Cognitive Work Analysis: Implications for Microworld Research on Human-Machine Interaction

Kim J. Vicente
Cognitive Engineering Laboratory
Department of Industrial Engineering
University of Toronto
Toronto, Ontario M5S 1A4
CANADA
Email: benfica@ie.utoronto.ca

ABSTRACT

Cognitive work analysis (CWA) has been developed to inform the design of complex human-machine systems. However, it can also be used to guide microworld research in a number of productive ways. This paper provides examples of how CWA was used to inform the design of experiments, the analysis of data, and the classification of performance measures in microworld research. In each case, CWA led to substantial benefits that may not have otherwise been enjoyed. Therefore, CWA is a very important tool that can lead to important insights that address several of the challenges associated with microworld research.

I. INTRODUCTION

Conducting cognitive engineering research is an extremely challenging endeavor. On the one hand, there is a strong need for research results to generalize to operational settings, otherwise the research is of little practical use. This requirement implies that experiments need to be conducted under conditions that are representative of actual work domains. On the other hand, there is also a strong need for defensible research results that can lead to a principled understanding of the factors that affect human performance in complex systems. This requirement implies that experiments need to be conducted under controlled conditions. To maximize representativeness, one could conduct studies with full-scope simulators, but experience has shown that it is very difficult to obtain defensible and statistically reliable results under such conditions, due to lack of experimental control. To maximize control, one can greatly simplify the phenomenon of interest, but this can lead to results which are not relevant to industrial systems. Some cognitive engineering researchers have suggested microworlds -- small-scale simulations of systems that are intended to be representative of industrial-scale complex human machine systems -- as a useful alternative to complement field studies and highly controlled laboratory studies. Their rationale is that microworlds are research vehicles that can lead to defensible results that are also potentially generalizable, thereby generating useful and informative findings.

Microworld research has been around for many years although it has only recently been labeled as such. During this time, researchers have commented that there are many methodological difficulties associated with such research (e.g., [2]), since many traditional laboratory methods cannot be meaningfully applied in such rich contexts. Other researchers have suggested that an a priori understanding of the system can help to overcome some of these problems (e.g., [5]). While some progress has been made in this area, there is no consensus as to how to conduct such an analysis and, perhaps more importantly, there are very few examples of how such an analysis can be, or has been, effectively used in experimentation.

The hypothesis of this paper is that a great deal can be gained by adopting cognitive work analysis (CWA) as a framework for conducting an a priori analysis of microworlds. Although CWA was originally proposed as a tool to inform the design of complex human-machine systems, it also has important implications for the conduct of research. While there is no widely accepted framework for conducting a CWA, we have found the taxonomy developed by Rasmussen [4] to be particularly useful due to its comprehensiveness and coherence. The remainder of this paper begins by providing a brief description of Rasmussen's framework for CWA. The paper then describes how CWA has served as an invaluable tool in conducting microworld research in our laboratory.

II. COGNITIVE WORK ANALYSIS

As mentioned, our research has been guided by Rasmussen's [4] taxonomy for CWA. The goal of this taxonomy is to identify the goal-relevant constraints that can shape the behavior of human operators in complex systems. Each type of behavior-shaping constraint provides a distinct layer of analysis. Ignoring organizational issues, the primary levels in this framework are: developing a representation of the

0-7803-2559-1/95 $4.00 © 1995 IEEE

functional structure of the work domain; identifying the decision activities associated with the different control tasks required in the domain; analyzing the various mental strategies and heuristics that can be used to perform each of the decision activities already identified; and identifying the competencies required of human operators. Several conceptual tools have been identified to deal with each of these levels of analysis, including: the abstraction hierarchy, the decision ladder, and the skills, rules, knowledge framework [4]. This framework for CWA has been applied in industry, most notably in the design of advanced control room designs for nuclear power plants (see [6]), but it has rarely been applied to microworld research in a comprehensive manner.

We have applied this framework to analyze the microworld shown in Figure 1. This microworld, known as DURESS (DUal REservoir System Simulation) II, is a thermal-hydraulic process simulation that is intended to be representative of industrial systems, albeit on a much smaller scale. The system consists of two redundant feedwater streams (FWSs) that can be configured to supply water to either, both, or neither of the two reservoirs. Each reservoir has associated with it an externally determined demand for water that can change over time. The system purposes are twofold: to keep each of the reservoirs at a prescribed temperature (40° C and 20° C), and to satisfy the current mass (water) output demand rates. To accomplish these goals, the subject has control over eight valves (VA, VA1, VA2, VO1, VB, VB1, VB2, and VO2), two pumps (PA and PB), and two heaters (HTR1 and HTR2). All of these components are governed by first order lag dynamics, with a time constant of 15 s for the heaters and 5 s for the remaining components. The system input temperature (T0), reservoir output temperatures (T1 and T2), and the volumes for both reservoirs (V1 and V2) are also displayed in Figure 1.

A comprehensive CWA of DURESS II can be found in Vicente and Pawlak [7]. This analysis consists of 4 phases: an abstraction hierarchy representation of the functional structure of the system; a decision ladder analysis identifying the decision tasks required for the various operating regimes, as well as the information needs associated with each of those decision tasks; an analysis of the mental strategies and heuristics that can be used for each of the significant decision tasks identified in the previous phase; and finally, an analysis, based on the skills, rules, knowledge taxonomy, of the operator competencies required for effective performance in the domain.

In the remainder of this paper, the role that this CWA has played in guiding microworld research with DURESS II will be described. Three applications of CWA will discussed:

1. design of experiments
2. analysis of data
3. classification of performance measures

Examples of how CWA can be applied to each of these issues will be provided, based on recent research conducted in our laboratory.

III. DESIGN OF EXPERIMENTS

Typical Issues

One of the challenges that arises in the conduct of microworld research is how compare performance on different trials in a meaningful way. Because of the richness of microworlds, task demands are broad in scope compared to those used in traditional experiments. As a result, one often finds that there can be a significant degree of variability in the difficulty associated with different trials in a microworld. This presents a problem for comparison of performance across trials.

Another issue that arises frequently in microworld research is the need to gradually increase the complexity of trials during the practice phase of an experiment. Microworld tasks can be very complex and, in fact, overwhelming to a novice. Consequently, it is important to present subjects with simpler trials first so that they can gradually become familiar with the domain, and only then introduce trials of greater complexity.

To deal with these typical issues, it would be useful to have an objective way of identifying the complexity of trials in a particular microworld.

Benefits of CWA

The CWA of DURESS II conducted by Vicente and Pawlak [7] led to a solution to these two problems. As part of the strategy analysis, a taxonomy of trials based on the range of feasible strategies was identified.

Referring to Figure 1, one of the important structural properties of DURESS II is that each FWS can only supply water at a maximum rate of 10 units/s of flow. This means that different demand setpoints (D1, D2) can be satisfied using three different configurations of the FWSs. First, if the sum of the output flow demands for the two reservoirs does not exceed 10 units/s (i.e., [D1+D2] < 10), then a single FWS can be used to feed both reservoirs. This simplifies the task since a smaller number of components can be used to control the system. Second, if the sum of the demands exceeds 10 units/s but neither of the demands exceeds 10 units/s, then a decoupling strategy can be adopted. That is, one can use one FWS to feed one reservoir and the other FWS to feed the other reservoir, keeping the crossover valves VA2 and VB1 closed. This makes control of the system simpler because there are no interactions between FWS, thereby eliminating the possibility that manipulating a control valve will have an undesired effect on the other reservoir. Third and finally, if one of the output flow demand setpoints is greater than 10 units/s, then that reservoir cannot be supplied by just one FWS. As a result, a many-to-many strategy must be adopted by opening at least one of the crossover valves (VA2 or VB1). This makes control of the system more complex

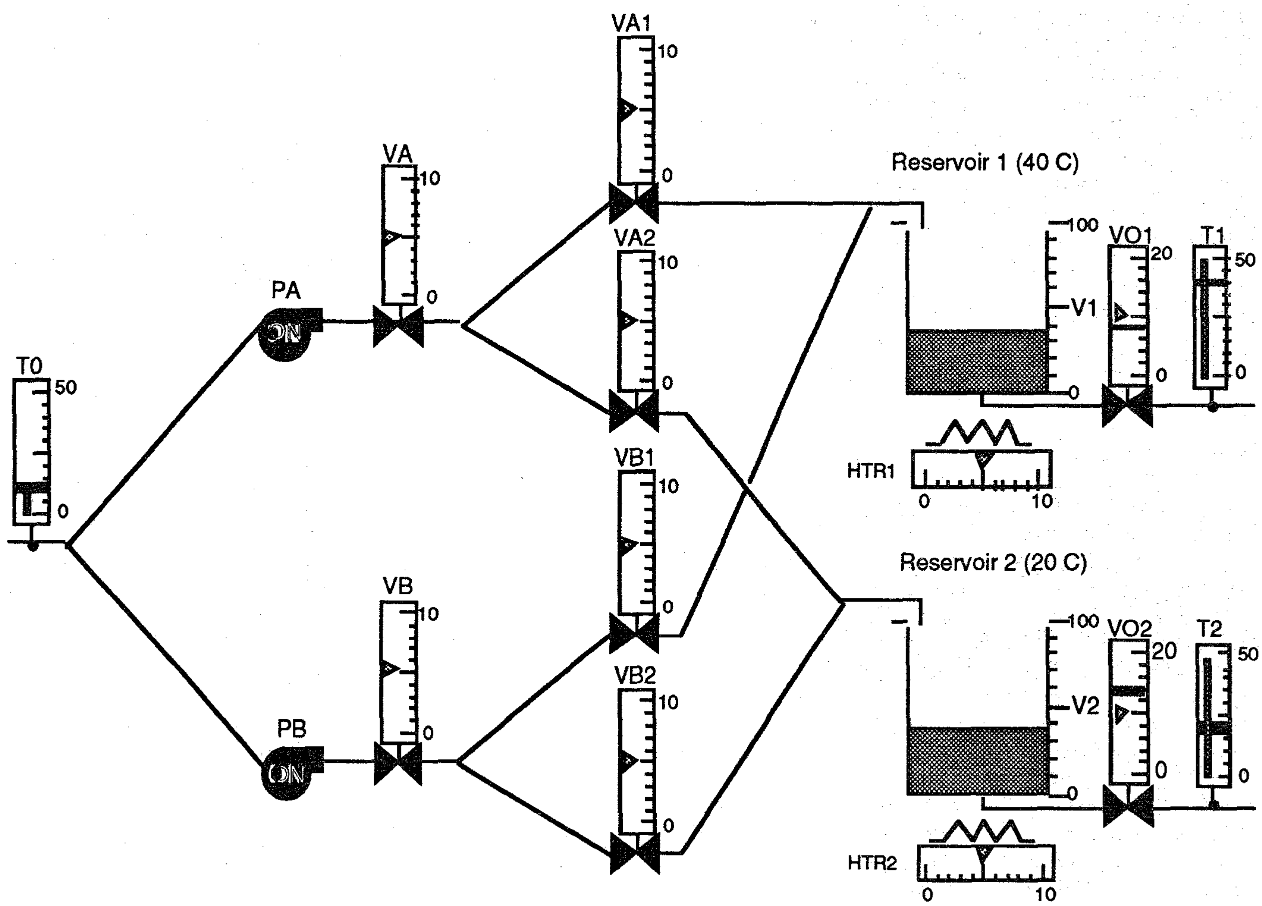

FIGURE 1. Diagram of physical structure of DURESS II microworld.

because adjusting the flow to one reservoir can affect the flow to the other because of the many-to-many mapping between FWSs and reservoirs.

This taxonomy shows the relationship between task demands (i.e., demand setpoints) and feasible strategies (and therefore, task complexity). This insight provided by CWA has been used in several experiments on DURESS II to address the issues mentioned at the beginning of this section. When it was important to compare performance across trials, an effort was made to match those trials by having the demands fall in the same category of taxonomy just described, thereby equating the complexity of the two trials in question. Similarly, when it was important to gradually increase the complexity of trials during practice, demand pairs that could be satisfied by simpler strategies were used for the initial practice trials. After subjects became familiar with the system, subsequent trials were introduced that were more complex in that they could not be solved using simplifying strategies.

Therefore, in both instances, the application of CWA to experimental design of microworld research provides the important benefit of giving the experimenter better control over the stimulus.

IV. ANALYSIS OF DATA

Typical Issues

Microworld research also leads to difficult challenges in terms of data analysis. For example, because of the richness of the tasks and the level of practice which subjects typically receive in such studies, coarse measures of performance rarely reveal differences between treatment groups. Also, because so many data are typically generated by microworld experiments, exhaustive analysis of the data is typically unrealistic because of the inordinate time involved. Both of these issues create a need for focusing the data analysis effort in a direction that is economic yet promises to be productive in terms of revealing differences between subject groups.

Benefits of CWA

The portion of the CWA of DURESS II that was directed at identifying the decision tasks required for effective system control led to an important insight for data analysis. The analysis revealed that the sensitivity of the heaters shown in Figure 1 depends on both the demand and volume. The former is a goal setpoint, and so subjects have no choice but to satisfy it if they are to perform the task. However, the choice of volume is a discretionary one. The CWA indicates that if subjects decide to stabilize the system at a low volume, then it will be more difficult to stabilize the temperature within the goal region because the temperature will be very sensitive to changes in heater setting.

This insight was applied to analyze the data from an experiment comparing the performance of two interfaces for DURESS II [3]. One of these interfaces represented the relationships between system more thoroughly than the other, and it was hypothesized that this would lead to improved performance. However, an analysis of typical measures of performance (e.g., task completion time) showed no difference between interfaces on normal (i.e., non-fault trials). The CWA suggested the hypothesis that the two interface groups might differ in terms of their steady state volumes, the rationale being that subjects with the interface showing system relationships should realize that low volumes should be avoided. An analysis of the steady state volumes for the two interface groups confirmed this hypothesis. Furthermore, it was also found that steady state volume was significantly correlated with the number of heater actions for low demand. This result was also motivated by the CWA, since one would expect that low volumes would lead to difficulties in stabilizing temperature, and therefore, an increased number of heater actions.

It is extremely unlikely that the experimenters would have had the idea of conducting these analyses were it note for the CWA conducted before the experiment. Therefore, one of the important benefits of CWA for microworld research is that it can lead to the identification of sensitive data analyses that can show differences between treatment groups that would otherwise not be observed.

V. CLASSIFICATION OF PERFORMANCE MEASURES

Typical Issues

Because there are typically many more degrees of freedom in microworld research than traditional research, there are many different types of performance measures that can be adopted. The question then becomes how does one measure performance in a study with a specific set of goals? Furthermore, if one decides to adopt a number of different measures, how can one organize the performance data in a meaningful and concise way?

Benefits of CWA

CWA can provide some help with this aspect of microworld research as well. In fact, each layer of constraint in Rasmussen's [4] framework for CWA defines a distinct class of performance measures. The first layer, a functional decomposition of the work domain, defines variables that identify the state of the work domain (e.g., state variables). The second layer, the decision activities for each domain task, defines product measures of performance that describe what subjects do (e.g., task completion time, number of errors, number of control actions). The third layer of constraints addressed by CWA, the strategies that can be used for each decision activity, defines process measures that describe how subjects do what they do (e.g., verbal protocols, action transition graphs, state-space diagrams). Finally, the fourth level, operator competencies, defines knowledge elicitation measures which describe what subjects know (e.g., categorization tasks, control recipes, transfer of training).

This categorization of performance measures was used to identify a set of measures that were used in a longitudinal experiment with DURESS II [1]. This led to several benefits. First, it allowed us to ensure that we had collected the performance measures that were required to answer all the questions of interest. For example, sometimes experimenters make claims about subjects' mental models without having any performance measures at the level of operator competencies. Such claims are simply indefensible. Second, it allowed us to associate a given set of measures with each specific experimental question of interest (e.g., knowledge elicitation tests with the question of what impact interface design had on operator competencies). This made the data analysis effort much more manageable, since the numerous performance measures could be segregated into meaningful groups.

Therefore, CWA also has important benefits for the classification of performance measures in microworld research. More specifically, it can help experimenters generate a comprehensive and coherent performance measurement plan that can, in turn, lead to deep insights and defensible inferences.

VI. CONCLUSIONS

The preliminary findings we have obtained to date suggest that CWA has a great deal of promise. If used widely in research on human-machine systems, it may lead to research findings that are both defensible and generalizable. This would represent an invaluable asset in the design of industrial complex systems, where because of the potential hazard involved, trial and error approaches to design are simply not acceptable.

CWA also has another benefit that has not been addressed in this paper. It can provide a common frame of reference for comparing different microworlds. A number of researchers have tackled identical research questions with a number of very different microworlds. However, several critical questions remain unanswered: How can one compare the demands placed by these systems in a psychologically relevant manner? How can one generalize and/or integrate research results across these various microworlds? To accumulate knowledge, researchers need to speak a common language. The framework for CWA developed by Rasmussen [4] is one such proposal. This paper has tried to outline the benefits of this approach so that other researchers may adopt it as well.

VII. ACKNOWLEDGEMENTS

This research has been sponsored by equipment and research grants from the Natural Sciences and Engineering Research Council of Canada, and by research contracts with the Japan Atomic Energy Research Institute (Dr. Fumiya Tanabe, contract monitor) and AECL Research (Jean-Yves Fiset, contract monitor).

VIII. REFERENCES

[1] K. Christoffersen, C. N. Hunter, and K. J. Vicente, "Cognitive 'Dipsticks': Knowledge Elicitation Techniques for Cognitive Engineering Research", (CEL 94-01), Toronto: University of Toronto, Cognitive Engineering Laboratory, 1994.

[2] N. Moray, P. Lootsteen, and J. Pajak, "Acquisition of Process Control Skills", *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. SMC-16, 1986, pp. 497-504.

[3] W. S. Pawlak, and K. J. Vicente, "Inducing Effective Control Through Ecological Interface Design". Manuscript submitted for publication, 1994.

[4] J. Rasmussen, *Information Processing and Human-Machine Interaction: An Approach to Cognitive Engineering*. North Holland: New York, 1986.

[5] P. M., Sanderson, A. G. Verhage, and R. B. Fuld, "State-Space and Verbal Protocol Methods for Studying the Human Operator in Process Control", *Ergonomics*, Vol. 32, 1989, pp. 1343-1372.

[6] K. J. Vicente, "Multilevel Interfaces for Power Plant Control Rooms I: An Integrative Review", *Nuclear Safety*, Vol. 33, 1992, pp. 381-397.

[7] K. J. Vicente, and W. S. Pawlak, "Cognitive Work Analysis for the DURESS II System", (CEL 94-03), Toronto: University of Toronto, Cognitive Engineering Laboratory,1994.