion of this process. While each domain had strengths and weaknesses, the group rapidly settled on the AHR as a useful problem for us to tackle.

This system (actually, a subsystem in the overall ethylene refining process) is relatively small (supported by only 2 schematic screens currently, as opposed to 4-5 screens for the pelletizer and nearly 200 for the full finishing end). While this may prove a disadvantage in the long run, its small size was virtually required to make our work manageable for our first year of research. Further, since the AHR is a component in the overall ethylene finishing system, the prospects for expanding the problem focus in the future (with the ability to integrate and scale up models built for this component) seem good. The next 'tier' of logical components in the process (involving the AHR, propylene refrigeration compressor, C2 splitter and the secondary demethanizer) comprises a relatively self-contained unit which is still only about half of the finishing end of the ethylene process. Similarly, interaction with the AHR involves a very small set of tasks or procedures—only one during normal operations-- although the decision about whether or not to execute that procedure is critical and difficult to make. While this may be less than ideal for the task modeling aspects of our research, further examination of the domain showed that the procedure is complex, involves extensive time pressure and may be unfamiliar to operators. Of the alternatives, the finishing end was deemed too large to make significant progress on during the first year of the program, and a focus on the Pelletizer involved larger problems in terms of the highly proprietary nature of the Sclair process and poorer availability of documentation and knowledgeable support for our knowledge acquisition activities.

Upsets in the AHR are the single most frequent cause of upsets in the overall ethylene process and down time for the AHR process costs roughly $1000/minute. Of upsets involving the AHR, roughly one third are caused by inappropriate initial decisions on the part of the operator (deciding not to go to flare when he should have), while another 50% are caused by poor execution of the flaring procedure. Furthermore, while an inappropriate flare decision (a false positive) can, *if well-executed*, cost 20 minutes of down time (= $20,000), even a well-executed false negative (deciding not to flare when you should have) will cost 4-6 hours of down time ($240,000 to $360,000). A poorly-executed false negative can easily cost double that amount. Thus, there are significant benefits to be obtained through *both* enabling better, more accurate initial decision making and through enabling better execution of the flaring procedure.

5.2 Description of the Acetylene Hydrogenation Reactor

The Acetylene Hydrogenation Reactor (AHR) is one subsystem used in NOVA's E1 ethylene processing facility. The AHR receives partly processed C2 feed which is composed mostly of ethane (C2H6) and ethylene (C2H4) with various trace elements, the most important of which is acetylene (C2H2). Further subsystems in the plant will separate the ethane and ethylene from the trace elements, but those processes are very sensitive to the presence of acetylene. The reason for the presence of the AHR is to remove this acetylene. The AHR does this by 'hydrogenating' it—that is, forcing it to undergo a chemical reaction

which adds an H2 molecule to each C2H2 to convert it to ethylene (C2H4). While the maximization of ethylene production is the overall goal of the E1 plant, the fact that slightly more ethylene is produced by hydrogenation of acetylene is incidental. Instead, the motivation for the removal of acetylene is that it enables the use of downstream processes to separate ethane from the existing ethylene. The AHR process also hydrogenates some of the existing ethylene, thereby turning it into ethane (C2H6). While this is not desirable, the impact on the overall quantity of ethylene and ethane produced is minimal. Instead, ethylene conversion to ethane is undesirable because it runs the risk of using up the available unattached hydrogen molecules, not leaving a sufficient quantity to accomplish the removal of the acetylene. The acceptable concentration of acetylene out of the AHR is less than 5 ppm.

The following is a summary of the AHR process used in NOVA's E1 facility. It may be helpful to cross reference this description with depiction of the AHR contained in Figure 6.9.

1. Raw natural gas enters the E1 facility and undergoes pyrolysis in a number of furnaces. This has the effect of converting some of the ethane and propane in the natural gas to ethylene and hydrogen. Other trace products are produced including carbon monoxide. Pyrolysis is not naturally a part of the AHR subsystem, since it occurs both temporally and geographically distant from the AHR but, for reasons that will become clear as the rest of the AHR process is described, the carbon monoxide present in the output of pyrolysis is critically important to the AHR. Thus, the AHR operator monitors and is given control over one aspect of pyrolysis which affects CO production—the addition of DiMethyle DiSulfide (DMDS) to the natural gas feed into the pyrolytic furnaces. The addition of DMDS reduces CO production--which is somewhat undesirable from the AHR operator's perspective, but it also reduces coke formation in the furnaces, which is desirable from the furnace operator's perspective.
2. Various processes which occur downstream from the pyrolysis furnaces, including the demethanizer (see section 7 below), separate and further process the gas mixture. By the time the gases enter the AHR system, they do so in two streams: one (called the feed stream, or the C2 stream) consists primarily of ethylene (C2H4) and ethane (C2H6) with trace amounts of acetylene (C2H2). The other consists primarily of H2 and CO. Each stream is driven by a pressure head produced by upstream compression equipment (K201), not a part of the AHR
3. The H2/CO stream is heated in a steam-driven heat exchanger (E413s) and then routed to an intersection with the C2 stream pipe.
4. The E1 facility is capable of sharing its hydrogen with NOVA's E2 facility, or of using E2's hydrogen if needed. E2's H2 can be routed into the E1 stream before or after heating in E413s, but E1's H2 can only be routed to E2 after heating. Differences in the content of H2 and CO in the streams will affect the reactions as described below.
5. Before it reaches this intersection and is mixed with the H2/CO stream, the C2 stream is heated twice. The first time is via the Reactor Cross Exchanger (E410) which uses hot effluent from the reactor (see below) to heat incoming, cooler C2 feed. The second is a steam-driven heat exchanger (E411).
6. 'Mixing' the two feed streams simply involves allowing them to intersect via a static turbulence inducer (SU-411). Following this, the mixed stream is allowed to flow into one of the two reactor vessels (the other is always off line and either undergoing regeneration or waiting to be put back on line).
7. The reactor vessels are currently filled with Dow Type-P Palladium catalyst which allow the following reactions to take place:
    - C2H2 + H2 → C2H4 + heat = "Acetylene Conversion"
    - C2H4 + H2 → C2H6 + heat = "Ethylene Conversion"
    - CO + 3H2 → CH4 + H20 = "CO reaction"
    - (with lots of heat and/or pressure) C2H4 → C + CH4 + lots of heat = "Ethylene Decomposition"
8. Acetylene conversion is desired. Ethylene conversion is undesired, but tolerable in small quantities. The CO reaction is used to regulate the other reactions as discussed below but it only operates within a narrow range and it produces undesirable side effects. Ethylene decomposition is highly undesirable and dangerous. Since it does not rely on the presence of hydrogen, reducing the H2/CO feed will not affect it. Instead pressure and/or heat must be reduced, and the quickest way to accomplish this is by venting to flare.
9. The catalyst has many weak and a few strong sites.

10. Precedence for 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
11. Thus, managing the reactor works as follows:
    - ensure that you've got enough CO in the reactor to occupy all of the strong sites
        - otherwise, ethylene will occupy those sites and be converted to ethane. This is both inefficient (you're trying to maximize ethylene content) AND dangerous—excess ethylene conversion can use up available H2 leaving none for acetylene resulting in "acetylene breakthrough" (getting too much acetylene in the AHR output).
    - Try to minimize CO so as to avoid occupying weak sites
        - CO on weak sites can mean not enough sites available for the acetylene reaction, thus, acetylene won't be fully converted and, again, you get breakthrough
    - Try to manage the ratio of H2/CO: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 won't be enough for total acetylene conversion, thus breakthrough
        - too much H2 (and/or too much heat) and, after all acetylene conversion, the last reaction (ethylene on weak sites) will occur and you'll get undesirable ethane.
12. Thus, CO is said to "improve selectivity of the catalyst" for the acetylene reaction.
13. Increased heat 'quickens' all reactions—that is, makes them more likely to occur. This increases the overall activity of the catalyst, but it 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, E411 and E413.
14. Increased pressure acts much like increased heat in making catalyst more active, but there is no convenient way to increase pressure in the reactor vessels. Decreasing pressure can be accomplished by routing feed or reacted product to flare.
15. All of the above reactions are stated as if they were absolute. They are not. They're stochastic.
16. Other reactions are possible given the presence of trace elements in the feed such as sulfur compounds, arsine, phosphine, halides and halogen. All of these have the effect of 'poisoning the catalyst'—that is, making it unreactive—but NOVA has never had these problems with the natural gas feed it uses in E1.
17. After reaction, the reacted product flow out of the reactors and downstream to the Reactor After Cooler (E412)—a heat exchanger driven by cool water. This cooler can be bypassed as well.
18. After E412, the reacted product stream can be diverted to E2, but is generally routed through the Reactor Cross Exchanger (E410) where it serves to heat the incoming C2 stream as described above. After E412, the reacted, cooled product stream proceeds out of the AHR subsystem to further refining (especially ethane separation) in the rest of the E1 facility.
19. Once the two input streams are mixed, they can be diverted to flare at many points in the AHR process. These include both before and from within the reactors, and before, from within or after E412. 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.

6. ADS Analysis of NOVA's E1 Acetylene Reactor

6.1 Overview of Analysis

Figure 6.1 shows an overview of, and will serve as a guide to, the ADS analysis of NOVA's E1 Acetylene Hydrogenation Reactor (AHR). We used five levels of abstraction (shown on the vertical axis of Figure 6.1) and three levels of decomposition (shown on the horizontal axis) in this analysis.

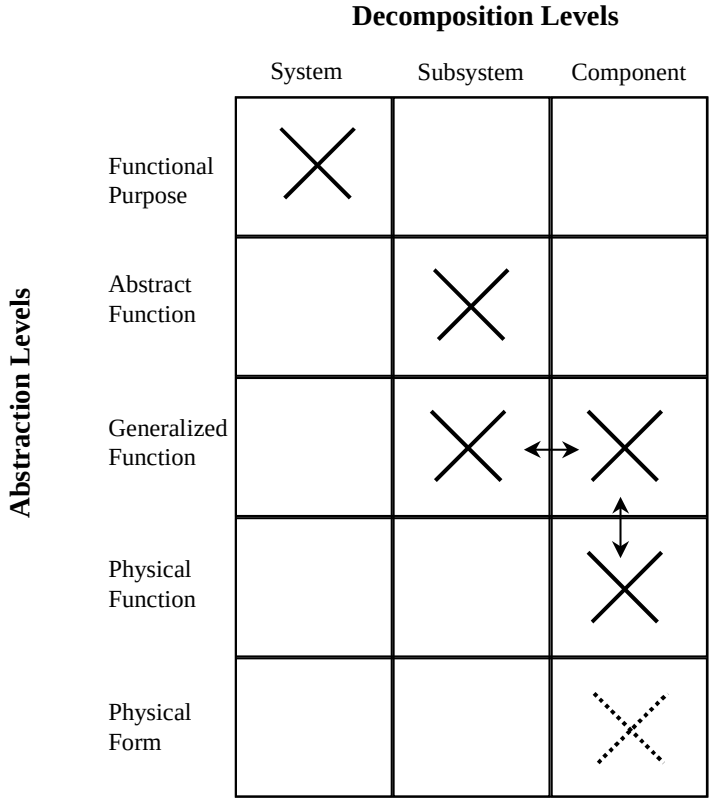

Figure 6.1. Matrix illustrating the Abstraction and Decomposition dimensions of the ADS for the E1 AHR. Xs indicate cells for which models have been prepared; arrows indicate prepared transition diagrams.

Each of the cells in the matrix shown in Figure 6.1 represents a specific model which could be constructed for the E1 acetylene reactor. The Xs in the cells of the matrix show which models were actually constructed and will be described below.

In practice, it has been repeatedly shown that constructing all of these models is rarely productive or efficient (Dinadis & Vicente, 1996; Bisantz & Vicente, 1994). The analyst must make judgements as s/he progresses as to the unique value or the novel information to be obtained by a new model for each cell. For example, we could construct a model of the functional purpose of each component of the AHR system, but generally speaking, the functional purpose of the components *is* the role they play in achieving the physical functions and generalized functions which lead to the overall functional purpose of the system. As long as we capture these relationships by building physical function and generalized function models at the component level, we will gain no new information by constructing a functional purpose model at that level as well.

Rasmussen (1985, see also Vicente, in press) collected a set of protocols of operators engaged in problem solving behavior with various systems. He mapped their verbal reports and actions against ADS models of the systems on which they were working and found that most problem solving activities fall along the diagonal from the Physical Form x Component cell to the Functional Purpose x System cell. This is, perhaps, not surprising. It is easier to think about high level (abstract) functions and purposes over the full set of equipment which accomplishes those functions and purposes, but when thinking about more concrete properties of the plant (that is, less abstract) such as equipment state, location, settings, etc. these properties are more easily considered at the level of specific components.

It is for these reasons that models were constructed for only 6 of the cells of the matrix in Figure 6.1, and that these cells generally fall along the top left to bottom right diagonal.

Finally, the double edged arrows in Figure 6.1 indicate that in addition to the models corresponding to the cells of the matrix, we have also provided "transition" diagrams. These diagrams will be discussed in more detail in section 6.4 below, but generally speaking, they are included to make explicit how entities at one level of the ADS are mapped into entities at another layer—for example, what components make up each

subcomponent, or what physical functions accomplish which generalized function, etc. Again, we have only provided those transition diagrams (as shown in Figure 6.1) which we believed to be most helpful; the others are implied.

6.2 Decomposition Dimension

As noted above, we have used three levels of decomposition in the Part-Whole dimension of this ADS analysis. In the direction of increasing decomposition, these are: system, subsystem and component.

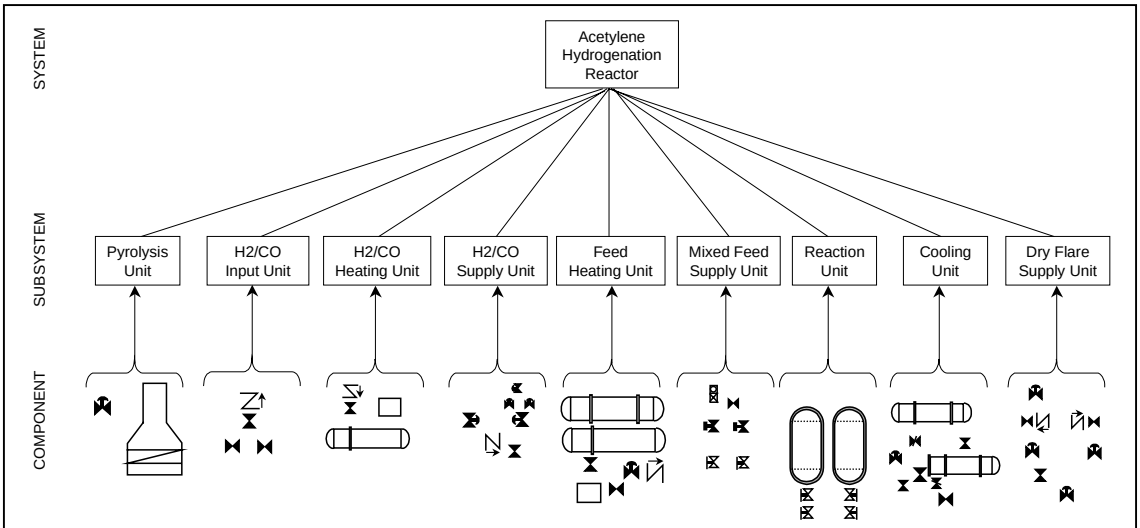

Figure 6.2. Part-whole decomposition of the AHR system.

Figure 6.2 shows this part-whole decomposition relationship graphically. There are 51 components[1] which have been aggregated into 9 subsystems which make up the overall AHR system. Subsystem boundaries are somewhat arbitrarily chosen, especially since operators on the E1 line tend to think of the AHR as a subsystem in the overall E1 facility and, while they do divide it into its components, they rarely create an intermediate level of aggregation between those components and the overall AHR. Nevertheless, it makes sense to create an intermediate level of decomposition to facilitate aggregation of information about the AHR system performance.

Table 6.1 lists the component names (taken from Figure 6.9) and the subsystems of which they are a part.

Table 6.1. Specific components which make up each subsystem.

Subsystem	Component
Pyrolysis Unit	FV135 H101-108
H2/CO Input Unit	VH2 VH3 VH4 CV1
H2/CO Heating Unit	E413 CV2 VS1

[1] Note that three components, E410 and VM12 and VM13 are part of two different subsystems: the Feed heating unit and the cooling unit, due to E410s dual role in system operations.

Subsystem	Component
	ST1199
H2/CO Supply Unit	TV440 CV6 VH5 FV413 SDV413 A SDV413 B SDV413 C
Feed Heating Unit	E410 E411 VM12 VM13 ST1052 TV410 CV3
Mixed Feed Supply Unit	SU411 VM1 MV410 MV411 VM4 VM5
Reaction Unit	R410A R410B VM6 VM7 VM8 VM9
Cooling Unit	E412 E410 VM10 VW1 VW2 VM14 VM12 VM13 HV41001
Dry Flare Supply Unit	PV441 CV4 VM2 PV410 A CV5 VM3 PV410 B VM11 PV412

6.3 Abstraction Dimension

As figure 6.1 shows, we have also used 5 levels of abstraction to describe the E1 AHR. As described in Section 4 above, these levels are linked by means-ends relationships with entities at the lower levels providing the 'means' by which 'ends' at the higher, more abstract levels may be achieved. We will provide general descriptions of each abstraction level here and provide specific descriptions of the models created in each cell in the following section.

Physical Form: The Physical Form layer is the least abstract or most concrete level of system description. At this level, emphasis is on the appearance and location of the physical entities which comprise the system. Connections between entities may be inferred from illustrations at this level, but they are more completely included at the Physical Function level.

Physical Function: At this level, we begin to abstract away from the reality of a particular physical entity and to talk about the function of that entity—albeit, still at a very concrete level. One indication of this is that names of entities are used for the first time at the Physical Function level—because we are now no longer referring to a specific physical entity (i.e., a particular piece of shaped metal), but rather to the entities *role* in the system (i.e., valve FV413). Whenever equipment can be manipulated to take on several different states (e.g., valve settings), the range of possible settings and the current setting are also represented at the Physical Function level. Especially at the component level, the Physical Function representation is similar to the traditional Piping and Instrumentation Diagram (P&ID) except that it does not include control or sensing instrumentation. Connections are explicitly included at the Physical Function level and represent the physical joining of entities.

Generalized Function: The Generalized Function layer represents the particular uses or roles that the physical functions are put to in the system as a whole. It no longer references the physical equipment specifically and, in fact, there may be a many to one or one to many relationship between Generalized Functions and Physical Functions. Typical concepts included at the Generalized Function layer are heat transfer, pressure propagation and the flow of commodities. We have decided, in the analysis reported here, to include chemical reactions at this level as well. Connections between nodes at the Generalized Function level represent causal connections, and the arrows associated with each connection indicate the direction of causal propagation. Note that double-headed arrows are entirely possible and indicate that each Generalized Function can affect the state of the other.

Abstract Function: The Abstract Function layer shows 'first principles' relationships pertinent to the domain. In industrial process control domains, this generally means mass and energy relationships in the work domain. Specific commodities and different types of heat transfer are no longer dealt with separately at this level, but are reduced to their role in the 'abstract functions' of mass and energy movement throughout the system. Connections at this level also reflect causality, as they did at the Generalized Function level.

Functional Purpose: The Functional Purpose layer is the most abstract of the five we have used. It states the overall purpose(s) for which the system has been designed. Functional purposes can generally be characterized by a simple statement of the value of a single variable. It is somewhat rare for even a small system in a complex, real-world setting to have only one purpose. Almost inevitably, the system may be achieving a specific work domain purpose (e.g., removing acetylene from C2 feed) but it will also be engaged in other functions such as safety, profitability, environmental protection, etc. We note that if these auxiliary purposes are not accomplished solely by the system being examined, but are instead a function of multiple systems in concert, then they ought not to be represented at the Functional Purpose level for the system under examination. With that caveat, we acknowledge that representing the interaction between multiple, and potentially conflicting purposes remains a topic for additional research.

6.4 ADS Cells

6.4.1 Physical Form x Component

The physical form layer of the abstraction space captures and represents information about the physical appearance and location of plant entities at the corresponding level of aggregation—in this case, the work domain components. Many prior uses of the ADS have involved modeling computer simulations. Since these have no objective physical form, this layer has frequently not been modeled.

Although we did not devote energy to capturing a full physical form representation for the E1 AHR system, we have obtained parts of it (hence, the dashed X in the Physical Form x Component cell in Figure 6.1). Figure 6.3 shows the general position and relative size of major AHR components in a 'God's eye view'.

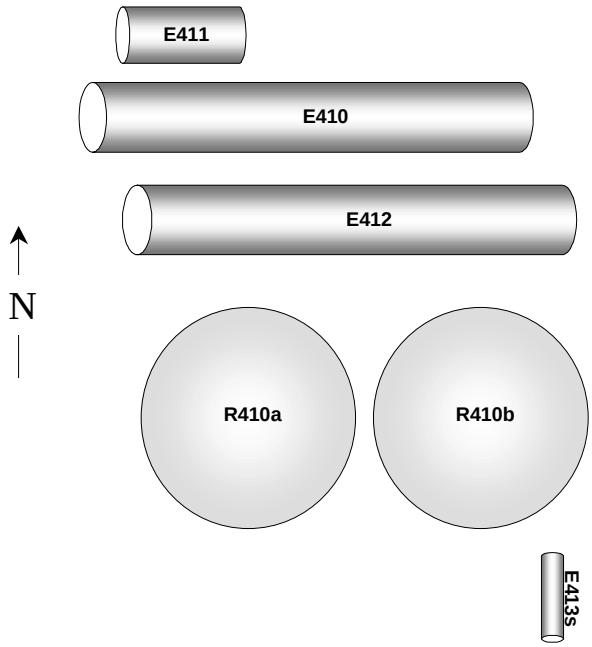

Figure 6.3. Physical Form x Component: Relative size and position of major E1 AHR components.

Figures 6.4-6.8 are photographs, taken at the E1 AHR site, showing the appearance of these same major work domain components. Figure 6.4 shows E411, the Reactor Feed Preheater and, in the background, the larger E410, the Reactor Heat Cross Exchanger. Figure 6.5 shows a more complete view of E410. Taken from the south, looking north, this view shows the relative size of E410 with the tip of E411 visible behind it. Figure 6.6 shows E412, the Reactor After Cooler, which is located to the south of E410 and about the same size. Figure 6.7 shows the much smaller E413, the Hydrogen Feed Heater, with our colleague Greg Jamieson standing nearby to show scale. Finally, Figure 6.8 shows the two reactor vessels themselves, R410 A and B, with E413 just visible in the foreground.

While these figures show portions of the physical form information for the AHR system, they are incomplete in a number of ways. First, as will be seen in section 6.4.2 below, while these figures show the location and appearance of the major components, there are many other components in the AHR system which are not depicted. Second, even for the components shown, while location and appearance are shown, connections are not shown (at least not clearly). Third, while appearance and location are shown for some components, what is depicted here is necessarily static and incomplete. Complete physical form information should include (for all components):

- The *current* appearance of all components (as might be provided by live video images or sight inspections).
- The accurate dimensions of all components (as might be provided by a blue print or CAD drawing), not just the relative sizes shown here.
- Where the component is located relative to known landmarks (something that might be provided by a site map).

Physical form information is important because it helps operators distinguish between the maze of equipment located in a complex system. Knowing where E413 is located and what it looks like will help an operator find it in

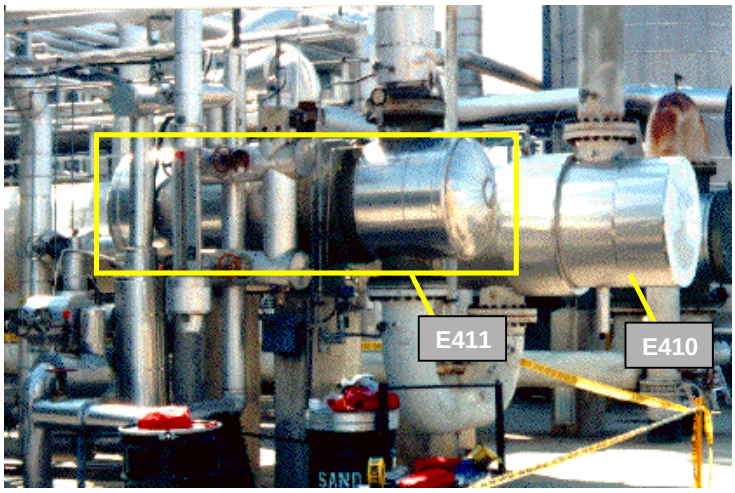

Figure 6.4. E411, the Reactor Feed Preheater and, behind it, the larger E410, the Reactor Heat Cross Exchanger.

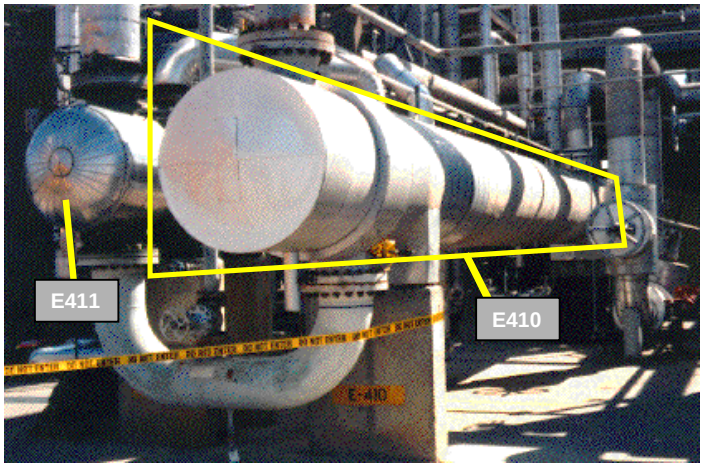

Figure 6.5. A view from the south, looking north, showing the full E410 and, behind it, the smaller E411.

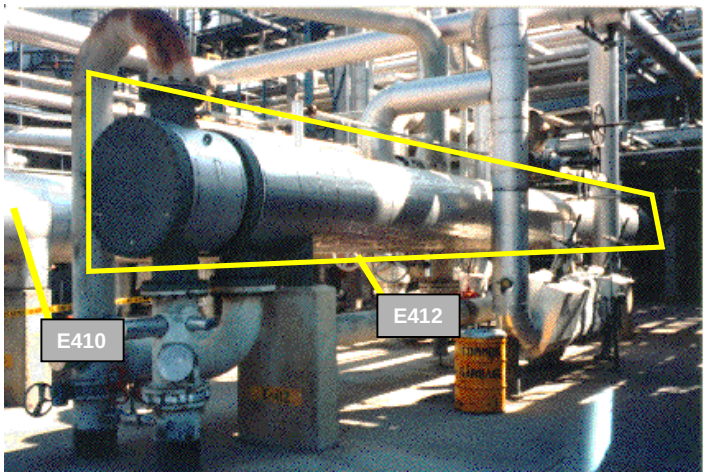

Figure 6.6. E412-- the Reactor After Cooler, located to the south of the similarly sized, E410.

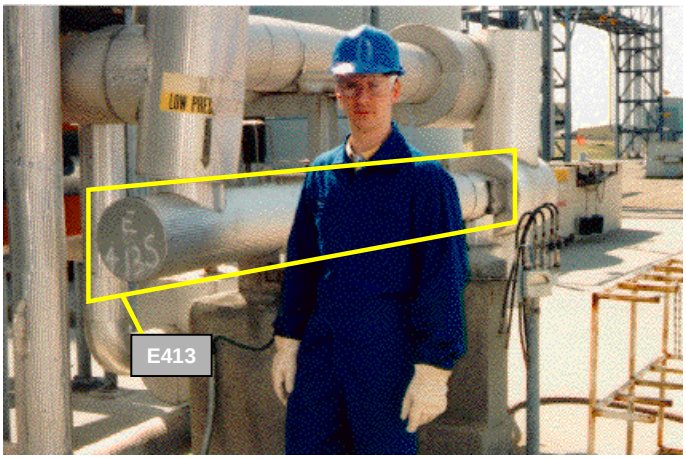

Figure 6.7. Appearance of E413, the Hydrogen Feed Heater.

the field if inspection or adjustments need to be made. Short of visiting it, know information about its location and appearance can help in causal reasoning. For example, just knowing the information conveyed in Figure 6.3 above enables one to predict that an explosion or fire in E411 would be much more likely to affect E410 than E413, since E413 is small, located far away and is protected by the bulk of R410 a and b. Finally, knowing current physical appearance (especially relative to past or expected appearance) is frequently the starting point for diagnostic reasoning.

6.4.2 Physical Function x Component

The Physical Function model for the components of the E1 AHR is shown in Figure 6.9. As discussed above in section 6.3, this figure should look much like a P&ID diagram for the AHR with the exception that instrumentation for control and sensing has been removed. What remains are the physical functions of the 'plant' itself. While the devices included in this diagram still retain something of the appearance of their real-world counterparts, this is primarily to help the user interpret the diagram. Relative size, location and actual appearance are not represented in this model, since they were included at the Physical Form x Component model above.

Constructing this model for a real-world system taught us just how complex such systems can be. The result in Figure 6.9 represents some compromises. First, while this diagram is significantly more detailed than the synoptic screens operator's currently deal with, it is somewhat less detailed than the P&ID diagrams for the E1 AHR. For example, NOVA's current board operator synoptic displays regularly

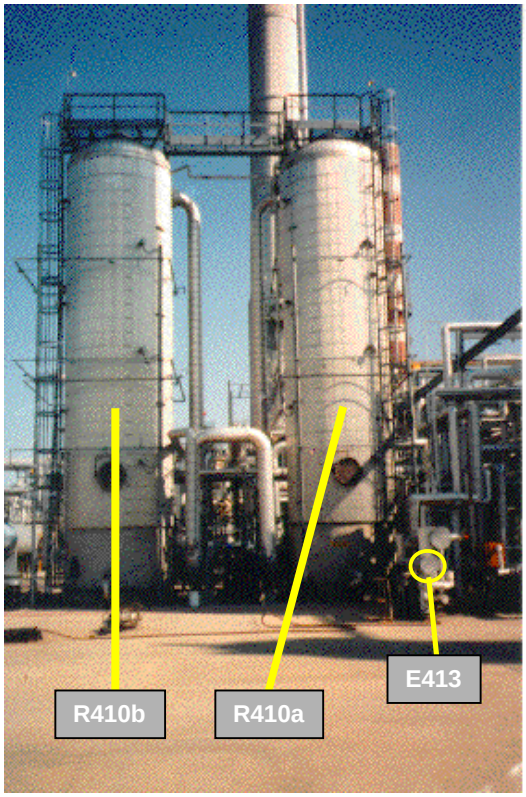

Figure 6.8. Reactor vessels R410 a and b, with E413 visible in the foreground.

omit the wealth of hand valves and check valves used throughout the system. These have generally been included in the Physical Function x Component representation. Similarly, the P&ID diagrams make it clear that most of the major components (and many of the valve sequences) have small drain valves attached to them. These have generally not been included in our model in the belief that they represent a finer level of decomposition than the component level we have chosen.

As always, deciding where to draw the boundaries for the 'system' is problematic and largely a matter of convenience and representational power. We chose not to include the product drying and green oil separation devices as a part of the AHR mainly as a matter of convenience. While many operators would tend to classify these devices as a part of the AHR, they do not have much impact on the reaction itself and instead impact downstream devices more heavily. Because they were conceptually separable, and because we needed to simplify the system under study, we chose not to include them. On the other hand, we chose to include the Pyrolysis furnaces and the DMDS feed to them as a part of the AHR system even though the are physically located far from the other AHR components. We made this choice because the composition of both the hydrocarbon feed and the H2/CO stream are critically determined by the performance of the furnaces and, furthermore, the AHR operator can influence these variables through the amount of DMDS added at that point. While this makes drawing boundaries around our 'system' geographically problematic, conceptually, it seemed best to include these components within the set of functions of the AHR.

One aspect of the abstraction performed at the Physical Function layer (over the Physical Form layer) is the assigning of names to components. We have tried to use NOVA's assigned names wherever possible. This proved problematic with regard to valves, however, for two reasons. First, NOVA does not (at least in the P&ID diagrams we had access to) assign names to hand controlled valves. Thus, we assigned names using the following convention: All valve names we assigned begin with a "V". The next letter refers to the type of material the valve is manipulating: S for Steam, W for water, H for the H2/CO stream, and M for the mixed H2/CO and hydrocarbon feed stream. Following this two letter designation, valves were numbered sequentially in the approximate order in which products passed through them in the AHR system. Thus, "VS1" is the first hand valve permitting steam flow in the AHR system; VM14 is the fourteenth valve whose role is permitting mixed feed flow.

The second reason that NOVA's chosen valve names are problematic is that they frequently refer to the type of control exercised on the valve. Thus, FV413 is a flow controlled valve, and PV412 is a pressure controlled valve. At a hardware level, however, if we do not consider the control approach (which is not appropriate to include in the ADS analysis), then all of these valves are physically the same. This situation is made clear by PV412 which is also referred to, in some NOVA documentation, as FV412 because there are two separate control loops which can activate the same physical device: one pressure based and the other based on flow. In the interests of maintaining familiarity for NOVA personnel and compatibility with NOVA documentation and displays, we have used the common NOVA names whenever they were available. It should be kept in mind, however, that we are designating the physical device, and not the controlling software, sensors, behavioral rules, etc.

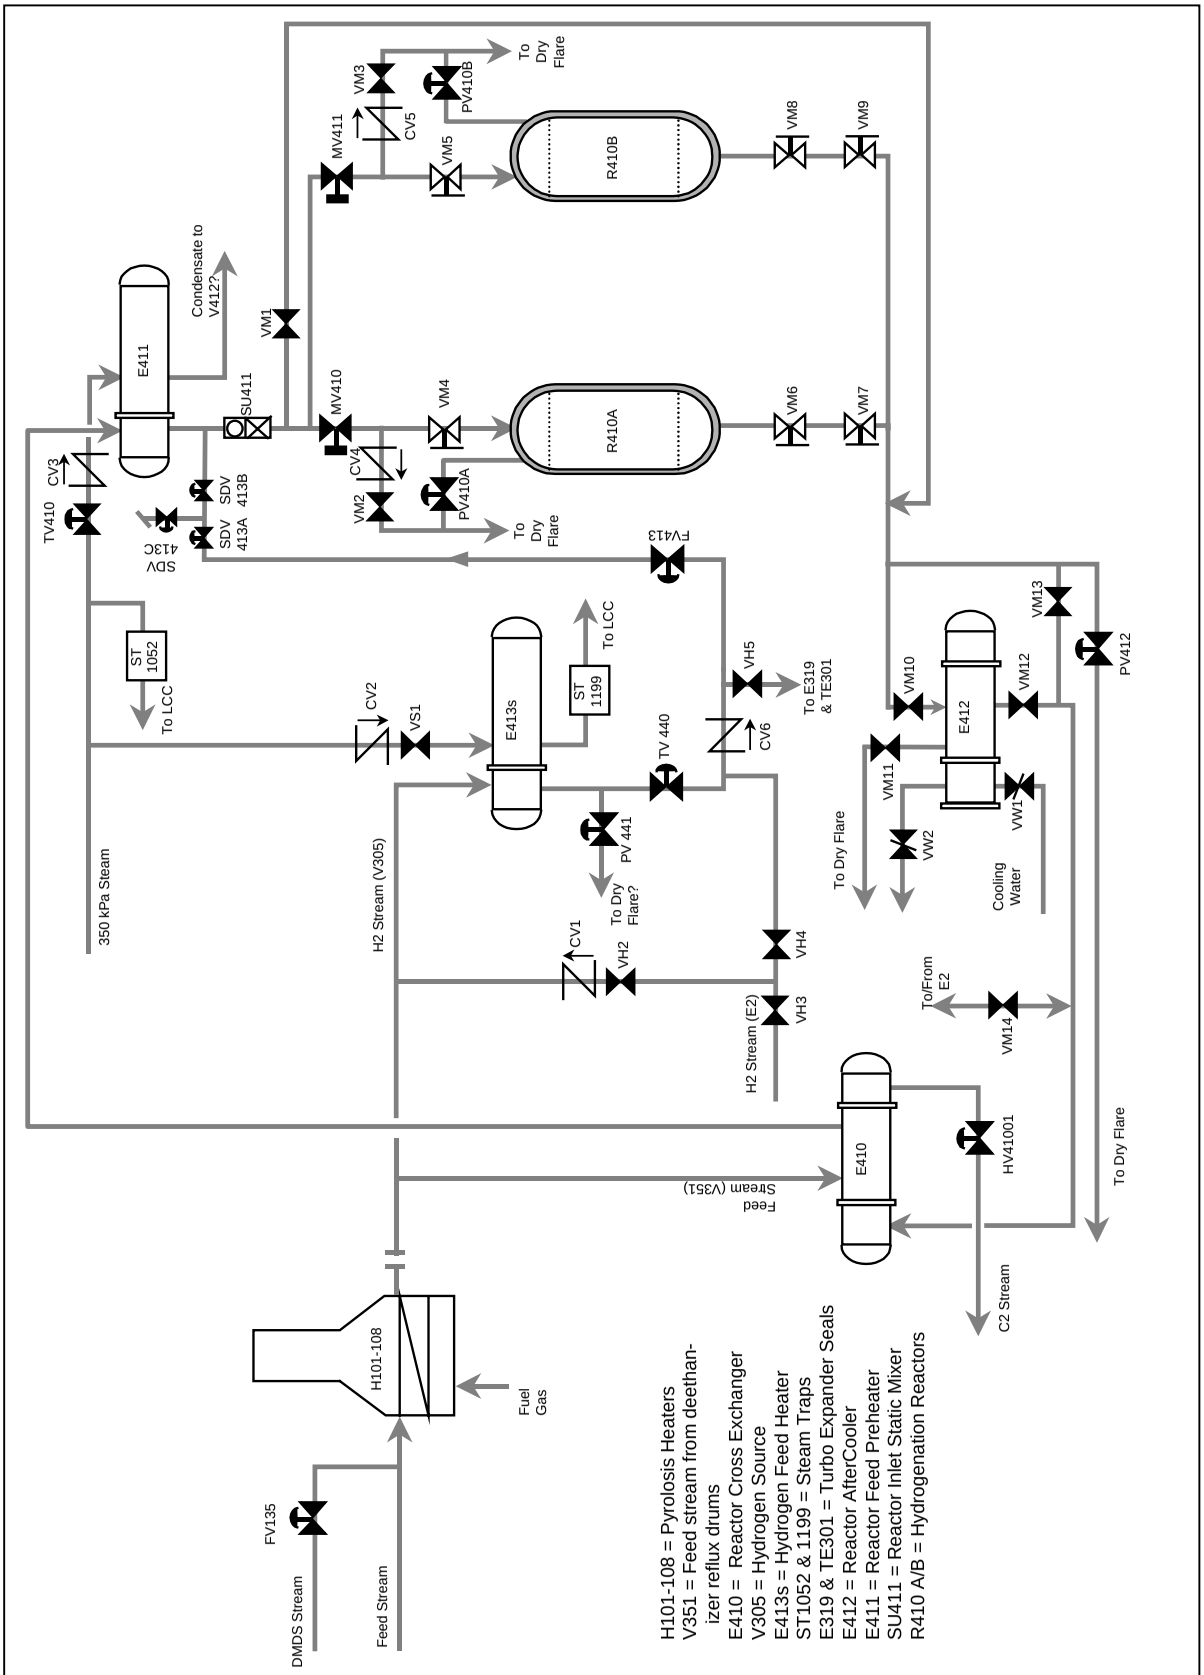

Figure 6.9. The Physical Function x Component level model.

To help remember device names, we have included a small key to prominent components in Figure 6.9.

Finally, recall that the connections illustrated at the Physical Function level are meant to represent physical connections. Thus, the gray lines in Figure 6.9 indicate physical connections (generally pipeline) between components. Arrows indicate *general* direction of flow, but special plant configurations and abnormal situations can upset these normal flows (except where check valves ensure flow in one direction only).

6.4.3 Generalized Function x Component

Figure 6.10 shows the Generalized Function x Component level model for the E1 AHR system. Generalized functions are at a level of abstraction higher than specific pieces of physical equipment, though they are of course, accomplished by them. Hence, in this model, there is no attempt to retain the shape or location of pieces of equipment. Instead, simple ovals are used to represent the general functions performed by the components of the system. Incidentally, the colored ovals and the use of wavy versus straight line connectors were included only to improve the comprehensibility of the very complex interactions around the chemical reactions. They have no other significance.

As is typical of petrochemical processes, generalized functions represented at this level usually include flow of materials and heat, pressures, and other general functions of plant equipment such as pyrolysis. Two aspects of this ADS are somewhat novel, however. These novel extensions of the ADS were worked out with Greg Jamieson and were also used in his analyses of a Fluidized Catalytic Cracking Unit (Jamieson, 1998; Jamieson and Vicente, 1998). They represent one innovation of this work which will be useful beyond this project itself.

The first novel extension is that flows of material are differentiated as to type. Other ADS analyses, especially those performed on the DURESS II system (e.g., Bisantz & Vicente, 1994; Vicente and Rasmussen, 1990), but also those for aviation applications Dinadis & Vicente, 1996), have generally not needed to worry about the composition of different elements within a flow. Those differences are obviously critical, however, in many petrochemical processing applications including the AHR. Thus, pyrolysis produces five different types of flows (C2H2, C2H4, C2H6, CO and H2) which are of interest for the AHR operator. These are separated into two streams and enter the Reactor by two different routes: the C2 products are, in fact, intermingled in the same pipeline flow, thus they are shown here as going through the same heat transfer processes. The CO/H2 stream is also routed through a single pipe, but it enters the AHR via a different route and is shown here undergoing a different set of heat transfer functions before the two streams are mixed and sent together to the reactor functions.

The second extension is the modeling of chemical reactions themselves. Again, prior uses of the ADS have not included chemical reactions, but they are key to the overall purpose of the AHR. We explored many different alternatives for representing chemical reactions within the ADS, but the most natural seems to be 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 others and how—information critical to understanding the AHR. 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, though at different rates. This fact is partially implied in Figure 6.9 by the fact that heat, in addition to the reaction products, flows out of the reactors.

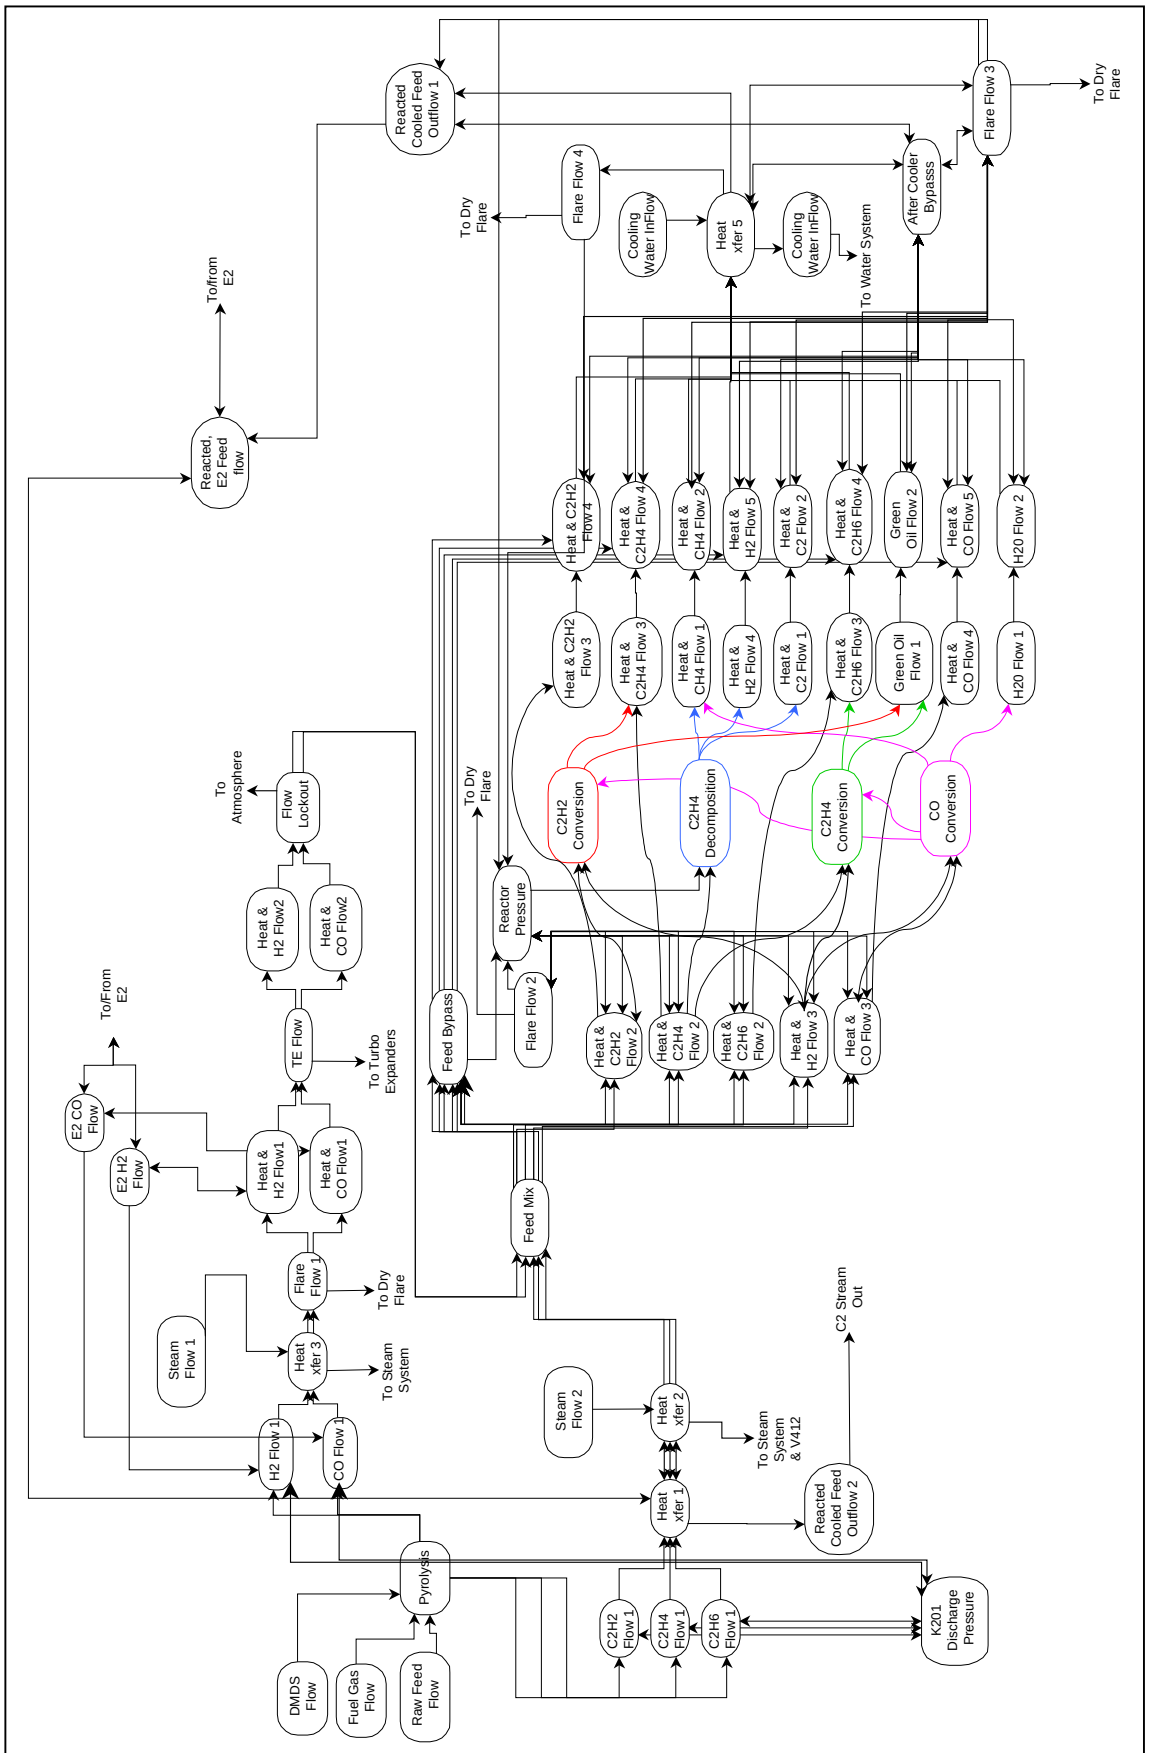

Figure 6.10. The Generalized Function x Component level model.

Connections depicted at the Generalized Function level represent causal links between the nodes and arrowheads indicate the direction of causality. For example, Figure 6.10 shows that DMDS flow, Fuel Gas Flow and Raw Feed Flow are all required to cause Pyrolysis. On the other hand, Pyrolysis is not required to cause any of these flows, but it is required for the flow of the five feed components described above: C2H2, C2H4, C2H6, CO and H2. Double-headed arrows indicate that causality can propagate in either direction—for example, E2 H2 flow can either cause Heat and H2 Flow 1, or it can be caused by it (reflecting the fact that E1 can either use hydrogen supplied by E2 or can supply hydrogen to it). By tracing causal connections around the reactions, it is possible to ascertain which feed components participate in which types of reactions and produce which products out. Several important facts can be deduced from these causal connections about the reactions. For example, (1) carbon monoxide can only participate in the CO Conversion reaction and does not directly participate in any of the other reactions, (2) the only way to get acetylene (C2H2) in the output of the reactors is by failing to convert it, and most critically, (3) that an H2 source is not required for there to be ethylene decomposition, but that adequate reactor pressure is required for this undesirable reaction.

The Generalized Function x Component and Physical Function x Component representations shown in Figures 6.9 and 6.10 are extremely complex. Much of this complexity is reduced in normal operations (and in Board Operator's displays) by eliminating consideration of little used routes, field-controlled hand valves not under the operator's control, etc. While that reduction of complexity for operations may be necessary, the fact that operators need to consult flow sheets and P&IDs implies that this information *is* critical—and least in some circumstances. It is important to include all possible functions in the ADS analysis, especially if it is to serve its role in aiding abnormal situations. As will be seen in section 6.4.5 below, some complexity reduction can also be achieved by aggregating components along the Part-Whole dimension and viewing Generalized Functions at the subsystem rather than the component level.

6.4.4 Physical Function to Generalized Function Transition Diagram at the Component Level

Figure 6.11 explicitly shows how Physical Functions (the specific roles played by physical objects in the system) map into Generalized Functions. This figure corresponds to the double-headed arrow which bridges the Physical Function x Component cell to Generalized Function x Component cell in Figure 6.1. The figure is arrayed in four rows for convenience. Within each row, there are a set of objects at the lower level linked to one or more nodes at the higher level. The objects at the lower level are taken from Figure 6.9 and are the Physical Function x Component entities in that figure. The objects at the higher level are taken from Figure 6.10 and are the Generalized Function x Component entities.

In addition to serving as an organizer and a check for the analyst (if all entities from both Generalized Function and Physical Function layers don't show up in this transition diagram, something is wrong), Figure 6.11 can be used to clearly show the utility or purpose of equipment in the AHR system and, conversely, show the equipment which supports AHR system functions. Since the transition from Physical Function to Generalized Function is one of abstraction along a means-ends dimension, we can interpret the transitions shown in this diagram as means-ends links. Thus, the set of Physical Function objects in each cluster can be said to be the 'means' and the Generalized Function nodes in each cluster are the "ends". Each cluster, therefore, represents the means available within the system to achieve a specific set of ends. For example, the means by which the generalized function 'Pyrolysis' can be achieved within the AHR system are the components H101-108. Some ends are accomplished by multiple means: Steam Flow 1 is accomplished by a combination of CV2, VS1 and ST1199. Similarly, some means accomplish multiple ends: TV440 and CV6 together accomplish both Heat and H2 Flow 1 and Heat and CO Flow 1 (not surprising, since they are valves regulating the combined flow of these two gases). The Reactor vessels R410A/B are the means for the chemical reactions, and for the pressure which makes them possible[2]. Whenever these sorts of multiple connections occur, they represent a potential for confusion by operators

[2] We considered representing the reactor catalyst beds, with their strong and weak catalyst sites, at a finer level of decomposition—a "sub-component" layer—but thus far have not found sufficient need for a representation this explicit.

who may forget the full set of means required to accomplish an end, or may overlook the fact that a given means affects more than the single end they are currently interested in.

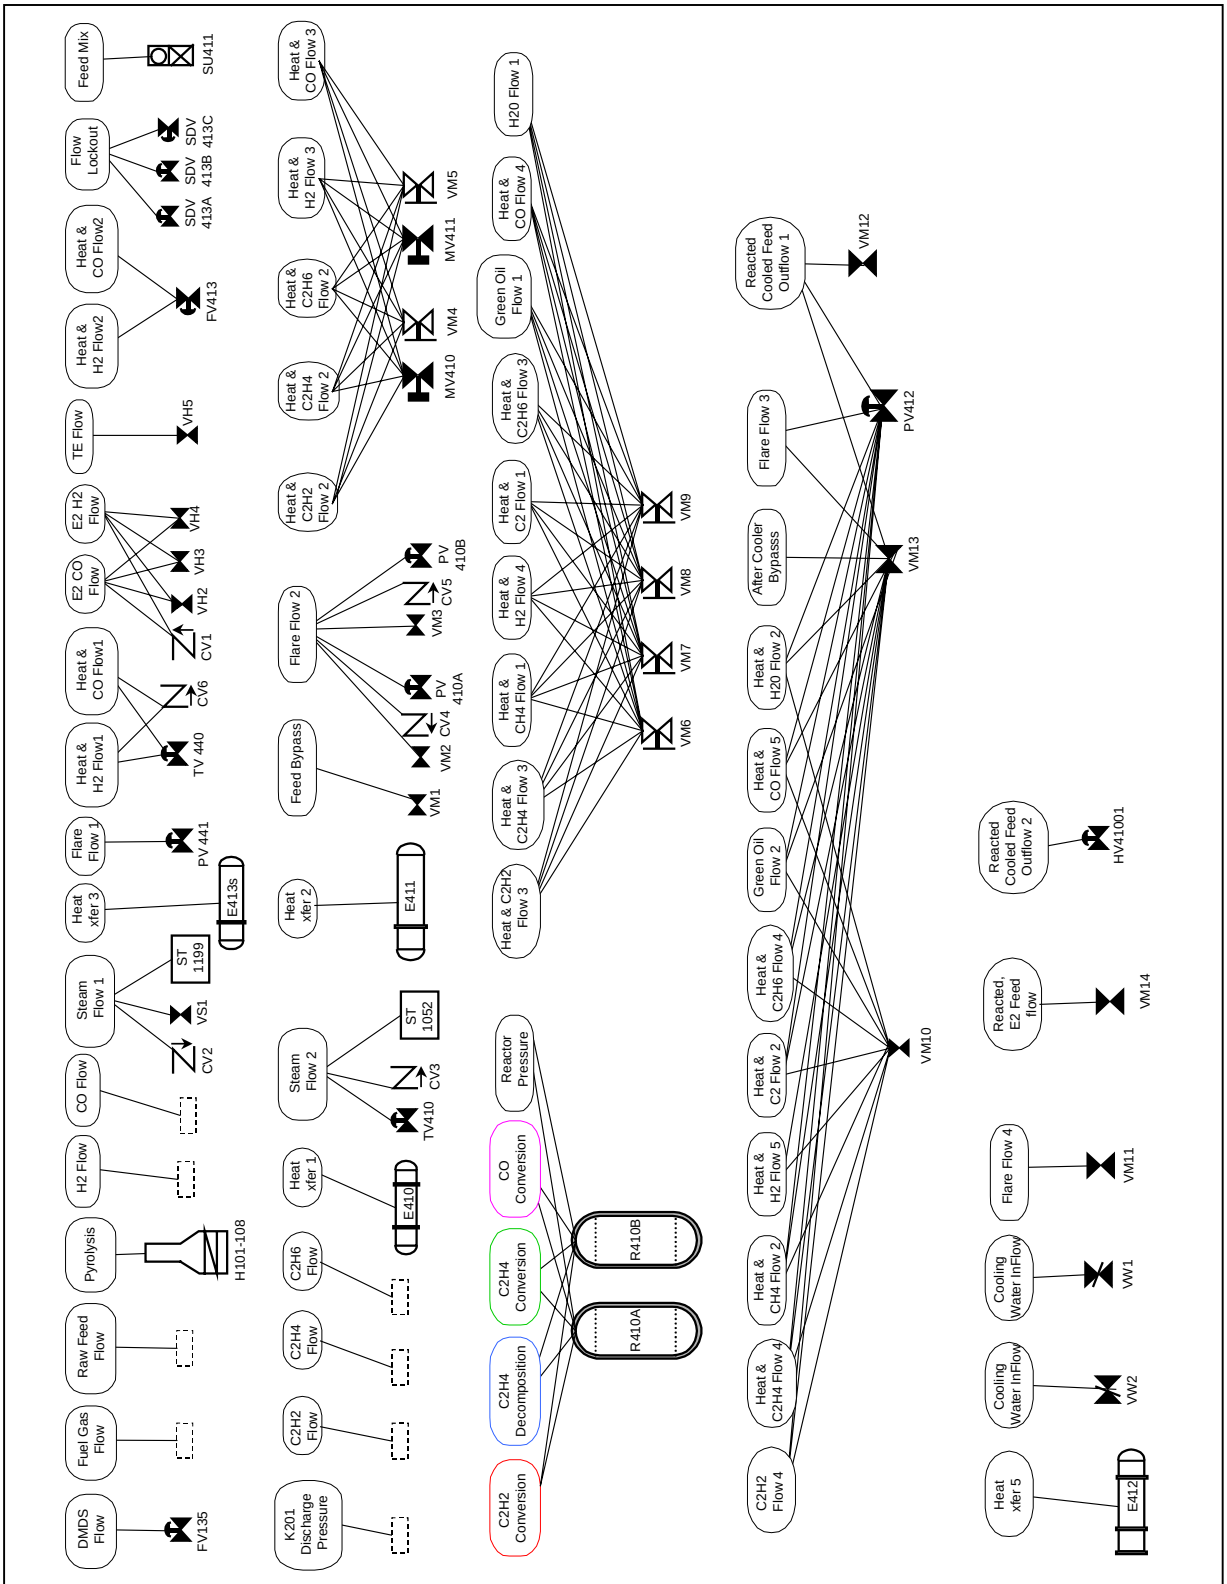

Figure 6.11. Transition diagram showing the links between Physical Functions and Generalized Functions at the Component Level.

Note that some Generalized Functions in Figure 6.11 have no corresponding means (this is represented by an empty, dashed box at the Physical Function level). This is not to imply that these function happen spontaneously, but rather than there are no direct means within the AHR system for achieving or affecting them. This is, of course, related to where we have drawn the boundaries of the AHR system, but the phenomenon of having important process variables be determined by upstream units outside one's control is a common one for industrial process operators. This technique clearly illustrates that the AHR operator has no means of taking action (with the AHR system) which will affect the flow or the pressure of C2 products into E410. Similarly, the E410 system provides no means for affecting the Raw Feed or Fuel Gas flow into the Pyrolysis furnaces, but there *is* a means to affect the addition of DMDS—and this is the operator's primary method of affecting the content of CO entering the AHR system.

6.4.5 Generalized Function x SubSystem

Figure 6.12 represents the first step along the Decomposition axis of Figure 6.1. In this model, we are no longer talking about component-level entities in the AHR system, but have instead moved up a level of aggregation and are now showing relationships between subsystem entities. On the other hand, because this is still a model at the Generalized Function level of abstraction, the same types of functions are described—flows of material and heat, general functions of the AHR system—but these are now described at a more aggregated level than they were in the Generalized Function x Component model. In fact, although we have not deemed this of high enough value to pursue, we could construct a Physical Function model of subsystem devices (drawing from the units described in Figure 6.2 above), and then construct a transition diagram across these subsystem levels similar to that shown for the Component cells in Figure 6.11 above.

As before, the arrows represent causal links between the nodes. Thus, we can tell that Pyrolysis affects the input of both the combined H2/CO stream and the C2 Feed Stream; and that Flare Flow can be achieved from either the Mixed Feed Input or the Product Cooling functions, but that it will affect them in turn.

Moving up a level of aggregation has greatly simplified the model shown in Figure 6.12 compared to that in Figure 6.10. This level of aggregation would place much fewer demands on user's attention, and in some ways it may seem more 'natural' to them. For example, the multiple constituents of each initial feed stream were somewhat artificially kept separate in the Generalized Function x Component layer—even though they exist in two aggregate gas mixtures. This decomposition is useful because, as we saw in Figure 6.10, it enabled us to talk about how the different chemical reactions took place. On the other hand, the level of aggregation in Figure 6.12 is useful for other purposes—such as monitoring general heat transfers or realizing that there *are* two and only two distinct sources of input to the mixed flow.

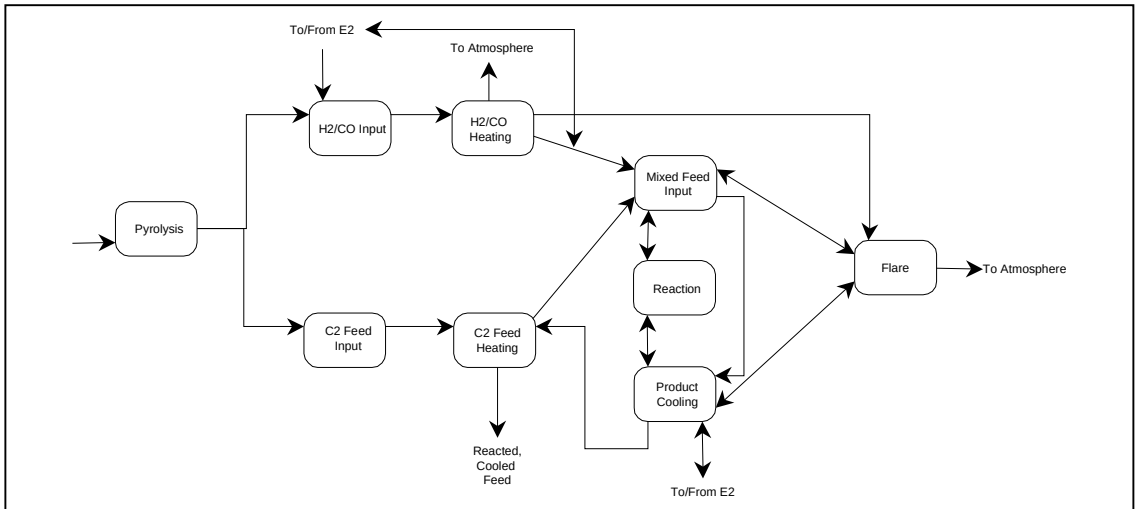

Figure 6.12. The Generalized Function x Subsystem model.

Maintaining the distinction between these levels of aggregation will help inform displays useful for both types of purposes.

6.4.6 Component to Subsystem Transition at the Generalized Function Level

Table 6.2 shows how Generalized Functions at the Component level (see Figure 6.10) have been aggregated to form Generalized Functions at the Subsystem level (see Figure 6.12). This table is the expansion of what is conveyed by the horizontal, double-headed arrow in Figure 6.1.

Note that even though we have moved up a layer of aggregation/decomposition, we have remained in the same abstraction layer. Thus, the concepts represented, generalized functions of the AHR system, remain the same. In the Generalized Function x Component model, concepts represented flows of material and heat, chemical reactions, etc. In the Generalized Function x Subsystem model, the same types of concepts are represented, but now they cover a more aggregated set. For example, instead of representing individual flows of reactants, aggregated flows such as the "H2/CO stream" or the "C2 Feed Stream" are now represented.

Table 6.2. Generalized Functions at the Component level which comprise Generalized Functions at the Subsystem level.

Generalized Function x Subsystem	Generalized Function x Component
Pyrolysis	DMDS Flow Fuel Gas Flow Raw Feed Flow Pyrolysis
H2/CO Input	K201 Discharge Pressure H2 Flow 1 CO Flow 1 E2 H2 Flow E2 CO Flow
H2/CO Heating	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
C2 Feed Input	K201 Discharge Pressure C2H2 Flow 1 C2H4 Flow 1 C2H6 Flow 1
C2 Feed Heating	Heat Transfer 1 Steam Flow 2 Heat Transfer 2
Mixed Feed Input	Feed Mix 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
Reaction	Reactor Pressure C2H2 Conversion C2H4 Conversion

Generalized Function x Subsystem	Generalized Function x Component
	CO 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 Green Oil Flow 1 H2O 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 Green Oil Flow 2 H2O Flow 2
Product Cooling	Cooling Water Inflow Cooling Water Outflow Heat Transfer 5 After Cooler Bypass Reacted Feed Outflow 1 Reacted E2 Feed Flow Heat Transfer 1 Reacted Feed Outflow 2
Flare Flow	Flare Flow 1, Flare Flow 2, Flare Flow 3, Flare Flow 4

6.4.7 Abstract Function x Subsystem

Abstract Functions are concerned with the 'first principles' or most basic representation appropriate to the domain. Abstract Functions in process control domains are generally concerned with mass and energy relationships in the plant—and the NOVA E1 AHR is no exception. The representation used to characterize Abstract Function relationships is adapted from Multi-level Flow Modeling (MFM—Lind, 1994). MFM uses six symbols to represent six highly abstract functions. These are presented and summarized in Figure 6.13.

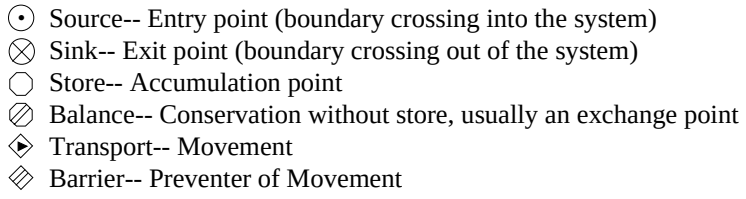

Figure 6.13. Symbols used in the MFM representation and their meanings.

Figures 6.14 and 6.15 move up one layer of abstraction from Figure 6.12, showing the Abstract Function x Subsystem model of the AHR system. The mass and energy models appear very similar—and this is to be expected. Most manipulations in the AHR process affect both the mass and energy relationships of the system as a whole. There only difference, however, is that the heat exchangers (E410-E413) show up as barriers in the Mass model and as in the Energy model. This reflects the nature of the heat exchanger technology. All heat exchangers on the E1 AHR line are variations on shell and tube mechanisms where hot material

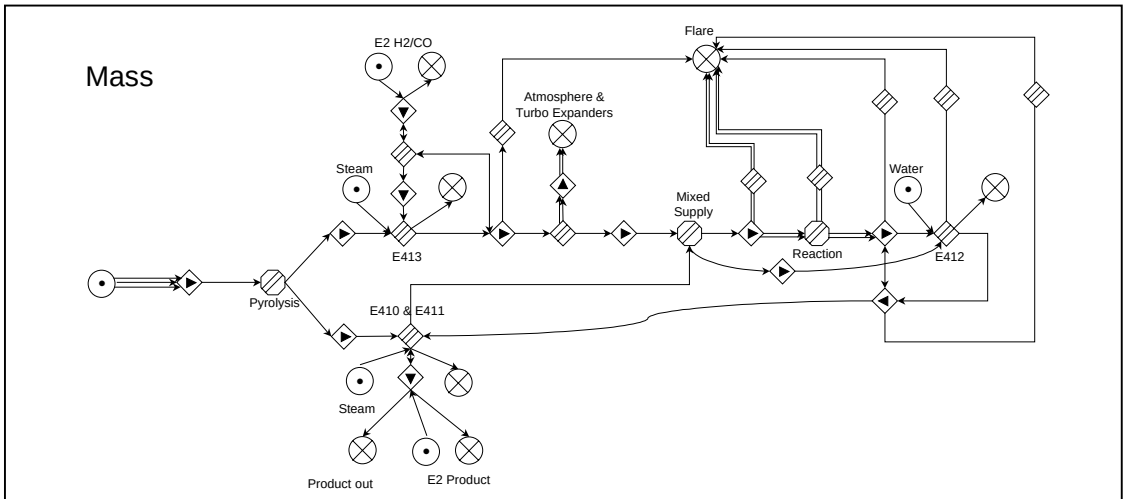

Figure 6.14. The Abstract Function x Subsystem Model of Mass Relationships.

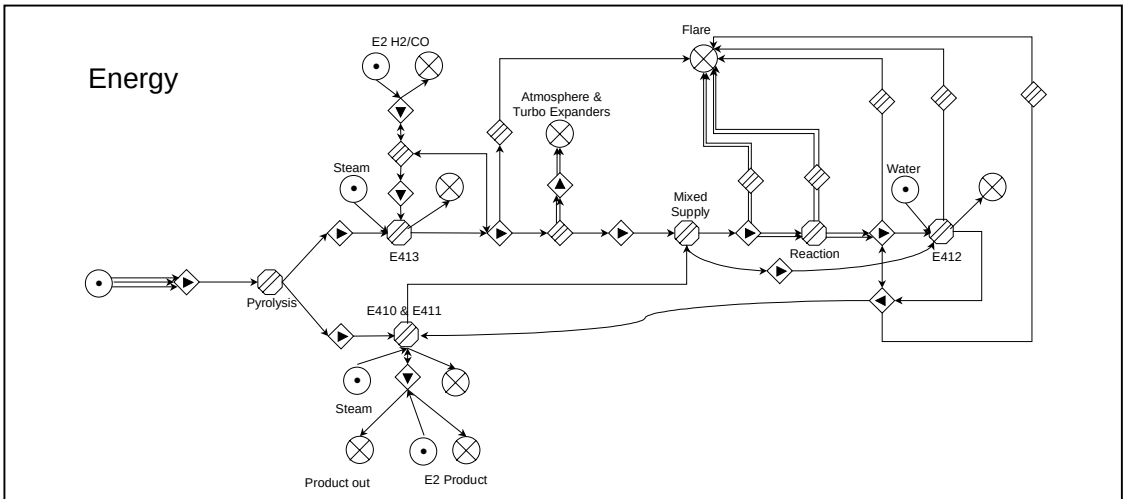

Figure 6.15. The Abstract Function x Subsystem Model of Energy Relationships.

(typically steam, but in E410 the hot product from the reactors is used) is passed through tubes in a shell through which a cool material is passing. Thus, while energy passes from the hot material to the cool material, the two materials never come in contact with each other. Thus, from the mass perspective, the shell and tube heat exchangers act as barriers, while from the energy perspective, they act as balances.

These models can be used to inform the designer or the operator about the behavior and capabilities of mass and energy in the system. For example, we can tell that it is possible to get rid of mass and energy from one input stream (the H2/CO stream) before it is mixed with the other because there are several sinks along this path, but this is not the case for the other stream. Also, we can see that both mass and energy should balance before and after the reactions take place as long as none is flared to atmosphere. Thus, even though the reactions are exothermic, overall mass and energy should balance across the reaction—albeit changed in form from molecular bonds to heat. Abstract Function level representations are particularly powerful in that they exploit conservation relationships that do not appear explicitly at any other level of the ADS.

6.4.8 Functional Purpose x System

The final cell of the ADS to be modeled in our analysis is the Functional Purpose abstraction of the full AHR system. This 'model' is a very simple statement of the purpose of the AHR as a whole as presented in Figure 6.16. In operations, the success of the system in achieving this purpose can be measured by a single variable—the amount of C2H2 remaining in the feed stream as it leaves the AHR.

The Functional Purpose of the Acetylene Hydrogenation System is to reduce the proportion of acetylene in the C2 product stream (normally to less than 5 ppm).

Figure 6.16. The Functional Purpose of the AHR system as a whole.

7. Test Case: Explaining an AHR upset incident

A NOVA memo, dated 30 May, 1985, details a 'reactor runaway' incident (involving ethylene decomposition) which had recently happened on the E1 ethylene site. This incident occurred during a shutdown of the E1 facility when unusual conditions in a system upstream of the AHR—the Demethanizer—failed to perform its normal function of separating hydrogen from the C2 Feed stream entering the AHR. The fact that this incident occurred during shut down, when many operational variables are not operating in their 'normal' ranges, and the fact that the presence of hydrogen in the C2 Feed was unexpected (perhaps even unprecedented) made this incident extremely difficult to diagnose and to manage. In some cases, as can be seen in the incident report quoted below and our analysis of it, actions which would normally serve to address AHR imbalances and bring it back into safe and normal states, caused further, dangerous reactions under the set of circumstances which prevailed.

As a test case of the AHR ADS models we have built, we will describe this incident below and then attempt to trace its causes, the structural states produced in the plant during the incident, the motivation of the operators as they attempted to address it, and the effects of their actions using the ADS model. If the ADS model can explain the incident, this is evidence that it is correct and reasonably complete. If it can be used to identify problem causes, to suggest courses of action and/or to predict their outcomes, then it is probable that a display built from the ADS analysis and conveying its information would also have helped in diagnosing and reacting to this incident.

Below, in Table 7.1, we first quote the sequence of events reported in the NOVA incident investigation (Cook and Poole, 1985). Then we will review the cause of the incident as reported by NOVA and retrace both the cause and the sequence of incident events with reference to the ADS analysis described above.

7.1 Sequence of Events in Reactor Runaway Incident

Table 7.1. Sequence of Events in NOVA E1 Reactor Runaway Incident.

Time	Events
11:26	▪ Reactor tripped on high temperature ▪ Reactor bypassed ▪ Reactor inlet MOV closed ▪ Hydrogen to reactor physically blocked in ▪ Temperatures increasing rapidly – ALERT sounded ▪ Temperature reached 186 degrees C – EVACUATE sounded ▪ Deluge on Reactors tripped ▪ Phone communications with E2, asking them to go on ALERT ▪ There was a request to close bypass, this was rejected by [name omitted] as plant EVACUATION alarm was sounded. ▪ MOV on inlet was open to provide flow through reactor to cool bed ▪ Temperature still increasing

	▪ MOV on inlet closed ▪ Depressuring to flare via PV412
11:36	▪ Temperatures still increasing at 5-6 degrees per punch to approximately 400 degrees from 200 degrees ▪ MOV on inlet opened again after temperature increase stabilized ▪ Temperature started to decrease ▪ Reactor bypass blocked in once safe for Operators to do so
11:54	▪ Temperature increasing on bed again ▪ MOV on inlet closed ▪ Temperatures decreasing ▪ MOV opened on inlet
12:00	▪ Temperatures down under 200 degrees
12:03	▪ Temperatures down under 100 degrees ▪ The ALL CLEAR was now sounded. All temperatures less than 32 degrees ▪ [Names omitted] noticed that the reactor was still on spec, therefore they realized that to react there was Hydrogen coming in with the feed.
12:46	▪ High temperature in on bed – 100 degrees ▪ Bed blocked in, reactor bypassed ▪ Nitrogen purge started on reactor ▪ The flare system was inspected after ALL CLEAR, damage was realized ▪ Suggestion to trip K201 during runaway was declined as to keep flow through reactor ▪ Reactor was pressure purged and flow purged to flare numerous times over remainder of shift and through the night shift.

7.2 Cause of Reactor Runaway Incident

The cause of the reactor runaway incident, as cited in the NOVA memo, was:

"No fractionation in demethanizer, resulting in concentrations of Hydrogen in feed to reactor."

7.3 Incident Analysis using AHR ADS

7.3.1 Analysis of Incident Cause

In this section, we will trace the sequence of events and relate them to the models from the ADS analysis. Events from the incident are repeated in italics. We begin with the incident cause and its implications below.

"No fractionation in demethanizer, resulting in concentrations of Hydrogen in feed to reactor."

The Demethanizer is a system within E1 that lies between the Pyrolysis furnaces and the rest of the AHR as described in sections 5.2 and 6. It lies in the gap illustrated in Figure 6.9 between the output of the pyrolysis furnaces and the separate H2/CO and C2 Feed Streams entering the rest of the AHR—in fact, the Demethanizer is primarily responsible for separating the H2/CO stream from the C2 Feed (and other products of pyrolysis) … when it is working properly. In this incident, due to abnormal conditions during shutdown, the Demethanizer was not working properly and allowed a proportion of hydrogen to proceed down stream in the C2 Feed stream.

While we have not analyzed the Demethanizer (and, therefore, can say little about how it works) the results of its functioning are clearly apparent in the Physical Function x Component (Fig. 6.9), Generalized Function x Component (Fig. 6.10) and Generalized Function x Subsystem (Fig. 6.12) models. We can trace the route of the contaminated C2 Feed using Figure 6.9—into E410 on to E411 and to SU411 where it was mixed with more H2, and then into one of the two reactors. The Generalized Function layer provides more interesting information—including some predictions about the results of this cause. In terms of the

generalized function model presented in Figure 6.10, there was an added, unusual H2 flow (in addition to the three C2 flows) into heat transfer 1 and on into the reactions. The description of the conversion reactions shown in Figure 6.10 predict that this added flow of H2 would increase the conversion of CO, C2H2 and C2H4, since more H2 enables more reactions to take place. Figure 6.10 also shows that the outcomes of these added reactions should be more C2H6 in the output than normal (due to a higher ratio of C2H4 conversion) but, presumably, no increase (and perhaps even some decrease) in the output of C2H4 since (a) virtually all C2H2 is converted to C2H4 even under normal circumstances, and (b) some C2H4 is being converted to C2H6. Most importantly, the knowledge about the reactions captured at the Generalized Function level shows that more conversion reactions will produce more heat—and if that heat rises sufficiently (200 degrees C is the notional threshold), ethylene decomposition will begin to become common, resulting in much higher heat production and self-sustaining reactions. Another implication visible from the Generalized Function x Component or x Subsystem figures is that the rising heat in the reactors and in the reactor effluent will be partially transferred to incoming C2 feed as it passes through E410 (the Reactor Cross Exchanger), thereby resulting in added heat flow back into the reactor via the incoming C2 Feed.

These energy relationships are even more visible at the Abstract Function x Subsystem level for energy relationships (Figure 6.15). Here we can see that increased energy in the system can be removed (up to some limit) by the energy sink associated with cooling water in E412, but that any additional energy will flow to the energy balance associated with the E410 Cross Exchanger from which it can either be routed out of the AHR system or will be routed back into the Mixed Feed Supply in the C2 Feed Stream.

The mass relationships at the Abstract Function x Subsystem level are perhaps less informative, but generally speaking, the Abstract Function level models are useful in rapidly identifying the source or effects of upsets. If sensitive enough measures could be used, they would have shown which equipment and processes were being affected and which were not in this incident. For example, there may have been an increase in mass flow (due to the unexpected hydrogen content) through E410 and E411 into the mixed feed supply and on into the reaction. Similarly, as will see below, after actions were taken on the H2/CO stream through E413, there was neither mass nor energy flow by that route—thus informing the operators earlier that the problem must lie elsewhere.

Finally, and critically, by tracing means-ends and part-whole links from the Physical Function x Component level to the Functional Purpose level we can see that the effect of excess hydrogen will *not* have an affect on the overall functional purpose of the system. If anything, excess hydrogen will result in an even higher proportion of acetylene conversion and, therefore, enable the system to continue to fulfill its purpose. This is what was observed in this incident.

7.3.2 Analysis of Incident Event Sequence

In this section, we step through events from the incident sequence (again, repeated here in italics) and discuss the ability of the ADS models to account for them.

- *Reactor tripped on high temperature*

The reactor trip is an automated system designed to shut off the H2 supply to the reactors via the SDV valves (SDV 413A, B and C) whenever temperatures in the reactors get dangerously high (at 110 degrees C). While the sensors and control logic are not modeled as a part of the ADS, the results of these actions are. The combined heat flows before and after the reactions (see Figure 6.10) are the temperature being monitored by the reactor trip function, and the result of the trip is to close SDV 413A and SDV 413B and open SDV 413C in order to vent H2 to atmosphere and prevent it being mixed with the C2 Feed stream and entering the reactions. The commanded settings of these valves appear at the Physical Function x Component level (Figure 6.9) and the expected effect on flow of H2/CO to Feed Mix appears at the Generalized Function x Component level (Figure 6.10). If the SDV valves were actually in their commanded settings (apparent at the Physical Form level), and if H2/CO Flow Lockout actually occurred (apparent at the Generalized Function level), then there should be no Heat and H2 Flow 3 into the reactors, no conversion reactions should take place and, thus, C2H2 flow out of the reactor should increase, but heat

flow within the reactor and its effluent should decrease. All of these concepts are reflected at the Generalized Function x Component layer.

Note in the discussion above that there may have been earlier indicators of the unexpected hydrogen source than the high reactor temperature trip. Chief among these is an added source of mass into the first heat transfer. Also, computation of the heat and reacted feed components on the basis of the known, expected source of H2 (the H2/CO stream measured at, say, H2 Flow 1 or Heat and H2 Flow 2 in Figure 6.10) would have shown no basis for the rise in reactor temperature.

- *Reactor bypassed*

Operators presumably opened valve VM1 (in Figure 6.9) to divert mixed feed flow around the reactors. Note that this is a field operated hand valve without board operator control, therefore knowledge of its appearance and location (Physical Form x Component attributes) were required for the field operator to find and manipulate it. If this valve successfully opened, it should have diverted some of the mixed feed flow around the reactor, affecting both the pressure and the quantity of reactants flowing into the reactor (see Generalized Function x Component model, Figure 6.10), thereby slowing the conversion reactions and lowering the heat produced. Note that eliminating the H2 in the mixed feed would have accomplished this alone, but would have taken longer. As long as the valves achieving feed input to the reactor remained open, however, feed would be shared along these two paths and, presumably, the reactions would continue (though at a reduced rate), hence operators closed the valves providing feed to the reactor in the next step below.

Note that the Generalized Function x Component and Physical Function x Component models imply that the operators could have achieved these same ends by diverting the mixed feed flow to flare either via VM2/VM3 or via PV410A/B (depending on which reactor was in use). This would seem to have been a preferred method of slowing the reactions, since routing unreacted feed around the reactors and downstream out of the AHR means failing the system's functional purpose—to provide acetylene-free feed. Flaring feed with acetylene content would seem preferable to routing it out of the AHR, but there may be other factors (such as holding tank or re-routing capabilities) in the downstream systems of which I am unaware.

- *Reactor inlet MOV closed*

Operators issued board commands to close the Motor Operated Valve (MOV) providing feed input to the reactor—either MV410 or MV411 depending on which reactor was in use. As discussed above, and given that they had previously opened VM1, this should route all flow through the reactor bypass and isolate the reactor, thereby more rapidly reducing conversion reactions and heat production.

- *Hydrogen to reactor physically blocked in*

Field operators physically ensured that the SDV valves were in their commanded positions. This amounts to Physical Form confirmation of settings modeled at the Physical Function level.

- *Temperatures increasing rapidly – ALERT sounded*

Despite the above actions which, according to reasoning from the Generalized Function level should have resulted in slowing reactions and decreasing temperatures in the reactor, temperatures continue to increase. This could have been due to a time lag, but for this magnitude of effect that would imply a significantly higher than expected concentration of hydrogen in the feed already in the reactors. That, we know in hindsight, was the case because greater than expected concentrations of H2 had been coming in with the C2 Feed stream. If time lags and excess H2 content in the reactor were not the culprits, it implies (again using Generalized Function reasoning, though now at the subsystem level) that some exothermic reactions are still taking place. That, in turn, probably implies that hydrogen is continuing to be input to the reactor. Moving down to the Physical Function x Component level, possible reasons for this might be faulty SDV valves (though these have already been field checked), or a faulty MOV valve (though there should have

been no H2 in the feed stream into the reactor if the SDV valves were working even if the MOV failed). Another possibility, this time reasoning from the Abstract Function level, is an unexpected source of energy into the reactor (causing temperature rises) or, reasoning at the Generalized Function level, an unexpected source of mixed feed into the reactor (causing more reactions) or an unexpected reaction occurring which was more exothermic than the normal conversion reactions.

- *Temperature reached 186 degrees C – EVACUATE sounded*

Temperatures continue to rise and are now nearing the point at which ethylene decomposition becomes common (information at the Generalized Function x Component level).

- *Deluge on Reactors tripped*

The deluge system is a part of a plant-wide safety system which sprays cooling water on hot components to remove excess heat. As such, it is not modeled in the ADS. It's effect, however, would appear at the Abstract Function level as a new energy sink (and mass barrier) associated with the reaction. The ends achieved are a reduction in energy which will mean a reduction in temperature.

- *Phone communications with E2, asking them to go on ALERT*

This was probably just good safety precautions, but it might have been motivated by the need to prepare E2 to take off spec reacted feed from the E1 Reactor Bypass. If so, the Physical Function to Generalized Function transition diagram (Figure 6.11) shows that someone would have had to manipulate VM14 to achieve Reacted E2 Feed Flow to E2.

- *There was a request to close bypass, this was rejected by [name omitted] as plant EVACUATION alarm was sounded.*

The motivation for this request is unclear, though the ADS suggests two possibilities. First, as discussed above, the fact that unreacted feed was being routed out of the AHR, violating its functional purpose, may have motivated a downstream system to ask the AHR personnel to stop it. Another possibility is suggested by the next set of events, where operators attempt to use the mixed feed input stream to cool the reactor bed. Physical Function and Generalized Function models show that having VM1 closed (and, therefore, the Feed Bypass function halted) would result in more feed flow into the reactor.

At any rate, the Physical Function to Generalized Function transition diagram (Figure 6.11) shows that the only means available to affect the Feed Bypass function is VM1. The Physical Function model shows that this is a hand valve which must be controlled from the field, and the Physical Form model would show that VM1 is located much too closely to the Reactor for safe field operations in these circumstances.

- *MOV on inlet was open to provide flow through reactor to cool bed*

Operators command a new, open setting for MV410 or 411, thereby allowing feed to flood back into the reactor. The fact that VM1 was still open (a Physical Function datum) and that there was still flow via the Feed Bypass function (a Generalized Function datum), means that some proportion of the feed is being routed via the bypass and not all of it is going into the reactor. Still, the objective was to allow the cooler, unreacted feed to drive down the temperature of the reactor bed. At the Abstract Function level, this appears as adding mass with lower energy to mass with higher energy in the reactor balance to produce a lower overall energy to mass ratio (and, thus, a lower temperature). If, in fact, hydrogen had been eliminated in the mixed feed flows, then this added feed should have produced no new reactions and no new heat beyond that of the feed itself (see Generalized Function x Component model, Figure 6.10) and therefore would have had the desired effect.

Instead, and still unknown to the operators, hydrogen had not been eliminated from the feed and thus, injecting this mixed feed into the reactors began another round of reactions as detailed in the Generalized Function x Component model in Figure 6.10.

- *Temperature still increasing*

This effect is predictable from our Generalized Function models, knowing that hydrogen was present in the feed. Given the operators' assumptions that H2 had been eliminated, this was exactly the opposite of the expected effect.

- *MOV on inlet closed*

Operators again commanded the closing of MV410 or MV411 to undo the effects of adding new feed.

- *Depressuring to flare via PV412*

Operators now command the opening of PV412, routing product after the reactors and after the reactor bypass to flare. The Generalized Function x Component model (Figure 6.10) shows this as Flare Flow 3 and indicates that it will affect Reactor Pressure, which in turn slows all conversion and decomposition reactions. The Abstract Function model shows, also, that any flow to flare acts as both a mass and energy sink. Hence, if the commanded setting of PV412 actually occurs and if product flows to flare via that route, the ADS models predict that pressure and temperature reductions should occur, reactions should slow or stop, and there should be a corresponding loss of Reacted, Cooled Feed Outflow 1 to the Product Cooling subsystem and beyond.

- *Temperatures still increasing at 5-6 degrees per punch to approximately 400 degrees from 200 degrees*
- *MOV on inlet opened again after temperature increase stabilized*
- *Temperature started to decrease*
- *Reactor bypass blocked in once safe for Operators to do so*
- *Temperature increasing on bed again*
- *MOV on inlet closed*
- *Temperatures decreasing*
- *MOV opened on inlet*
- *Temperatures down under 200 degrees*
- *Temperatures down under 100 degrees*
- *The ALL CLEAR was now sounded. All temperatures less than 32 degrees*
- *[Names omitted] noticed that the reactor was still on spec, therefore they realized that to react there was Hydrogen coming in with the feed.*

This series of events, which occurred over 27 minutes as opposed to 10 for the prior set of events, is not entirely explained by the ADS as currently configured. It is quite likely that lag effects account for some of the recorded observations. For example, the first recorded event after routing product to flare (an event that, according to the ADS, should have *reduced* temperatures) is that temperatures continued to increase rapidly. The very next recorded events, however, imply that they had stabilized and, subsequently, began to decrease. Our ADS identifies the need for modeling lag effects (at the Generalized Function level), but we don't have good models included in it currently. That, combined with the fact that specific data about time of events and rates of flow are not included in the incident report, make it difficult to tell whether the model is inaccurate or whether the sequence of events reported, with their associated lags, is misleading.

It is also probable that ethylene decomposition, which was responsible for the higher temperatures (certainly those over 200 degrees C) was halted by the depressurization. Our Generalized Function model clearly predicts this. Once this source of heat was stopped, and associated energy and mass were routed out of the system via the flare, the addition of new, cool, hydrogen-rich feed would have started the conversion reactions again, but would not initially have been sufficient to restart decomposition. Heat buildup would occur more slowly and the flow of cool feed might well result in a net reduction of temperature in the reactor bed for a while. Again, the lack of exact models, including time lags, and the lack of exact incident data make an explanation difficult.

Blocking in the reactor bypass (that is, ensuring that VM1 is closed by field inspection and manipulation) could now be done as the reactor reached a safer state. The effects of closing this flow route (as shown in the Generalized Function and Abstract Function models) was to increase the flow through the reactor. Since the incoming feed was cool but hydrogen-rich, initial cooling would take place, followed by temperature rises as the exothermic conversion reactions (see Figure 6.10) began to take place. This was observed, and the response was to close the inlet valve thereby stopping flow into the reactor (see Figure 6.11). Existing hydrogen-rich feed in the reactor should have continued to undergo conversion reactions at this point, resulting in increasing heat and the lack of pressured flow through the reactor should have kept the heated product in the vessel. Instead, decreasing heat is reported. Again, time lags, cool down after depressurization and lack of precise data may be the fault. For example, consider the following scenario: a thorough depressuring of the reactor associated with a loss of much of the reactants to flare, followed by gradual input of small quantities of reactants, followed by stopping all inputs. Our Generalized Function model predicts that at temperatures and pressures below those sufficient for ethylene decomposition, the reactants would quickly undergo their reactions, producing some heat and excess ethane and eliminating virtually all C2H2, but would then stop reacting and the reactor would begin to cool gradually (more rapidly, if the deluge water is still being applied). This is approximately what was observed, but it makes assumptions about quantities and conditions that are not recorded in the incident report.

When the MOV is opened again, cool but hydrogen-rich flow is re-established through the reactor. Temperatures continue to decline over the next 10 minutes until they are below 32 degrees C. (Normal inlet temperatures for the Reactor are 40-60 degrees C with a 20-25 degree rise over the bed assuming normal C2H2 and H2 concentrations). Again, while the Generalized Function model leads us to expect conversion reactions in the reactor under these conditions, producing heat, an overall drop in temperatures might still be plausible if flows were small and the reactor bed was cooling from a much higher temperature. More importantly, if deluge water were still being used to cool the reactor vessel, then the Abstract Function model predicts that this added energy sink would more rapidly and completely contribute to the loss of heat energy in the system. Similarly, the incident report does not record how long flow to flare was maintained. Continuous depressurization to flare (see below) would have resulted in a rapid and complete flow of feed through a very low pressure reactor, effectively halting most or all reactions and cooling the beds to the temperature of the feed flowing over them.

If temperatures are ambiguous indications of the presence of conversion reactions, confounded by time lags and a host of external forces, then the presence or absence of C2H2 in the feed outflow from the reactor is a somewhat clearer indication. At 12:03, operators evaluated the functional purpose of the AHR system and found, surprisingly, that it was still being met. Tracing back through the levels of our models, we see what the operators knew: that the only means for removing acetylene from C2 feed available in the AHR system (assuming it is in there in the first place) is by acetylene conversion reaction—and the only means available to cause that Generalized Function, is the presence of C2H2 and H2 in the reactor vessel. In spite of the fact that they had taken steps to eliminate the only known source of H2, acetylene was being removed. Therefore, acetylene conversion must be going on. Therefore, hydrogen must be present. This bit of detective work is well supported by the Generalized Function x Component model (see Figure 6.10).

Presumably, between 12:03 when operators noticed that acetylene conversion was still going on, and 12:46 when temperatures in the reactor again reached critical levels, plant personnel were busy trying to detect where the hydrogen was coming from. Physical Form and Physical Function models would be helpful at this level, suggesting the possibility of leakage in the SDV valves (and several potential repair routes by diverting H2 flow earlier in the H2/CO input stream). Once that possibility had been eliminated, the next most likely one might be that H2 was somehow entering the reactor via alternate mass input routes—specifically, via the only other existing route in the system: the C2 Feed Input stream. There are, in fact other alternatives—especially other input lines to the reactor which are part of the Reactor Regeneration system—but these were not considered part of the AHR and were not included in our ADS models.

- *High temperature in on bed – 100 degrees*

Over the next 43 minutes, hydrogen-rich feed continues to enter the reactor. As shown in our Generalized Function models, these components cause conversion reactions which inevitably produce heat. As this

process continues, heat builds in the reactor and temperatures increase. This process is exacerbated by the feed back loop through E410 and, if deluge water had been removed (probable after temperatures reached normal, safe levels), the increase would be much more rapid.

- *Bed blocked in, reactor bypassed*

Operators again take the step of opening VM1 and activating the Feed Bypass function. Even if they had not yet learned that H2 was entering via the C2 Feed stream due to a failure in the Demethanizer, they were aware that prior (and more normal) steps of blocking out the H2/CO input stream (via the SDV valves or other, upstream possibilities) were inadequate because the H2 was coming from somewhere else. Hence, the safer course was to remove the reactants from the reactor bed altogether. The only ways of achieving this (see Figure 6.10) are by continuous flaring (which is inefficient, hard on equipment as noted below and, perhaps, environmentally unsound) or by routing feed around the reactor. They chose the later course.

- *Nitrogen purge started on reactor*

Nitrogen purge is generally a part of the reactor regeneration system and, thus, has not been included in our model. It is unclear from this report whether the operators were beginning a reactor regeneration (something which would normally be done during a shut down anyway), or were simply using the capabilities of another system to aid in the management of the AHR. In either event, they were using the inert, non-reactive nitrogen to stop all reactions (and thereby remove uncertainty) in the AHR.

- *The flare system was inspected after ALL CLEAR, damage was realized*

This could, perhaps, be evidence of the excessive use of the flare which would, in turn, imply that the flare system had been used throughout the 27 minute time period described above when reactor behavior was not fully predicted by our ADS models.

- *Suggestion to trip K201 during runaway was declined as to keep flow through reactor*

This comment appears to be out of sequence. By this point, the reactor runaway was past. Our Generalized Function x Component model (Figure 6.10) shows that K201 discharge pressure provides the pressure which drives flows throughout the AHR. Tripping it would have stopped flows into the reactor, but this could have (and was) accomplished by other means including Feed Bypass via VM1, routing to flare via VM2 or VM3 and closing of MV410 or MV411. More importantly, removing K201 pressure would have made one of the operators' more effective strategies impossible: using the flow of cool C2 Feed (caused by K201 Discharge Pressure, see Figure 6.10) to reduce the temperature in the reactor bed.

- *Reactor was pressure purged and flow purged to flare numerous times over remainder of shift and through the night shift.*

It is unclear whether this final comment refers to purging (i.e., with Nitrogen) to remove contaminants including excess hydrogen, or purging via flare in response to essentially unmanageable increases in temperatures due to excessive conversion reactions when they tried to back to normal operations (using hydrogen-rich C2 Feed), or as a precursor to the shut down operations which were underway when the incident occurred.

7.4 Lessons Learned

Generally speaking, the ADS performed on NOVA's E1 AHR did a good job of providing the deep knowledge of AHR functional structure required to understand the nature of this complex and unanticipated incident. It is worth noting that Chris Miller, the author who performed this ADS analysis, is a psychologist by training with no particular background in engineering, chemistry, physics or natural gas processing. As such, it is a testimony to the power of the ADS that it structured his investigation of the AHR system sufficiently to enable this level of understanding.

There are several areas for improvement, however. As noted repeatedly above, our models are either incomplete or inaccurate with regards to their qualitative predictions about temperature effects following the reactor runaway. The lack of understanding (on the part of the analyst) of the effects of time lags on material and heat flow, and their combined effects on temperature, is a problem which should be addressed. The ADS does show us, however, the need for understanding such effects particularly in the links between Physical Functions (such as valve settings) and corresponding Generalized Functions (such as flow through the valve), and between Generalized Functions (the feedback loop formed by heat and the conversion reactions). Having said that, efforts to use ADS analysis techniques have rarely, to date, dealt with complicated temporal effects and this is an application area where future research could be done.

Finally, the above incident analysis shows the strengths and weaknesses of an ADS performed on a small part of a larger system. The causes (unusual Demethanizer operations), many of the effects (high heat an acetylene content on systems downstream from the AHR), and at least some of the management strategies (use of deluge water) in this incident all lay outside of the AHR ADS we performed. While their effects generally were understandable within the AHR models we had created (and this is evidence of the ADS's robustness), lack of access to knowledge about those other systems would limit the AHR operator's ability to predict or manage upsets. Bigger, more encompassing models will be increasingly (perhaps exponentially) difficult for analysts to create or manage, but they will provide increasingly (perhaps exponentially?) better knowledge about the constraints and capabilities of the whole system with which the operator must interact.

A somewhat subtle, but significant point about the model should be stressed. This incident is certainly not a part of the normal operating procedures for E1—nor was it a part of the system documentation we worked with in constructing the ADS models. It is obvious, for example, that the Generalized Function x Component level (see Figure 6.10) does not include an 'H2 Flow' function into Heat Transfer 1 as a part of the C2 Feed Input subsystem (see Figure 6.12). *Nevertheless*, in spite of their failure to include this unexpected Generalized Function, our suite of models was still able to detect associated effects of the problem(s)—increased mass input, unexpectedly high heat production in the reactors, increased ethane production)—and to provide many valid predictions about their effects and recommendations for how to deal with them. Vicente has labeled this robustness of the ADS in the face of unexpected contexts 'event independence' (Vicente and Tanabe, 1993) and it is a common side effect of capturing the 'deep knowledge' about system function and structure which the ADS strives to do.

8. Display Requirements Derived from the ADS Analysis

With the same approach we used in Miller & Vicente, (1998b), it is possible to derive a set of display requirements from the ADS analysis presented in Section 6. This section contains that requirements list. Note that we do not claim that all of these requirements must be met via the same display, nor must they all be presented simultaneously, nor even that these are the only display requirements that should be taken into account in creating displays for the AHR (indeed, Miller and Vicente, 1998b suggests that there are task-based requirements which are not a part of this list). Rather, these should be regarded as a set of requirements for information that the user will need to access in some fashion in order to interact with and control the AHR successfully in the range of situations s/he may encounter.

There would appear to be substantial redundancy in this list. In most cases, especially under 'normal operations', it is not necessary to have all of this information. For example, if one knows the setting of a valve, flow under normal conditions is easily deducible. There are two points to be made about this redundancy. First, even in 'purely' redundant circumstances (ones in which some information is derivable from other information under all possible circumstances), one tenet of human factors is that making information explicit makes the user's job easier. If information is obtainable via direct perception rather than by complex, error-prone, and workload intensive cognitive processes, the user will be more likely to get it correctly and with less work. Second, much of what appears 'redundant' in normal operating circumstances, is not redundant under abnormal circumstances. This is precisely what enables detection and diagnosis in those circumstances. In our simple example above, valve setting and flow are redundant as long as adequate pressure and correct valve operation occur (and discounting response lags). If, on the

basis of that redundancy, we chose to omit display of one or the other of the values, the operator would be incapable of knowing whether the valve was operating correctly. Thus, while the list below may contain redundant displays, extreme care should be taken in deciding to omit any of this information from a final display or display set.

1. All physical components of the AHR (as identified by the Physical Form level of the ADS) should be represented. These are:

1. H101-108	18. PV410 B	35. VM8
2. E410	19. PV412	36. VM9
3. E411	20. HV41001	37. VM10
4. E412	21. MV410	38. VM11
5. E413	22. MV411	39. VM12
6. R410 A	23. VS1	40. VM13
7. R410 B	24. VH2	41. VM14
8. SU411	25. VH3	42. VW1
9. FV135	26. VH4	43. VW2
10. PV441	27. VH5	44. CV1
11. TV440	28. VM1	45. CV2
12. FV413	29. VM2	46. CV3
13. TV410	30. VM3	47. CV4
14. SDV413 A	31. VM4	48. CV5
15. SDV 413 B	32. VM5	49. CV6
16. SDV 413 C	33. VM6	50. ST1052
17. PV410 A	34. VM7	51. ST1199

2. Information about the appearance and location of physical components listed in 1 should be included—ideally, both as a reference (i.e., 'normal') state and a current state.
3. All physical functions of the AHR components (as identified by the Physical Function x Component cell of the ADS) should be represented. That is, the state of the components identified above should be illustrated. These are:

1. H101-108 temperature	18. PV412 setting	35. VM9 setting
2. E410 shell and tube side temperatures	19. HV41001 setting	36. VM10 setting
3. E411 condensate level, shell and tube side temperatures	20. MV410 setting	37. VM11 setting
4. E412 shell and tube side temperatures	21. MV411 setting	38. VM12 setting
5. E413s condensate level, shell and tube side temperatures	22. VS1 setting	39. VM13 setting
6. R410 A temperature profile, pressure, days in service	23. VH2 setting	40. VM14 setting
7. R410 B temperature profile, pressure, days in service	24. VH3 setting	41. VW1 setting
8. FV135 setting	25. VH4 setting	42. VW2 setting
9. PV441 setting	26. VH5 setting	43. CV1 direction
10. TV440 setting	27. VM1 setting	44. CV2 direction
11. FV413 setting	28. VM2 setting	45. CV3 direction
12. TV410 setting	29. VM3 setting	46. CV4 direction
13. SDV413 A setting	30. VM4 setting	47. CV5 direction

14. SDV 413 B setting	31. VM5 setting	48. CV6 direction
15. SDV 413 C setting	32. VM6 setting	49. ST1052 level
16. PV410 A setting	33. VM7 setting	50. ST1199 level
17. PV410 B setting	34. VM8 setting	

4. In addition to actual state, the range of possible states for each component listed in 3 should also be shown.
5. Physical connections between physical functions should be shown. Physical connections are illustrated in Figure 6.10. Among the connection information which should be included is:
   a) The raw feed stream and the fuel gas stream enter the pyrolysis furnace via different paths.
   b) DMDS is added to the raw feed stream (achieved by FV135) before that stream reaches the furnaces.
   c) The C2 feed stream and the H2/CO stream are split after the pyrolysis furnaces, but before reaching the remaining AHR components. Thus, these streams enter the AHR separately.
   d) Two different sources of H2/CO stream are available to the AHR—the stream from the pyrolysis furnaces and the E2 stream.
   e) The E1 H2/CO stream enters E413 directly, with no limiting valve available.
   f) The E2 H2/CO stream can be routed into the E1 H2/CO stream before it reaches E413 by appropriate settings on VH2-VH4.
   g) The E2 H2/CO stream can also be routed into the H2/CO stream for the AHR *after* (that is, bypassing) E413 by appropriate settings on VH2-VH4.
   h) The E1 H2/CO stream can be routed to E2 after it has passed through E413 by appropriate settings on VH2-VH4, PV441 and TV440 but it cannot be routed E2 before E413 due to CV1.
   i) The H2/CO stream can be diverted to flare immediately after E413 by means of PV441.
   j) The H2/CO stream can be diverted to E319 and TE301 (the Turbo Expanders) after E413 and TV440, but before FV413 by means of appropriate settings on VH5.
   k) CV6 is positioned immediately before VH5 to prevent flow of H2/CO back into E413.
   l) FV413 is located before the SDEV valves and before the point where H2/CO and C2 feed mix (at SU411).
   m) SDV A and SDV B are located on either side of a pipe venting to atmosphere which is controlled by SDV C. All three of these are located immediately before SU411.
   n) SU411 is located immediately after the connections of the H2/CO and C2 feed streams.
   o) The C2 feed stream enters the tube side of E410 directly, with no limiting valve available.
   p) From E410, the C2 stream is connected to shell side of E411, again with no limiting valve available.
   q) After E411, the C2 stream intersects the H2/CO stream immediately before SU411. Again, there is no limiting valve, bypass or diversion path available.
   r) Steam from the same source supplies both E411 and E413, but the pathway splits before entering either heat exchanger.
   s) VS1 and CV2 are connected before E413 and can affect the flow of steam to it, and CV2 prevents steam flow from it.
   t) ST1199 is connected after E413 and can affect the level of condensate in it.
   u) TV410, CV3 and ST1052 are located before E411 (ST1052 on a diverging path) and can affect the flow of steam into it, and CV3 prevents steam flow from it.
   v) Steam and condensate flow out of E411 to V412.
   w) 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.
   x) The two reactors (R410 A and B) are located side by side, but there is no direct connection between them. Flow from reactor to reactor can be achieved by upstream or downstream valves, however.
   y) MV410 and VM4 are located on a pipeline entering R410A, while VM6 and VM7 are located on a pipeline leaving that reactor. R410 B has a similar configuration using MV411, VM5 before and VM8 and VM9 after.
   z) There are connections to flare before the reactors and VM2 lies on this connection before R410 A while VM3 lies on the connection before R410 B. CV4 lies on the connection before VM2 to

prevent backflow from the flare to the pipeline, while CV5 performs a similar function before VM3 on R410 B.

aa) There are also connections to the flare from the reactors themselves. PV410A lies on the connection from within R410A and PV410B lies on the connection from R410B.
bb) Although separate pipelines exit the reactors, they connect after VM6-VM9.
cc) 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).
dd) The tube side of E412 is connected to the flare (VM11 lies on this connection) or to a pipeline leading to the shell side of E410 (VM12 lies on this connection).
ee) Cooling water enters the shell side of E412 via a pipeline which contains VW1 and leaves via a pipeline which containes VW2.
ff) The connection from E412 to E410 intersects a connection bringing reacted feed to or from E2. VM14 lies on this connection.
gg) The connection from E412 enters and exits the shell side of E410 where it leaves the AHR. HV41001 lies on the exiting connection.

6. All generalized functions from the AHR ADS should be illustrated at the component level. That is, the actual flow and storage of heat and water should be represented distinctly from the settings described in #3 above and the possible connections described in #5. In practice, this means the following values should be represented:
   a) The flow of DMDS into the pyrolysis furnaces
   b) The flow of raw feed into the pyrolysis furnaces
   c) The flow of fuel gas into the pyrolysis furnaces
   d) The flow of H2 out of the pyrolysis furnaces[3]
   e) The flow of CO out of the pyrolysis furnaces
   f) The flow of C2H2 out of the pyrolysis furnaces
   g) The flow of C2H4 out of the pyrolysis furnaces
   h) The flow of C2H6 out of the pyrolysis furnaces
   i) The flow of H2 entering the AHR at E413
   j) The flow of CO entering the AHR at E413
   k) The flow of steam into E413
   l) The flow of heat into the H2/CO stream in E413
   m) The flow H2/CO to flare
   n) The flow of heat and H2 to or from E2.
   o) The flow of heat and CO to or from E2.
   p) The flow of heat and H2/CO to the Turbo Expanders.
   q) The flow of heat and H2/CO through the SDV valves and/or into the atmosphere.
   r) The flow of heat and H2/CO through SU411.
   s) The discharge pressure of K201.
   t) The flow and pressure of C2H2 entering the AHR at E410
   u) The flow and pressure of C2H4 entering the AHR at E410
   v) The flow and pressure of C2H6 entering the AHR at E410
   w) The flow of heat from the reacted feed stream into the C2 feed stream at E410.
   x) The flow of steam into E411.
   y) The flow of heat into the C2 feed stream at E411.
   z) The flow of heat and C2 feed through SU411.
   aa) The flow of heat and mixed feed through the reactor bypass.
   bb) The flow of heat and mixed feed to flare after SU411 but before the reactor.
   cc) The flow of heat and CO into the reactor
   dd) The flow of heat and H2 into the reactor.

[3] It is unclear to me where, exactly, each of the five products of interest (H2, CO, C2H2, C2H4, and C2H6) are separated from the outflow of pyrolysis. In principle, the AHR operator should have access to *predictions* about the proportions of these feed components that are going to reach the AHR as early as they can reliably be made. I believe that CO content can be reliably predicted immediately after pyrolysis, but I am not sure about the other components.

ee) The flow of heat and C2H2 into the reactor
ff) The flow of heat and C2H4 into the reactor
gg) The flow of heat and C2H6 into the reactor
hh) Pressure in the reactor
ii) The rate of C2H2 conversion
jj) The rate of C2H4 conversion
kk) The rate (or presence) of C2H4 decomposition
ll) The rate of CO conversion
mm) The flow of heat and CO out of the reactor
nn) The flow of heat and H2 out of the reactor.
oo) The flow of heat and C2H2 out of the reactor
pp) The flow of heat and C2H4 out of the reactor
qq) The flow of heat and C2H6 out of the reactor
rr) The flow of heat and CH4 out of the reactor
ss) The flow of heat and C2 out of the reactor
tt) The flow of heat and green oil out of the reactor
uu) The flow of heat and H2O out of the reactor
vv) The flow of heat and CO after reactor effluents and reactor bypass join.
ww) The flow of heat and H2 after reactor effluents and reactor bypass join.
xx) The flow of heat and C2H2 after reactor effluents and reactor bypass join.
yy) The flow of heat and C2H4 after reactor effluents and reactor bypass join.
zz) The flow of heat and C2H6 after reactor effluents and reactor bypass join.
aaa) The flow of heat and CH4 after reactor effluents and reactor bypass join.
bbb) The flow of heat and C2 after reactor effluents and reactor bypass join.
ccc) The flow of heat and green oil after reactor effluents and reactor bypass join.
ddd) The flow of heat and H2O after reactor effluents and reactor bypass join.
eee) The flow of reacted feed around E412 and toward E410.
fff) The flow of reacted feed to flare before E412.
ggg) The flow of reacted feed into E412.
hhh) The flow of reacted feed out of E412 toward E410.
iii) The flow of reacted feed from E412 to flare.
jjj) The flow of cooling water into and out of E412.
kkk) The flow of heat from the reacted feed in E412
lll) The flow of reacted, cooled feed to or from E2.
mmm) The flow reacted, re-cooled feed (after E410) to the rest of the plant.

7. In addition to actual, current value for each of the items in 6, the range of possible values for each component should also be shown.
8. The fact that the value of each component's generalized function can be effected by changes in that component's physical function state should also be clearly illustrated either through the appearance or the behavior of the interface (or both). These include the following facts:
   a) The effect of valve setting on flow through the valve for all valves.
   b) The effect of all check valve orientations and pressure resistances on direction of flow.
   c) The effect of the settings of VH3 and VH4 (and, secondarily, VH5, PV441, TV440, FV413 and the SDV valves) on flow of H2/CO through VH2.
   d) The effect of the setting of PV441 on flow of H2/CO through TV440 (and, secondarily, on flow through VH2-VH4).
   e) The effect of the setting of TV440 on flow of H2/CO through CV6 and downstream, and on flow through VH2-VH4.
   f) The effect of the setting of VH4 and VH2 on flow of H2/CO through VH3.
   g) The effect of the setting of VH5 on flow of H2/CO through TV440 and/or VH4 and downstream to FV413.
   h) The effect of the setting of FV413 on the flow of H2/CO through CV6, VH5 and downstream to the SDV valves.
   i) The effect of the setting of each of the SDV valves on the flow through the remaining two. Especially the fact that flow through SDV413B to the pipeline junction and SU441 is dependent on SDV413C being closed.

j) The fact that pressure from K201 drives the flow of C2 and H2/CO products throughout the AHR, but that there is no physical function setting for affecting that pressure.
k) The fact that the level of condensate in ST1199 (along with the heat of the steam) affects the heat transfer from steam to H2/CO feed in E413.
l) The fact that the level of condensate in ST1052 affects the flow of steam through TV410.
m) 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.
n) The effect of the setting of VM1 on flow of mixed feed through MV410 and MV411.
o) The effect of the settings of flare valves (VM2, VM3, PV410A and PV410B), secondary reactor input valves (VM4 and VM5) and reactor outlet valves (VM6-VM9) have on flow through the primary reactor input valves: MV410 and MV411.
p) The effect of the setting of the reactor flare valves (VM2, VM3, PV410A and PV410B) on flow through VM4 and VM5 into the reactors.
q) The effect of the settings of the pairs of reactor outlet valves (VM6 and VM7, and VM8 and VM9) have on flow through the reactors, and on each other.
r) The effect of the settings in the components in one reactor chain have on flow through the other.
s) The effect of the setting of PV412 on flow of reacted feed through VM10 and VM13.
t) The effect of the setting of VM13 on flow of reacted feed through VM10 and VM12.
u) The effect of the setting of VM10 on flow through VM13 and through E412 to VM12.
v) The effect of the setting of VM11 on flow through VM10, VM12 and E412.
w) The effect of the setting of VM12 on flow through VM13 and VM10.
x) The effect of the setting of VW1 on flow of cooling water through E412 and on to VW2.
y) The effect of the setting of VW2 on flow of cooling water through VW1 and E412.
z) The fact that the settings of VW1 and VW2 affect the rate of heat exchange in E412.
aa) The effect of the setting of VM14 on flow of cooled, reacted feed through VM12, VM13 and into E410.
bb) 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.

9. Connections between generalized component functions should be illustrated in the interface indicating direction of affectivity or causality. The fact that the value of each generalized function can be affected by changes in other components' generalized function states should also be clearly illustrated either through the appearance or the behavior of the interface (or both). More complex examples of these interactions include:
a) The effect of the rate of DMDS flow on CO flow out of pyrolysis.
b) The fact that K201 discharge pressure affects flows throughout the AHR.
c) The fact that diverting H2/CO flow to or from E2 means affecting the heat input to either Heat Transfer 1 or downstream to feed mixing.
d) Similarly, diverting H2/CO flow to flare means affecting the heat and H2 and CO flows downstream.
e) Ditto for diversion to the Turbo Expanders.
f) Ditto for Flow lockout and diversion to atmosphere.
g) The rate of heat transferred in Heat Transfer 1 can affect the rate transferred in Heat Transfer 2 (and vice versa).
h) 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).
i) Ditto for routing mixed feed to flare before the reactors.
j) The effect of reactor pressure on [especially] Ethylene decomposition (higher pressure lowers the temperature required for it to occur).
k) The effect of acetylene conversion on acetylene flow out of the reactor. (It will eliminate it. The only way C2H2 can appear after the reactors (assuming the bypass is closed) is if it was not all converted in the reactors.) Ditto for CO and CO conversion.
l) The effect of CO conversion on H2O flow out of the reactor (it is the only way to produce it; more CO conversion means more H2O) and on CH4 flow (it is one of two ways to produce it, the other is C2H4 decomposition).

m) 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).
n) The effect of Acetylene and Ethylene conversion on green oil production (they are the only ways to produce it).
o) The effect of the availability of H2 on CO conversion, C2H2 conversion and C2H4 conversion—it is required for all three and limits them if not available.
p) CO Conversion affects the availability of H2 (and the availability of catalyst sites, though we have not chosen to model that level of decomposition), and therefore affects the incidence or rate of C2H2 and C2H4 conversion.
q) 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.
r) The effect of all reactions on heat flow out of the reactor: all are exothermic and happen more rapidly (or at lower pressures) with increased heat. This is especially true of ethylene decomposition. Thus, there is the potential for self perpetuating reactions if heat is not removed from the reactor.
s) The effect of flare flow after the reactor (at either Flare flow 3 or 4) on reactor pressure, heat and reactant availability. It also, of course, reduces the availability of heat and reacted feed downstream.
t) The effect of bypassing the reactor aftercooler on the heat flow to E2 or to Heat Transfer 1.
u) The effect of heat and feed flow to Heat Transfer 1 on the rate of heat transfer to the incoming C2 Feed.

10. Generalized functions should also be represented at the Subsystem level, and the interface should make it clear that the subsystems consist of their appropriate components (as listed in Table 6.1 for physical components, and in Table 6.2 for the relationship between Generalized Functions at the Component level and those at the Subsystem level).
11. Connections between generalized subsystem functions should also be illustrated indicating direction of affectivity or causality either through the appearance or the behavior of the interface (or both). These include the following facts:
    a) That pyrolysis affects both H2/CO input and C2 feed input streams.
    b) That H2/CO can either come from Pyrolysis or be brought in from E2.
    c) H2/CO feed can also be sent from E1 to E2 either before or after heating.
    d) That H2/CO heating heats only the H2/CO input stream (regardless of where it comes from). It has no effect on the C2 Feed stream.
    e) That C2 Feed can only be sent to the Mixed Feed Input unit after heating, but that, by use of E2 feed, it is possible to for the H2/CO feed to bypass its heating unit.
    f) That mixed feed can either go to the reactors, to flare, or can bypass the reactors and go directly to product cooling. Each route affects the availability of mixed feed to send along the other routes.
    g) Mixed, reacted feed can also be sent to flare.
    h) Mixed, reacted and partially cooled feed can be sent to E2, or it can be brought in from E2 into the Product Cooling subsystem.
    i) Mixed, reacted and cooled feed plays a role in heating the C2 feed stream before it is routed out of the AHR.
12. All abstract functions identified in the ADS should be represented in the interface. Current values and ranges for these entities should be represented at the subsystem level. These are illustrated in Figures 6.14 and 6.15. In all there are 6 mass sources, 8 mass sinks, 11 mass barriers and 3 mass balances, as opposed to 6 energy sources, 8 energy sinks, 8 energy barriers and 6 energy balances.
13. The fact that the value of each subsystem's abstract function can be affected by changes in the subsystem's generalized function (or lower level functions) should also be clearly illustrated in the interface either through the appearance or the behavior of the interface (or both). This includes the following facts:
    a) Pyrolysis affects both mass and energy input streams.
    b) Mass and energy can either come from or go to E2 via two routes.
    c) The Flare subsystem causes an out flow of mass and energy wherever it is initiated from.

d) Heating and cooling functions (H2/CO heating, C2 Feed heating and Product Cooling) cause a transfer of energy, but not of mass to the main mass flow in the system (paired sources and sinks for steam and cooling water cause their mass to exit and leave the AHR quickly, and barriers prevent their mixing with product—under normal circumstances).
e) The reaction subsystem is a mass and energy balance—that is, it neither produces nor removes energy or mass.
f) Ditto for feed mixing.
g) Ditto for bypassing the reaction.
h) The balance of energy across the Product Cooling and C2 Feed Heating subsystems (which share the same Generalized Component Function Heat Transfer 1) means that energy lost from the flow from E412 must be absorbed by the C2 feed flow to the mixed supply, and vice versa.
i) There are no inventories for mass or energy in this system.

14. Connections between abstract subsystem functions should be illustrated indicating direction of affectivity or causality. The fact that the value of each abstract function can be affected by changes in other abstract function states should also be clearly illustrated either through the appearance or the behavior of the interface (or both). These include the following facts:
    a) Mass and energy sources increase mass and energy in the system, but are constrained by the capacity of the system and of the sinks.
    b) Mass and energy sinks decrease mass and energy in the system and are constrained by source flow.
    c) The transport routes to E2 can either act as sources or sinks for mass and energy.
    d) The heat transfer devices (E410-12) act as mass barriers, but energy balances.
    e) The flow of energy either through or around the reactions will affect the energy available to be balanced at E412, and E410.
    f) The energy flows through the product cooling subsystem (E410 in particular) form a closed loop. Thus, increased energy out of the balance at E412 will increase the energy entering the balance at E410, which will either increase the energy re-entering the mixed feed supply balance or increase the energy exiting the AHR via the sinks or both.
15. Functional purposes uncovered in the ADS analysis should be represented at the system level and included in the interface. These are:
    a) Target C2H2 content upon leaving the AHR system.
16. In addition to these targets, actual current values for functional purposes should also be shown. These are:
    a) Current C2H2 content upon leaving the AHR system.
17. The fact that the value of each subsystem's functional purposes can be affected by changes in the subsystem's abstract functions (or lower level functions) should also be clearly illustrated in the interface either through the appearance or the behavior of the interface (or both). This includes the following facts:
    a) That changes in mass into and out of the reaction balance can affect the functional purpose.
    b) The fact that use of a mass sink (specifically, the flare can prevent bad product from flowing out of the AHR.

9. Comparison of ADS-Derived Display Requirements to Current AHR Displays

An obvious next step in the using the display requirements generated in the previous section would be to compare them to NOVA's current set of displays for the AHR to determine their strengths and weaknesses and then suggest design modifications to enhance NOVA's displays. While we have not yet taken those steps, the display requirements list provided in section 8 puts either us, or NOVA personnel, in a good position to do so.

In other work (e.g., Vicente and Rasmussen, 1990; Dinadis and Vicente, 1996; Jamieson, 1998), researchers using ADS analysis techniques have proceeded from the analysis directly to designing a display 'from scratch' without making an explicit attempt to integrate that display into ongoing operations or other existing displays. It might be more effective for real-world fielding of such a display, especially when the ADS has only been performed for a small part of the overall work domain, to instead attempt to find those

aspects of the current displays that need to be modified. This approach, presuming that modifications are not extreme, would result in greater familiarity to existing operators, and therefore less need to retrain them.

We have not yet taken those steps for several reasons. First and foremost is that the current available level of effort did not support them. Second is the fact that we have not yet done a sufficiently detailed analysis of the existing E1 AHR displays to fully understand them or evaluate their strengths and weaknesses against the requirements described above. Finally, the overall goal of the UMP research project is to generate modeling techniques and interfaces that integrate display requirements from both task-based and work domain-based analyses. This report provides only half of those requirements. Future work will address task-based analyses and the integration of requirements from both approaches into a single interface design.

10. Conclusions and Lessons Learned

The chief lesson learned from this effort was that the ADS *can* be done for NOVA ethylene refining work domains. Although these domains pose novel challenges rarely encountered by users of the ADS in the past (such as chemical reactions, geographically distributed plant operations and complex temporal interactions), the ADS can be adapted to analyze, structure and interpret the domain. Further, using ADS techniques offers valuable insights about the work domain including: display requirements, an explanatory structure for incidents, a critique of existing displays, a framework for training 'deep knowledge,' etc. Since this work, along with Jamieson (1998) is one of only two projects to date applying the ADS to realistic petrochemical processing operations, this finding is important and encouraging for users in those domains.

This work also serves as an illustration of the utility of the ADS as a knowledge acquisition tool. Most of this analysis was performed by a psychologist, untrained in the supporting knowledge of engineering, chemical and physical reactions, etc. that are required to understand the domain—and a complete novice at the specifics of NOVA's ethylene operations. In many ways, we were using the ADS as a guide to training ourselves—to acquiring the knowledge necessary to understand the operation of the AHR (hence the ability to explain and understand the incident analyzed in section 7 was a particularly good test of how well we had done). In this regard, then, the ADS can be very helpful in guiding the search to obtaining deep knowledge about a work domain—even for a novice.

The ADS provides a straightforward identification of the information required to monitor current AHR behavior, understand its implications, diagnose upsets and take actions. These are illustrated in section 8. While these are not yet a display, a knowledgeable designer could use these requirements to construct a display him- or herself. This offers an intriguing suggestion: that it might be easier and more efficient for a person knowledgeable in the operation of the work domain to construct the ADS and then use it as a vehicle to transfer domain knowledge to an interface designer. A 'knowledgeable person' could either have general knowledge (e.g., a chemical engineer who understands conversion reactions) or specific knowledge about the work domain (e.g., an engineer from NOVA's E1 facility itself)—with the later approach offering the best fit, of course. Engineer and display designer could then, perhaps, jointly review the ADS to ensure understanding on the part of the display designer and to construct a display requirements list, and then the display designer could work with that list to construct the interface. Of course, this approach would require that the engineer (or some other member of the design team) to have knowledge about how to use the ADS modeling technique itself.

Such efficiencies are likely to become increasingly important if the ADS is to be used as a common tool in real-world interface design. The level of effort required to perform this analysis of a portion of the overall E1 process (approximately three person months) may be difficult for industry to justify. Another useful approach to keeping costs and time required for ADS construction down, is currently under investigation by the U.S. Air Force Labs at Wright-Patterson Air Force Base: the development of a support tool to both speed model construction and help manage the complexity of the models which are developed. Such tools

will be a welcome addition to help improve the accuracy and reduce the cost of ADS construction. Ultimately, though, one must ask whether there is really any other justifiable alternative? If the 'deep knowledge' about work domain constraints and capabilities are not available in the display itself (which implies that they must have been available and understood by the designer of the display), then they will need to be worked out in real time by operators—perhaps in the midst of an unexpected upset. Are there more efficient ways of gaining (whether by a display designer or a novice operator) a thorough knowledge of the constraints, capabilities, functions and behaviors of the system than the ADS approach? Task analysis, the chief alternate contender, while it does offer insights not offered by ADS, does not offer the deep knowledge that both system operators and display designers need.

On a more specific level, this effort has, for the first time, used the ADS to represent and reason about chemical reactions. 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, aggregating them in a more holistic 'reaction' subsystem along the part-whole dimension—is an important, but natural, addition to the ADS framework. If our approach proves robust, this will be a valuable addition to the repertoire of ADS techniques and will extend the range of applicability for the ADS into a host of chemical, petrochemical and even biological realms where it has not been applied to date.

We note, from this effort, the difficulty of defining system boundaries and of knowing how much to include at each level in the analysis in a highly complex and tightly-coupled, real-world system. The decision to include aspects of the pyrolysis system violated the geographical organization of the AHR and even operators mental organization of the system, but it provided the analysts with the ability to represent an important manipulation of AHR behavior—the role of DMDS in the pyrolysis feedstock. On the other hand, we chose not to represent how feedstock is separated into C2 components and H2/CO feed after pyrolysis, and this prevented us from fully understanding the causes of the incident analyzed in section 7. Similarly, we separated out the regeneration system and the product drying and green oil removal systems from the AHR, even though operators tend to think of these as all portions of the AHR. This decision simplifies our models, and would simplify displays created from those models, but it also had the effect of reducing our knowledge about possible sources of feed or hydrogen, possible effects of actions within the AHR on downstream units, and possible courses of action that might be taken to address the problem studied in section 7. While most of the above refers to the 'breadth' of the system modeled, the problem also exists in the 'depth' dimension. We chose to model hand valves and bypass routes, but did not model drain valves. This decision similarly simplified our models, but did so only by editing out some knowledge that might be important to designers or operators in some future incident.

These problems are endemic to any system analysis or design effort. The ADS offers an advantage, however, in that it provides a powerful structure—hierarchical in two dimensions—for organizing information from very detailed to very aggregated and very concrete to very abstract. We saw the power of this dual hierarchical organization in simplifying the Generalized Function x Component model to the Generalized Function x Subsystem model. We also suspect that this transition provides a powerful *set of ways* for operators to think about the domain: reasoning about flows of individual reactants at the component level vs. about collected flows at the subsystem level; reasoning about individual reactions and the conditions under which they occur at the component level and reasoning about the 'selectivity' of 'the reaction' at the subsystem level. If either broader or deeper information about the AHR system is needed, the ADS shows us where and how to incorporate it.

And finally, as can be seen in section 8, we have used the ADS to produce a large quantity of detailed display requirements—each with a rationale for inclusion. The generation of this list and its availability for use as a design aid and checklist is an improvement over ad hoc approaches to interface design. The explicit system-based rationale for the completeness of this list and the inclusion of each element on it is an improvement over the single path trajectories, with their associated display requirements, provided by most task-based analytic techniques.

But several questions remain to be answered before we can ascertain the ultimate quality of these requirements and the ADS approach to generating them:

- *Unique*--to what degree are these requirements similar to or different from requirements identified using more tradition means by NOVA's own display designers?
- *Complete*--To what degree is this the full set of requirements? Are there others necessary?
- *Useful*—Will such a display, constructed from these requirements, produce superior performance from the human-machine system in which it is used?

We have reason to suspect, from prior research, that these requirements will be useful, but an assessment of their true quality must await future work.

11. References

Bisantz, A. and Vicente, K. (1994). Making the abstraction hierarchy concrete. International Journal of human-computer studies 40:83-117.

Cook, D. and Poole, A. (1985). "Minutes of Reactor Runaway Incident Investigation", Alberta Gas Ethylene Company, Ltd., File No. E1-0-1.11.6-85. May 30.

Dinadis, N. and Vicente, K. (1996). Application of Ecological Interface Design to Aviation. CEL 96-07. Toronto; Canada. Cognitive Engineering Laboratory, University of Toronto.

Gibson, J. J., & Crooks, L. E. (1938). A theoretical field-analysis of automobile-driving. American Journal of Psychology, 51, 453-471.

Jamieson, G. (1998). Ecological Interface Design for Petrochemical Processing Applications. Technical Report CEL 98-04. Toronto; Canada. Cognitive Engineering Laboratory, University of Toronto.

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

Miller, C. and Vicente, K. (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." Technical Report CEL 98-07. Toronto; Canada. Cognitive Engineering Laboratory, University of Toronto.

Miller, C. and Vicente, K. (1998b). "Comparative Analysis of Display Requirements Generated via Task-Based and Work Domain-based Analyses: A Test Case using DURESS II." Technical Report CEL 98-08. Toronto; Canada. Cognitive Engineering Laboratory, University of Toronto.

Rassmussen, J. (1985). The role of hierarchical knowledge representation in decisoin making and system management. IEEE Transactions on Systems, Man and Cybernetics, 15, pp. 234-243.

Vicente, K. J. (1996). Improving dynamic decision making in complex systems through ecological interface design: A research overview. System Dynamics Review, 12, 251-279.

Vicente, K. J. (in press). Cognitive work analysis: Towards safe, productive, and healthy computer-based work. Erlbaum: Mahwah, NJ.

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

Vicente, K. and Rasmussen, J. (1990). The ecology of human-machine systems II: Mediating 'direct perception' in complex work domains. Ecological Psychology, 2(3). Pp. 207-249.

Vicente, K. and Tanabe, F. (1993). Event-indepent assessment of operator information requirements: Providing support for unanticipated events. In *Proceedings of the Topical Meeting on Nuclear Plant Insrumentation, Control, and Man-Machine Interface Technologies*. (LaGrange Park, IL; American Nuclear Society). 389-393.