Int. J. Human-Computer Studies (2002) **56,** 569–596
doi:10.1006/ijhc.1006
Available online at http://www.idealibrary.com on IDEAL®

Work domain analysis and sensors I: principles and simple example

Dal Vernon C. Reising†
Department of Mechanical and Industrial Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, USA

Penelope M. Sanderson
ARC Key Centre for Human Factors and Applied Cognitive Psychology, The University of Queensland, St Lucia, QLD, Australia. E-mail: psanderson@humanfactors.uq.edu.au

(Received 15 June 1999 and accepted in revised form 30 March 2002)

In this paper we establish a foundation for understanding the instrumentation needs of complex dynamic systems if ecological interface design (EID)-based interfaces are to be robust in the face of instrumentation failures. EID-based interfaces often include configural displays which reveal the higher-order properties of complex systems. However, concerns have been expressed that such displays might be misleading when instrumentation is unreliable or unavailable. Rasmussen's abstraction hierarchy (AH) formalism can be extended to include representations of sensors near the functions or properties about which they provide information, resulting in what we call a "sensor-annotated abstraction hierarchy". Sensor-annotated AHs help the analyst determine the impact of different instrumentation engineering policies on higher-order system information by showing how the data provided from individual sensors propagates within and across levels of abstraction in the AH. The use of sensor-annotated AHs with a configural display is illustrated with a simple water reservoir example. We argue that if EID is to be effectively employed in the design of interfaces for complex systems, then the information needs of the human operator need to be considered at the earliest stages of system development while instrumentation requirements are being formulated. In this way, Rasmussen's AH promotes a formative approach to instrumentation engineering.

KEYWORDS: work domain analysis; sensors; instrumentation; abstraction hierarchy; ecological interface design.

"What can be sensed forms a fundamental limiting feature of displays. This limiting feature is not always given the emphasis it deserves. [...] A measuring instrument which indicated directly and accurately its position relative to the earth's surface would render obsolete the entire science of navigation as we know it. This illustrates clearly the limits imposed on navigation systems by lack of a sensing instrument. The whole system of celestial and radio

†Present address: Honeywell Laboratories, Honeywell Int., Inc., 3660 Technology Drive, Minneapolis, MN 55418. E-mail: dalvernon.reising@honeywell.com.

1071-5819/02/$ - see front matter

fixes, the use of compass information, inertial and other "dead reckoning" systems, are all substitutes for what we would like to sense directly and display but do not yet have the sensing means for. The initial step in considering the design of the displays for a particular manual control system is to analyze the information the operator would really like to have and to consider the sensing instruments available to obtain it for him. Too often this analysis is not carried out, and it is assumed that the operator requires information that, for instance, has been displayed on similar systems in the past ..." Kelley (1968, p. 90–91).

1. Introduction

Kelley's (1968) words suggest the need for a use-centered approach to display design for dynamic systems that is as important today for supervisory control as it was 30 years ago for manual control. Kelley makes the important point that human performance will be limited by displays and displays, in turn, will be limited by what can be physically sensed about a system or process.[‡] Therefore, during system development cognitive engineers should made important contributions to the process of identifying sensor and instrumentation requirements to ensure that the human operator has information that will adequately support the kind of performance needed.

However, the question of what should be displayed to the human operator, and therefore what needs to be sensed, has become more complex since Kelley made his comment. First, human operators are now usually supervisors of processes that are largely automated. Human operators are expected to be able to step in and take over wholly or in part from the automation if the system encounters unexpected conditions or enters a state unanticipated by designers (Rasmussen, 1986; Bainbridge, 1983). Second, it is increasingly realized that displays should reflect not just what individual operators wish to see, as Kelley (1968) suggests, but what has been determined from formal analyses to be *necessary* for operators to see in order to exercise adequate control (Rasmussen, 1999; Vicente, 1999). Research indicates that operators need information not only about the kind of straight-forward physical plant functioning that instrumentation usually senses, but also about higher-order properties of plant functioning (Woods & Roth, 1988; Rasmussen, Pejtersen & Goodstein, 1994; Vicente, 1999). Higher-order properties include relations such as ratios, rates of change, and progress towards goals and functional abstractions such as heat exchange, mass and energy. Such higher-order properties are usually not directly sensed, but instead may be derived from multiple sensed values combined mathematically with each other, from sensed values combined mathematically with constants (such as density and specific heat), or from both. Empirical research suggests that a human operator supplied with information about higher-order properties is better equipped to cope with any unanticipated variability encountered by a system (Vicente, in press).

Our goal in this paper is to demonstrate that the abstraction hierarchy (AH) formalism developed by Rasmussen (1979, 1986, 1988, see also Rasmussen *et al.*, 1994)

[‡] Note that global positioning systems (GPS) have not solved the problem identified by Kelley (1968) as they do not directly sense "position relative to the earth's surface" but instead estimate it with a (potentially) extremely high level of accuracy using a combination of satellites, microwave signals, atomic clocks, etc. For further information see http://www.colorado.edu/geography/gcraft/notes/gps/gps_f.html.

may be a powerful tool in determining the instrumentation needed to support human operators when they must handle unanticipated variability. The AH is a key element in the approach to display design known as ecological interface design (EID) (Vicente & Rasmussen, 1990, 1992). The AH helps analysts determine the information that human operators need if they are to help the system behave in a new, desired way. It can also be used to pinpoint the instrumentation and/or computation that will provide the information. Finally, the AH may be a powerful tool for tracing the impact of inadequate instrumentation on how information is generated. It may help us determine the impact of inadequate instrumentation on an EID-based interface, and therefore on human–system performance.

This paper is the first of two papers that explore the use of the AH to illustrate instrumentation needs when designing interfaces for complex dynamic systems. We discuss these issues primarily in the context of the EID approach to display design, but the message extends to other approaches as well. In Part I (the present paper), we provide a background to EID as an approach to advancing the reliability of human–system integration and we cover some basic details of how sensors are designed and how they can fail. We then review some of the concerns that have already been raised about the vulnerability of ecological interfaces to sensor unavailability or unreliability. Then we introduce Rasmussen *et al.*'s (1994) work domain analysis (WDA)—in particular the AH formalism—to perform an analysis of how sensor information is used to calculate the higher-level variables that are needed when carrying out EID. We introduce "sensor-annotated AHs" that represent this information visually. We use sensor-annotated AHs to determine the impact of different kinds of instrumentation configurations on configural displays, which are integrative graphical elements that may form part of an EID-based interface. Finally, we work through a simple example with a water reservoir.

In Part II (the accompanying paper) we perform a fuller analysis with a simulated pasteurization plant. Our goal in Part II is to show that the techniques developed here extend to a more complex system. The Pasteurizer II microworld simulation described in Part II is supporting a program of analysis and experimentation on the impact of different instrumentation policies on human performance with EID-based interfaces.

In what follows we are not attempting to replace existing methods for performing instrumentation engineering. Instead, we wish to see if the analytic tools of EID might help cognitive engineers and instrumentation engineers communicate more effectively when considering the information that human operators might need if they are to preserve system functioning in the face of the unexpected.

2. Ecological interface design

In this section, we provide a brief description of EID as an approach to supporting the human operator. We distinguish EID from other approaches to interface design, and we summarize EID's successes to date and challenges to overcome. Finally, we develop a suggestion by Vicente (1999) that ecological interfaces help human operators handle

unanticipated situations by providing what is known in control theory as analytic redundancy (Frank, 1990; Sha, Rajkumar & Gagliardi, 1996; Seto, Krogh, Sha & Chutinan, 1998).

EID is an approach to the design of displays for complex dynamic processes that has been under development for the last decade (Rasmussen & Vicente, 1989; Vicente & Rasmussen, 1990, 1992; Rasmussen *et al.*, 1994; Vicente, 1996; 1999; in press). EID is based on the insight that most major incidents and accidents occur when human operators encounter conditions unanticipated by systems designers. Therefore, the main goal of EID is to provide principles for the design of interfaces that will help human operators in unanticipated conditions, while at the same time preserving their ability to exercise normal control and to handle anticipated abnormalities.

EID is based on two ideas from Rasmussen's cognitive work analysis (CWA) (Rasmussen *et al.*, 1994; Vicente, 1999). First, the domain of work (roughly, the system or process) should be analysed as a structural means–ends hierarchy. This involves using the AH formalism, which is the key analytic template for WDA. Second, human cognitive work should be supported at the most appropriate level of cognitive control. This involves using the skills–rules–knowledge (SRK)-based behavior distinction. Three design principles emerge for EID from the SRK distinction, described by Vicente (in press) as follows.

1. Skill-based behavior—workers should be able to act directly on the interface.
2. Rule-based behavior—there should be a consistent one-to-one mapping between the work domain constraints and the perceptual information in the interface.
3. Knowledge-based behavior—the interface should represent the work domain in the form of an abstraction hierarchy to serve as an externalized mental model for problem solving. (Vicente, in press, p. 4).

In the sections that follow we provide more detail about the AH formalism and the SRK distinction, noting their significance for EID. Finally, we discuss semantic mapping, which is one means by which the EID design principles can be instantiated.

2.1. WORK DOMAIN ANALYSIS

WDA is a description of the structure and functioning of the domain of work, independent of events, tasks, strategies, or actors. WDA therefore describes the "field" upon which activity will take place (Rasmussen *et al.*, 1994; Vicente, 1999). WDA is therefore distinctly different from task analysis, which usually specifies certain goals, tasks, events, and may even specify strategies and actors. The difference between WDA and various forms of task analysis have been discussed in detail in Vicente (1999, Chapter 7).

The analytic tool usually used to perform WDA is a two-dimensional framework in which the work domain is described at different levels of abstraction (i.e. in different analytic languages) and at different levels of decomposition (i.e. structural aggregation).

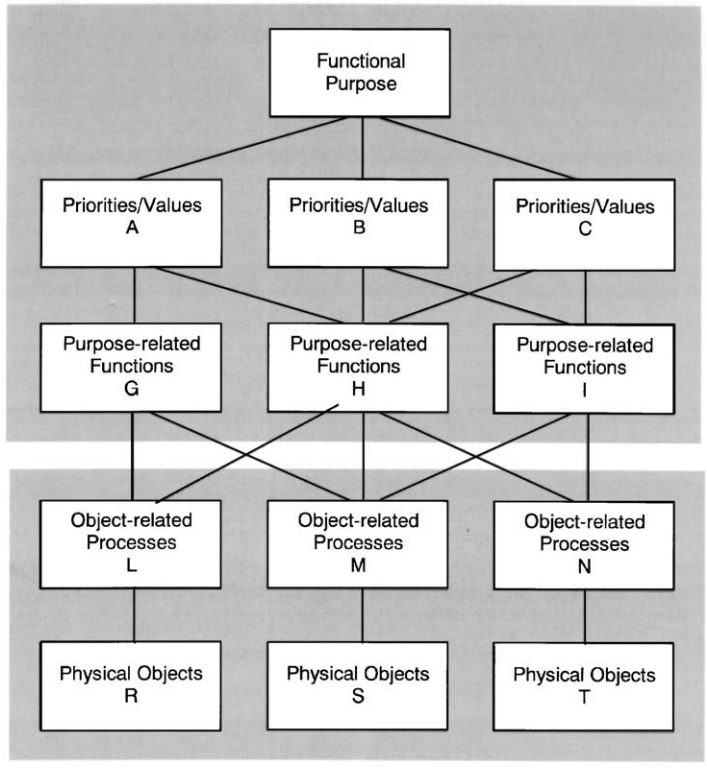

FIGURE 1. "Lattice-like" representation of the abstraction dimension of work domain analysis, with some sample means–ends relations shown. Upper part shows purposive relations and lower part physical relations.

In this paper we will work principally with the dimension of abstraction, which is shown in Figure 1 in a lattice format. The five levels of abstraction[§] include the following.

- Functional Purpose—the reason or purpose that the system or domain exists.
- Priorities/values—higher-order properties, priorities and values of the work domain that are preserved, conserved, maximized or minimized, such as the conservation of mass and energy.
- Purpose-related Functions—the functions, devoid of physical instantiation, that must be present for the functional purpose of the work domain to be fulfilled.

[§]The labels of the levels of abstraction can change across domains and to highlight particular meanings. Labels often used are Functional Purpose, Abstract Function, Generalized Function, Physical Function, and Physical Form. The labels used in the present paper reflect some of Rasmussens's more recent thinking about the nature of the abstraction hierarchy (Rasmussen, pers. comm., February, 1998) and we have chosen to adopt them. See Reising (1999) for more details of the origin and rationale for the labels being used for the abstraction hierarchy levels in this paper.

- Object-related Processes—the functional properties of the physical elements such as vats, pumps, heaters and so on, without necessary reference to their particular purpose in the work domain in question.
- Physical Objects—enumeration and description of the literal hardware form of devices and instruments, and their configuration.

At each level in the abstraction hierarchy the whole system is described, but with a different language of description. The nodes connected across different levels of abstraction have a means–ends or why–how relation to each other. From a node at a given point, nodes above indicate why the device, function, or purpose is there (ends) whereas nodes below indicate how the device, function or purpose is implemented (means). Lines in Figure 1 between nodes indicate specific means–ends relations that might be relevant for specific system functions. Not shown in Figure 1 are topographic links that show flow of function or information within a level of abstraction.

The decomposition dimension describes the work domain at various levels of granularity, from the whole work domain or system, to subsystems, to individual components. Levels of decomposition have a part–whole relation to each other which is clearly distinct from the means–ends relation.

An ecological interface supports operator reasoning at different levels of abstraction by showing directly the system's governing constraints within and across levels. An ecological interface will therefore show the current system state in the context of a model of *proper functioning* of that system. The ecological interface should reveal all physical and functional boundaries and constraints so that any deviation from proper functioning will stand out. In this manner, human operators can more quickly see where the system is operating in relation to its physical and functional boundaries and constraints, and can plan appropriate action. As a result, the operator is better supported for dealing with both routine and non-routine events, including unanticipated system states (Rasmussen & Vicente, 1989; Vicente & Rasmussen, 1990, 1992) and the whole human–machine system is more robust.

2.2. COGNITIVE CONTROL AND SRK

One of the goals of EID is to simultaneously support the skill-, rule- and knowledge-based levels of cognitive control. Skill-based behavior is the automatic sensory-motor actions that are in response to space–time signals from the environment. Rule-based behavior can be accounted for by if–then rules in response to signs from the environment, where the rules are either stored in memory or externalized in operating procedures. Knowledge-based behavior describes the operator's cognitive activity when problem solving has to be actively engaged. In this situation, the operator cannot rely on skill- or rule-based reasoning as signals and signs for action are missing—instead the information from the environment is in the form of symbols whose significance must be worked out before action can happen. The operator formulates a target state based on the system's purpose, assesses the environmental conditions, and forms a course of action to reach the target state.

The skill-, rule- and knowledge-based behavior distinctions involve increasing amounts of "executive" cognitive control and workload. Therefore, EID proposes that

a display or interface should not force the level of cognitive control to higher than the minimum required for a task (Vicente & Rasmussen, 1990, 1992). The idea is to promote effective human operator action by letting operators take full advantage of their perceptual capabilities where possible, while at the same time providing support for more cognitively loading diagnostic and planning activities.

2.3. SEMANTIC MAPPING

Given an analysis of the work domain and requirements for cognitive control, the display designer has necessary (but not always sufficient) information about the system parameters and properties that should be displayed to human operators. The designer seeks an ecological interface that will reveal the first principles of operation of the work domain in question and that will reveal possibilities for action and inherent constraints within the system as operators pursue system goals [see Rasmussen *et al*. (1994) and Vicente (1996, 1999) for examples of this process].

One powerful way of achieving the above goal is to use analytical geometry to map the system's governing constraints—both physical and purposive—to the dynamic behavior of graphical forms so that operators can directly perceive system state and action alternatives. This process has been called semantic mapping (Woods, 1991; Bennett & Flach, 1992; Bennett, Flach & Nagy, 1997). Further discussions of semantic mapping are available in Hansen (1995) and Reising and Sanderson (in press).

Figure 2 provides an illustration of the logic behind semantic mapping. At left is a geometric form—a rectangle—whose area is the product of its width and height. The width/height/area form is a faithful geometric representation of the mathematical relationship between two numbers and their product. At right of Figure 2 is a physical relationship that can be described as a relationship between two numbers and their product, and so can be faithfully represented with the width/height/area form. The result is a so-called "configural" display that shows the contributing elements (width and height) the whole (area) plus the physical boundaries of operation. Because configural displays show higher-level properties as well as the values of individual

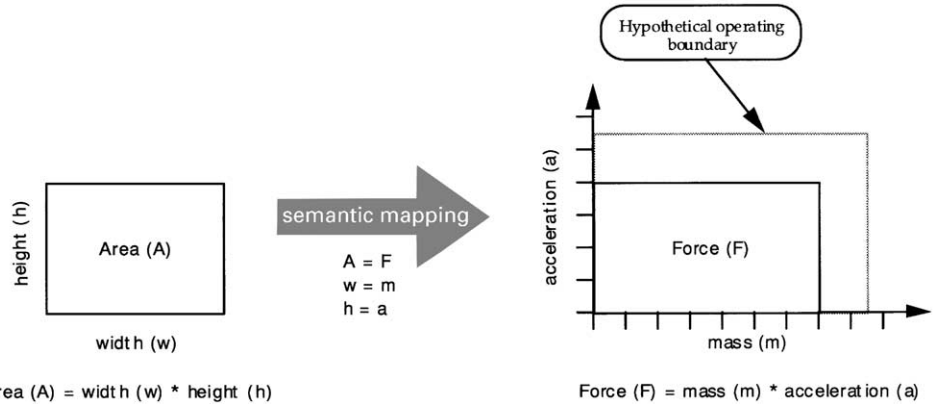

FIGURE 2. Illustration of the principle of semantic mapping between a geometric form (left) and a physical relationship of interest to generate a configural display (right) with boundaries of operation marked.

system elements, they are often used as building blocks in ecological interfaces. However, as work by Ham and Yoon (2001, in press) has shown, the principles behind EID work even without a strongly graphical approach to semantic mapping.

2.4. DISTINCTION BETWEEN EID AND NON-EID APPROACHES

Before proceeding with a discussion of empirical support for EID and remaining issues, it may be helpful to briefly note similarities and differences with other interface development methods. EID is not unique in advocating a thorough analysis of the system the human operator must work with. Many researchers—particularly those in the European tradition—also advocate such an approach [see for example work by Sundstrom (1993, 1997), Hollnagel (1993, 1998), Bainbridge (1991) and Johannsen (1992) amongst others]. Indeed, the AH component of EID builds on a tradition to which these and other researchers contributed significantly (e.g. Lind, 1981). Moreover, EID is not unique in advocating the use of visualizations of system state that reduce cognitive workload because many others have explored this (Duncan, Praetorius & Milne, 1989; Woods, 1991; Bainbridge, 1991).

EID differs from non-EID approaches in two ways: (1) it focuses on supporting human operator in the face of unanticipated variability and (2) it focuses on analysing the domain of work independently of tasks, events, strategies or actors. Many other approaches focus on identifying events and contingencies for which interfaces should support operators, and tasks that operators must perform (Reed, 1992; Wells, 1997). Although such approaches are an essential part of interface design, they are targeted at supporting human operators under conditions different from those that EID targets.

We noted previously that because the AH is a model of proper functioning of a system, showing all physical and functional constraints and boundaries, any deviation from proper functioning due to a failure will be evident. Lind (1981) discusses this property with respect to the use of mass and energy conservation laws to help detect system failures and Vicente (1999) relates it to the concept of analytic redundancy in control theory (Frank, 1990). The general principle is shown in the top part of Figure 3. The AH functions as a kind of "observer" to the system in question, the constraints it embodies demonstrating proper functioning and providing an interpretive framework for the human operator (rather than automation) to detect any deviation. Moreover, one can view the control capabilities the human operator gets by using the AH representation as a control backup that has broader capabilities than those provided by a controller specific to anticipated events, tasks and states. This is shown in the lower part of Figure 3. This is a further form of analytic redundancy analogous to that used to preserve effective control over systems undergoing on-line upgrades (Seto *et al.*, 2001). Overall, the concept of analytic redundancy helps explain how EID differs from other approaches to interface design.

2.5. EMPIRICAL SUPPORT FOR EID

EID is proving to be a powerful way of supporting human operators of complex dynamic processes, especially when operators are faced with unanticipated events or unexpected variability. Lee, Kinghorn and Sanquist (1995) and Vicente (1996, 2002)

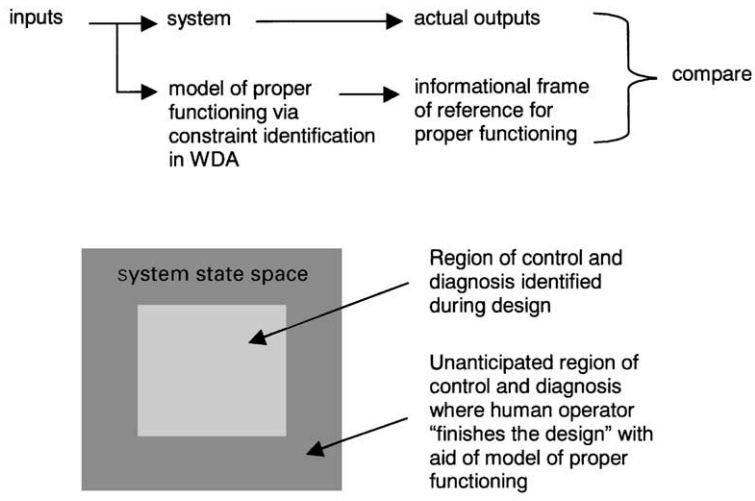

FIGURE 3. Illustration of the principle of analytic redundancy. Top diagram shows WDA as a set of constraints on outputs given proper functioning. Lower diagram shows that the region of control and diagnosis identified during design is contained with the model of proper functioning identified via WDA constraints.

have reviewed the growing number of studies in this area and noted further research needs. Rather than repeat those finding in detail, in the following sections we point to the most robust findings and highlight shortcomings that are evident. In particular, we focus on the problem of sensor and instrumentation engineering.

2.5.1. Demonstrations of EID's effectiveness. Results of studies using Vicente's DURESS simulation show that interfaces designed under EID principles support better human operator performance under unanticipated conditions, faster and more accurate fault detection and diagnosis, less variable control performance, and better declarative knowledge about functional properties of a system, than interfaces not designed using EID principles (for a full summary and review of studies with DURESS see Vicente, 1996). Although concerns have been expressed about the adequacy of experimental contrasts between displays in these experiments (Maddox, 1996; Reising, 1999) the accumulation of evidence with DURESS suggests that EID-based interfaces help human operators cope with unanticipated variability better than non-EID interfaces. At the same time, EID-based interfaces support human operators at least as well as non-EID interfaces when handling normal control and anticipated forms of abnormality (Christofferson, Hunter & Vicente, 1996; Pawlak & Vicente, 1996). Further studies have found that the variability of performance during normal operation was less when subjects were using the EID-based interface—this is a clear benefit of an EID-based interface (Yu, Lau, Vicente & Carter, 1998).

Subsequent experiments using systems other than DURESS have provided additional evidence of EID's effectiveness and have addressed some concerns (for a full review see Vicente, in press). For example, Reising and Sanderson (2000*a*) have used a simulated pasteurization plant to show that EID supports sensor and system fault detection effectively only if the underlying instrumentation is adequate—an issue

that will be detailed in this and the accompanying paper. Ham and Yoon (2001*a, b*) have used a simulated nuclear power plant to show that performance at handling abnormalities— especially unanticipated abnormalities—improved as information from adjacent levels in the AH was added to an interface, even when simple bar graph displays rather than highly configural graphics were used.

In recent years, EID studies have extended from relatively simple process simulations such as the above, to development of prototype EID interfaces and proof-of-concept studies for large-scale systems such as real-life chemical and energy process control (Olsson & Lee, 1994; Dinadis & Vicente, 1996; Reising, Sanderson, Jones, Moray & Rasmussen, 1998; Burns, 2000*a*, *b*), aviation (Dinadis & Vicente, 1999), network management (Burns, Barsalou, Handler, Kuo & Harrison, 2000), and medicine (Sharp & Helmicki, 1998). Furthermore, some EID prototype interfaces have been implemented in full-scope simulators (Itoh, Sakuma & Monta, 1995; Yamaguchi & Tanabe, 2000). Further examples are provided in Vicente (1999, 2002). All these efforts represent various points in the necessarily conservative process to be taken when moving a new interface concept from the laboratory to full implementation in high-risk industries and therefore cannot be expected to have happened much earlier.

2.5.2. Equivocal or problematic results with EID. Overall, there is still some way to go before EID will be a well-defined engineering approach to human—machine interface design that comes with assurances about the quality of human performance nd predictions of exactly where display superiority should be found. Some of the previous shortcomings are starting to be addressed. For example, EID is already being applied to complex, real-world systems, as outlined in the previous section. The problem of handling the very large number of variables found in real-world systems through spatial and temporal integration is being handled by Burns (2000*a*, *b*). Clear design principles for moving from the results of an AH and SRK analysis to an effective visual representation in an interface are starting to be articulated (Rasmussen, 1999; Reising & Sanderson, 2002; Liu, Nakata & Furuta, 2002).

Of particular concern for the present research, however, is the shortcoming that designers and researchers do not know how effective ecological interfaces are when critical sensors are unreliable (Anyakora & Lees, 1974; Vicente & Burns, 1995) or, worse, unavailable, as is often the case in industry (Hayter, 1996; Beltracchi, 1998*a*) and particularly in the medical domain (Sharp & Helmicki, 1998; Hajdukiewicz, Vicente, Doyle, Milgram & Burns, 2001). Vicente has recently labeled this issue as still a high priority issue to resolve for EID, stating:

> Many important issues have yet to be addressed, let alone solved. Some of these issues, such as sensor failure, may turn out to be "show stoppers" (Vicente, 2002).

Most performance tests of EID-based interfaces have used fundamentally reliable data because the tests have been performed in laboratory or simulator environments. The concern of the present research is whether sensor presence/absence, unreliability or failure will seriously compromise the effectiveness of EID-based interfaces, which can depend heavily on displaying information derived mathematically from lower-level sensor information in a graphical manner. Clearly, an EID-based interface cannot provide effective analytic redundancy of the kinds shown in Figure 3 if the interface is

ineffectively showing the higher-order constraints governing the operation of a work domain. In the section that follows, challenges to EID from sensor and instrumentation issues are outlined that point to the need for further analytic work.

2.6. CHALLENGES TO EID FROM INSTRUMENTAITON

Challenges to EID from instrumentation come from limitations in the sensor technology, the need to derive or borrow some values to obtain higher-order information and the impact of sensor inadequacies on displays in an interface. The key issue is that when sensor failures are difficult to detect, they make it more difficult for the operator to distinguish sensor failures (requiring subsequent maintenance or recalibration) from true systems failures (requiring a corrective response).

2.6.1. Sensor technology limitations. Two kinds of instrumentation limitations are sensor accessibility (Beltracchi, 1998*a*; Hayter, 1996; Lindsay, 1990; Vicente & Rasmussen, 1992) and sensor reliability (Vicente & Rasmussen, 1992). Most research on EID has assumed that the variables and constraints that should be displayed will be measurable and will be reliable (although for discussions of potential problems see Vicente & Rasmussen, 1992; Reising & Sanderson, 1996; Sharp & Helmicki, 1998; Jamieson, 1998; Hadjukiewicz *et al*., 2001). Most empirical tests of ecological interfaces have used fundamentally reliable sensor data, even when system failures have been tested. This is partly because evaluations have been performed with digital simulations.

2.6.2. Derivations based on limited sensor information. EID advocates displaying higher-order variables and relationships for a system, such as mass and mass balance, and energy and energy balance. However, these variables and relationships have to be derived from lower-order variables such as volume, temperature, density and specific heat (Reising & Sanderson, 1996). This may create difficulties if such lower-order information is unreliable or if it must itself be estimated. Moreover, if there are few sensors in the system, information about some parts of the system may need to be borrowed from sensors at adjacent, but separate, parts of the system.

2.6.3. Impact of sensor failures on interface. Little has been written about designing EID-based interfaces that will be robust in the face of instrumentation failures, other than to highlight the possible ambiguities that could arise (Moray, Lee, Vicente, Jones & Rasmussen, 1994; Reising & Sanderson, 1996, 1998; Vicente *et al*., 1996). Because EID-based interfaces represent information at all levels of the AH, they will usually depend heavily on information derived mathematically from lower-level sensor information. If that information is wrong, then it may be reflected in the interface in a manner that makes it difficult to work out what has happened. For example, the human operator may find it difficult to distinguish sensor failures and system failures (physical plant malfunctions that have functional manifestations, such as leaks, blockages, failures or valves, pumps, and heaters, and all failures of electronic circuitry independent of sensors). Using the approach outlined in the present paper, Reising and Sanderson (2000*a*, *b*) have provided some first empirical evidence of the impact of inadequate vs. adequate instrumentation on an EID-based interface.

Given that these aspects of sensor technology do exist, in this paper we address the effectiveness of EID-based interfaces in the context of the following general questions.

- When would we expect the performance of operators using the configural displays typical of EID to be compromised if the sensors upon which those displays are based fail, or are unreliable?
- Is sensor failure detection by human operators improved when the higher level information suggested by the EID framework is provided?
- How badly is operator performance compromised if higher-order variables are derived not from dedicated sensors that measure raw data directly, but rather from higher-order variables that are themselves are derived from distally measured raw data?

3. Sensor unreliability and unavailability

If we are to design EID-based interfaces that are robust over sensor failures, then we need to understand something of the different forms of sensor inadequacy. We also need to understand current instrumentation and control engineering methods for handling sensor inadequacies.

3.1. SENSOR INADEQUACIES

For every property of an engineered system, a sensor may be *present* or *absent*. In what follows we discuss the inadequacies that can occur in each case.

3.1.1. Sensor present problems. If a sensor is present, either it may give a reading that is *accurate* within expected normal variation, it may give an *unreliable* reading, or its reading may be *unavailable* (i.e. the sensor has failed outright). The distinction between unreliability and unavailability is fluid, depending upon when the operator decides that a reading carries no information about its parameter. For example, if a reading drops to zero, sits at a constant value, drifts or behaves erratically, then the operator may treat it as unavailable. However, if the reading either varies within a too-narrow range, shows high-frequency disturbance or suddenly becomes displaced, then some information value may be preserved and the operator may treat the reading simply as unreliable (cf. Anyakora & Lees, 1974).

If the operator knows the failure modes of a sensor, then sensor output can carry considerable information. A sensor failure can originate at either the transduction from a physical event to a (usually) digital signal, the transmission of the digital signal from source to display, or the transformation of the signal to a digital or analog display format (Jovic, 1992; Johnson, 1993). Digital readouts and digitally based displays (whether alphanumeric or graphical) have long been known to have special problems. Relatively early in the conversion of control rooms to digital technology, Anyakora and Lees (1974) noted that "... some care is required in the design of the [CRT] display if it is to be a facility which is truly equivalent to the recorder with respect to malfunction detection. For example, the operator learns much from the noise on a chart record and this may undergo modification on a CRT" (p. 248). More recently, Stubler and O'Hara

(1996) note that digital sensors can lock up with no obvious indication of trouble. With an analog system a needle will usually fall to zero, whereas with digital sensors there may be no obvious indication of trouble because the reading may "freeze" at a value within normal range. Stubler and O'Hara also note that lockups can happen after a command has gone out but before it takes place physically, which results in a mismatch between the displayed command, the physical state of the system, and the sensor reading. Knowing that such new failure modes exist does little to reduce the complexity facing the operator when trying to assess evidence for system vs. sensor failures.

3.1.2. Sensor absent problems. A sensor may be absent because the property is inherently unmeasurable, because the technology does not apply to measure it, or for practical reasons related to the measurement process.

First, the sensor may be absent because the property is not inherently measurable (e.g. entropy, enthalpy[¶] and so on). One real-world example comes from research on direct perception displays for nuclear power plant (NPP) control (Moray *et al*., 1994). An ideal variable to display in the case of NPP control would be the departure from nucleate boiling ratio (DNBR) (Reising & Sanderson, 1996). However, to measure DNBR directly, a sensor would have to measure the ratio of steam bubbles to water pockets along the surface of each fuel rod in the reactor core. Beltracchi (2000) has also recently suggested that DNBR would be a useful parameter to communicate to plant operators. However, he points out that "There is no indication of how much steam is being produced through nucleate boiling, thus there is no way to determine how much water to replace" (Beltracchi, 1998*b*, p. 13).

Second, the sensor may be absent because the technology required to measure the property does not exist. Another real-world example from the NPP control is monitoring the neutron flux of the reactor core. Although sensors are located around the core, there are no sensors in the center of the core to provide continuous, accurate indications of the neutron flux profiles throughout the various quadrants of the reactor core. Because of the harsh conditions of the nuclear core, no sensors presently exist that could give accurate, sustained measures of the variable. Presently, sensors that measure the flux are dropped down into the core at regular (weekly or monthly) intervals, from which neutron flux profiles are generated and updated (Dr Barclay Jones, pers. comm., January 25, 1996).

Third, the sensor may be absent because the property is measurable but remains unmeasured for reasons of cost, standardization or concern about the impact of the measurement process (for examples from a medical domain: see Sharp & Helmicki, 1998, p. 352).

3.2. INSTRUMENTATION AND CONTROL ENGINEERING METHODS TO HANDLE SENSOR INADEQUACIES

The preceding problems are well understood in the area of instrumentation and control engineering. Because any assessment of plant status starts with sensor data, sensor failures must be distinguished from system failures and the true value of the sensed

[¶] See Beltracchi (1998*a*, p. 8).

variable determined. This is especially important when sensor information helps to drive automatic controllers.

For this reason, engineers have developed sophisticated methods for determining the validity of sensor data. Most methods exploit different forms of redundancy, which Lee (1994) has classified grossly as spatial and temporal redundancy. *Spatial redundancy* includes replication, functional and analytic redundancy (Sha *et al*., 1996). Replication redundancy usually provides an identical sensor at the same sensing point. Given the same inputs, the output of the two sensors should be identical. Functional redundancy works just as replication redundancy does, but with different engineering to reduce the chance that a shared design flaw might make both sensors fail at the same time. Various forms of cross-calibration and voting schemes determine the final value accepted (Gotcher & Burroni, 1993; Hashemain, Riner, Bunch & Petersen, 1993). Finally, analytic redundancy is based on a model of the system that might use quite different inputs and produce quite different outputs from the sensor itself, but that allows the sensor value to be validated (Deutsch, Ornedo & Lindsay, 1983; Clark & Campbell, 1982; Lee, 1994; Sha *et al*., 1996; Dorr, Kratz, Ragot, Loisy & Germain, 1997). Sensor validation methods using *temporal redundancy* depend on the analysis of sensor values over time, using filtering techniques (Massoumnia, 1986) and Bayesian approaches (Dragoni, Giorgini & Pant, 1998) for example. Pattern recognition techniques (Griebenow, Hansen & Sudduth, 1995) and neural nets (Himmelblau & Bhalodia, 1995) have also been used for sensor validation. Overall, such advances have made it possible to consider the development of "smart sensors" that will be self-calibrating and self-diagnosing (Doyle, Garrison, Johnson & Smith, 1998).

Our goal in this paper is not to propose a new form of sensor data validation or a new way to estimate state variable values in the absence of sensor information. We are not concerned with the information an automatic controller requires, but instead the information the human operator requires to handle unanticipated variability. Our goal is to establish analytic tools that indicate the information an EID-based interface must display and to identify when sensor inadequacies might threaten the integrity of that information.

As Stubler and O'Hara (1996) have noted, sensor validation techniques are not guaranteed to be perfect, even though they are very powerful. Therefore, the human will sometimes have to detect sensor error, particularly in situations unanticipated by the sensor validation techniques on hand. Anyakora and Lees (1974) have noted that humans use redundant information from *a priori* expectations about sensed values, past signals from sensors, readings from duplicate sensors, readings from other sensors that are related in known ways to the sensors in question, or the topographic position of an sensor. Our point is that humans might also exploit the analytic redundancy provided by an AH-based representation to discriminate sensor from system failure. In the section that follows we explore the prospects for such an approach.

4. Effects of sensors on ecological interfaces

In this section we review in more detail what is known about the interaction between instrumentation and EID-based interfaces, looking at both field and laboratory studies.

4.1. MEASURABLE AND AVAILABLE DATA IN THE FIELD

EID assumes that variables and constraints to be displayed are measurable and available. An EID-based interface requires certain key physical properties of a system to be sensed or derived so that they can be displayed (Vicente & Rasmussen, 1992). The consequences of any shortcoming have been brought home nicely by Hayter (1996) who analysed the potential for an EID-based interface to be installed in an existing nuclear power plant for boiler level control. He found that there was insufficient instrumentation in the plant to support the calculation of mass or energy—specifically, there were inadequate flow sensors and flow calculations in the plant. Hayter concluded that "the existing field instrumentation [...] is not sufficient to allow the direct implementation of an ecological interface design based display in the suite of displays provided by the new plant display system" and recommended that in future system upgrades flow sensors and associated cabling should be installed.

Likewise, Lindsay (1990) reports on the development of a direct perception display for an experimental breeder reactor, "which, *to the extent that instrumentation is available*, represents a top down design with the thesis that a nuclear power plant is a heat engine" (p. 266; emphasis added). Therefore, as Reising and Sanderson (1996) have argued, the needs of the human operator may contribute new requirements to instrumentation engineering during plant design or redesign.

4.2. EFFECTS OF SENSOR SHORTCOMINGS ON ECOLOGICAL INTERFACES

To date there has been no systematic study of the effects of sensor unreliability or failure on the visual coherence of EID-based interfaces. All performance tests of EID-based interfaces have used fundamentally reliable data because the tests have been performed in laboratory or simulator environments. However, Vicente and Rasmussen (1992) noted that "empirical research is needed to determine how robust performance with an interface based on the abstraction hierarchy is with respect to these sources of uncertainty" (p. 600). Moreover, in their study of the Rankine Cycle display, Vicente *et al.* (1996) acknowledged that sensor failures could conceivably compromise the geometric form and therefore the information value of such displays. As they state, "the effect of a failed sensor on operators' understanding and control of the plant is unknown. [...] when one data point fails, the lines showing the relations between that point and other points connected to it are distorted [...] it is not known what effect this might have on operators' abilities to interpret the remaining information on the display" (p. 520). A further example with the Pasteurizer II simulation is provided in Reising and Sanderson (2002). These concerns relate to the reliability of software underlying displays—how displays should behave when unusual values are sent to them so that operators can distinguish sensor from system failures.

4.3. DISCRIMINATING SENSOR AND SYSTEM FAILURES WITH ECOLOGICAL INTERFACES

Proponents of EID would predict that EID-based interfaces should lead to faster and more accurate discrimination of sensor failures, and there is indirect evidence to show this. Beltracchi's (1987) Rankine Cycle display of the thermodynamic heat engine cycle in nuclear power plants has been evaluated against more conventional displays (Moray

et al., 1994; Vicente *et al*., 1996). Two failure detection scenarios involved sensor failures—one a sudden drop to zero and the other an exponential drift to zero. There was an overall superiority of the Rankine Cycle display for failure detection, with sensor failure detection seeing the same kind of boost as system failure detection.

Simulator-based studies have occasionally used sensor failures alongside system failures as singular events to be detected, rather than as a persistent or emerging state of the system to be handled while other tasks are performed. In such studies there was no conscious analysis of how detectable the sensor failures would be, using heuristics similar to those proposed by Anyakora and Lees (1974), under different instrumentation design policies. Therefore, we do not know the effect of EID-based interfaces on operators' ability to detect sensor failures or how operators control the system with an EID-based interface when there are sensor failures.

The empirical evaluation of the Rankine Cycle display (Moray *et al*., 1994; Vicente *et al*., 1996) and the EID-based interface for Pasteurizer II (Reising & Sanderson, 2000*a*, *b*) are probably the only studies to date that have evaluated the effectiveness of EID-based displays for diagnosing sensor failure. Previous studies specifically examining EID (Bisantz & Vicente, 1994; Vicente, Christoffersen, & Pereklita, 1995; Christoffersen *et al*., 1996, 1997; Pawlak & Vicente, 1996) and other advanced displays (e.g. Bennett, Toms & Woods, 1993; Edlund & Lewis, 1994; Gillie & Berry, 1994) have almost exclusively used system failures. This observation is not meant to be a criticism of those studies, however. They are part of a long tradition of research into the human operator in process control which has focused exclusively on system failures when examining fault diagnosis (e.g. Rasmussen & Rouse, 1981; Morris & Rouse, 1985; Duncan, 1987; Moray & Rotenberg, 1989; Lee & Moray, 1992, 1994).

5. A simple example

A very simple water reservoir example helps us illustrate the use of the AH to show how sensors are placed and information derived (cf. Borer, 1991, p. 146, 161; Johnson, 1993, p. 485). Such a water reservoir might be a part of a sub-system in petrochemical refineries, conventional or nuclear power stations, and so on (see Figure 4). Normally, an AH would not be developed for such a small component of a system. However, by choosing a simple yet fully worked example we are better able to convey the purpose and power of sensor-annotated AHs as a tool for EID and the reader is better prepared for the more complex pasteurization example in Part II (Reising & Sanderson, 2002).

Assume that the purpose of the reservoir sub-system shown in Figure 4 is to ensure that the next sub-system, whatever it is, receives a supply of fluid at a constant volume flow rate. As the reservoir sub-system is currently configured, a constant volume flow rate at the output pipe would be achieved by maintaining a constant volume in the vat (Borer, 1991, p. 146). The volume of the vat can be controlled by the valve on the input pipe. An AH analysis of this sub-system is presented in Figure 5.

In general, EID would advocate the use of a configural display that supported the direct perception of higher-order properties and means–ends relations in the reservoir sub-system. A configural display can be adapted from the EID-based interface for DURESS (Vicente, 1991) (see Figure 6). The meter in Figure 6 labeled V_A corresponds

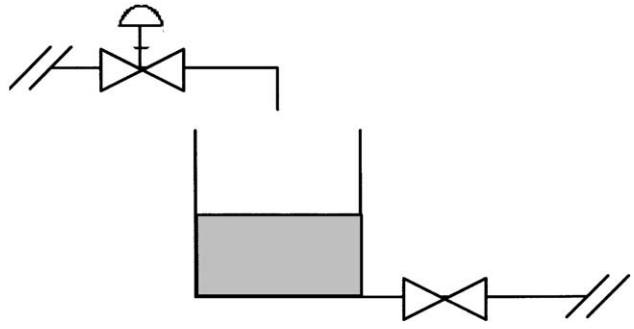

FIGURE 4. A common "system" used in instrumentation and control texts, stripped of any transducers and automatic controllers.

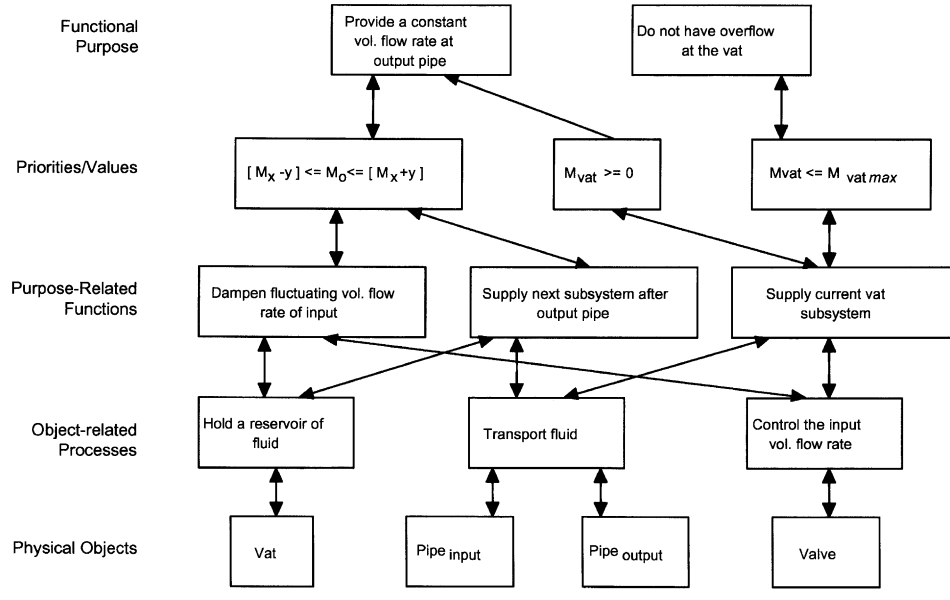

FIGURE 5. The AH analysis showing the means–ends relations for the reservoir sub-system presented in Figure 4.

to the setting of the valve on the input pipe in Figure 4 that controls the volume flow rate for the sub-system. The volume of the vat has been transformed to mass (labeled M_V in Figure 6) and is displayed in the center of the "funnel" display. The output flow rate is expressed as mass flow rate (labeled M_o in Figure 6) and has the target region that should be met, per the first node at the Values & Priority Measures level of the AH presented in Figure 5. The other two constraints at the Values & Priority Measures level are also indicated in Figure 6.

Under normal control engineering practice, the system would probably be instrumented as presented in Figure 7. By instrumenting the system for cascade

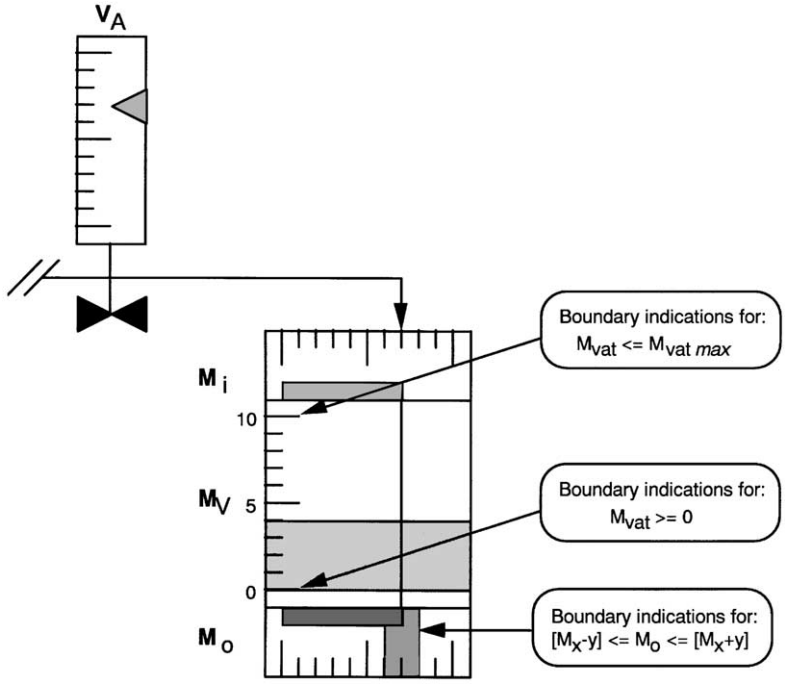

FIGURE 6. The configural display meant to capture the means–ends relations presented in Figure 5 for the reservoir sub-system.

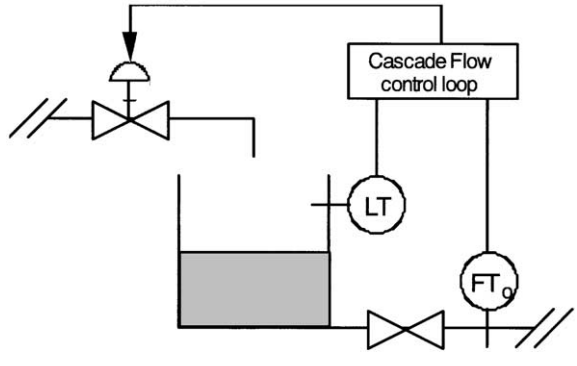

FIGURE 7. The reservoir sub-system with necessary transducers and cascade control loop (see Johnson, 1993, p. 486) (LT = Level Transducer; FT = flow transducer).

control (see Johnson, 1993, p. 486), the sub-system maintains a constant volume level in the vat such that the pressure head maintains a constant volume flow rate while not overflowing the vat—given that no sensor failures or system failures arise. Given this instrumentation configuration, the information at each level of the AH that should be displayed to the operator can be assessed (see Figure 8).

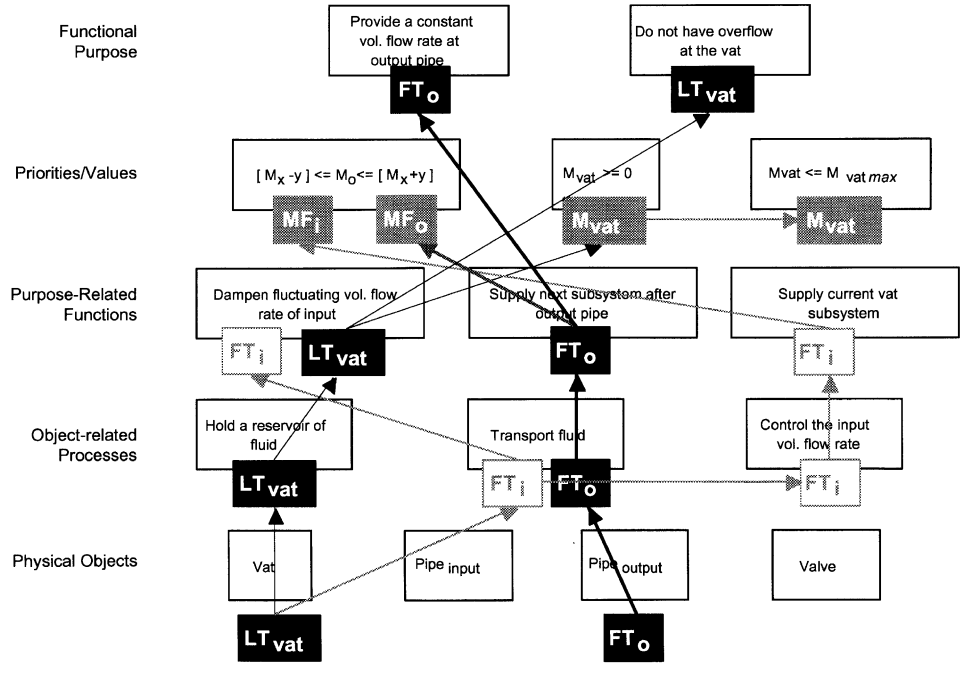

FIGURE 8. The sensor-annotated AH for the reservoir sub-system in Figure 7, showing the "location" of sensors at the Physical Object level, and the propagation of information from those sensors up through the AH. Black boxes show good information and dark gray boxes correspond to information that is normally derived. White boxes represent information that, if needed, must be derived by other means.

Note that certain "nodes" of information are supported directly by the dedicated sensors in the sub-system (see the black rectangles in Figure 8). However, other information must be derived. First, the mass flow rate for the output leg (MF_o in the gray box in Figure 8) and the mass in the vat (M_{vat} in the gray box in Figure 8) are derived by multiplying the measured variable by the density of the fluid. Second, information about the volume flow rate and mass flow rate of the input pipe (MF_i in the white box in Figure 8) would have to be derived from the change in vat volume per unit time. These values are needed to support the configural display—specifically the indication of the mass flow rate value (M_i in Figure 6) as well as the indication of equal mass flow rate into and out of the vat (the vertical line connecting M_i and M_o). Note that the robustness or "reliability" of the visual graphics in the configural display holds as long as no sensor failure occurs.

Given all the above, we can determine what will happen to the information to be displayed as well as to the visual representation of the display when a sensor does fail. For example, the level transducer indicating the volume of the vat might start to drift upwards. We could examine the propagation of the failure in the sensor-annotated AH in Figure 9. The mass level, M_{vat}, and the derived mass flow rate for the input pipe, MF_i, will increase as the level transducer drifts upwards.

The effect of the level transducer failure on the configural display is presented in Figure 10 where it is clear that M_i and M_v are increasing, even though V_A and M_o are

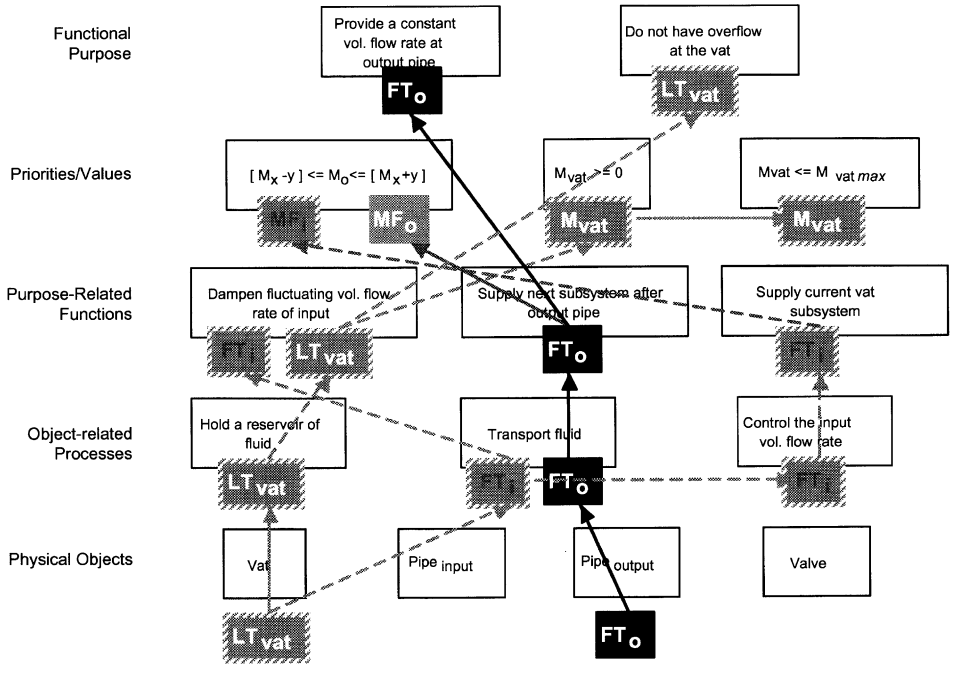

FIGURE 9. The annotated AH showing the propagation of compromised information (gray boxes with striped borders connected by dotted lines) caused by a faulty level transducer at the vat. Black boxes are good information and white boxes represent distally derived information.

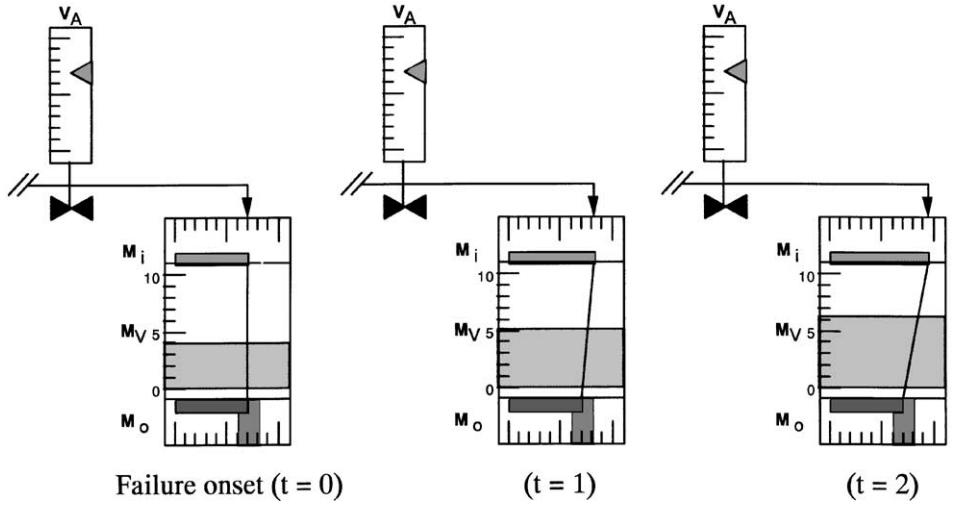

FIGURE 10. The effect of a drifting level transducer at the vat on the configural display when there is no measure of input flow rate. M_i appears to increase even though valve setting V_A is constant, because M_i is derived from an F_i already derived from the level transducer.

unaffected. An operator should still be able to meet the primary functional purpose. The sub-system maintains a constant volume flow rate to the next sub-system, because of the accurate performance of the flow transducer. However, the operator may try to prevent the vat from filling up by closing the valve V_A a little. This would have a negative effect on maintaining the constant flow rate at the output because it would effect the pressure head in the vat even though there was no actual system state change.

Therefore, based on the AH analysis and the information requirements of the configural display, a cognitive engineer could recommend that a flow transducer be placed on the input pipe to ensure more reliable communication of system state to the operator (see Figure 11). With an input flow transducer, less derivation is needed. Moreover, the consequence of any one sensor failure has been isolated from the other "paths" of information to be displayed (see Figure 12). If the same transducer failure were now to occur, the mass of the vat would still "increase" but the mass flow rate of the input pipe would not change. In this case, the operator might be more inclined to attribute the change in the mass of the vat to a faulty sensor and to correctly make no control action.

In general, we can distinguish between two forms of inadequacy that emerge from the AH when the instrumentation is underspecified: *topographic inadequacy* and *derivational inadequacy*. Such inadequacies are known by instrumentation engineers, but they have a particular significance in the context of EID.

Topographic inadequacy involves having impoverished sensor information within a level of abstraction. Specifically, the value of a variable measured by a sensor at a specific

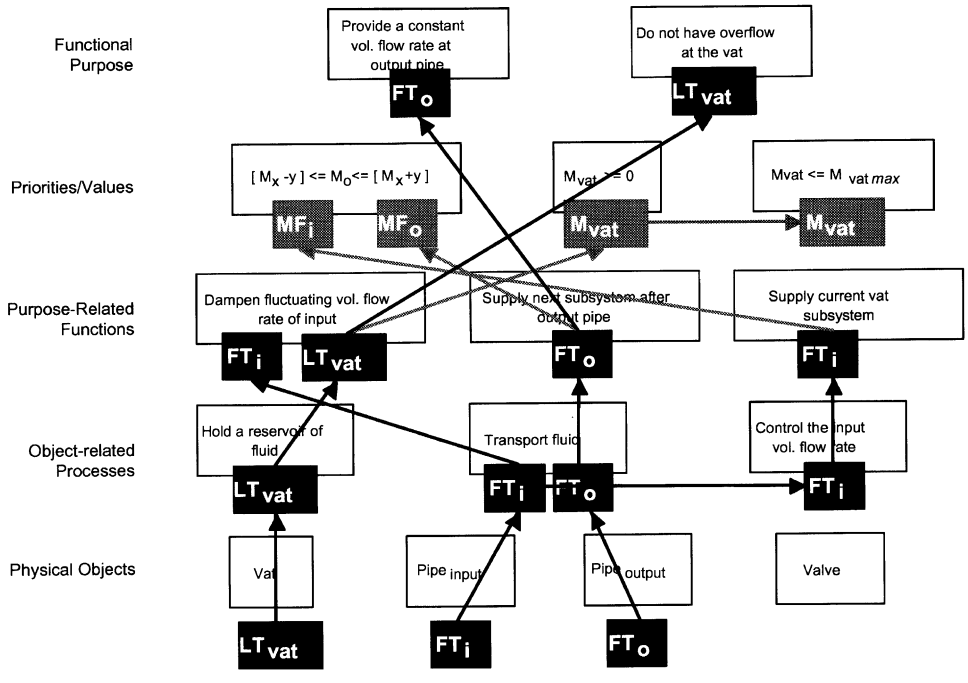

FIGURE 11. The annotated AH showing the propagation of information after a second flow transducer was added.

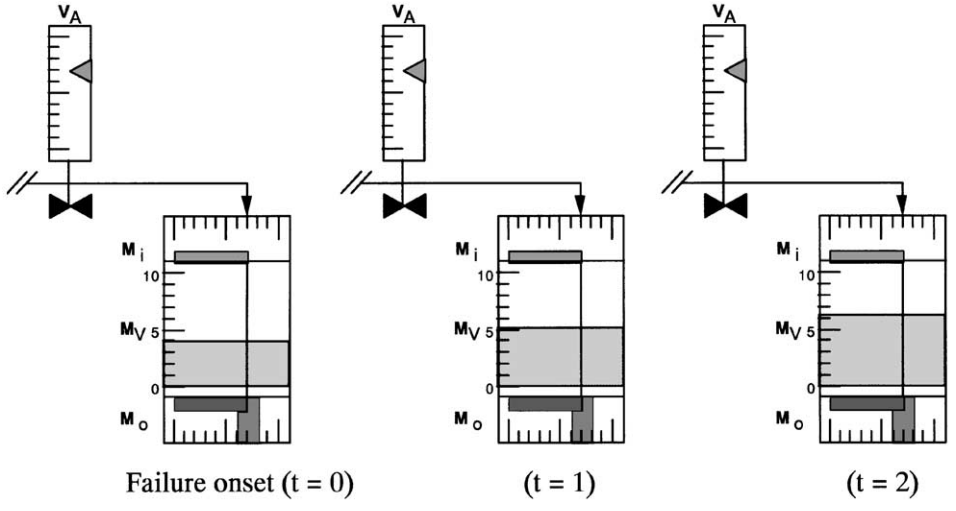

FIGURE 12. The effect a drifting level transducer at the vat on the configural display with a second flow transducer assigned to the input pipe. M_i no longer rises as vat level rises, because M_i is now derived from the directly sensed F_i.

location is taken to stand for the same variable at other locations. In the water reservoir example, topographic inadequacy is initially evident in the sensing of flow at Physical Object level (see Figures 8 and 10) where a flow transducer is placed on the output but not the input leg. When the level transducer drifts upwards, because there is no independent input flow transducer, it is assumed that input flow rate is increasing. Topographic inadequacy is usually seen in more complex systems than the simple water reservoir discussed here, and will be illustrated further in Part II (Reising & Sanderson, 2002).

Derivational inadequacy involves relying on impoverished sensor information when performing derivations for higher levels of abstraction. Specifically, the derivation uses topographically borrowed values or values that are themselves derived, rather than directly sensed values. This is seen in the water reservoir example when the mass input value is derived from an independent volume level sensor on the vat, rather than relying on a volume flow rate sensor on the input pipe.

In summary, the simple water reservoir example shows the ambiguity that results when topographical and derivational inadequacies affect information used in a display. Rather than making the deviation from proper functioning easy to see, the display makes the deviation hard to see. Additional information, here gained through an additional sensor, provides a better basis for detecting a sensor failure.

6. Conclusion

In this paper we have argued that Rasmussen's AH formalism (1979, 1986, 1998; see also Rasmussen *et al.*, 1994) can be extended and annotated to specify the instrumentation needed to supply information needed by human operators when handling unanticipated variability. EID advocates the use of the AH to identify information that should appear

in an EID-based interface if the interface is to provide analytic redundancy. However, the validity of such information is constrained by the availability of sensors and reliability of derivations. Moreover, the impact of invalid data might be amplified in displays that use configural graphics. Sensor-annotated AHs help analysts consider the information human operators must have to provide the kind of analytically redundant control needed to handle unanticipated variability. Because such information has implications for sensor and instrumentation design, the needs of the human controller must be considered early enough in system development to influence sensor and instrumentation specifications (Reising & Sanderson, 1996). In this way, sensor-annotated AHs may extend EID's formative approach to interface design (Sanderson, 1998; Vicente, 1999).

In the quote at the start of this paper, Kelley (1968) claimed that "What can be sensed forms a fundamental limiting feature of displays" (p. 90). Sensor-annotated AHs show that what *is* sensed (which is often limited by what *can be* sensed) sometimes leads to information inadequacy and misleading displays. Kelley also claimed that "The initial step in considering the design of the displays for a particular manual control system is to analyze the information the operator would really like to have and to consider the sensing instruments available to obtain it for him" (p. 91).

Our analysis goes beyond Kelley's suggestion in three ways. First, the logic behind EID is to show the human operator not just the information he or she would like to have, but information that a formal analysis shows he or she *needs* to have in order to have the kind of analytically redundant model of proper functioning that supports human operator adaptation in the face of unanticipated variability. Second, we argue that analysing the information the operator needs should be the initial step not just when designing displays, but also when drawing up the engineering specifications of the system. Third, the information displayed to the operator comes not just from sensors, as Kelley implies, but also from derivations performed on data coming from one or more sensors. Sensor-annotated AHs show very clearly how borrowing sensor information within system topography, and passing topographically inadequate information to higher-order derivations, can lead to problems of interpretation for the human operator because of derivational inadequacy. Sensor-annotated AHs may put cognitive engineers in a stronger position to indicate the instrumentation required to support the human operator when he or she is required to adapt to unanticipated variability.

In putting forward these ideas, we do not suggest that control engineers and instrumentation engineers have failed in any way. The extraordinary achievements of the 20th century in automatic control and intelligent fault diagnosis bear witness to the effectiveness of such efforts. Instead, we suggest that supporting the human operator with information based on all levels of the AH, and building reliable visual representations of such information, adds some new criteria to instrumentation engineering that might be considered alongside other criteria.

The water reservoir example in this paper shows that sensor-annotated AHs can be effective tools in analysing the effect of instrumentation on the availability of information at different levels of abstraction. Sensor-annotated AHs help to predict the impact of sensor failures on ecological interfaces, and thereby the ability of the operator to distinguish true system failures from sensor failures. Reising and Sanderson (2002) present a more detailed example of the same approach and they contrast the effects of different instrumentation on EID- and non-EID-based interfaces.

References

Anyakora, S. N. & Lees, F. P. (1974). Detection of instrument malfunction by the process operator. In E. Edwards & F. P. Lees, Eds. *The Human Operator in Process Control*, pp. 238–248. London: Taylor & Francis.

Bainbridge, L. (1991). Multiplexed VDT display systems: a framework for good practice. In G. R. S. Weir & J. L. Alty, Eds. *Human–Computer Interaction and Complex Systems*. New York: Academic Press.

Bainbridge, L. (1993). Ironies of automation. *Automatica*, **19,** 775–779.

Beltracchi, L. (1987). A direct manipulation interface for water-based Rankine Cycle heat engines. *IEEE Transactions on Systems, Man, and Cybernetics*, **17,** 478–487.

Beltracchi, L. (1998*a*). *Display Concept for Monitoring Process Variables, Functions and Cycles*. Unpublished manuscript.

Beltracchi, L. (1998*b*). Encoding a model-based display with dynamic data. Unpublished manuscript.

Beltracchi, L. (2000). Encoding a model-based display with dynamic data. *Proceedings of the IEA2000/HFES2000 Congress*. Vol 3, pp. 579–582. Santa Monica, CA, HFES.

Bennett, K. & Flach, J. M. (1992). Graphical displays: implications for divided attention, focused attention, and problem-solving. *Human Factors*, **34,** 923–935.

Bennett, K., Flach, J. M. & Nagy, A. (1997). Visual displays. In G. Salvendy, Ed. *Handbook of Human Factors and Ergonomics*. New York: Wiley.

Bennett, K. B., Toms, M. L. & Woods, D. D. (1993). Emergent features and graphical elements: designing more effective configural displays. *Human Factors*, **35,** 71–97.

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

Borer, J. (1991). *Microprocessors in Process Control*. New York: Elsevier Science Publishing.

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

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

Burns, C. M., Barsalou, E., Handler, C., Kuo, J. & Harrigan, K. (2000). A work domain analysis for network management. *Proceedings of the IEA2000/HFES2000 Congress*. Vol. 1, pp. 469–472. Santa Monica, CA: HFES.

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

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

Clark, R. N. & Campbell, B. (1982). Instrument fault detection in a pressurized water reactor pressurizer. *Nuclear Technology*, **56,** 23–32.

Dinadis, N. & Vicente, K. J. (1996). Ecological interface design for a power plant feedwater system. *IEEE Transactions on Nuclear Science*, **43,** 266–277.

Dinadis, N. & Vicente, K. J. (1999). Designing functional visualizations for aircraft system status displays. *International Journal of Aviation Psychology*, **9,** 241–270.

Dorr, R., Kratz, F., Ragot, J., Loisy, F. & Germain, L. (1997). Detection, isolation and identification of sensor faults in nuclear power plants. *IEEE Transactions on Control Systems Technology*, **5**, 42–60.

Deutsch, O. L., Ornedo, R. S. & Lindsay, R. W. (1983). Implementation of real-time signal validation at EBR-II. *Transactions of the American Nuclear Society*, **45,** 660–661.

Doyle, F. J., Garrison, A., Johnson, B. & Smith, W. D. (1998). *Process Measurement and Control: Industry Needs*. NSF/NIST Process Measurement and Control Workshop Report. National Institute of Standards and Technology, USA.

Dragoni, A. F., Giorgini, P. & Panti, M. (1998). Self-monitoring distributed monitoring systems for nuclear power plants. In J. Mira, A. P. del Pobil & M. Ali, Eds. *Methodology*

and Tools in Knowledge-Based Systems, Vol. 1. Lecture Notes in Computer Science, No. 1415. Berlin: Springer-Verlag.

Duncan, K. D. (1987). Fault diagnosis training for advanced continuous process installations. In J. Rasmussen, K. Duncan & J. Leplat, Eds. *New Technology and Human Error*, pp. 209–221. New York: John Wiley & Sons.

Duncan, K. D., Praetorius, N. & Milne, A. B. (1989). Flow displays of complex plant processes for fault diagnosis. In E. D. Megaw, Ed. *Contemporary Ergonomics*, pp. 199–206. London: Taylor & Francis.

Edlund, C. & Lewis, M. (1994). Comparing ecologically constrained and conventional displays in control of a simple steam plant. *Proceedings of the Human Factors and Ergonomics Society 38th Annual Meeting*, pp. 486–490. Santa Monica, CA: Human Factors and Ergonomics Society.

Frank, P. M. (1990). Fault diagnosis in dynamic systems using analytical and knowledge-based redundancy—a survey and some new results. *Automatica*, **25,** 459–474.

Gillie, T. & Berry, D. (1994). Object displays and control of dynamic systems. *Ergonomics*, **37,** 1885–1903.

Gotcher, R. & Burroni, R. J. (1993). Reactor coolant system resistance temperature detector cross comparison and possible sources of measurement error. *Instrumentation, Controls, and Automation in the Power Industry. Proceedings of the 38th Power Instrumentation Symposium*, Vol. 36, pp. 65–79. Research Triangle Park, NC: Instrument Society of America.

Griebenow, R. D., Hansen, E. J. & Sudduth, A. L. (1995). Applied pattern recognition for plant monitoring and data validation. *Instrumentation, Controls, and Automation in the Power Industry. Proceedings of the 38th Power Instrumentation Symposium*, Vol. 38, pp. 21–30. Research Triangle Park, NC: Instrument Society of America.

Hajdukiewicz, J. R., Doyle, J., Milgram, P., Vicente, K. J. & Burns, C. M. (1998). Modeling a medical environment: an onotology for integrated medical informatics design. *International Journal of Medical Informatics*, **62,** 79–99.

Hajdukiewicz, J. R., Vicente, K. J., Doyle, D. J., Milgram, P., & Burns, C. M. (2001). Modeling a medical environment: An onotology for integrated medical informatics design. *International Journal of Medical Informatics*, **62,** 79–99.

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

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

Hansen, J. P. (1995). Representation of system invariants by optical invariants in configural displays for process control. In P. Hancock, J. Flach, J. Caird & K. Vicente, Eds. *Local Applications of the Ecological Approach to Human–Machine Systems*, Vol. 2, pp. 208–233. Hillsdale, NJ: Lawrence Erlbaum Associates.

Hashemain, H. M., Riner, J. L., Bunch, C. E. & Petersen, K. M. (1993). RTD cross calibration in pressurized water reactors. *Instrumentation, Controls, and Automation in the Power Industry. Proceedings of the 36th Power Instrumentation Symposium*, Vol. 36, pp. 81–101. Research Triangle Park, NC: Instrument Society of America.

Hayter, D. (1996). *An Application of Ecological Interface Design: Boiler Level Control*. Unpublished manuscript. Cognitive Engineering Laboratory, Department of Mechanical and Industrial Engineering, University of Toronto, Canada.

Himmelblau, D. & Bhalodia, M. (1995). On line sensor validation of single sensor using artificial neural networks. *Proceedings of the 1995 American Control Conference*, Vol. 1, pp. 766–770.

Holbert, K. E. & Upadhyaya, B. R. (1990). An integrated signal validation system for nuclear power plants. *Nuclear Technology*, **92,** 411–427.

Hollnagel, E. (1993). *Human Reliability Analysis: Context and Control*. London: Academic Press.

Hollnagel, E. (1998). *Cognitive Reliability and Error Analysis Method — CREAM*. Oxford: Elsevier Science.

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

Jamieson, G. A. (1998). *Ecological Interface Design for Petrochemical Processing Application*. Unpublished Master of Applied Science thesis. University of Toronto, Toronto, Canada.

Johannsen, G. (1992). Towards a new quality of automation in complex man-machine systems. *Automatica*, **28,** 355–373.

Johnson, C. (1993). *Process Control Instrumentation Technology* (4th edn). New York: John Wiley & Sons.

Jovic, F. (1992). *Process Control Systems: Principles of Design, Operation, and Interfacing* (2nd edn). London: Chapman & Hall.

Kelley, C. R. (1968). *Manual and Automatic Control*. New York: Wiley.

Lee, S. C. (1994). Sensor value validation based on systematic exploration of the sensor redundancy for fault diagnosis KBS. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-24, 594–605.

Lee, J. D., & Moray, N. (1992). Trust, control strategies and allocation of function in human–machine systems. *Ergonomics*, **35,** 1243–1270.

Lee, J. D. & Moray, N. (1994). Trust, self-confidence, and operators' adaptation to automation. *International Journal of Human–Computer Studies*, **40,** 153–184.

Lee, J. D., Kinghorn, R. A. & Sanquist, T. F. (1995). *Review of Ecological Interface Design Research: Applications of the Design Philosophy and Results of Empirical Evaluations*. Battelle Research Center Technical Report, Seattle, WA, April.

Lind, M. (1981). The use of flow models for automated plant diagnosis. In J. Rasmussen & W. Rouse, Eds. *Human Detection and Diagnosis of System Failures*, pp. 411–432. New York: Plenum.

Lindsay, R. W. (1990). A display to support knowledge based behavior. *Proceedings of the ANS Topical Meeting on Advances in Human Factors Research on Man–Computer Interactions: Nuclear and Beyond*, pp. 266–270. LaGrange Park, IL: ANS.

Liu, Q., Nakata, K., & Furuta, K. (2002). Display design of process systems based on functional modeling. *Cognition, Technology, & Work*, **4**, 48–63.

Maddox, M. (1996). Critique of "A longitudinal study of the effects of ecological interface design on skill acquisition" by Christoffersen, Hunter, and Vicente. *Human Factors*, **38,** 542–545.

Massoumnia, M. A. (1986). A geometric approach to the synthesis of failure detection filters. *IEEE Transactions on Automatic Control*, AC-9, 839–846.

Molloy, R. & Parasuraman, R. (1996). Monitoring an automated system for a single failure: vigilance and task complexity effects. *Human Factors*, **38,** 311–322.

Moray, N., Lee, J., Vicente, K. J., Jones, B. G. & Rasmussen, J. (1994). A direct perception interface for nuclear power plants. *Proceedings of the Human Factors and Ergonomics Society 38th Annual Meeting*, pp. 481–485. Santa Monica, CA: Human Factors and Ergonomics Society.

Moray, N. & Rotenberg, I. (1989). Fault management in process control: eye movements and action. *Ergonomics*, **32,** 1319–1342.

Morris, N. M. & Rouse, W. B. (1985). The effects of type of knowledge upon human problem solving in a process control task. *IEEE Transactions on Systems, Man, and Cybernetics*, **15,** 698–707.

Olsson, G. & Lee, P. L. (1994). Effective interfaces for process operators—a prototype. *The Journal of Process Control*, **4,** 99–107.

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

Rasmussen, J. (1979). *On the Structure of Knowledge—A Morphology of Mental Models in a Man–Machine System Context*, Risø-M-2192, Risø National Laboratory Electronics Department, Roskilde, Denmark.

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

Rasmussen, J. (1988). A cognitive engineering approach to the modeling of decision making and its organization in: Process control, emergency management, CAD/CAM, office systems,

and library systems. In W. B. Rouse, (Ed.), *Advances in man-machine systems research* (pp. 165–243). Greenwich, CT: JAI Press.

Rasmussen, J. (1998). *Ecological Interface Design for Complex Systems: An Example: SEAD—UAV Systems*. Final Report, EOARD-Contract: F61708-97-W0211, HURECON, Denmark, April.

Rasmussen, J. (1999). Ecological interface sign for reliable human–machine systems. *International Journal of Aviation Psychology*, **9,** 203–223.

Rasmussen, J. & Rouse, W. B. (Eds.). (1981). *Human Detection and Diagnosis of System Failures*. New York: Plenum Press.

Rasmussen, J., Pejtersen, A. M. & Goodstein, L. P. (1994). *Cognitive Systems Engineering*. New York: John Wiley & Sons.

Rasmussen, J. & Vicente, K. J. (1989). Coping with human errors through system design: implications for ecological interface design. *International Journal of Man–Machine Studies*, **31,** 517–534.

Reed, J. (1992). A plant local panel review. In B. Kirwan & L. K. Ainsworth, Eds. *A Guide to Task Analysis*, pp. 267–288. London: Taylor & Francis.

Reising, D. C. (1999). *The Impact of Instrumentation Location and Reliability on the Performance of Operators Using an Ecological Interface for Process Control*. Unpublished Doctoral Dissertation. Department of Mechanical and Industrial Engineering, University of Illinois at Urbana-Champaign, IL.

Reising, D. C. & Sanderson, P. M. (1996). Work domain analysis of a pasteurization plant: building an abstraction hierarchy representation. *Proceedings of the 40th Annual Meeting of the Human Factors and Ergonomics Society*, pp. 293–297. Santa Monica, CA: Human Factors and Ergonomics Society.

Reising, D. C., & Sanderson, P. M. (1998). Designing displays under Ecological Interface Design: Towards operationalizing semantic mapping. In *Proceedings of the 42nd Annual Meeting of the Human Factors and Ergonomics Society* (pp. 372–376). Santa Monica, CA: Human Factors and Ergonomics Society.

Reising, D. C. & Sanderson, P. (2000*a*). Testing the impact of instrumentation location and reliability on ecological interface design: control performance. *Proceedings of the Joint Meeting of the Human Factors and Ergonomics Society and the International Ergonomics Association (IEA2000/HFES2000)*, Vol. 1, pp. 124–127. Santa Monica, CA: HFES.

Reising, D. C. & Sanderson, P. (2000*b*). Testing the impact of instrumentation location and reliability on ecological interface design: fault diagnosis performance. *Proceedings of the Joint Meeting of the Human Factors and Ergonomics Society and the International Ergonomics Association (IEA2000/HFES2000)*, Vol. 3, pp. 591–594. Santa Monica, CA: HFES.

Reising, D. C. & Sanderson, P. (2002). Work domain analysis and sensors II: pasteurizer II case study. *International Journal of Human–Computer Interaction*, X-ref: S1071-5819(02)00048-4.

Reising, D. C. & Sanderson, P. (in press). Ecological interface design for pasteurizer II: a process description of semantic mapping. *Human Factors*, 44.

Reising, D. C., Sanderson, P. M., Jones, B. G., Moray, N. & Rasmussen, J. (1998). A direct perception display for rule-based behavior: supporting power plant startup with a "lattice" display. *Proceedings of the 42nd Annual Meeting of the Human Factors and Ergonomics Society*, pp. 224–228. Santa Monica, CA: Human Factors and Ergonomics Society.

Sanderson, P. M. (1998). Cognitive work analysis and the analysis, design, and evaluation of human-computer interactive systems. *Proceedings of the Australian/New Zealand Conference on Computer–Human Interaction (OzCHI98)*, pp. 220–227. Los Alamitos: IEEE Computer Society.

Seto, D., Krogh, B., Sha, L. & Chutinan, A. (1998). Dynamic control system upgrade using the Simplex architecture. *Proceedings of the American Control Conference*, pp. 3504–3508, Philadelphia, PA, 24–26 June.

Sha, L., Rajkumar, R. & Gagliardi, M. (1996). Evolving dependable real-time systems. *Proceedings of the 1999 IEEE Aerospace Applications Conference*, pp. 335–346. Aspen, CO: IEEE New York.

Sharp, T. D. & Helmicki, A. J. (1998). The application of the ecological interface design approach to neonatal intensive care medicine. *Proceedings of the 42nd Annual Meeting of the Human Factors and Ergonomics Society*, pp. 350–354. Santa Monica, CA: Human Factors and Ergonomics Society.

Stubler, W. F. & O'Hara, J. M. (1996). Human factors challenges for advanced process control. *Proceedings of the Human Factors and Ergonomics Society 40th Annual Meeting*, pp. 992–996. Santa Monica, CA: Human Factors and Ergonomics Society.

Sundstrom, G. (1993). User modelling for graphical design in complex dynamic environments: concepts and prototype implementations. *International Journal of Man–Machine Studies*, **38,** 567–586.

Sundstrom, G. (1997). Designing support contexts: helping operators to generate and use knowledge. *Control Engineering Practice*, **5,** 375–381.

Vicente, K. J. (1991). *Supporting Knowledge-Based Behavior Through Ecological Interface Design*. Unpublished Doctoral Thesis, University of Illinois, Urbana, IL.

Vicente, K. J. (1996). Improving dynamic decision making in complex systems through EID: A research overview. *System Dynamics Review*, **12,** 251–279.

Vicente, K. J. (1999). *Cognitive Work Analysis: Towards Safe, Productive, and Healthy Computer-based Work*. Mahwah, NJ: Lawrence Erlbaum Associates.

Vicente, K. J. (2002). Ecological interface design: progress and challenges. *Human Factors*, **44,** 62–78.

Vicente, K. J. & Burns, C. M. (1995). *A Field Study of Operator Cognitive Monitoring at Pickering Nuclear Generating Station-B.* Technical Report CEL 95-04, Cognitive Engineering Laboratory, University of Toronto, Toronto.

Vicente, K., Christoffersen, K. & Pereklita, A. (1995). Supporting operator problem solving through ecological interface design. *IEEE Transactions on Systems, Man, and Cybernetics*, **25,** 229–244.

Vicente, K. J., Moray, N., Lee, J. D., Rasmussen, J., Jones, B. G., Brock, R. & Djemil, T. (1996). Evaluation of a Rankine cycle display for nuclear power plant monitoring and diagnosis. *Human Factors*, **38,** 506–521.

Vicente, K. J. & Rasmussen, J. (1990). The ecology of human–machine systems II: mediating "direct perception" in complex work domains. *Ecological Psychology*, **2,** 207–249.

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

Wells, (1997). *Major Hazards and Their Management*. Rugby, UK: Institute of Chemical Engineers.

Woods, D. D. & Roth, E. M. (1988). Aiding human performance II: from cognitive analysis to support systems. *Le Travail Human*, **51,** 139–172.

Woods, D. D. (1991). The cognitive engineering of problem representations. In G. R. S. Weir & J. L. Alty, Eds. *Human–Computer Interaction and Complex Systems*, pp. 169–188. London, UK: Academic Press.

Yamaguchi, Y. & Tanabe, F. (2000). Creation of interface system for nuclear reactor operation: practical implication of implementing EID concept on large complex system. *Proceedings of the Joint Meeting of the Human Factors and Ergonomics Society and the International Ergonomics Association (IEA2000/HFES2000)*, Vol. 1, pp. 571–574. Santa Monica, CA: HFES.

Yu, X., Lau, E., Vicente, K. J. & Carter, M. W. (1998). Advancing performance measurement in cognitive engineering: the abstraction hierarchy as a framework for dynamical systems analysis. *Proceedings of the Human Factors and Ergonomics Society 42nd Annual Meeting*. pp. 359–363.

Paper accepted for publication by Associate Editor, G. Sundstrom.