Cognitive Work Analysis for the DURESS II System

Kim J. Vicente and William S. Pawlak

CEL 94-03

Cognitive Engineering Laboratory University of Toronto Department of Industrial Engineering
4 Taddle Creek Rd. Toronto, Ontario, Canada M5S1A4
phone: (416) 978-0881 email: benfica@ie.utoronto.ca fax: (416) 978-3453

PREFACE

This technical report descibes a cognitive work analysis of the DURESS II system. It is intended to be a resource for all research projects with this system, since it can be helpful for the design of experiments, selection of dependent variables, data analysis, and to provide a general context with which to understand subjects' behaviour. The report only represents our current understanding, and is not intended to be a definitive analysis. It is expected that the report will be expanded and revised as more research is conducted on DURESS II. This version is dated 20/5/94 and is based on analyses conducted by Vicente (1987), Vicente and Rasmussen (1990), Pawlak and Burns (1993), and Bisantz and Vicente (1994).

TABLE OF CONTENTS

LIST OF TABLES AND FIGURES

1. INTRODUCTION

This report will describe the CWA that was performed on the DURESS II system. Most of this analysis was completed before the beginning of the experiment. Additional information was obtained during data analysis and was therefore also included. This CWA proved to be invaluable to the entire project as it not only aided in designing the P+F interface but, more importantly, it helped define both the experimental conditions and the analysis measures used. A thorough CWA can be used as a constant reference source for a particular system. If this analysis is performed correctly, the resulting work defines exactly what can and cannot be performed with or on that particular system. DURESS II is constrained by certain physical laws and functional relationships that govern the possibilities of system states and operator control. The CWA revealed the various layers of constraints and relationships that govern how the system functions, what actions are possible and/or meaningful, and what measurement methods would be useful for analysis.

This particular CWA is based on the four phase methodology proposed by Rasmussen (1986). The first section deals with the work domain representation, which identifies system constraints on operators' behaviour. The section on operator control tasks is used to identify which tasks/decision activities need to be performed in each of four phases of the system operation (start-up, normal operation, fault detection and compensation, and shut-down). The mental strategy analysis is a discussion of the different methods by which an operator can perform the previously identified tasks/decision activities. Finally, a description of the levels of cognitive control used in controlling the system, indicating which types of knowledge, skills, and mental competencies are involved in performing system tasks, is presented. This particular CWA is based on previous analyses conducted by Vicente (1987), Vicente and Rasmussen (1990), Pawlak and Burns (1993), and Bisantz and Vicente (1994).

2. WORK DOMAIN REPRESENTATION

2.1. The Abstraction Hierarchy

The abstraction hierarchy representation of a work domain is used to map the field in which operators perform their actions. Thus, the abstraction hierarchy does not attempt to define what needs to be done, but rather, it defines the constraints of the work domain

that will limit what actions an operator can and cannot perform. If all the possible functional relationships within a work domain are represented in the abstraction hierarchy, then any action the operator performs on the system within that domain can be traced through that hierarchy. It is important to realize that the abstraction hierarchy is intended to represent the set of goal-relevant constraints governing the operation of the controlled system. As a result, it does not contain representations of any specific system events or operator tasks. An event-dependent representation of a work domain cannot, by definition, provide support for unanticipated events (Vicente and Tanabe, 1993). Hence, an event-independent representation is needed. The abstraction hierarchy discussion presented in this section is adopted from Bisantz and Vicente (1994).

2.2. The Abstraction Hierarchy for DURESS II

Figure 1 shows the five levels of abstraction that have been found to be useful for describing technical systems (Rasmussen, 1985). These levels are as follows: Functional Purpose, Abstract Function, Generalized Function, Physical Function, and Physical Form. This figure presents this hierarchy as well as a conceptually orthogonal part-whole decomposition of the DURESS II system. This part-whole hierarchy allows reasoning through different levels of system decomposition in addition to the different levels of abstraction. Note that not all of the 'boxes' in Figure 1 are being used to describe the system. Results from other studies have shown that operators, at higher levels of abstraction, think of a system at a coarser level of description (Vicente, 1992). Conversely, at lower levels of abstraction, operators think about the system at finer levels of decomposition. Therefore, only the representations identified in Figure 1 have been developed for DURESS II. Figures 2, 3, and 4 show the means-end, part-whole, and topological links between objects in the system representation, respectively.

Beginning with the part-whole dimension, three levels of resolution were selected: component, subsystem, and system. The objects at the component level of decomposition are the pumps, valves, heaters, and reservoirs. At the next level, these components are aggregated into meaningful subsystems. Thus, the objects are now transport subsystems, storage subsystems, and heating subsystems. Finally, at the system level, the entire system is described as a single whole. Part-whole links are shown in Figure 2.

The abstraction hierarchy, which is orthogonal to the part-whole dimension, consists of the five previously defined levels of description. These are shown in Figure 3, and are described below.

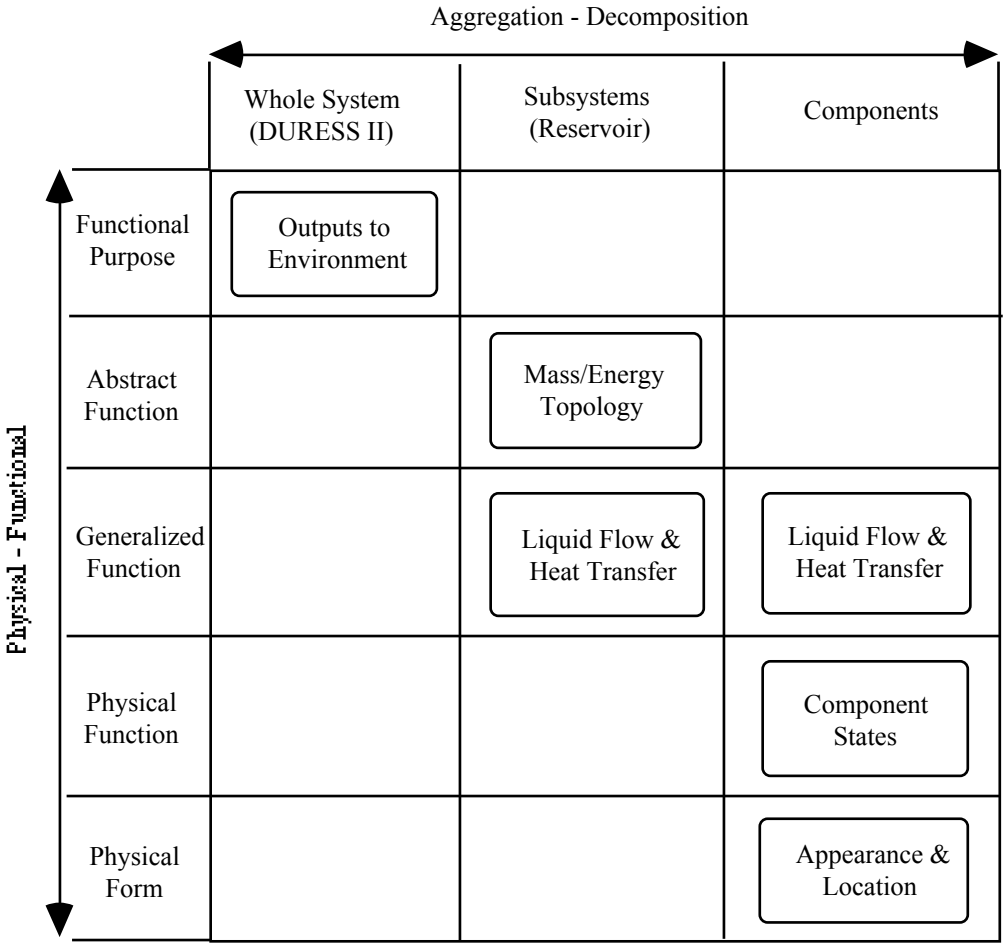

Figure 1. Means-end/part-whole space for DURESS II (Vicente, 1991)

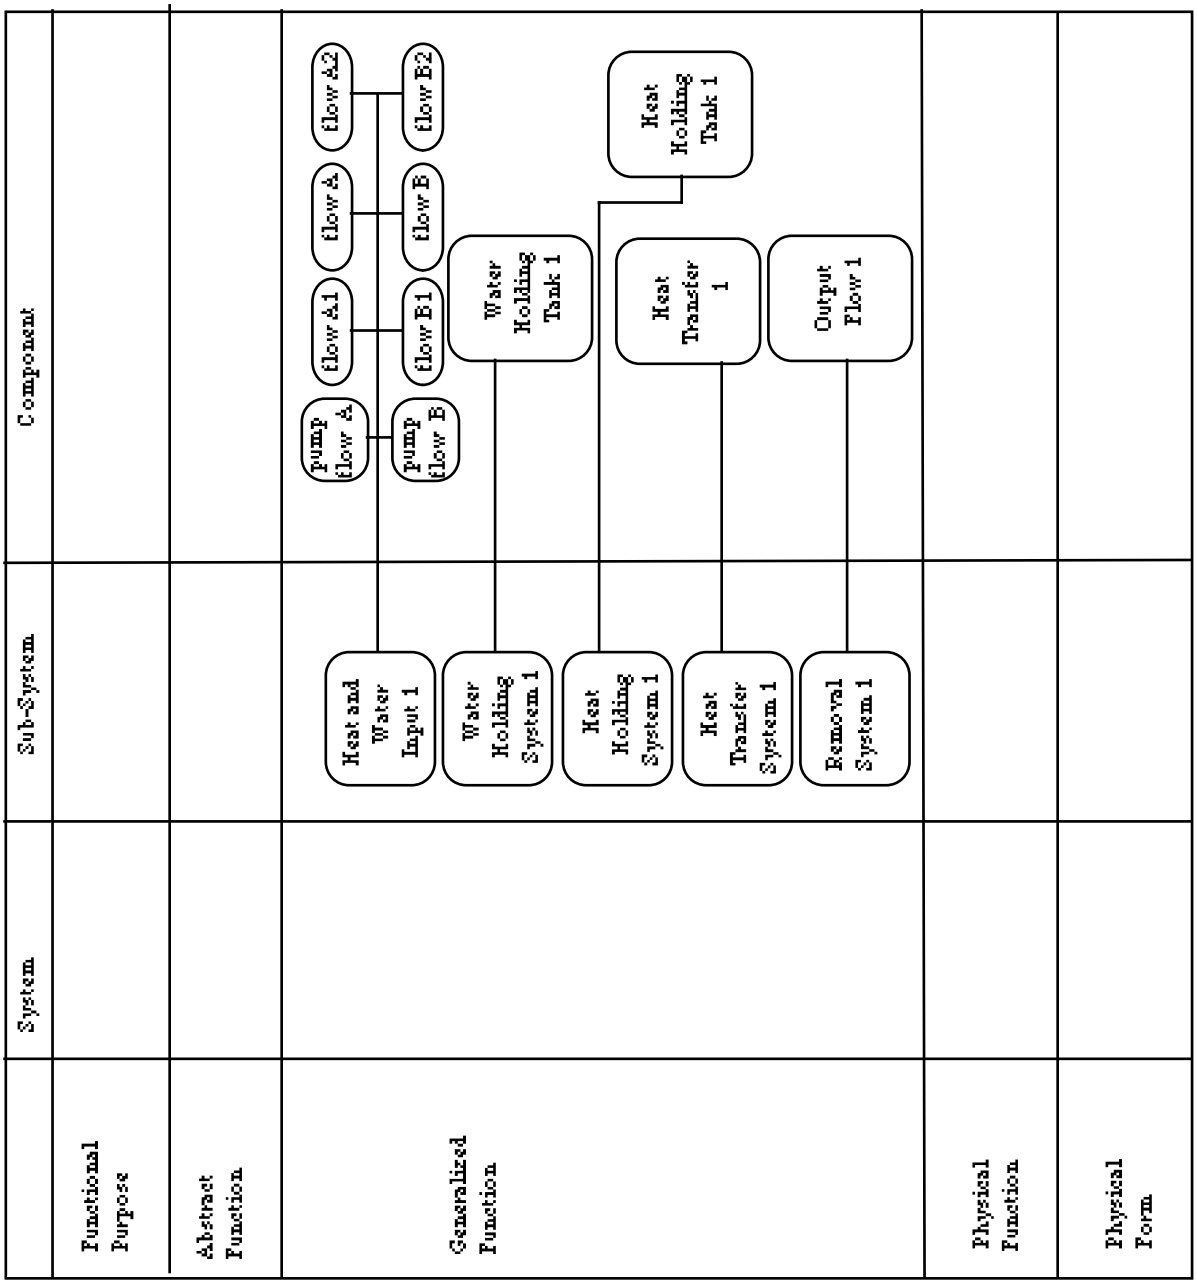

Figure 2. Means-end links in DURESS II (Bisantz and Vicente, 1994)

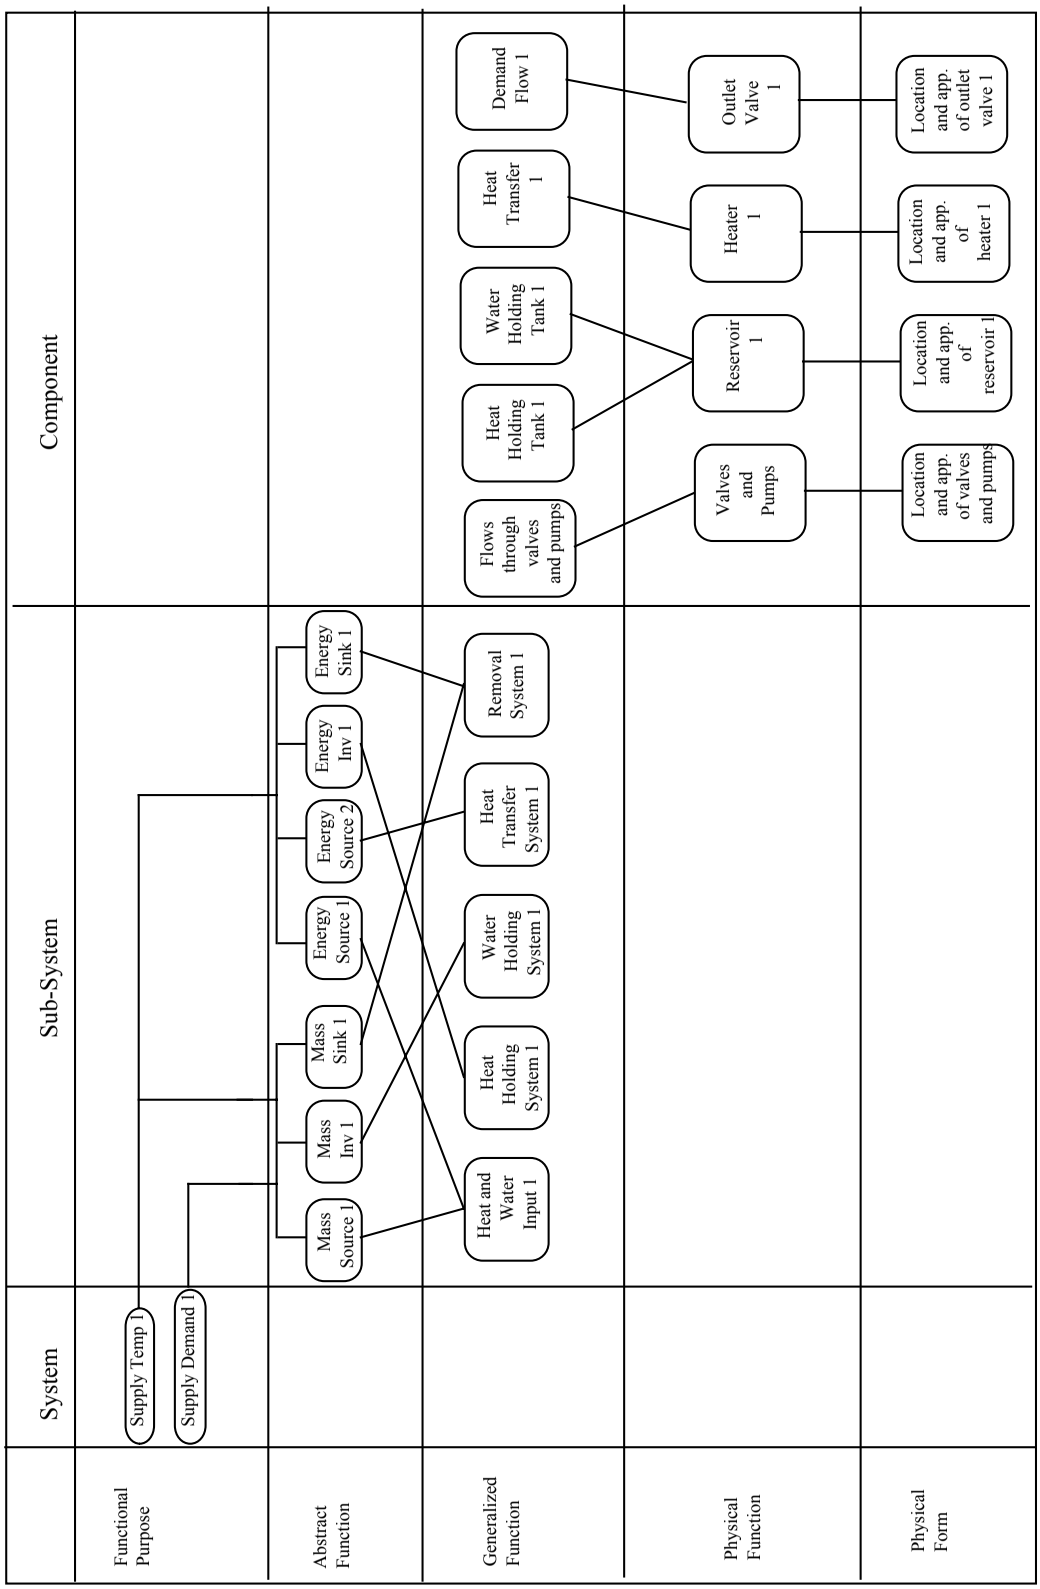

Figure 3. Part-whole links in DURESS II (Bisantz and Vicente, 1994)

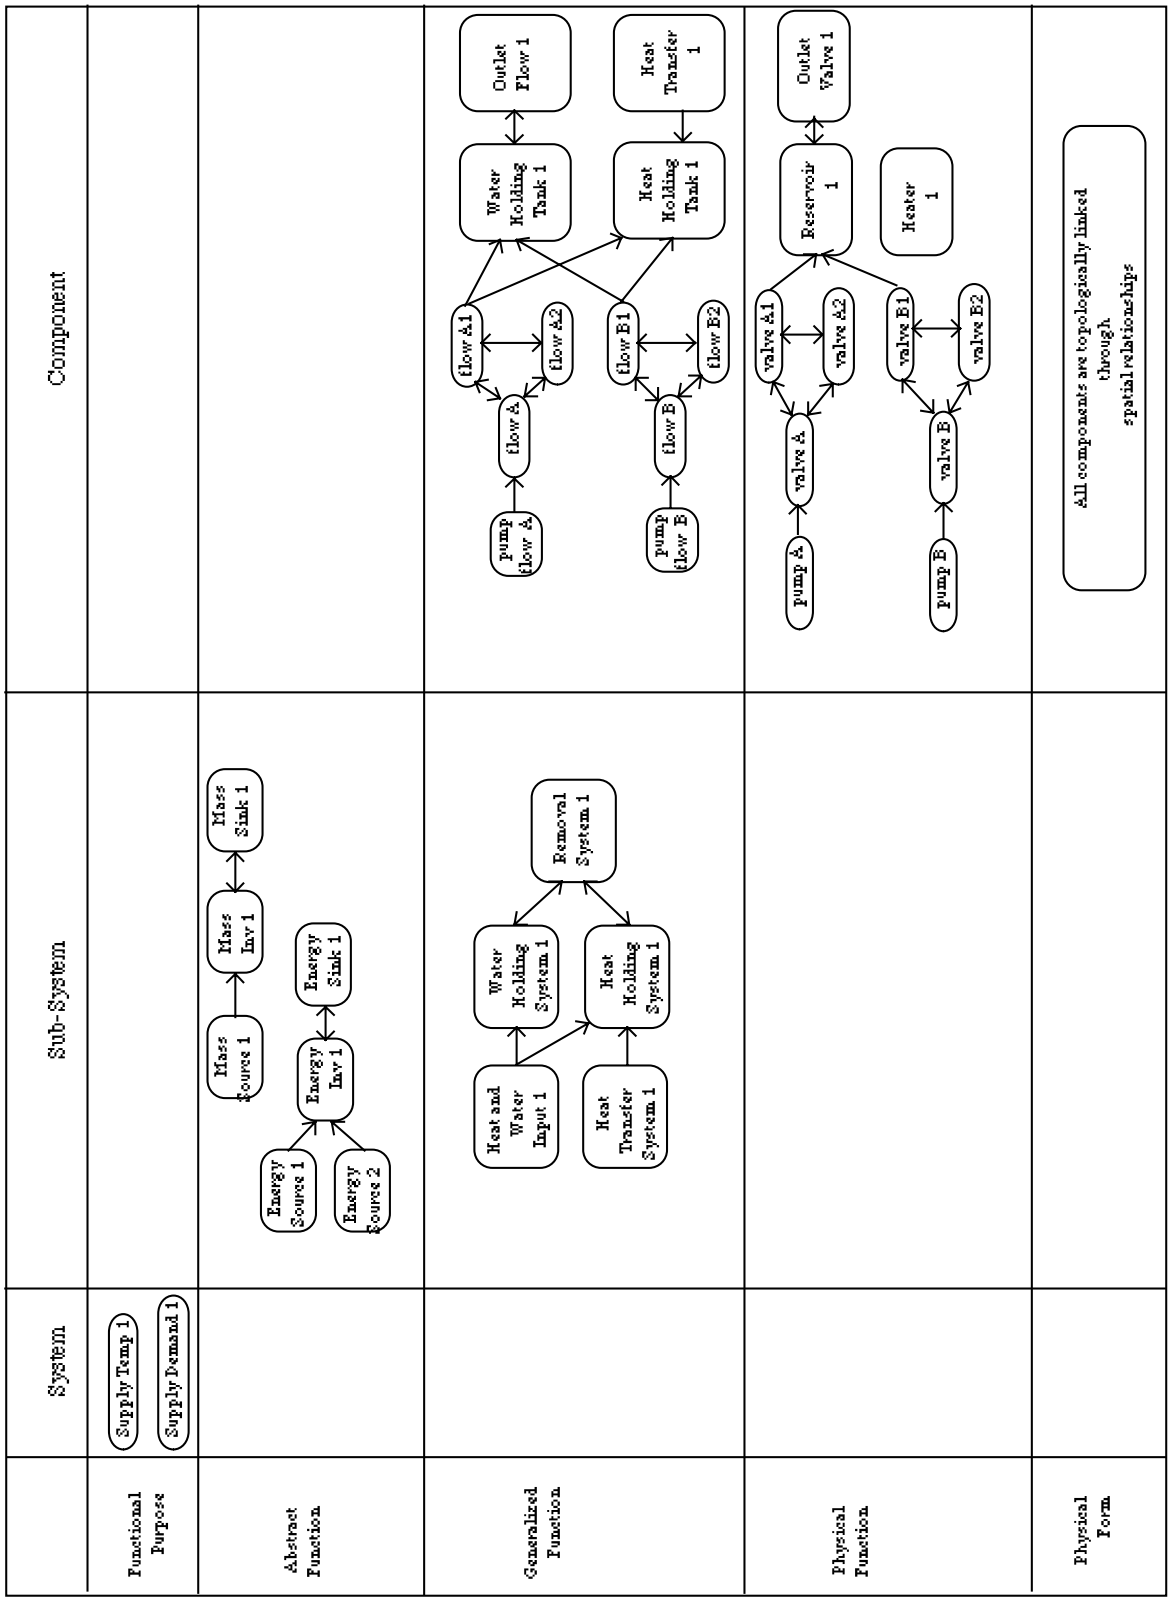

Figure 4. Topological links in DURESS II (Bisantz and Vicente, 1994)

Functional purpose. Objects at this level of abstraction correspond to system goals, and therefore are appropriately described at the system level of the part-whole decomposition. There are four goals in this system: Keep the water at the setpoint temperature for each reservoir (two goals), and keep enough water in each reservoir to keep up with the current demand flow rate (two goals).

Abstract function. This level can be described in terms of the flow of water and energy which balances the conservation of mass and energy for each storage subsystem. In addition to shifting downward in abstraction from the Functional Purpose level, this corresponds to a decomposition from the system to sub-system level (see Figure 1). As shown in Figures 3 and 4, each subsystem has one mass and energy store (the reservoirs), one source of mass (input water), two sources of energy (input water and the heater), and one sink of mass and energy (the output valve). Topological links at this level, shown in Figure 4, indicate the flows of mass and energy through the subsystems.

Generalized function. Flows and storage of heat and water are described at this level of abstraction. At the subsystem level of decomposition (see Figures 3 and 4), the rate of flow of water and heat transfer from the input stream, rate of heat transfer from the heating system, storage of heat and storage of water in the reservoirs, and rate of removal of heat and water due to demand are described for both subsystems. A further decomposition to the component level, shown in Figure 2, allows the description of the rate of heat transfer and water flow through each valve and pump, as well as the rate of heat transfer from the heater, storage of heat and water in the reservoir, and rate of removal of heat and water due to demand. For both the subsystem and component descriptions, the topological links, shown in Figure 4, indicate the flows of water and heat through the components.

Physical function. The states of system components are described at this level of abstraction. Because only individual components have measurable states in this system, the descriptions are at the component level of decomposition. The settings of valves, pumps, and heaters are described, along with the volume and temperature in the reservoir. Topological links at this level indicate physical connections between components (see Figure 4).

Physical form. At this level, the appearance, condition, location, and anatomical configuration of each component are described. The topological links reflect spatial relationships between components.

The next section of this CWA examines the types of decisions operators will have to make during system operation for normal conditions (start-up, tuning, and shut-down)

and in fault situations. This is accomplished through the use of Rasmussen's (1976) decision ladder.

3. CONTROL TASKS

The decision ladder diagram (see Figure 5) is a ladder of abstraction with the left-hand leg representing the analysis of a situation and the other, descending leg representing the planning and execution of a suitable action. The short cuts, or shunts, indicate the skipping of steps in the sequence (Rasmussen, 1976).

The operation of the DURESS II system has been broken down into four major modes: System start-up, normal or steady state operation, system shut-down, and fault diagnosis and correction. Each mode is a general classification of control, or subtasks, for the various operation states of this system, which are defined below:

- Start-up deals with initiating the feedwater system and the meeting of initial goal demands from shut-down state.

- Normal operation is the maintaining of a safe level of operation and the achieving of the present set of demands imposed on the system operator. It does not deal with problems or irregularities that occur when the system deviates from normal operating conditions or reacts differently from what was intended by the operator when a set of control actions was implemented.

- Shut-down contains the control tasks which concern the meeting of a specific goal state, in this case it is a "zero-state" which is to be achieved.

- Fault management includes the large set of possible control actions when dealing with a system that is deviating significantly from normal operation. The type of control actions required will be determined by the severity of the disturbance. This disturbance will either trigger a set of predefined control actions (typical disturbances) or an analysis of both the system and the disturbance so as to formulate a new set of control actions (non-typical, unpredictable disturbances).

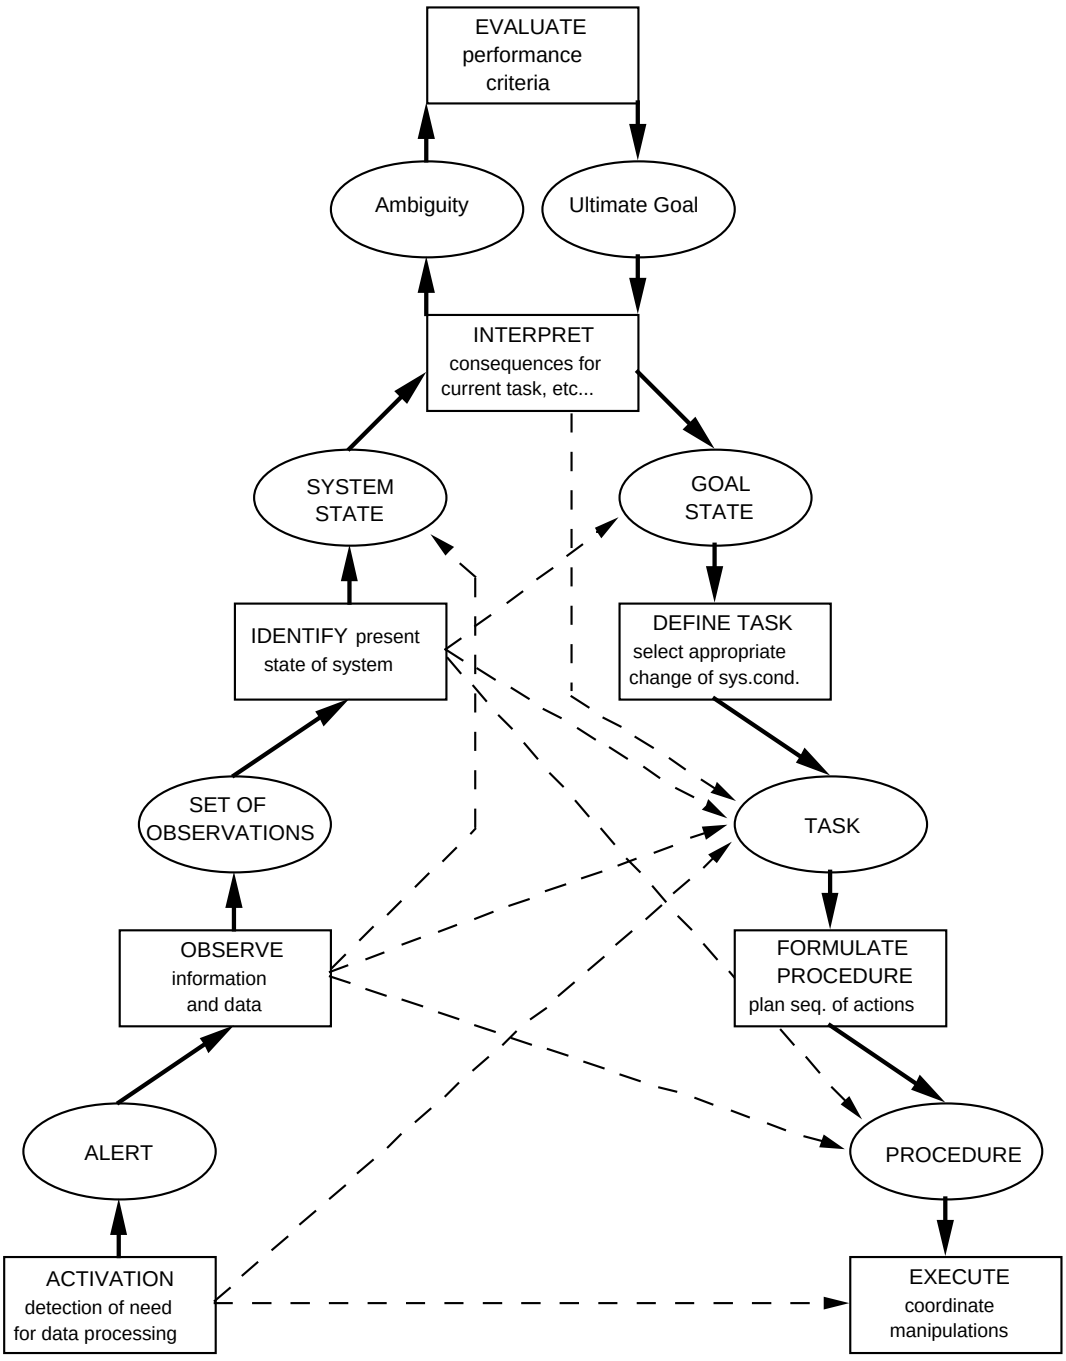

Figure 5. Rasmussen's (1976) decision ladder

The four DURESS II operating modes, with respect to their representation using the decision ladder, are discussed below These representations follow a "path" around the ladder, which indicate the operation of the system, for that particular mode.

3.1. System Start-Up

The decision ladder for system start-up is represented in Figure 6 on the following page.

1. *Goal State.* The trajectory around the ladder begins at knowledge of the goal state, since operators will already know that they need to start-up the system. The goal state is given to the operator in terms of demands (D1, T1 and D2, T2).

2. *Goal State Iteration.* The highest level goal state is to meet the demand flows and temperatures. There are, however, several levels of subgoals. For example, one subgoal particular to system start-up is to initiate the flow of water into the system. Another subgoal is to initiate the heat transfer from the heater to the reservoir. The operators must meet these subgoals first, before meeting the highest level goal. The operators must carry out this decomposition of subgoals in order to be able to define the meaningful tasks necessary in order to meet the highest goal. This decomposition of goals into subgoals has been indicated by a loop indicating an iteration on goal state.

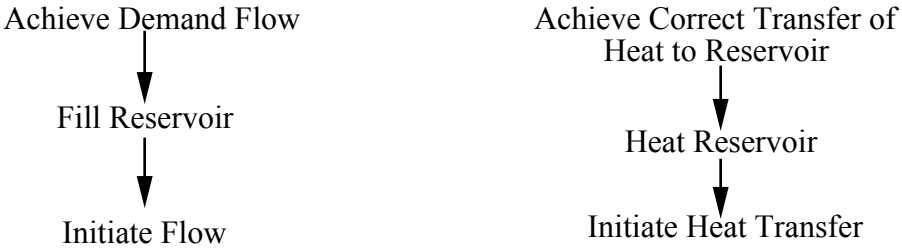

3. *Define Task.* From the identification of each lower level subgoal, operators must define their task in order to meet each subgoal. In this case, the operators decide that the feedwater streams need to be configured and the heat will need to be established.

4. *Task.* The operators will now know the task(s) that need to be performed. As in the above example, water flow and heat transfer both need to be initiated.

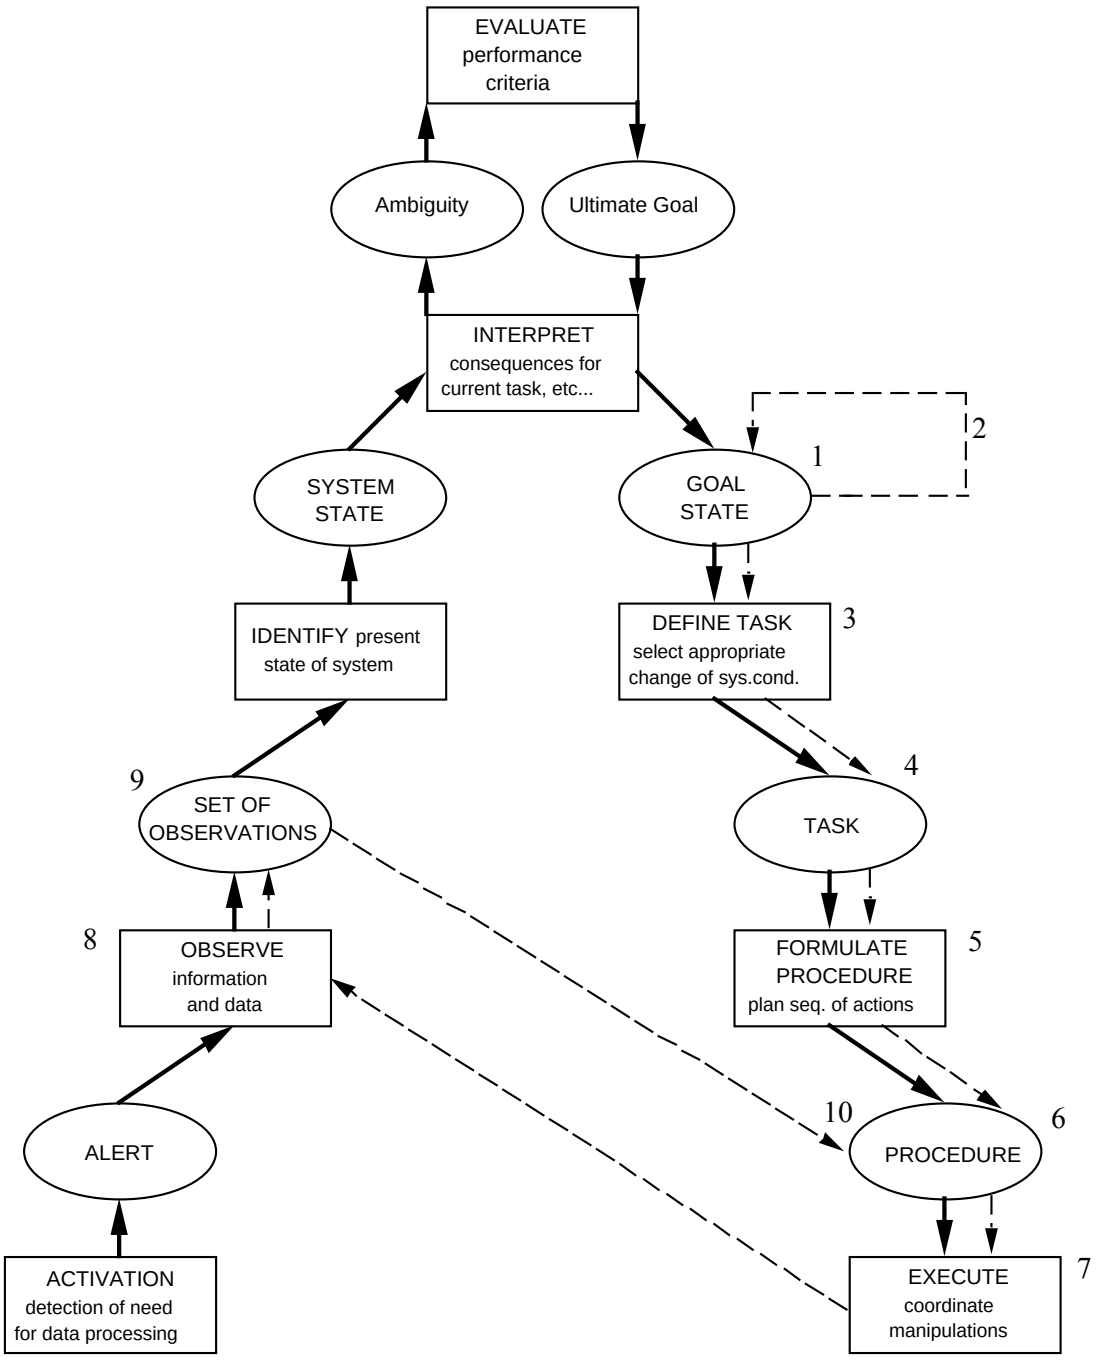

Figure 6. Decision ladder for start-up

5. *Formulate Procedure.* From the task, the operators must formulate a procedure or sequence of actions to accomplish the specific task. Using the example of initiating water flow into the system, an appropriate procedure might be: 1) Open valves, 2) turn on pumps, and after flow has been achieved, 3) set the valves to the settings needed to meet the next higher level goal (meet demand flows). For the initiation of heat transfer task, an appropriate procedure might be: 1) Wait for the reservoir to contain some amount water, 2) turn on the heaters, and 3) set the heater to the setting needed to meet the demand temperatures.

6. *Procedure.* The operators now have a procedure (or set of procedures) to follow in order to achieve the goal state. This basic procedure for start-up is to open the valves, turn on the pumps, ensure that water is going into the reservoirs, and then to turn on the heaters.

7. *Execute.* Execution involves coordinating the above manipulations required by the procedure. In this case, at a very crude level, the operators need to establish flow first, and then heat the reservoirs.

8. *Observe.* After executing the appropriate actions, the operators will observe the response of the system to ensure that the procedure performed was the correct one to obtain the desired goal or sub-goals.

9. *Set of Observations.* After observing the system, operators will have a set of observations which they use to compare, once again, to the overall (sub)goals desired and to check whether or not these (sub)goals have been achieved.

10. *Procedure.* After the goal state has been reached, the operators need to reduce the input flow and heaters so that the goal states are not passed. The interface feedback helps to formulate a procedure to determine the setting of the valves and heaters. After this is completed, the operators can continue by fine-tuning the system (through roughly the same steps) in order to meet the prescribed demands.

It should be noted that steps 2-4 may be skipped by expert operators. Once the goal state is realized, expert operators may know how to formulate the proper procedure to reach this goal state, ignoring the need to concentrate on the various lower-level, nested subgoals.

3.2. Normal Operation

The decision ladder for normal operation is represented in Figure 7 on the following page.

1. *Observe*. The operators have to monitor the system (in terms of the goal state) to see if there are any mismatches between the goal state (the system is observed in terms of the goal state) and the present state of the system.

2. *Set of Observations*. The operators now have a set of observations which (in this case) indicate that there is a difference between the needed goal state and the current state of the system.

3. *Define Task*. The operators now have to define the task(s) that will bring them back to the desired goal state. Once again this process is iterative, as each possible sub-task needs to be defined. As before, the operators will define the task as reconfiguring the feedwater streams and/or reconfiguring the heaters (depending on the type of mismatch between actual system state and the goal state).

4. *Task*. The operators now know that their task is to change the valves and heater settings (as needed) in order to reach the goal state.

5. *Formulate Procedure*. Once again, the operators decide the valve and heater settings needed in order to obtain the desired goal state.

6. *Procedure*. The operators now have a procedure (or set of procedures) to follow in order to achieve the goal state. Here, the procedure could be "lowering the setting on heater A by 2". This procedure follows this form for all valves, pumps, and heaters, as necessary.

7. *Execute*. The operators now carry out the procedure(s).

8. *Observe*. Once again the operators observe the system to monitor the time when the system is once again at the goal state and the valves, pumps, etc. can be cut back (in order to not 'overshoot' the goal state).

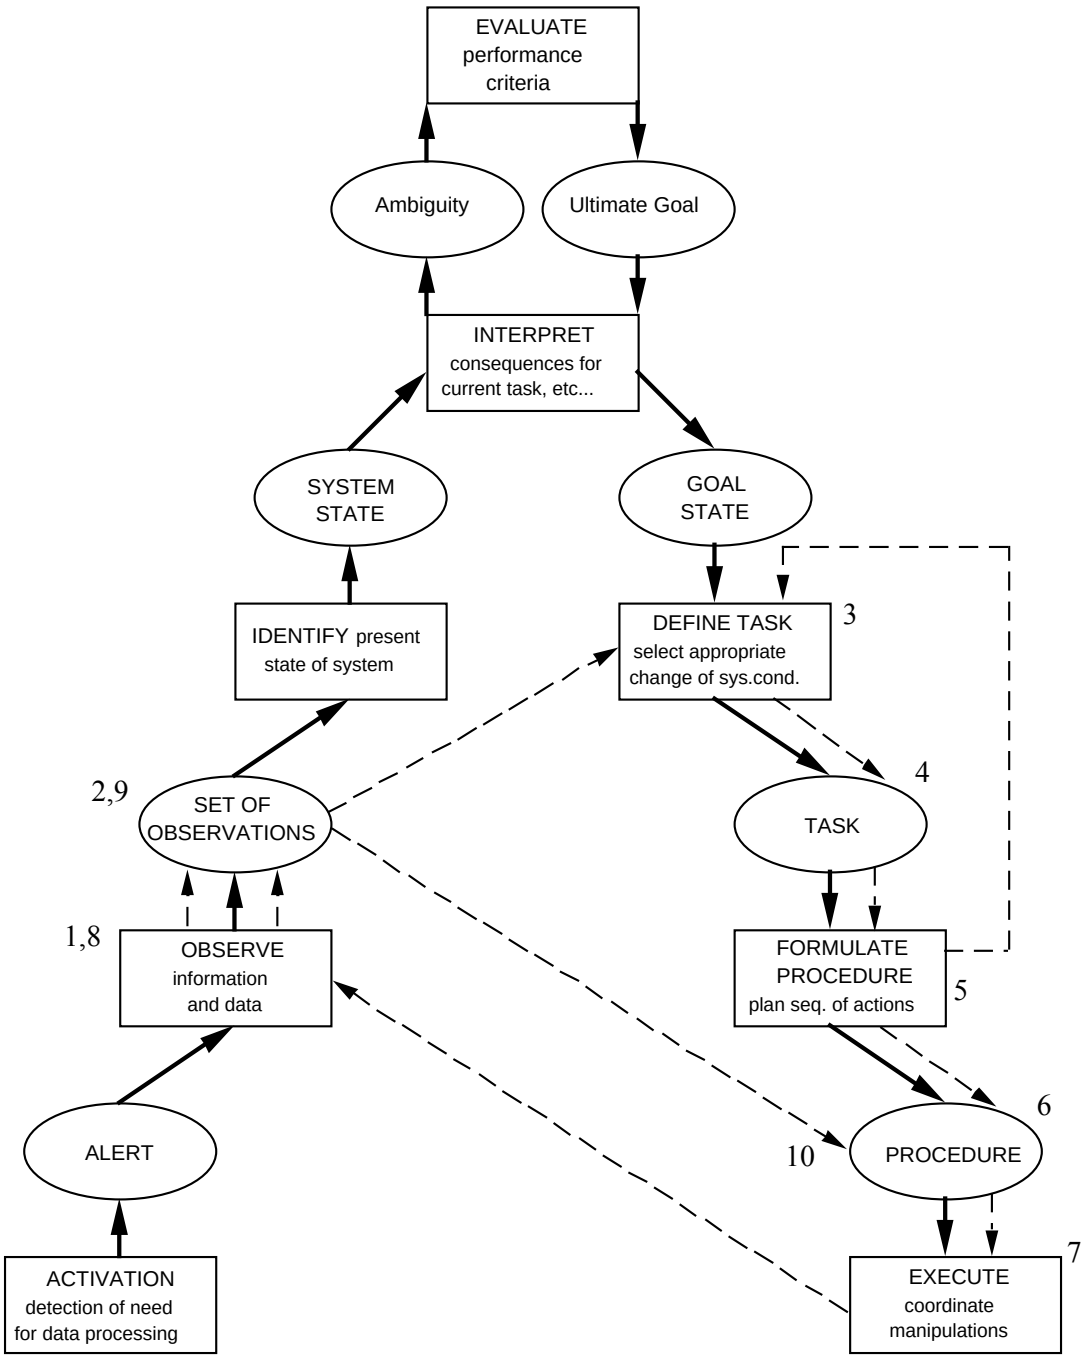

Figure 7. Decision ladder for normal operation

9. *Set of Observations.* After observing the system, operators will have a set of observations which they use to compare, once again to the overall goal and subgoals desired and to check whether or not these goals have been achieved.

10. *Procedure.* After the goal state has been reached, once again the operators need to reduce the input flows and heaters so that the goal state is not passed. The feedback provided by the interface helps to formulate a procedure to determine the setting of the valves and heaters. After this is completed, the operators can continue by fine-tuning the system (through roughly the same steps) in order to meet the prescribed demands.

It should be noted that steps 3-5 may be skipped by expert operators. Once the operators realize that the system needs (for example, tuning) attention, they may know the proper procedure needed to once again reach the goal state, ignoring the need to concentrate on the various lower-level, nested subgoals.

3.3. System Shut-down

The decision ladder for system shut-down is represented in Figure 8 on the following page.

1. *Goal State* The decision ladder for this operational mode begins with goal state. The operators begin this decision chain knowing that the system must be shut down

2. *Task.* The operators already know, by definition, that the task is to eliminate flow and eliminate the heat to the reservoir. Therefore, there is no real need to define the task.

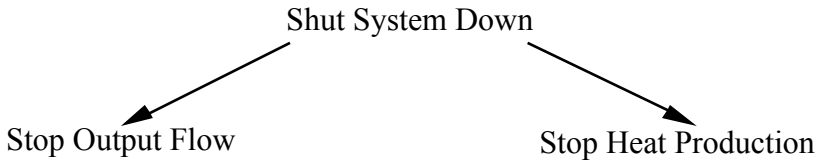

3. *Formulate Procedure*. The operators will then formulate a procedure which will bring the system towards the goal state (shut down). For example: 1) stop heat production (turn heaters off) and 2) stop flow (turn pumps off, open valves, allow reservoir to drain). Note that the first part of this procedure, stop heat production, should be done before the second half, stop flow.

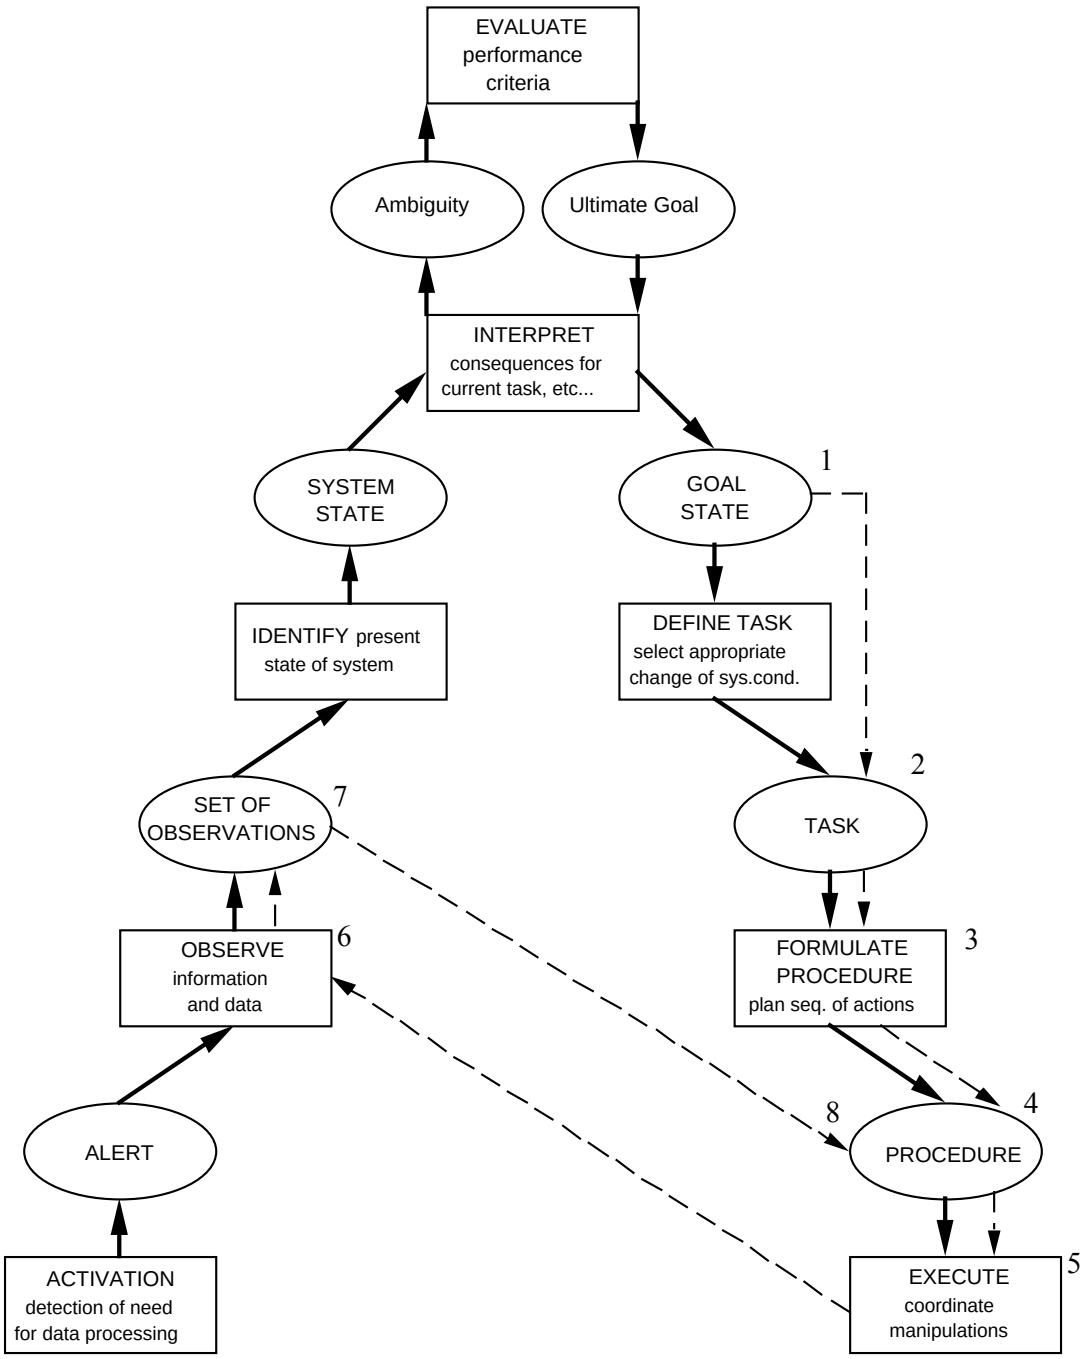

Figure 8. Decision ladder for system shut-down

4. *Procedure*. The operators would now have a procedure to follow for shutting down the system.

5. *Execute*. The operators now execute the procedure.

6. *Observe*. Finally, the operators would once again check to make sure that the system is responding in the appropriate manner (i.e. moving towards the goal state of shut-down).

7.*Set of Observations.* After observing the system, operators will have a set of observations which they use to compare, once again to the overall goal desired and to check whether or not these goals have been achieved.

8. *Procedure.* The operators need to monitor the system based on feedback, to ensure that the goal state will be reached. The operators continue by fine-tuning the system (through roughly the same steps) in order to meet the goal state.

It should be noted that steps 2 and 3 may be skipped by expert operators. Once the operators realize that the system needs to be shut-down, they may immediately know the proper procedure needed complete this task, ignoring the need to concentrate on the various lower-level, nested subgoals.

3.4. Fault Management

The decision ladder for Fault Management is represented in Figure 9 on the following page.

1. *Activation*. The operators will detect a need for action because a fault of some sort is occurring and the system is not at goal state. This may be indicated by the interface, and if the interface design is sufficient, the operators will not need to continually observe the system, but it will be readily apparent that a fault has occurred.

2. *Set of Observations*. The resulting observations will indicate to the operators that a fault has indeed occurred. It is important to note here that one of several things could occur, depending on the type of fault occurring. If for example, the fault is one in which the operators do not need to know (immediately) *why* the fault is occurring, but only need

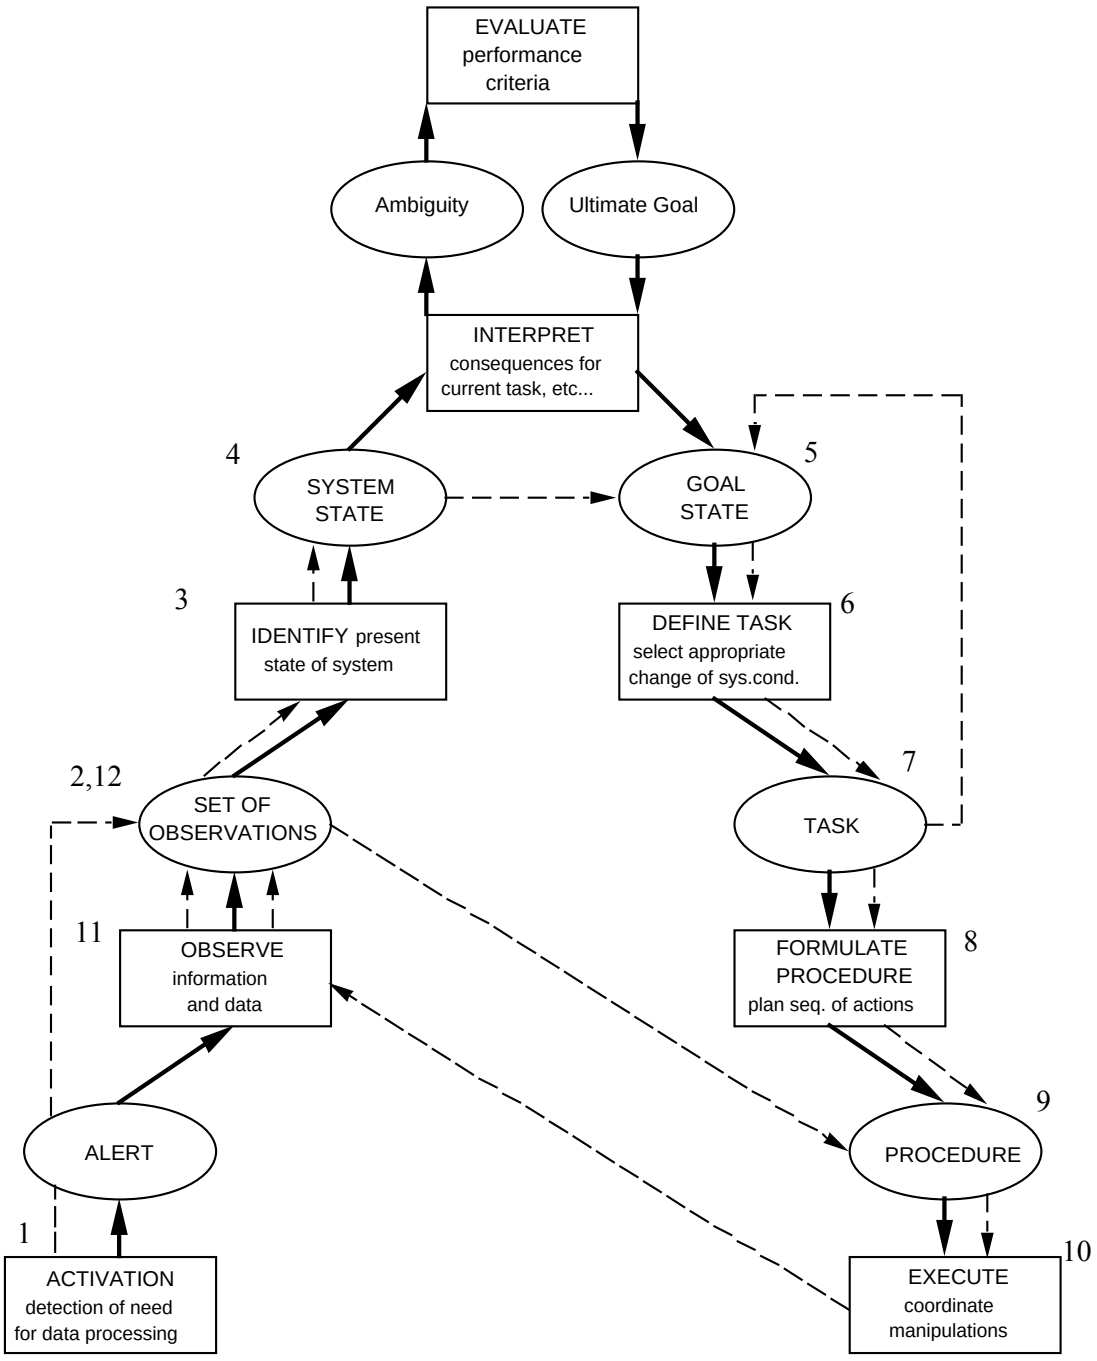

Figure 9. Decision ladder for fault management

to know that the fault has occurred and how to compensate for it, then the operators may only need to go directly to 'task' to compensate, saving the analysis of the fault for a later time. An example of this would be a leak in a reservoir. In this case, the operators may only need to make sure that the input rate is increased so that the water level in the tank remains constant and does not deplete or overflow the surrounding containment chamber. This would result in the quick formulation of a procedure (increase input water rate--decrease outlet valve setting, etc.). This procedure would then be carried out and later, the operators can attempt to figure out why the fault occurred.

3. *Identify*. Based on the set of observations taken, the operators need to identify the state of the system. In this case, the operators have identified that a fault is occurring and will attempt to understand exactly what that fault may be and why it has happened. Moving to this part of the ladder is one way to manage a fault. Rasmussen (1981) defines this method as root-cause problem solving. In this case, the operator attempts to find out the cause of the fault before any action is taken.

4. *System State*. The operators now have an understanding of the system state, the type of fault that occurred, and why that fault occurred. This system state can then be compared to the goal state to determine the extent of the fault.

5. *Goal State.* The operators examine the system and compare it to the previously known goal state.

6. *Define Task.* The operators now need to define the task(s) necessary to compensate for the fault. This process of defining the task in terms of the goals and subgoals of the system is recursive, as above.

7. *Task*. The operators, now understanding system state, would have the task of reconfiguring the input feedwater streams (and possibly the outlet valve) to compensate for the leak in the reservoir.

8. *Formulate Procedure*. The operators now need to figure out how to change the valve settings, etc. to compensate for the fault so that the goal state can be re-established.

9. *Procedure*. The operators now have a procedure to compensate for the reservoir leak. In this case the operators realize that the outlet valve setting cannot be changed because

the goal state must be maintained. Therefore, the input flowrate must be increased to compensate for the leak.

10. *Execute*. Once the consequences of the procedure are deemed acceptable, the operators can execute the procedure.

11. *Observe*. The operators once again monitor the process to make sure that the fault is being corrected.

12. *Set of Observations.* After observing the system, operators will have a set of observations which they use to compare, once again to the overall goal and subgoals desired and to check whether or not these goals have been achieved and the fault has been managed.

13. *Procedure.* The feedback from these observations enable the operators to continue by fine-tuning the system (through roughly the same steps) in order to ensure that the fault is being compensated for and that the system is meeting the prescribed demands.

It is important to note some of the different shortcuts in the decision ladder that different operators may take. In the above discussion, the operators spent time trying to figure out what the fault was and why that fault happened. Another method of fault management that was mentioned was one of simple compensation. In this case, the operators would skip steps 3 and 4. After making the system observations, the operators can immediately compare these observations to the goal state, realize that a fault has occurred, and attempt compensation. Here, the reasons *why* the fault occurred is of no consequence, for the moment. The immediate problem is to compensate for the fault (in this case, a reservoir leak) to minimize losses. Once the fault has been compensated for, then the operators can try to determine the cause of the fault.

Once again, expert operators may not even need to compare the system observations with the goal state and they can proceed directly to task or procedure. In all of these ladders, the more an operator knows about the system, the fewer number of ladder rungs will have to be "visited" on their path to successful control.

4. MENTAL STRATEGY ANALYSIS

The abstraction hierarchy for the DURESS II system helped to develop an understanding of the system itself -- what goals must be accomplished and what means are available to reach those goals. The decision ladder analyses enabled the identification of the decision activities associated with the four modes of system operation. This section of the CWA presents an analysis of possible control strategies that can be used for controlling DURESS II under both normal and abnormal conditions.

The strategies presented can be broken down into three main categories, which are considered to be the main information processing activities identified in the decision ladder analysis. The first category of strategies that operators can use for control are identified as planning strategies. These are typically used for both start-up and tuning tasks. These strategies are:

- Input flow strategies. Strategies for achieving desired flows into the reservoirs. These strategies involve the different configurations of the valves and pumps of the feedwater streams.
- Reservoir flow strategies. Strategies for achieving the demand flows out of the reservoir. These strategies involve the setting of different volume levels in the reservoirs (when and how to stabilize the volume levels).
- Reservoir heating strategies. Strategies for achieving the required demand temperature. These strategies involve the manipulation of different heater settings, different tank volumes, and different input and output flows to control output temperature.

The second category of strategies are specific to the shut-down task. These strategies involve two different modes of system control, defined by whether or not operators create a symmetrical structure within the system to complete the task. Finally, the third category of control strategies are diagnostic strategies. These are specifically useful in abnormal situations, and are thus identified as fault management strategies. They include strategies for determining the cause of an observed fault and for compensating for that fault.

Note, however, that these strategies are inter-constrained. For example, adopting a certain reservoir flow strategy may limit which input flow strategies are feasible and which reservoir heating strategies can be used. The current independent breakdown, however, allows for comparison between different strategies within each class and for relating them directly to decision ladder activities. For each set of the strategies discussed

below, the potential advantages and disadvantages of each individual strategy are outlined. The ability of operators to meet new demand levels and to handle faults has been considered when postulating these advantages and disadvantages.

4.1. Input Flow Strategies

There are three main strategies for configuring DURESS II to provide the appropriate amount of flow into the reservoirs. These strategies are identified according to the configuration of the six input valves in the two feedwater streams. The strategies are:

- Single FWS Configuration - Uses only 1 of the feedwater streams
- Decoupled FWS Configuration - Uses a subset of the valves in each stream
- Full FWS Configuration - Uses all six input valves

In addition, for any one of these strategies, the operator may choose to fully open the primary valves (VA and VB). Control would then be directed towards the secondary valves (VA1, VA2, VB1, and VB2). Doing so reduces the complexity of configuring the valves to maintain the desired input flowrate. This will be further explained as a "valve complexity reduction strategy," following the discussion of the three strategies listed above.

Table 1 lists the strategies that are available to meet the system goals under different setpoint conditions. Note that slight differences (e.g., the difference between closing one valve instead of another) may exist within each of the above strategies, and these will be indicated as necessary. Below, the three strategies are discussed along with the advantages and disadvantages of each. Also included are considerations for changes in setpoints and fault detection and management.

Single FWS configuration. With this strategy, the operators use only one of the FWSs to obtain the desired input flow rate. This system configuration is easier for the operator to manipulate, especially in the case where the main valve (VA or VB) is fully opened, as mentioned above. Using only 1 FWS, there is only one unique solution to solving the input flowrate equation. However, the major problem with this strategy is that it cannot be used when ($D1 + D2 > 10$). In this case, the operators would need to change strategies and bring on-line the other feedwater stream and then reconfigure the streams. This may involve recalculating algebraic equations to meet the desired flow

$(D_1 + D_2) \leq 10$	$[10 < (D_1 + D_2) \leq 20]$ and $(D_1$ and $D_2 \leq 10)$	$[(D_1 + D_2) \leq 20]$ and $(D_1$ or $D_2 > 10)$
Single FWS	-------	-------
Decoupled FWS	Decoupled FWS	-------
Full System (3 or 4)	Full System (3 or 4)	Full System (3 or 4)

Table 1. Table of the classes of setpoints for system goals, and the operating strategies that will work under those conditions.

rate. If this change of strategy is required, operators may have the problem of dealing with the time lags associated with bringing an entire feedwater system on-line. Also, they must then incorporate the new valves that are now active in the second feedwater stream into the calculations necessary to meet the demands. If a fault occurs while operating under this strategy, the problems may be overwhelming. Time lags, and the fact that the operators need to solve equations, are just the beginning. Also consider the fact that the feedwater stream that is off-line may contain a fault. Since this is unknown to the operators when they attempt to bring that system on-line, there could be additional complexities.

Decoupled FWS configuration. In this strategy, both feedwater streams are being used, but only a subset of the valves are being employed. For example, VA2 and VB1 are turned off, thereby having the operator use VA, VA1 in feedwater stream 1 and VB, VB2 in feedwater stream 2. Thus, each FWS supplies water to a separate reservoir. This makes it easier to control the system since it is possible to treat each reservoir independently, without worrying about interactions. Again, if VA and VB are fully opened, the operators need only manipulate the remaining two valves to control the system. In this case, there is only 1 unique solution for each configuration of valves employed. This strategy can continually meet maximum system demands ($D1 + D2 \leq 20$) as long as $D1$ and $D2 \leq 10$.

A major problem with this strategy is the case when ($D1$ or $D2 > 10$). Since 10 units per second is the maximum flowrate that a FWS can achieve, the operators would need to bring the other valve in the other feedwater stream on-line to add the additional input flow to the reservoir. Due to the interaction in the system, this may affect the input flowrate to the other reservoir. Again, in this strategy, operators have the problem of dealing with time lags. Also, they must incorporate the new valves that become active in the FWS into the calculations necessary to meet the demands. If a fault occurs while operating under this strategy, the problems may once again be overwhelming. The fact that half of the valves in each feedwater stream are off-line may involve time lags if there should be changes in setpoints or faults. To bring a valve on-line, operators will need to deal with those time lags and as well as the introduction of this new variable into the equations they must solve. If operators are not familiar with solving these new equations, problems may occur.

Full FWS configuration. With this strategy, the operators manipulate either all three valves in one FWS and two valves in the other, or all three valves in both FWSs to

maintain the correct input flows into the reservoirs. In this configuration, the system is highly flexible. If there is a fault or a new demand, all valves and pumps are on-line and available for control. As well, this is the only strategy that can continually meet the maximum flow into the tanks, if necessary (D1 + D2 = 20), and the only strategy that can meet the following conditions: (D1 + D2 ≤ 20) and (D1 or D2 > 10).

Being the most complex of the strategies, a large change in the demand setpoints may result in the operators having to reconfigure the valve settings. This may prove to be difficult to accomplish due to the possible calculations involved. However, in the previous strategies, not all of the components are on-line. In this strategy, since all of the components are already on-line, effort (in terms of starting up new components) is reduced. Therefore, both the system and the operators are capable of handling the new setpoints, once the correct valve configurations have been established. Small changes in setpoints may only require small changes to the valve settings.

The Full FWS configuration involves all three input valves in each subsystem to obtain the proper input flow rate. This is the most difficult strategy to adopt because there are an infinite number of solutions to the problem. This is due to the fact that two equations (flowrate into each reservoir) must be solved for six unknown variables. To solve these equations, operators must understand how the flow rate of one valve affects and is affected by the settings of the other valves within the two feedwater streams. In addition, the operators must also understand how the two feedwater streams are coupled with the two reservoirs. Undoubtedly, this can be a cognitively demanding strategy.

Valve complexity reduction strategy. As previously mentioned, this strategy can be used in conjunction with any of the above strategies. The operator sets the primary valves (VA and VB) fully open and controls the system using only the secondary valves (VA1, VA2, VB1, and VB2). Doing so requires operators to only monitor the secondary valves instead of both the primary and the secondary valves. Using this strategy, understanding what flow levels are entering the reservoir only requires a summation of the input flows, not computing the weighted sum of the flows based on the configuration.

However, there may still be an infinite number of solutions with this strategy, especially if the operator is using it in conjunction with the Full FWS configuration. Additional problems with this strategy are realized in the event of a fault, especially if one of the secondary valves becomes stuck while fully open. For example, if VA1 were to break open, the operator would need to know that VA can be manipulated in place of VA1 and still meet the desired flow rate. This may be difficult for operators to understand because of the relationship between the flows through valves VA and VA2

(i.e., a setting on VA would affect the flow through VA2 as well as through VA1, assuming VA2 > VA). It should be noted that, in the case of this error, to achieve the initial input flowrate to R1 could require manipulations of the second FWS stream as well. In this case, it might be worthwhile for the operators to consider changing strategies and/or possibly shutting down the broken feedwater system.

Table 2 presents the three original strategies (the Full FWS strategy was broken down into three and four valve settings according to the number of secondary valves being used) along with the constraints that govern the usage of each. Also, a list of criteria used to compare the utility of each strategy for certain situations is presented. Based on this, it is evident that the optimal strategy depends upon the operators' criteria, their level of system knowledge, and the current demands.

4.2. Reservoir Flow Strategies

Through the decision ladder analysis, it was determined that, for the various modes of system behaviour, operators will need to determine how to achieve the demand flowrates. These particular strategies deal with achieving the demand output from the reservoir using different combinations of input and output flows for the reservoirs. The three strategies will be discussed below, along with the potential benefits and disadvantages of each. Three different strategies for meeting these demanded flowrates are presented below. These strategies are:

- Shut-off Flow, Full or Partially Full Reservoir
- Constant Flow, Constant Reservoir Volume
- Constant Flow, Increasing/Decreasing Reservoir Volume

Shut-off flow, full or partially full tank. With this "batch production" strategy, the operators stop meeting demand, fill the reservoir, and then heat the water to the appropriate temperature. When the water in the reservoir is heated to the correct temperature, the operators then meet the demand by draining the reservoir. Partially filling the reservoir is essentially the same strategy, but with a faster cycle. Using this method, it is easy to meet the demand flow and to achieve the demand temperature since there is only one energy source to consider. This strategy can meet all potential demand levels for a limited period of time.

Criterion	Single FWS	Decoupled FWS	Three Valve	Full Valve
Cannot be used for	$(D_1 + D_2) > 10$	D_1 or $D_2 > 10$	$(D_1 + D_2) > 20$	$(D_1 + D_2) > 20$
Likelihood that a change in demand results in a need to reconfigure system	Likely	Not as likely but possible	Zero	Zero
Facility with dealing with a major fault	Most cumbersome	Less cumbersome	Even less cumbersome	Least Cumbersome
Nature of resulting coupling between reservoirs	Two-way	None	One-way	Two-Way
Nature of mapping between valve flow rates and reservoir input flow rate	1:1 for both reservoirs	1:1 for both reservoirs	1:1 for one, 2:1 for the other	2:1 for both

Table 2. Criteria that govern the usage of control strategies.

There are several disadvantages, however, to adopting this strategy. The most important of which is the loss of productivity while the reservoir is filling, which would probably not be tolerated in a real industrial system. Along with this are the added costs of continually starting and stopping the system. A further disadvantage with this strategy is realized in the case when the reservoir is completely filled. If the operators overheat the water in the reservoir, they cannot compensate by adding additional water as coolant, as this could overflow the tank. Therefore, to use this strategy, operators must understand the heater/reservoir dynamics extremely well and know exactly when to turn the heater off.

Constant flow, constant reservoir volume. With this strategy, operators keep the reservoir volume constant and match the input flow to the demand (output) flow. The volume of water in the reservoir affects how long it will take for the water to reach its steady state temperature. The volume of water that the operator maintains in the reservoir can be considered a buffer. Here, it is very easy to achieve the demand flow and this demand level can be maintained indefinitely. This system never needs to be shut down and there is no further waiting for a reservoir to fill. As well, changes in demand flow are relatively easy to achieve. However, one problem with this method is that not all possible classes of demand flows can be attained. Whenever the two demand flows sum to more than 20 units, at least one reservoir cannot be operated by this strategy. In this case, operators would have to use one of the other two strategies.

Constant flow, increasing/decreasing reservoir volume. With this strategy, the initial input flow is either greater or less than the output flow. The volume of water in the reservoir, therefore, is constantly increasing or decreasing, respectively. Using this strategy, all flow levels can be achieved for a limited period of time. However, eventually the reservoir will near the point of overflow or depletion and will therefore need to be emptied or refilled as required.

It would be best to recommend that the operator adopt the constant flow, constant reservoir volume strategy. From a productivity standpoint, this is the most desirable strategy. Further, this strategy can be quite easily used in conjunction with the constant flow, increasing/decreasing reservoir volume strategy. Doing so allows the operators to continually meet the demands as the reservoirs fill (during start-up) and empty (during shut-down).

4.3. Reservoir Heating Strategies

Each reservoir must supply water not only at the various demanded flow rates, but also at the two different temperature setpoints. Analytically, there are five classes of strategies for heating the water in the reservoirs to the target temperatures. Each strategy varies in terms of the time they take to reach the desired temperature and the amount of control required from the operator. Note that these strategies can be used in conjunction with each other (e.g., open-loop followed by fine tuning via feedback control). The strategies are:

- Proportional Feedback Control (PFB)
- Derivative Feedback Control (DFB)
- Open-Loop Control (OL)
- Multi-Variate Control (MV)

The two main factors that determine how quickly the temperature setpoint for a reservoir can be attained and maintained are 1) the sensitivity of the heater around the goal area, and 2) the rate at which the temperature changes. For all of the following heater control strategies, operators must realize that the tolerance around the temperature goal state (the width of the green goal area mapped onto the E_{inv} scale for the P+F interface) is affected by the steady-state volume (V_{ss}). The larger the V_{ss}, the wider this goal area tolerance, and thus the easier it will be to maintain the temperature within the goal area. The rate at which the Temp/E_{inv} changes (the slope on the energy balance graphic in the P+F interface) is affected by the throughput of the reservoir (externally set according to demand) and the V_{ss}. In this case, the rate of change is slower for larger throughputs, and slower for larger V_{ss} values. It must be noted that the effects of V_{ss} in terms of tolerance and rate of change of E_{inv} are coupled, and therefore, there is no trade-off. Thus, there should be an invariant relationship between the given throughput (demand) and the V_{ss} (chosen by the operator) if behaviour is adaptive.

Figure 10 (a-e) illustrates the five strategies in terms of operator control inputs (heater/valve settings) and system response (temperature). Figure 10e is a special case of feedback control that will be further discussed. For all of the graphs, the solid horizontal line represents the goal temperature, the solid sloping line represents the current temperature, and the dashed line represents operator control inputs. The horizontal axis of each of these graphs is time.

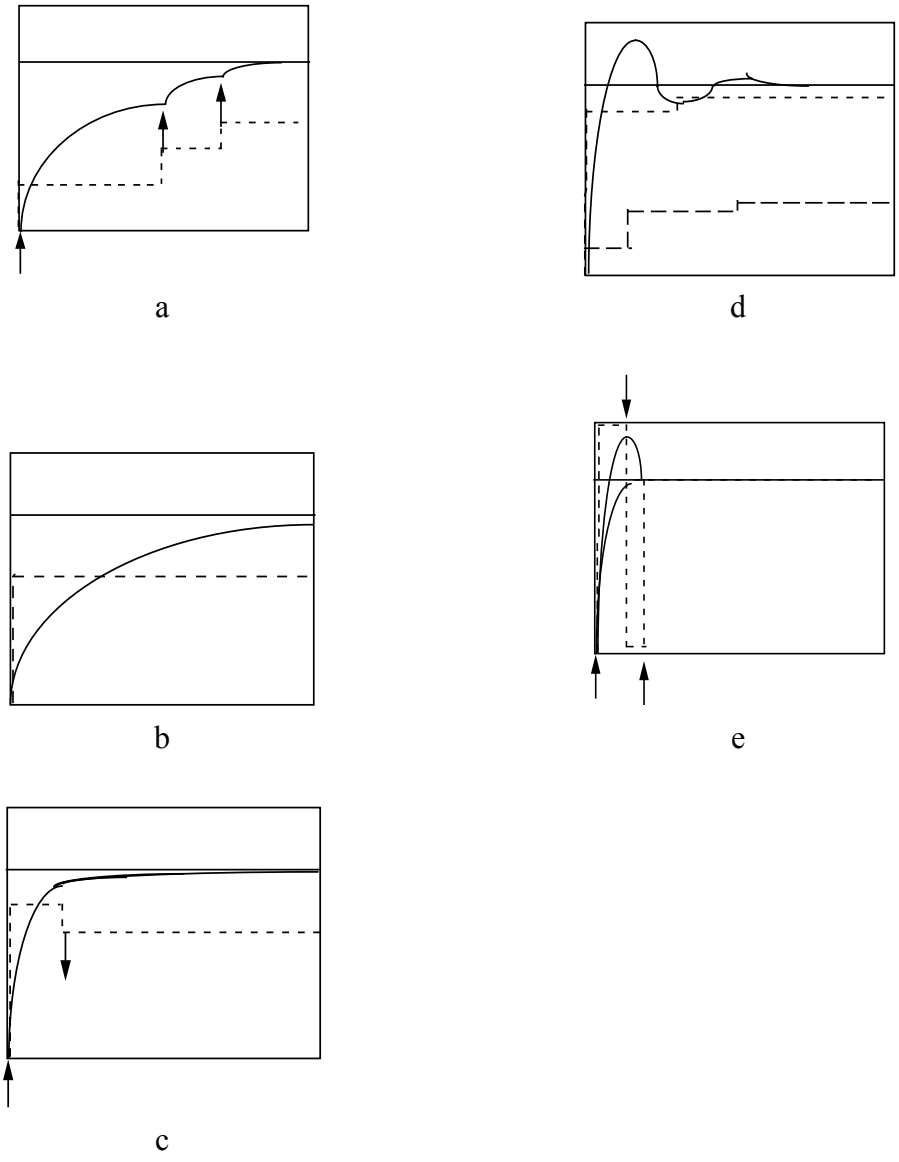

Figure 10. Heating control strategies

For the following definitions, these four abbreviations below will be used to classify the control strategies into their distinctive groups.

H	Heater control action. Can be either ¬ or ¯.
T <> T_{GOAL}	Indicates current temperature (compared with the goal temperatures of 20˚C and 40˚C).
T	Indicates current direction of temperature change. Can be either ¬ or ¯.
Mass	Indicates using system valves (input or output) to affect temperature. Can be either ¬ or ¯.

Proportional feedback control (PFB). Figure 10a illustrates the proportional feedback control strategy. This method of control is based on the current state of the system. It is a reactionary control input for when the current system state does not equal the goal state. Because this strategy is not adapted to the dynamics of the environment, it can lead to instability in control (e.g., oscillations). If the following statements are true, the operator will be using the Feedback Control strategy.

if ($T > T_{GOAL}$) and (H¬) or if ($T < T_{GOAL}$) and (H ¯)

Using this strategy, the operators adjust the heater setting in increments, waiting to see if the optimal temperature has been reached. If this temperature has not been reached, they need to readjust the heater setting again, and again wait to see if the current temperature meets the goal. Presumably, operators might make increasingly fine adjustments as the goal temperature is being reached.

While the accuracy to which the desired temperature setting may be achieved is a little lower than other strategies (because the operator has to guess, to some degree, the optimal heater setting), once this goal is achieved, the operator can leave the setting alone. The best use of this strategy is if there is a fairly long delay between action by the operator and a reaction from the system.

Model-based, open loop. This heating strategy is illustrated in Figure 10b. This strategy requires the operators to examine both the given demand flow rate and the incoming temperature of the water and to use this information for computing the heater

setting. This strategy conforms to the constraints of the system, and does not try to overrule them, which leads to stable control.

An advantage of this strategy is that operators do not have to further monitor the system after the change has been made. Therefore, they can direct their attention to other tasks if necessary. However, a problem with this strategy is that is does not necessarily achieve the desired temperature as quickly as possible. It requires the system to respond to the new setting and change until a constant temperature plateau has been reached. If there are any disturbances within the system, the desired energy level of the reservoir may not be reached. Also, using this strategy requires the operators to either solve complicated computations or to have a lot of experience controlling the system before the system dynamics can be internalized.

Derivative feedback control (DFB). Figure 10c illustrates the derivative feedback control strategy. In this strategy, the dynamics of the system (reactive forces) are exploited and used for more flexible and stable control. Thus, if the following conditions are true, the operators will be exhibiting anticipatory behaviour.

if ($T\neg$), ($T > T_{GOAL}$), and ($H\ \overline{)}$ or if ($T\ \overline{)}$, ($T < T_{GOAL}$), and ($H\neg$)

The operators begin control as in the proportional feedback strategy, but can actually use a very high setting at first, gradually lowering the setting as the temperature of the water in the reservoir is approaching the goal temperature to achieve a faster response time. This requires the operators to anticipate the time it will take for the temperature to reach the goal and to act accordingly. While it may take a while for the operators to internalize the dynamics of the system in terms of heater time lags, this strategy may be very successful and economical when employed.

Multi-Variate Control (MV). Figure 10d illustrates this strategy. In this figure, the upper dashed line represents the heater control movements, while the lower dashed line represents the flow increases. In this strategy, mass is used either as an information variable (as in the first case below), or as a control variable. The first condition below is usually only met when a change in demand setpoints has occurred. For example, if the setpoints suddenly increased greatly, the operators would most likely increase the amount of flow both into and out of the reservoir. The increased inflow of 10˚ C water, which keeps the reservoir from completely draining, and the increased outflow needed to meet the demand will decrease the temperature of the water in the reservoir. Thus, even

though the reservoir temperature was greater than the goal temperature and was also increasing, changing the flowrates could ultimately cause the temperature to fall below the goal temperature. To avoid this, the operators can increase the heater setting first, before changing the flowrates, in an attempt to maintain the system temperature within the goal range. The second two conditions require the operators to change the mass or flowrate of the reservoir to affect the temperature. This is usually done when the operators are having an extreme amount of trouble controlling the system and want to quickly change the temperature of the reservoir.

if ($T\ \bar{}$), ($T > T_{GOAL}$), and ($H\ \bar{}$) or if ($T\ \bar{}$), ($T > T_{GOAL}$), and ($Mass\ \bar{}$)
if ($T\neg$), ($T < T_{GOAL}$), and ($Mass\neg$)

One problem with this strategy is that it is cognitively more complex to control than the others. In this case, there is an additional degree of freedom which must be managed to achieve the desired temperature. Also, this strategy may require volumes which are approximately in the middle of the reservoir to allow for the increased/decreased mass without overflowing/emptying the reservoir. To use this strategy successfully, the operators must have a keen sense of the dynamics of the system. Otherwise, the operator-reservoir system could behave unstably, exhibiting larger and larger oscillations. Also note that the first example ($T\ \bar{}$), ($T > T_{GOAL}$), and ($H\ \bar{}$) is useful only when there is a change in demand setpoints.

Bang-Bang Control: A Special Case of Feedback. As seen in Figure 10e, this strategy involves turning the heaters on full (dashed line), which increases temperature at the fastest rate (upper solid line). When the temperature (lower solid line) is nearing the goal state (solid horizontal line), the operators must turn off the heater. Once the temperature is within the goal region, the operators turn on the heater to the point where energy flowing into the reservoir equals the energy flowing out of the reservoir (EI = EO).

This strategy is the fastest for attaining goal temperature and is best explained with respect to the P+F interface. The operators turn the heater on full and when the temperature gets near the goal area, the heater is turned off. The heat transfer rate will be greater than needed at this point, and temperature will begin to slowly decrease its positive rate of change. Turning the heater off reduces heat transfer at the fastest rate. Just before EI = EO (energy is stabilized), and the temperature is still within the goal area, the heater is turned on to the point where the heater setting matches the current heat

transfer rate displayed on the interface. This will establish a correct temperature, with a balance between energy input and energy output, resulting in a vertical line in the energy balance representation for the P+F interface.

Operators using the P interface are greatly disadvantaged when trying to use this strategy because they are not provided with the visual feedback needed to perform this correctly. In fact, it is extremely difficult to use this strategy as the heat transfer rate and the energy gradient are unknown to the operators. If the operators are not very good at anticipating when the temperature will reach the goal area, this strategy may involve minor adjustments to the heater (using one of the other strategies).

Which strategy is used will depend on how well the operators understand the dynamics of the system. This understanding may come from operator experience or by effective interface design. Obviously, the bang-bang strategy is the most effective, but it is not very easy to adopt with the P interface.

4.4. Shut-down Strategies

For shut-down , two different control strategies have been identified. One of these strategies consists of two phases, the first of which is an initial proactive structuring of the system in which like components, for each reservoir or feedwater stream, are approximately matched in values. For example, if a subject were to close the output valve VO2, he would also close valve VO1. In the second phase of this strategy, control actions are concurrently performed on each pair of like components, regardless of the state of the individual reservoir or feedwater stream. An advantage of this control strategy is increased cognitive economy. Using this strategy, a subject may not have to keep track of both configurations, thus reducing the visual sampling and memory load needed for control. The parallel actions only require the subject to calculate the appropriate component values once for each set of like components. Therefore, the proactive creation of structure reduces subsequent control demands. One potential disadvantage of this strategy is that the extra control action on the second component, in order to match like components, may be unnecessary. Another potential disadvantage may be realized when the initial reservoir states are not equal (e.g., volume for R1 is high and volume for R2 is low). In this case, depleting both reservoirs at the same rate would lead to a system failure, as R2 may completely empty with the heater still in operation.

For the other control strategy, no initialization phase is present. Control actions are independently tailored to each individual reservoir state and as a result, the values for

like components pertaining to each reservoir or feedwater stream do not match. In addition, control actions on like components do not happen at the same time. This type of control strategy may be viewed as more efficient because control is tailored to the current system state. However, there is an increased cognitive load, and possibly an increased visual sampling rate required of the subject.

4.5. Fault Management Strategies

There are two main types of strategies available for fault management. The first strategy involves root-cause diagnosis. Using this strategy, operators attempt to understand what the fault is and why it has occurred, before trying to compensate. Using the second strategy, operators simply compensate for the observed fault first, with little regard for the reasons why it occurred. After compensation, then the operators may attempt to diagnose the cause of the fault.

A root-cause diagnostic strategy involves examining the system to gain an understanding of what has caused the fault, before compensation begins. This search can be either symptomatic (such as hypothesis-test, pattern recognition, etc.) or topographical (searching along the various breakdowns of the system (means-end, part-whole, topological) for relationships between either physical or functional variables). For more information on this type of fault management strategy, see Rasmussen (1981 and 1986).

The second type of fault management strategy involves compensating for the fault first, and leaving the diagnosis for a later time. To explain this type of strategy, consider the case when a valve becomes blocked. The operators can see that the reservoir level is decreasing, since a blocked valve results in no flow. If this output valve is still open, the water in the reservoir would continue to leave the system. The operators can simply reconfigure other input valves so that the input into the reservoir equals the output, thereby compensating for the fault. After this is accomplished, and the system is once again at steady-state, then the operators can then attempt to understand the fault event, in terms of the affected system constraints, to determine the reasons why the fault may have occurred.

5. LEVELS OF COGNITIVE CONTROL

The objective of this part of the CWA is to determine what competencies operators need to control the DURESS II system. Rasmussen's (1983) taxonomy for

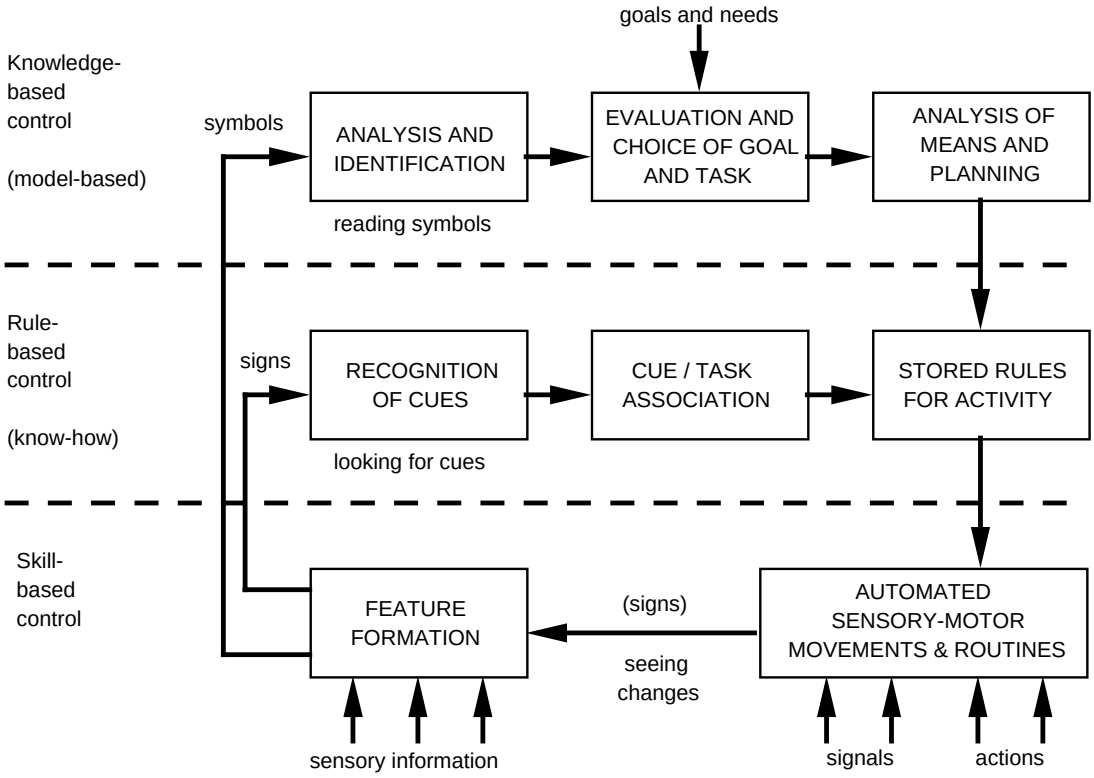

Figure 11. Information flow in SRK framework (from Rasmussen, 1983)

levels of cognitive control will be used to conduct this analysis. Operators can control the system by using the three levels of cognitive control specified in this taxonomy. Figure 11 on the following page shows the flow of information between these three levels of control. Note, however, that this figure does not present any information that is specific to DURESS II, but merely serves as a basis for understanding the relationships between the three levels of cognitive control. Presented below is a discussion of these three levels of control (SBB, RBB, and KBB) and how they apply to DURESS II.

5.1. Skill-Based Behaviour

The level of SBB represents the highly skilled sensory-motor control performance directed by automated patterns of movements. Sequences of such sub-routines will be controlled by stored rules and/or activated by what Rasmussen (1983) refers to as signals. The sensory-motor control patterns at the SBB level will only be concerned with the manipulation of items on the control panel, or interface surface--which are the signals (of system state) presented to the operators. Thus, the only operator requirement for SBB is the ability to use a computer mouse to "point and click" on the components that are being manipulated.

5.2. Rule-Based Behaviour

This level of control involves the operators using the changes in the display as signs that a certain control procedure or set of rules are to be used. The interface should provide signs that operators can use as cues for the selection of an appropriate action. The action alternatives consist of a set of operating procedures and routine control strategies that the operators can be trained to execute when certain cues occur. An example of this is the sign that the system needs to be shut-down. The "sign" in this case is the two reservoir demand levels changing to zero. At this sign, the operators should know that the system needs to be shut-down and that certain shut-down procedures should be followed. These procedures are typically learned either through experience or training.

5.3. Knowledge-Based Behaviour

For the effective utilization of the abstract reasoning that is typical of knowledge-based behavior, the operators must have good mental models of the processes. In order to support this level of cognitive control, the system model must be represented at several different conceptual levels representing components, physical processes, or general functions (see the abstraction hierarchy discussion above) depending on the situation or

system state (Rasmussen, 1983). Thus, the operators must have a correct and complete mental representation of the abstraction hierarchy for DURESS II, to effectively control the system at this level. Figure 12 shows the relationships between system process, the interface, and an operator's mental model.

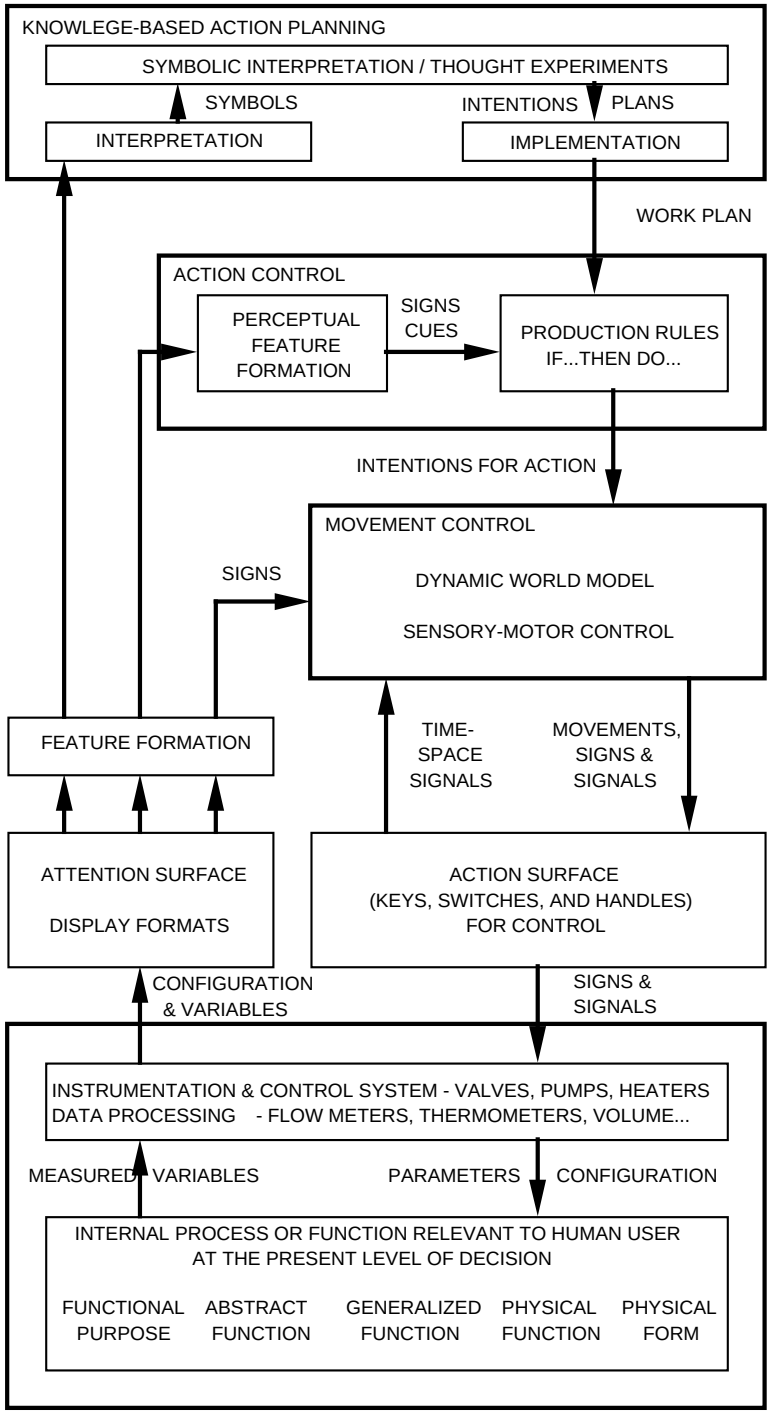

Figure 12. The mapping between process, interface, and operator mental model. (from Vicente and Rasmussen, 1988)

REFERENCES

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

Pawlak, W. S. and Burns, C. M. (1993). Cognitive work analysis of the DURESS system. Unpublished manuscript, University of Toronto, Toronto, ON.

Rasmussen, J. (1976). Outlines of a hybrid model of the process plant operator. In T. B. Sheridan and G. Johannsen (Eds.), Monitoring behavior and supervisory control (pp. 371-383). New York: Plenum.

Rasmussen, J. (1981). Models of mental strategies in process plant diagnosis. In J. Rasmussen and W. B. Rouse (Eds.), Human detection and diagnosis of system failures (pp. 241-258). New York: Plenum.

Rasmussen, J. (1983). Skills, rules, knowledge: Signals, signs, and symbols, and other distinctions in human performance models. IEEE Transactions on Systems, Man, and Cybernetics, SMC-13, 257-267.

Rasmussen, J. (1985). The role of hierarchical knowledge representation in decisionmaking and system management. IEEE Transactions on Systems, Man, and cybernetics, SMC-15, 234-243.

Rasmussen, J. (1986). Information processing and human machine interaction: An approach to cognitive engineering. New York: North-Holland.

Vicente, K. J. (1987). DURESS: A domain for cognitive engineering research. Unpublished manuscript, Risø National Laboratory, Roskilde, Denmark.

Vicente, K. J. (1991). Supporting knowledge-based behaviour through ecological interface design. Unpublished doctoral dissertation, University of Illinois at Urbana-Champaign, Urbana, IL.

Vicente, K. J. (1992). Multilevel interfaces for power plant control rooms I: An integrative review. Nuclear Safety, 33, 381-397.

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