ECOLOGICAL PSYCHOLOGY, *16*(4), 255–285

What Does Computer-Mediated Control of a Thermal-Hydraulic System Have to Do With Moving Your Jaw to Speak? Evidence for Synergies in Process Control

John R. Hajdukiewicz
ACS Advanced Technology
Knowledge Systems Lab, Honeywell
Minneapolis, MN

Kim J. Vicente
Cognitive Engineering Laboratory
Department of Mechanical & Industrial Engineering
University of Toronto
Toronto, Canada

Coordination phenomena can take many diverse forms, but ecological psychologists have focused primarily on understanding human motor control. In this article we report an experiment on human–machine coordination that was designed to replicate and extend an early experiment on human jaw movement during speech production that provided initial evidence of synergies. Participants controlled a thermal-hydraulic process simulation for about 1 hr per weekday for approximately 1 month. Half of the participants used a human–computer interface that presented predominantly lower level physical (P) information, whereas the other half used an interface that presented higher level functional (P+F) information as well. During the last block of trials, local perturbations were introduced by increasing the time constant of a particular component per trial by a factor of 20. The component perturbations had less impact on the performance of the P+F participants than the P, and

Requests for reprints should be sent to Kim J. Vicente, Department of Mechanical & Industrial Engineering, University of Toronto, 5 King's College Road, Toronto, Ontario, Canada M5S 3G8. E-mail: vicente@mie.utoronto.ca

this effect was mostly localized to components that had alternative degrees of freedom for control. Most important, the P+F participants exhibited more evidence of higher level control than the P, providing some initial evidence for synergies in process control. These findings suggest that it may be possible to develop a general unified theory of coordination that subsumes motor control and human–machine interaction as special cases.

The word *coordination* means many things to many people, but in ecological psychology, it almost always refers to the complex phenomenon of goal-directed human movement. Students of motor coordination have found the concept of *synergies*—task-specific functional units consisting of an assembly of component structures—to play an important role in their research (Bernstein, 1996; Turvey & Carello, 1996). For example, Kelso, Tuller, Vatikiotis-Bateson, and Fowler (1984) examined the coordination of jaw muscles during speech and found evidence for synergies: Participants were able to demonstrate robust control in the face of novelty and change, and this control was achieved by keeping higher level goals constant and by modifying lower level processes in a way that was specific to both the task goal and the nature of the perturbation. Subsequent work in motor control has extended this line of thinking in a number of significant ways—for example, by mathematically modeling the form of particular task synergies and by investigating the development of synergies via learning (e.g., Kelso, 1995).

The conjecture behind the research presented here is that coordination may be a general systems phenomenon that transcends human motor control (Hajdukiewicz & Vicente, 2002; Yu, Lau, Vicente, & Carter, 2002). Under this view, mechanical, biological, human–machine, and perhaps even social coordination may be different manifestations of the same underlying phenomenon, all of which could be explained in terms of synergies, albeit of different forms. This article pursues this strategy by investigating computer-mediated human–machine coordination in process control. Our experiment was modeled after that of Kelso et al. (1984) so that we could determine whether mediated human–machine coordination performance showed any of the hierarchical structure and control across levels seen in motor coordination. Participants were asked to control a thermal-hydraulic process simulation for about 1 hr per weekday over 1 month. After they were experienced, various types of disturbances were introduced to examine participants' abilities to use a variety of means to maintain the same goal states. To test for the plausibility of synergies, we also introduced a novel manipulation that was not investigated by Kelso et al. and that, indeed, cannot be easily adopted in motor control. The information presented to participants was varied by using two different human–computer interfaces for the same process—one that provided mainly low-level physical information and another that provided higher level functional information as well. Because synergies are said to rely on the existence of higher level invariants, we expected that the interface with higher level functional information would be more likely to lead to successful adaptation to local conditions observed in motor coordination than the interface with strictly low-level information (Vicente & Rasmussen, 1990).

Many ecological psychologists are accustomed to dealing with abstractions (e.g., based on dynamical systems theory), so they believe that the concept of synergies has a wide scope of applicability. However, this community has focused primarily on investigating motor coordination and thus has not explored the generalizability of synergies to many other forms of coordination such as mediated human–machine interaction. Social coordination has been investigated, but again of motor movement (e.g., Schmidt, Christianson, Carello, & Baron, 1994). Therefore, there is no empirical basis for believing that computer-mediated control of a thermal-hydraulic process simulation would have anything at all to do with unmediated jaw movements during speech. But if coordination is indeed a general systems phenomenon, then we would expect to see the patterns observed by Kelso et al. (1984) in our experiment as well—at least with the interface presenting higher level functional information—providing some initial empirical evidence for synergies in process control.

The potential contributions of this research are not restricted to ecological psychology, however. As far as we know, the concept of synergies has not been adopted before in our home discipline—cognitive engineering. Therefore, the extension of this theoretical construct to a new applied phenomenon may also have implications for research on how to design more effective human–computer interfaces in complex sociotechnical systems.

THEORY AND BACKGROUND

A number of concepts from cognitive engineering (Vicente, 1999) need to be introduced to explain how we were able to operationalize the concept of synergies outside of motor control. The first of these is the *abstraction hierarchy* (Rasmussen, 1985), a framework for modeling complex work domains. The lowest level in the abstraction hierarchy is the set of physical components that can be acted on in the work domain, and the higher levels are the functions and purposes that need to be satisfied. The links connecting adjacent levels of the abstraction hierarchy identify the available structural degrees of freedom that need to be coordinated to achieve effective performance. From this perspective, the control of a generic complex system can be viewed as a coordination problem.

Consistency in control can conceivably be observed at any level of the abstraction hierarchy. For example, as shown in Figure 1a, *higher level control* focuses on higher level functions by using any feasible lower level components in a context-specific way. In this case, the higher level functions and outcome (i.e., product) are consistent across situations, whereas the lower level actions vary from one situation to the next. In contrast, as shown in Figure 1b, *lower level control* focuses on a specific way of using components in a context-independent way, much like a

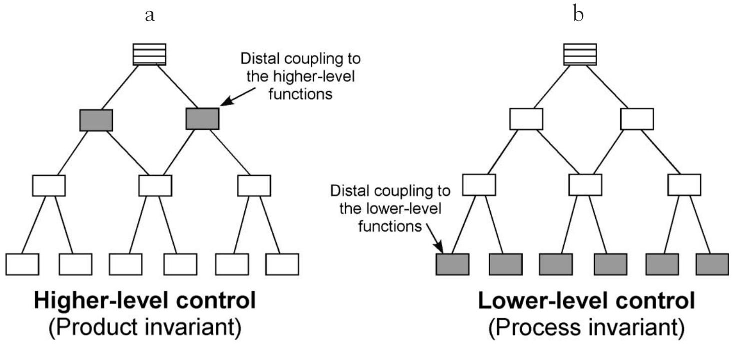

FIGURE 1 Differences between (a) higher level and (b) lower level control (Hajdukiewicz, 2001).

rote recipe. In this case, the lower level actions (i.e., process) are consistent across situations.

Figure 2 shows how these two different types of control would manifest themselves in the face of a hypothetical perturbation to the dynamics of a single component marked "X." The top part of Figure 2 shows that just before the change in component dynamics occurs, both types of control are using the same component (marked with circles). The bottom part of Figure 2 shows what will happen with the two levels of control after the change in component dynamics occurs. With higher level control, the focus is on supporting the higher level functions. A change in component dynamics affects a particular higher level function. An alternative component (marked with a circle) that supports the same higher level function can be used instead to adapt to the unanticipated change in component dynamics. In this case, the consistency in control is at the level of the higher level function, not the lower level actions. This type of adaptation to changes in context would provide evidence for synergies. In contrast, with lower level control, the focus is on low-level actions. Accordingly, the component that was used before the perturbation is used after the perturbation as well, without consideration for the effects on the higher level functions, because the rote process that is used to act on components has not changed. In this case the consistency in control is at the level of component actions, not higher level function. This empirical signature would be evidence against synergies and in favor of something akin to a motor program.

At least two other preconditions must be satisfied for the higher level control, typical of synergies, to be feasible (Hajdukiewicz, 2001). First, there must be structural redundancy in the work domain. If there is only one component for achieving a particular function, then it seems that it would be very difficult to adapt in the face of a perturbation to that component. Second, there must be higher level information to specify the state of affairs so that people can resolve the available degrees of freedom perceptually (as opposed to analytically, using mental heuristics or derivation).

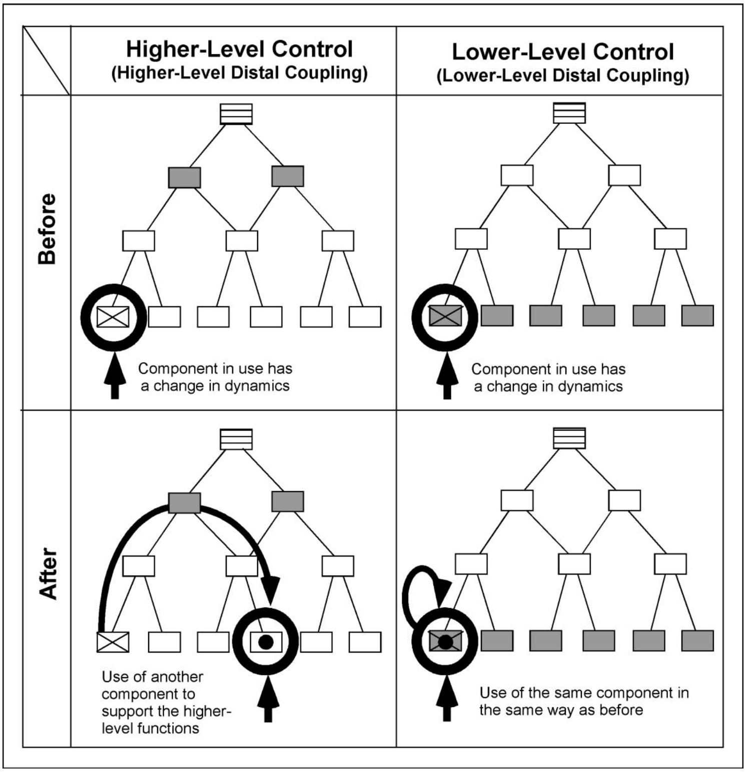

FIGURE 2 A hypothetical example of how higher and lower level control interact with a local change in component dynamics (marked as "X"; Hajdukiewicz, 2001).

Thermal-Hydraulic Process Simulation

To investigate these issues empirically, we used DURESS II, a thermal-hydraulic microworld simulation (Pawlak & Vicente, 1996; Vicente & Rasmussen, 1990). The physical structure of DURESS II is shown in Figure 3 with associated component capacities. DURESS II consists of two feedwater systems that supply water to two reservoirs. Human operators have control over 8 valves (6 input valves: VA, VB, VA1, VA2, VB1, and VB2; and 2 output valves: VO1 and VO2), 2 pumps (PA and PB), and 2 heaters (HTR1 and HTR2), resulting in a total of 12 degrees of freedom. Participants were required to achieve the dual goals of satisfying external, dynamic demands for water (D1g and D2g) while also maintaining each of the reservoirs at their respective temperature goals (T1g and T2g).

Figure 4a provides an outline of the work domain representation that was developed for DURESS II (Bisantz & Vicente, 1994; Vicente & Rasmussen, 1990). There are three levels of resolution in this space connected by part–whole links (System, Subsystem, and Component). Also, there are five levels of abstraction

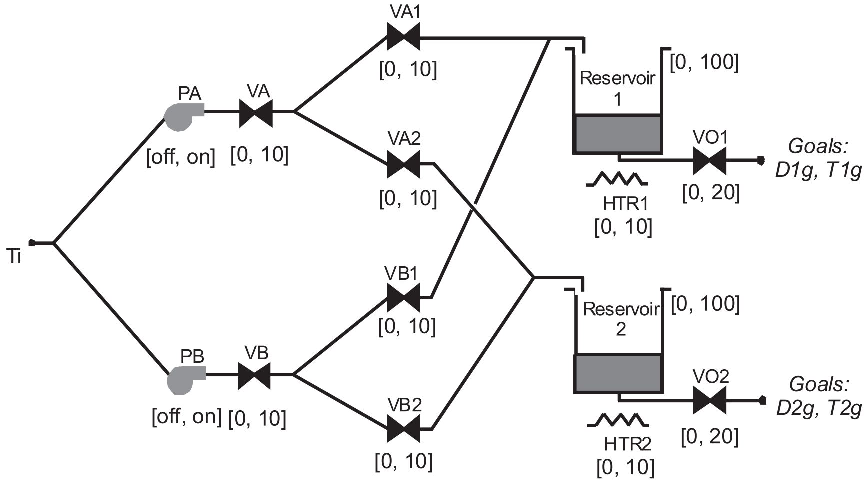

FIGURE 3 A schematic diagram of the DURESS II process control microworld with capacities shown.

connected by structural means–ends links (Physical Form, Physical Function, Generalized Function, Abstract Function, and Functional Purpose). The bottom level of Physical Form is not used here because it refers to the physical location and appearance of the work domain, features that are not particularly meaningful in a microworld simulation like DURESS II.

Figure 4a shows that the abstraction and part–whole dimensions, although conceptually orthogonal, are coupled in practice. At higher levels of abstraction (e.g., Functional Purpose), participants tend to think of the work domain at a coarse level of resolution (e.g., System), whereas at lower levels of abstraction (e.g., Physical Function), participants tend to think of the work domain at a detailed level of resolution (e.g., Component). Therefore, certain cells in the space are not very meaningful (e.g., Functional Purpose/Component). In the specific case of DURESS II shown in Figure 4a, four cells were identified as being useful for the purposes of process control. Each of these cells contains a different representation of the same work domain. Figure 4b lists the variables associated with each representation and specific functional (means–ends) links between adjacent levels of abstraction (Hajdukiewicz, 2000).

Two interfaces were developed for DURESS II (Pawlak & Vicente, 1996; Vicente & Rasmussen, 1990). The P (physical) interface, shown in Figure 5a, represents only lower level physical information in the abstraction hierarchy representation. The P+F (physical and functional) interface, shown in Figure 5b, also includes higher level function information (e.g., flow rates, heat transfer rates, and mass and energy balances), thereby including all of the abstraction hierarchy representation. Functional relationships are shown using emergent feature graphics so that participants can become perceptually attuned to the interface.

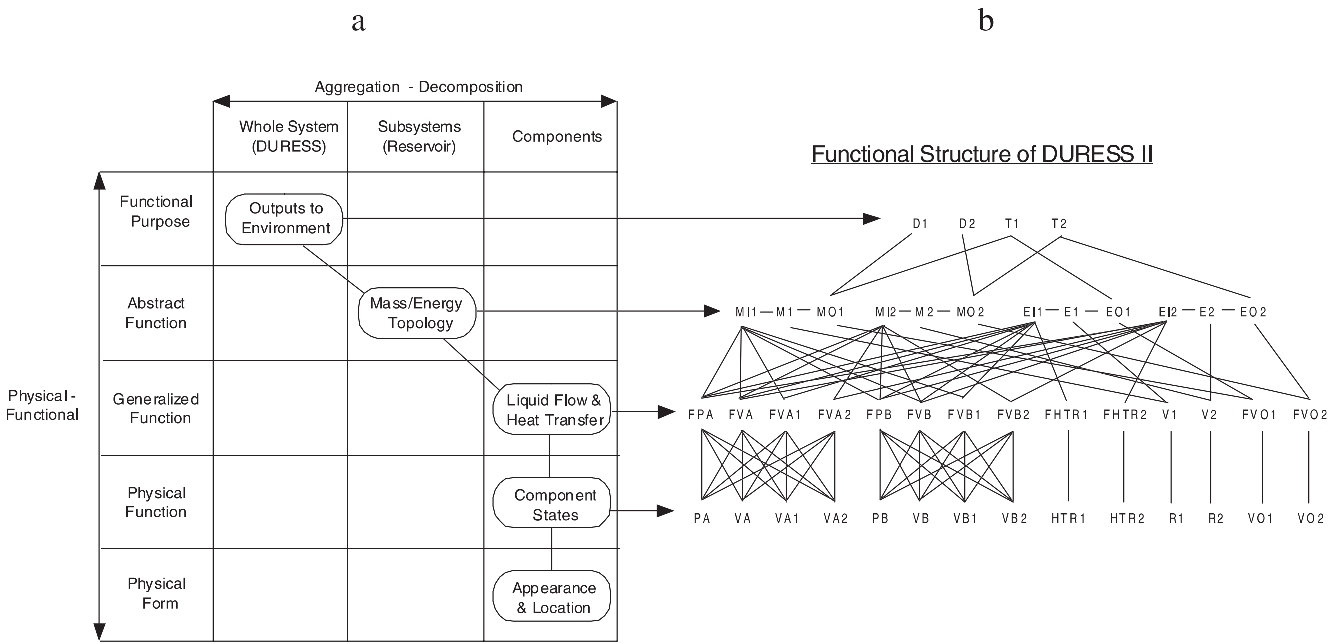

FIGURE 4 Work domain analysis of DURESS II. (a) Adapted from “Making the abstraction hierarchy concrete,” by A.M. Bisantz and K. J. Vicente, 1992, *International Journal of Human-Computer Studies, 40,* p. 90. Copyright 1994 by *International Journal of Human-Computer Studies*. Adapted with permission from Elsevier. (b) Adapted from *Proceedings of the Conference on Human Factors in Computing Systems (CHI 2000)*, by J. R. Hajdukiewicz, 2000. Copyright 2000 by J. R. Hajdukiewicz. Adapted with permission.

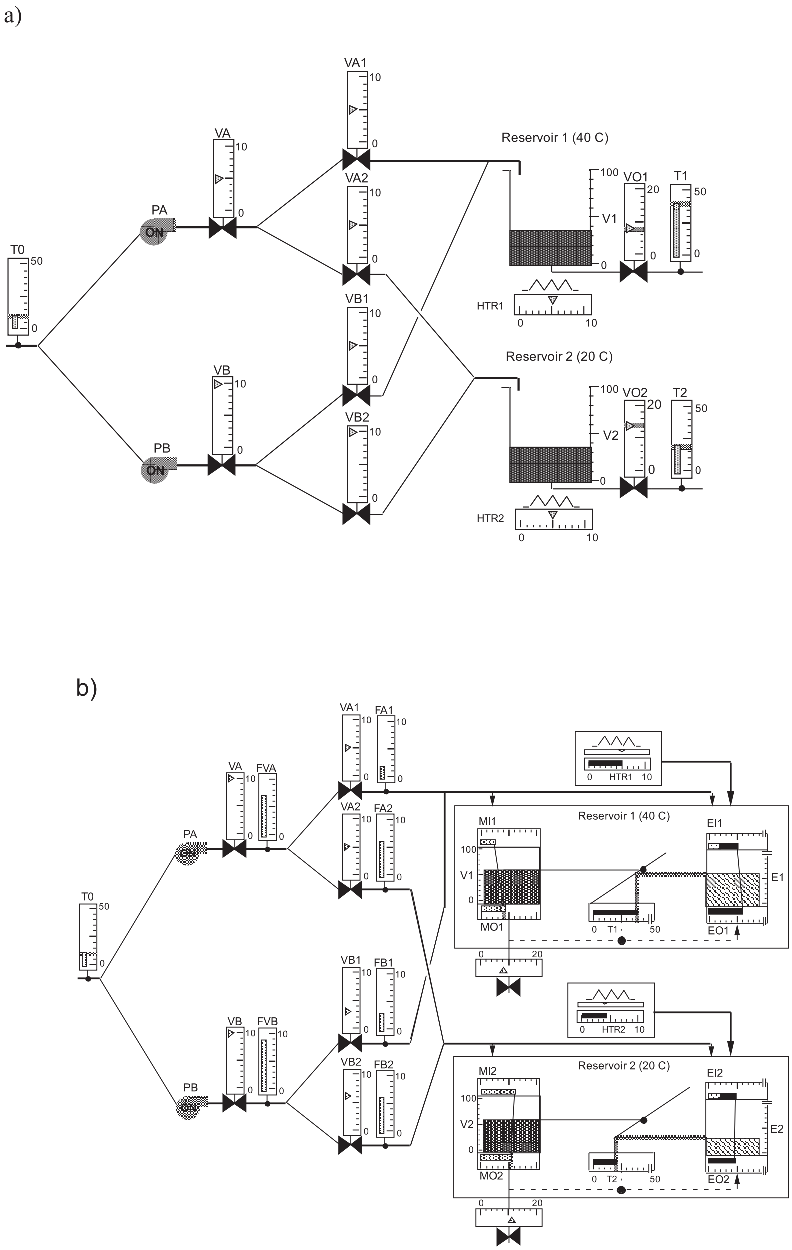

FIGURE 5 (a) P interface and (b) P+F interface. Reprinted from *International Journal of Human-Computer Studies, 44,* W. S. Pawlak & K. J. Vicente, "Inducing Effective Operator Control Through Ecological Interface Design,", pp. 633–688, Copyright (1996), with permission from Elsevier.

The predicted impact of structural degrees of freedom on coordination can be illustrated in the context of the P and P+F interfaces for DURESS II. The input valves provide structural redundancy in the supply of mass flow into the two reservoirs. For example, two valves can be used to supply water to each reservoir (VA1 and VB1 for Reservoir 1, VA2 and VB2 for Reservoir 2), assuming the pumps, VA, and VB are on and open. This redundancy provides flexibility in choosing components to support the higher level functions and maintaining the same higher level trajectory over time. Also, to some extent the input valves can serve as indirect means for cooling water. Thus, the use of synergies is possible because different valves can be used to support higher level functions, resulting in higher level control.

Consider the specific example of a switch in feedwater configuration strategy at the onset of an input valve perturbation (Figure 6). Three strategies can be used for feedwater configuration: single, decoupled, and full (Vicente, 1999). The single strategy uses only one pump and the valves connected to it to fill both reservoirs. The decoupled strategy uses both pumps and their connected valves to feed different reservoirs. The full strategy uses both pumps and their connected valves to feed both reservoirs. If a participant initially used a decoupled strategy and one of the input valves (VB2 in Figure 6) suddenly became slow to respond, a switch to a single or full strategy while minimizing the use of the affected valve would be evidence for higher level control because the same higher level functions (MI1 and MI2 in Figure 6) would be maintained using different input valves (VA1 and VA2 for the single strategy; VA1, VA2, and VB1 for the full strategy).

In contrast, in the case of the heating function and the output demand function, there is only one structural degree of freedom—the heater and the output valve, respectively. Note, however, that there are dynamic degrees of freedom that

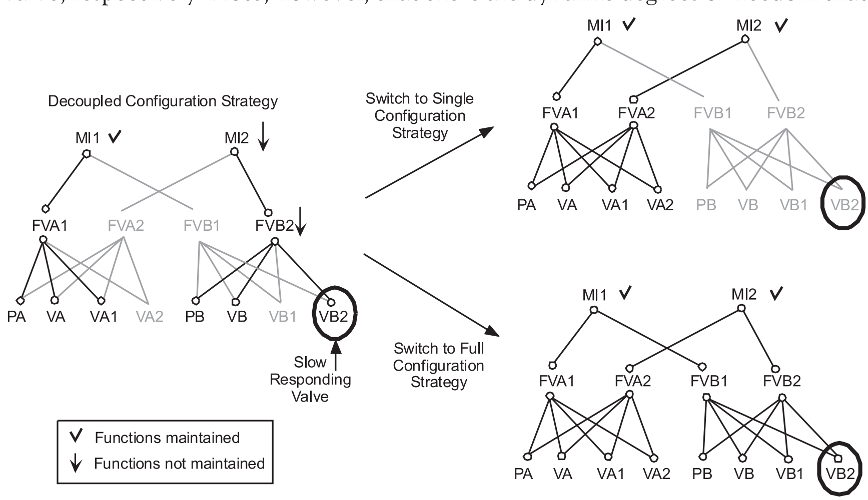

FIGURE 6 Evidence for higher level control through feedwater configuration strategy switching (Hajdukiewicz, 2001).

could potentially be exploited, providing a basis for higher level control. For example, the output valve can be set at its maximum setting to increase the rate at which the demand flowrate increases, and then that setting can be reduced when the flowrate approaches its desired value. Such a strategy exploits the first-order dynamics of the components to achieve steady state more quickly.

METHOD

This study investigated the impact of local perturbations in component dynamics on level of control and the ability of participants to succeed in adapting to change, as well as the impact of structural characteristics of the work domain on adaptation to the local perturbations (i.e., single vs. multiple structural degrees of freedom).

Hypotheses

If higher level functional information provides better support for participant adaptation to change and novelty, then one would expect that performance with the P+F interface would be less affected by the local perturbations than that with the P. However, this advantage is minimized when there are fewer structural degrees of freedom and less flexibility in compensating for local changes. Therefore, perturbations of the input valves with multiple structural degrees of freedom are expected to have a greater impact on the performance advantage of the P+F over the P interface compared to perturbations of heaters or output valves, both of which have only a single structural degree of freedom. Furthermore, for those participants who are able to adapt successfully with the P+F interface, one would expect that they would rely on higher level control to do so.

Experimental Design

A repeated measures mixed design was adopted with interface (P or P+F) as a between-participants variable and trial type (normal and perturbation) as a within-participant variable. The study consisted of 80 trials divided into 4 matched blocks of 20 trials. The first 3 blocks were the same (with one exception noted following) and provided an opportunity for participants to learn how to control the microworld under normal conditions. During this phase, the time constants were 5 sec for pumps and valves and 15 sec for heaters. During the last block, local changes to the component time constants were introduced. The time constants of specific components were randomly changed, one per trial, by multiplying their normal value by a factor of 20 (i.e., 100 sec for pumps and values and 300 sec for heaters). The perturbed components were the inflow valves VA1 and VB2, heaters HTR1 and HTR2, and output valve VO2; the input valves had multiple structural degrees of freedom, and the heaters and output valve had a single structural degree of freedom. This change made the affected component very slow to respond. Finally, the early trials had two local perturbations (VA1 in Trial 2, HTR1 in Trial

4) that were matched with the same local perturbations in late trials (VA1 in Trial 62, HTR1 in Trial 64).

Participants

There were eight paid participants in each interface group. All of them were university engineering students who were primarily selected based on their willingness to participate, their degree of relevant formal training (i.e., 1–3 courses in fluid mechanics and thermodynamics), and their cognitive style (see following).

Apparatus

The interactive, dynamic simulation (DURESS II) was programmed on an SGI computer with an IRIX operating system. Information was presented on a 19 in. high-resolution color graphics monitor. Control actions were input using a computer mouse. Each time an input was made, the state of all of the simulation variables was automatically logged for data analysis.

Task

For each trial, participants were required to start up the process and achieve steady state by satisfying the four target conditions simultaneously (i.e., two outflow demands and two temperature goals) for 5 consecutive min within a 30-min time limit.

Procedure

In the introductory session, participants read an explanation of the purposes of the experiment and filled out a consent form and an initial questionnaire. They then completed the Spy Ring History test (Pask & Scott, 1972) so that their cognitive style could be assessed. A *serialist* cognitive style reflects a propensity to focus on isolated facts, whereas a *holist* cognitive style reflects a propensity to focus on integration of relationships across facts. A *versatile* cognitive style reflects an ability to adopt either a serialist or a holist style to suit the situation. Previous research with DURESS II showed that this measure is a significant predictor of individual differences in task performance (Torenvliet, Jamieson, & Vicente, 2000), so the Spy Ring History test results were used here to assign participants to the two interface groups, thereby controlling for cognitive style.

Participants were then given a tutorial on the components of the microworld (independent of the interface). After completing the tutorial, each participant was given a brief exercise based on the tutorial. Next, the specific interface assigned to each participant (P or P+F) was introduced; the instructions described only the elements of the display, not the functioning of the simulation.

In each subsequent session, participants performed several trials requiring them to control the simulation. Each session lasted approximately 1 hr per day. Each participant operated the simulation for approximately 25 hr in total, resulting in approximately 400 total participant-hours of data collection. Before their first trial, and immediately after both the 60th and 80th trials, participants were asked to write out a control recipe—a knowledge elicitation measure that tries to identify participants' control strategies (see following). At the end of the experiment, each participant was debriefed individually.

Dependent Variables

The main measure of performance was trial completion time. This measure was defined as the time to reach a steady state condition as defined by the task previously described starting from the initial shutdown state (i.e., all pumps, values, and heaters were off, and the reservoirs were empty).

The degree to which participants' responses on perturbation trials exhibited successful adaptation to the disturbances (and thus evidence for synergies) was assessed using a set of across-level state patterns that we describe in detail in the Appendix. A different set of criteria was developed for each type of perturbation, reflecting the context-specific nature of synergies.

A knowledge elicitation measure—a control recipe (Irmer & Reason, 1991)—was also administered. Participants were asked to write down a set of instructions describing how they controlled the process. These instructions were to be sufficiently detailed to allow someone who had never seen the process before to control it in the same manner as the participant.

Data Analysis

Wherever meaningful, statistical significance was assessed using 95% confidence intervals because these provide the same information as traditional null hypothesis tests, such as analysis of variance, as well as additional important information (see Loftus, 1993, for details). If the confidence interval bars for two means do not overlap, then the difference between means is statistically significant at the $p < .05$ level. Also, when appropriate for small sample sizes, the nonparametric Fisher test for exact probability was performed, and statistical significance was assessed at the $p < .05$ level. Data were analyzed using several statistical software packages (Matlab, Microsoft Excel, and SAS).

RESULTS

Trial Completion Time

Figure 7 shows the effect of local perturbations on average trial completion times for each interface group. For the nonperturbation trials in the last block of 20 trials, the P+F interface was not significantly faster than the P interface. More important, for the perturbations trials affecting components with multiple structural degrees of freedom (i.e., the input valves VA1 and VB2), there were large and statistically significant differences between the interface groups; these perturbations negatively impacted the performance of the P group to a higher degree than P+F.

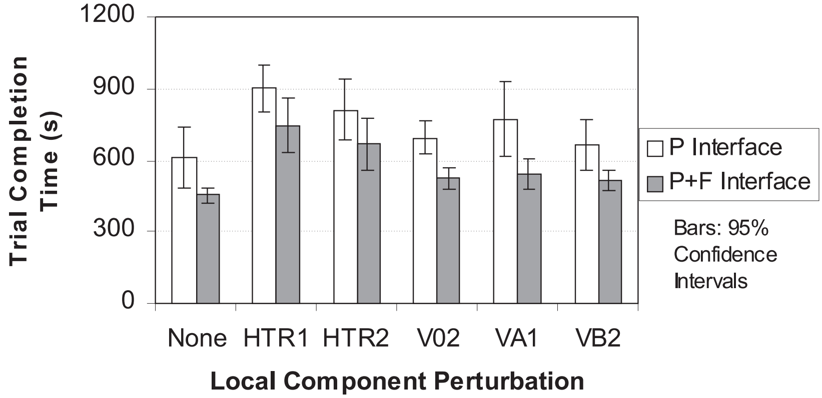

FIGURE 7 Effect of local perturbation on trial completion time.

In contrast, for the perturbations affecting components with a single structural degree of freedom, the results were mixed. Perturbations at the heaters (HTR1 and HTR2) showed small and nonsignificant differences between the two interface groups; both interface groups were equally and negatively impacted by the heater perturbations. Perturbations at the output valve (VO2) showed larger and significant differences between the two groups; the P group was disrupted more than the P+F by the output valve perturbations.

State Patterns

In these analyses, evidence for higher level control was assessed by comparing qualitative patterns in the data with reference patterns indicating level of control for each perturbation (see Appendix).

Output valve perturbation (VO2). Figure 8 shows the number of participants who exhibited the pattern of higher level control for the VO2 perturbation (see Figure A1, column B). For the first perturbation, three of the eight P+F participants exhibited this behavior compared with one of the eight P participants; this difference was not significant (nonparametric Fisher test for small sample size: $p =$ .569 exact, two-tailed). In later VO2 perturbation trials, seven of the eight P+F participants exhibited this behavior compared with none of the eight P participants; this difference was significant (nonparametric Fisher test for small sample size: $p = .001$ exact, two-tailed). For trials that did not have this perturbation, none of the participants from both groups showed this behavior.

These results provide evidence that almost all P+F participants were using higher level control for VO2 perturbations because they adjusted their actions in a context-specific fashion to support the higher level variable of mass outflow while achieving the target conditions in a timely manner. Almost none of the P participants exhibited this behavior.

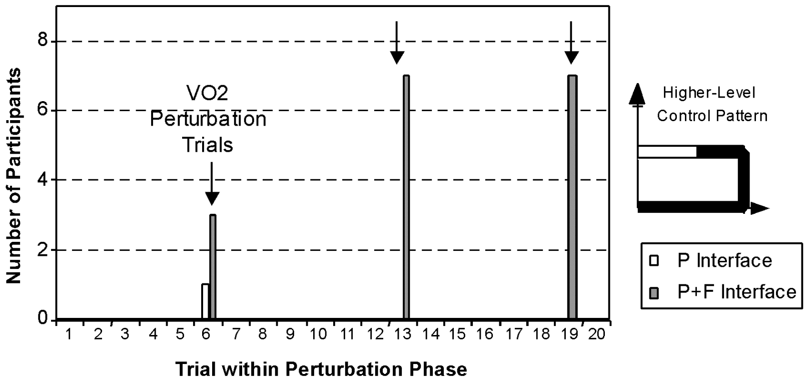

FIGURE 8 Number of participants exhibiting higher level control behavior for VO2 perturbations based on the criteria for higher level control in Figure A1.

Heater perturbations (HTR1 and HTR2). Figure 9 shows the number of participants exhibiting the patterns for higher level control for each heater (see Figure A1, column B). Both P and P+F participants had evidence of higher level control at the perturbation trials of HTR1 (Trials 4, 9, and 14 in Block 4) and HTR2 (Trials 5, 11, and 17 in Block 4), but not on most other trials. P+F generally had more participants exhibiting this behavior overall compared with P; however, the differences were not significant using the nonparametric Fisher test.

The results suggest that most participants in both groups were trying to use higher level control during the heater perturbations. P participants had information about heater setting and temperature only, whereas P+F participants additionally had functional information (heat flow and energy). This fact may explain the slight difference between the two groups in terms of number of people exhibiting this behavior. Also, there were three P participants who derived heuristics to determine the steady-state heater settings based on the outflow requirements (i.e., target regions). These participants used these heuristics as reference points for controlling temperature; they exploited the dynamics of the heater settings to increase temperature quickly and then placed the heater setting to the steady-state value when the temperature was in the target region. This resulted in a higher level control pattern for heater perturbations.

Input valve perturbations (VA1 and VB2). Figure 10 reports the number of participants who exhibited higher level control by either exploiting the input valve dynamics (criteria from Figure A1, column B) or by not using the perturbed input valve. As expected, the P+F group had more participants who exhibited the higher level control patterns at the perturbation trials of VA1 (Trials 2, 10, and 16 in Block 4) and VB2 (Trials 7, 12, and 18 in Block 4) compared with P. This suggests that more P+F participants were using higher level control during the input

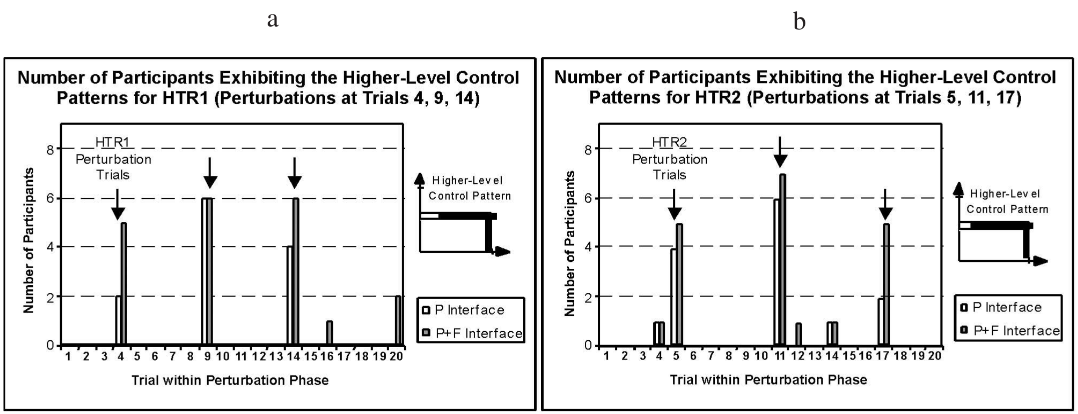

FIGURE 9 Heater utilization pattern for perturbation trials (a) HTR1 and (b) HTR2 based on the criteria for higher level control in Figure A1.

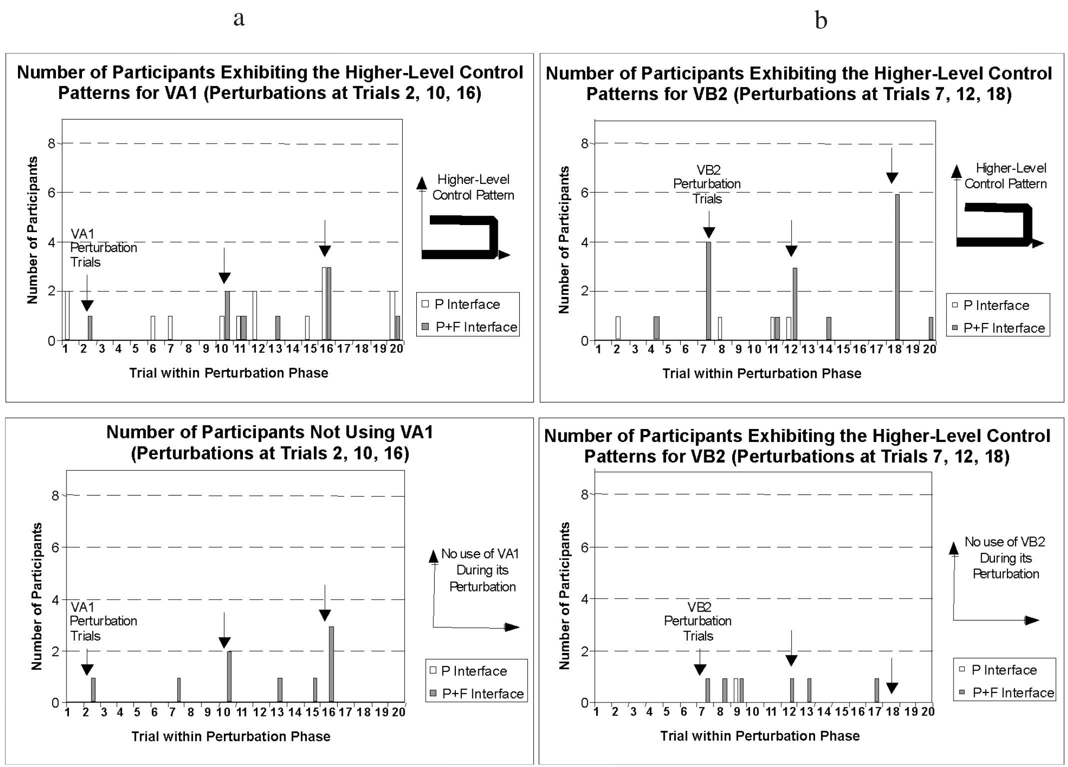

FIGURE 10 Number of participants who either exploited the component dynamics (top) or did not use the perturbed valve (bottom) for (a) VA1 and (b) VB2 perturbation trials based on the criteria for higher level control in Figure A1.

valve perturbation trials than P based on exploiting component dynamics or by not using the perturbed valve.

Figure 11 shows the number of participants who exhibited higher level control by switching feedwater configuration strategies based on the criteria from Figure A2. As expected, the P+F group had more participants who switched configuration strategies at the perturbation trials of VA1 and VB2 compared with P. This suggests that more P+F participants were using higher level control during the input valve perturbation trials than P based on switching feedwater configuration strategies.

Integrating the results from Figures 10 and 11, the number of participants who exhibited higher level control by either exploiting the input valve dynamics or switching configuration strategies is shown in Figure 12 for the input valve perturbation trials. These trials provided the most flexibility in reconfiguring DURESS II because there were multiple ways of supporting the higher level function of mass flow into the reservoirs. As expected, most of the P+F participants showed patterns consistent with higher level control, whereas most of the P participants did not. Using the nonparametric Fisher exact probability test (two-tailed), the differences between P and P+F were significant for Trials 67 (p = .007) and 70 (p = .041) and not significant for Trials 62 (p = .119), 72 (p = .132), 76 (p = .608), and 78 (p = .132). The results suggest that more P+F participants used higher level control during input valve perturbations than P. One plausible reason is that the P+F interface displays the functional information of flow rate through the valves and mass inflow to the reservoirs. P+F participants may perceive more readily any

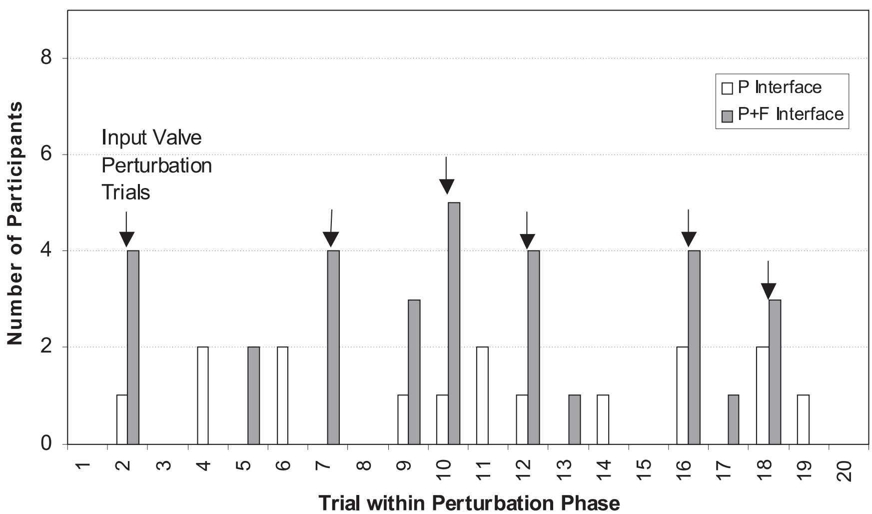

FIGURE 11 Number of participants who switched configuration strategies based on the criteria for higher level control in Figure A3 for input valve perturbation trials.

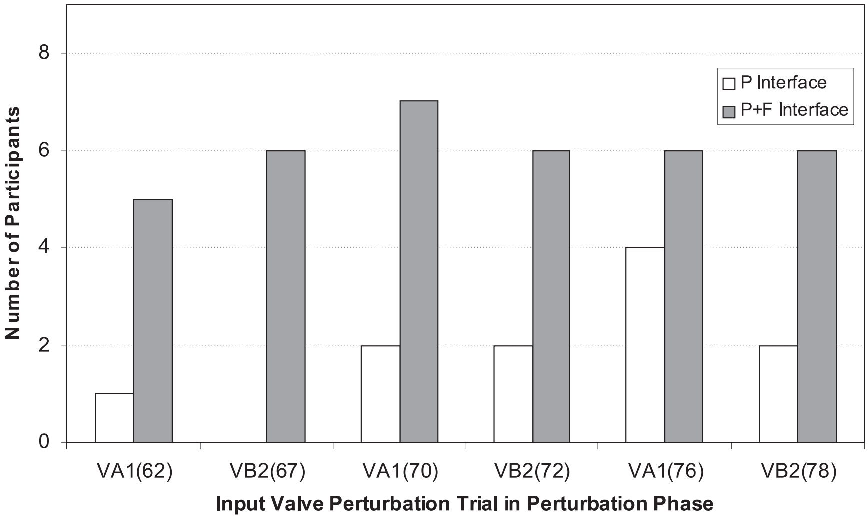

FIGURE 12 Number of participants who either exploited the component dynamics or switched feedwater configuration strategies.

changes to the valve dynamics and depressed higher level functions (e.g., mass inflow) and more successful opportunities to control DURESS II to compensate for this deficiency (e.g., use of cross-flow valves).

Overall. Figure 13 shows an integrated view of the evidence for higher level control in Block 4 state patterns. The results from individual trials for each participant were aggregated; normal trials were separated from perturbation trials. State patterns for each trial were assessed in terms of the criteria for higher level control described in the Appendix. If the state pattern showed evidence for higher level control using any of these criteria, the trial was counted as exhibiting higher level control. This procedure was followed for all trials in Block 4 for each participant. The percentage of trials with evidence for higher level control was calculated for (a) normal trials (i.e., Trials 1, 3, 8, 15, and 20 in Block 4) and (b) perturbation trials (i.e., all other trials in Block 4). The averages and confidence intervals for each interface group are shown in Figure 13.

For normal trials, both interface groups had a low average percentage of trials with evidence for higher level control (not statistically significant), as expected. For perturbation trials, the P+F group had a significantly higher average percentage of trials with evidence for higher level control compared with P. This suggests that the P+F interface is better at inducing higher level control than P for local perturbations.

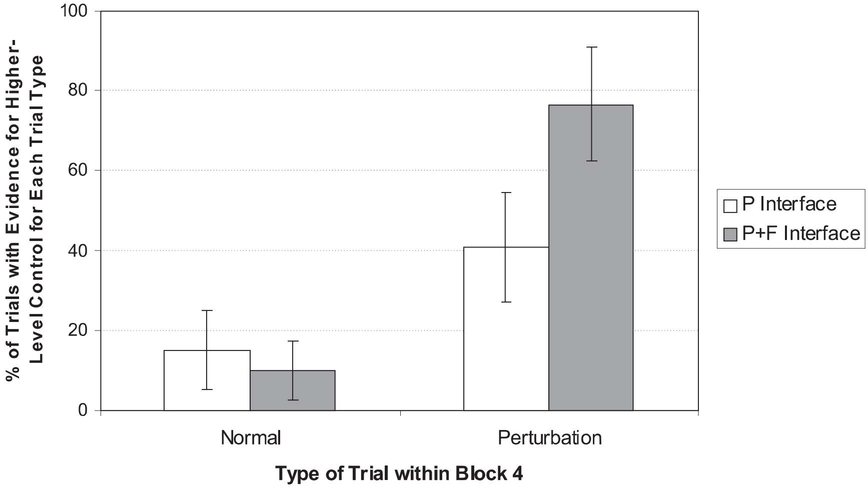

FIGURE 13 Percentage of perturbation trials with evidence for higher level control based on state patterns for specific perturbations.

Control Recipes

Control recipes provide a subjective assessment of level of control by requesting participants to write out the strategy they use for controlling the microworld. Three measures were used to make this assessment in the perturbation phase: (a) reported conditions where actions were dependent on perturbation context (e.g., "if VA1 doesn't work properly, use VB1 to put water into the reservoir 1"), (b) any statements of actions supporting higher level variables (e.g., "use input valves and heaters to make energy in = energy out"), and (c) any statements with specific lower level relationships (e.g., "set VA1 = VO1").

In the last block of the learning phase no participants mentioned any conditions resulting in action changes based on the responsiveness of the work domain components (i.e., local perturbations), as expected, because there were no perturbations. In the perturbation phase, five of the eight P participants and seven of the eight P+F participants reported adjusting their actions based on the effects of the perturbations (e.g., "using a cross-flow strategy when a valve was very slow to respond"); however, the difference between P and P+F was not significant (nonparametric Fisher test for small sample size: $p = .569$ exact, two-tailed). Examples of recipes reporting actions based on the perturbation context are given in Figure 14 for a P+F and a P participant. Of these participants, only one of five P participants correctly identified the source of the perturbation, compared with seven of seven P+F participants; this difference was sta-

a	*b*
P+F Participant	P Participant
A slow component was indicated by a slow response—that is, not heating up quickly or slow response to inflow or outflow. *Flow slow inflow valve:* Try to bypass it by using other flow patterns—for example, if VA1 is slow, use VB1 as source for Tank 1 and VA2 as source for Tank 2. *For slow outflow valve (VO1 or VO2):* Put valve at maximum (20) until the desired outflow is reached, then move valve to that position. For tank inflow, slowly bring inflow matching current outflow with a tank level approximately 20–40. Have heater on and adjust. *For slow heater:* Get steady state flow with water level at 20–40 in tank. Put heater to maximum (10) until it is approximately reading desired output temperature. Then move heater setting to that heat flow spot. This is the most difficult to control. Using water input to control temperature just messes things up. My assumption is that you should leave the flow alone and slowly adjust the dysfunctional heater.	*Special situations:* Heater malfunction and reservoir malfunction—malfunction referring to not responding, filling up too fast, or emptying out too fast (reservoir case). In both cases, do basic setup, then do the following: *Heater case:* The temperature does not rise fast enough—that is, it does not respond to the heater setting. In this case, set heater to maximum (10) and wait until the indicated temperature (red line) is about 1–2 notches (1 = if VO# is low, 2 = if VO# is high), then set HTR to the basic setup level. *Reservoir case:* The reservoir is filling up too fast or emptying too fast. There are two ways you can approach this situation. (a) Use both input valves to lower (or increase) the input flowrate until reservoir stabilizes. (b) Use only one valve to lower (or increase) until reservoir stabilizes. Once the reservoir stabilizes, it will always go back to the basic way of behaving, so once that starts to happen, return all valves altered to the basic setup to keep the reservoir stable. Also note the temperature in case it changes with the input flowrate change.

FIGURE 14 Part of a control recipe for (a) P+F participant and (b) P participant.

tistically significant (nonparametric Fisher test for small sample size: $p = .010$ exact, two-tailed).

Figure 15 shows the number of participants who reported the perturbation conditions as a cue or referent for driving different actions or sequences of actions. Most P+F participants provided recipes for all three classes of perturbations (input valves, heaters, and output valve). In contrast, P participants were split in providing recipes for two conditions (heaters and reservoir/feedwater streams), one routine, or no routines specific to particular perturbations.

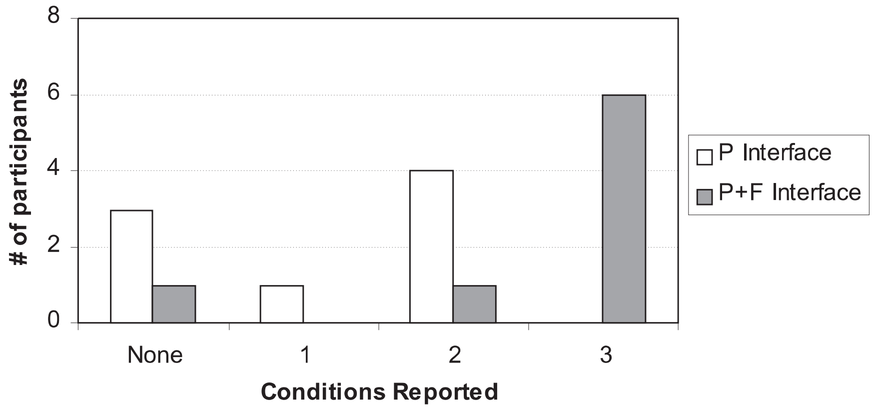

FIGURE 15 Number of participants reporting to adjust actions based on perturbation conditions.

DISCUSSION

Participants who successfully adapted to the perturbations based on trial completion times showed patterns reflecting higher level control through the state patterns and control recipes. As specific local component perturbations occurred, patterns of state variables across abstraction hierarchy levels that were consistent with higher level control (e.g., switching configuration strategies during an input valve perturbation) provided evidence that participants were coupled to the higher level structure of the work domain. Control recipes provided subjective evidence for higher level control after the perturbation phase. Recipes that reported using higher level functions to control the work domain and adjusting strategies for acting on components based on the perturbation context provided evidence for higher level control. These results show that higher level control was associated with successful adaptation to the local perturbations.

One key factor for higher level control is the requirement for multiple work domain affordances or opportunities to meet the work domain purposes. A work domain that has a greater number of structural degrees of freedom makes possible, and provides greater flexibility for, higher level control than a domain with fewer or no structural degrees of freedom; this makes successful adaptation possible. With multiple degrees of freedom, there are numerous opportunities to align with the work domain purposes. As such, using different but functionally equivalent components can compensate for a perturbation to one component. If there are fewer structural degrees of freedom, then there are fewer opportunities to compensate for the perturbations and to maintain the higher level functions.

In this study, successful adaptation through higher level control was limited by the structural characteristics of the work domain. Perturbations at components with multiple structural degrees of freedoms (i.e., input valves) resulted in a performance advantage for the P+F group over the P compared with the components with a single structural degree of freedom (i.e., heaters). In particular, the smallest proportional differences in trial completion times between the P+F group and the P occurred with heater perturbations. The largest differences occurred with input valve perturbations. The differences with the input valve perturbations could be explained by the flexible ways of invoking higher level control, made possible by the existence of multiple structural degrees of freedom. One exception to this pattern occurred with the output valve perturbations (single structural degree of freedom), resulting in a performance advantage for the P+F group over the P. The reason for this exception may be attributed to the lack of goal variable information for the P group.

A second key factor for higher level control is the mapping of functional information about the work domain onto higher level invariant forms in the display, facilitating the strong coupling between participant and affordances that is necessary for coordinated action. In this study, the two interface groups (P and P+F) had different amounts of functional information. The display of higher level functional information with the P+F interface resulted in more instances of higher level control and success in adapting to component perturbations than the P interface (see also Hajdukiewicz & Vicente, 2002). Moreover, the local perturbations had a proportionately larger negative impact on the P group compared with P+F in terms of trial completion times. In addition, more P+F participants, as compared with P participants, had state patterns that provided evidence for higher level control compared with P. Finally, more P+F participants had control recipes that reported the use of higher level functions and action changes based on perturbation context (both criteria for higher level control) as perturbations were introduced. These results suggest that higher level invariants are useful but not essential for higher level control and successful adaptation. Exceptions can occur when participants compensate for the lack of information by developing heuristics through analytical reasoning, as a few P participants did.

The results of this study can be compared with the jaw perturbation studies in motor control conducted by Kelso et al. (1984). In the domain of speech motor control, anatomical structures in the jaw and mouth link functionally to the purposes of speech production. Kelso et al. conducted three experiments introducing unexpected jaw perturbations as the participants uttered specific sounds. In the first two experiments, different parts of the jaw and mouth (e.g., upper lip, lower lip, and tongue) were examined when the same perturbation was introduced for different speech contexts. The results showed that different parts of the jaw and mouth were used based on the task context for the same jaw perturbation. In the third experiment, the jaw perturbations were varied for the same speech task. The results from this experiment showed that different parts of the mouth (e.g., upper lip actions) were used based on the perturbation for the same speech task. These experiments provided evidence for the use of synergies (or higher level control). Despite the changes in perturbation and task context, participants were able to produce the same articulatory patterns as in normal speech; the participants were able to perceive the disturbances and use different parts of the jaw and mouth to achieve the same speech outcome.

In the domain of DURESS II, numerous components link functionally to the purposes of meeting outflow and temperature demands (refer to the abstraction hierarchy analysis in Figure 4). In this study, different local perturbations were introduced for the same task context (i.e., DURESS II startup). Despite these changes, most P+F participants were able to adapt to the perturbations by using components in different ways (e.g., switching feedwater configuration strategy) to achieve generally the same outcome. Evidence for higher level control was shown using various measures (state patterns and control recipes; see previous section). Most of the P+F participants were able to perceive the dynamics of the work domain and opportunities for higher level control, resulting in success in adapting to change. For the most part, P participants did not exhibit this behavior, primarily because the higher level functional information was not provided in the interface.

The results of these two sets of studies from very different domains—speech motor control and thermal-hydraulic process control—show interesting parallels. The third jaw perturbation experiment is the closest when comparing the results with the DURESS II studies. Perturbations for both domains occurred at the lower component levels. For speech motor control, the perturbations were introduced by changing the constraints on jaw movements; for DURESS II, the perturbations were introduced by changing component time constants. With these perturbations, the same general results were found in both sets of studies: Higher level control occurred as lower level perturbations were introduced. This finding is important given the strong differences in tasks and environment characteristics; without the theoretical light shed by the concept of synergies, there is very little reason to believe that these two very different domains would be related in any way. The similarity in results suggests that the pursuit of a generalized systems theory of coordination is worthwhile.

The use of synergies as a theoretical construct to study process control has important implications not just for basic research in ecological psychology, but also for applied research in cognitive engineering. In terms of methodology, the hypothesis that synergies could be relevant to mediated human–machine interaction led us to develop a new set of quantitative measures of performance based on the abstraction hierarchy. These measures have several advantages over the performance measures that are typically used in cognitive engineering (Yu et al., 2002). In terms of experimental design, the concept of synergies also allowed us to make trial-specific predictions about how performance should vary as a function of the component being perturbed rather than the more typical, but less rigorous, aggregate prediction that one interface should be better than another overall (Vicente & Torenvliet, 2001). Similarly, we were able to predict that performance with the

P+F interface should be disrupted more by perturbations to display elements that represent higher level functions than to those that represent lower level physical information (Hajdukiewicz & Vicente, 2002). Collectively, these innovations allowed us to test a cognitive engineering framework—ecological interface design—in a more sophisticated and rigorous way than in the past. Finally, the adoption of synergies has also allowed us to situate cognitive engineering research on human–machine coordination with research in other disciplines, such as ecological psychology, that have also benefitted from the concept of synergies, something that does not appear to have been done before.

CONCLUSIONS

The study presented in this article resulted in two novel and significant contributions to the ecological psychology literature. First, this appears to be the first study to systematically investigate synergies (or higher level control) in human–computer interaction and process control. Second, given the parallel with results from motor control, this study may provide a contribution toward a generalized theory of coordination.

Despite these contributions, this research has limitations that motivate several research topics. First, only one class of perturbations was investigated (i.e., local perturbations to the work domain). Hajdukiewicz and Vicente (2002) systematically investigated the impact of two other types of perturbations (global perturbations of the work domain and perturbations of interface form), but it is not known how the results from either of these studies generalize to other types of perturbations. Second, the microworld, although representative, was limited in scale (i.e., very few components compared with industrial systems), so the generalizability of the results to larger scale systems remains to be evaluated. Third, the domain chosen was process control, which could limit the extent to which the results generalize to other work domains with different types of constraints. Fourth, just as Kelso et al. (1984) did not identify the synergies developed by participants in their studies, the exact form of the synergies used in this task were also not identified. The results suggest the exploitation of mass and energy balance constraints, but further work is needed to model the synergies mathematically.

Investigating these additional research issues would be of interest, not only to the applied goal of shaping human–machine coordination in practice, but also to the fundamental goal of understanding generalized coordination in theory.

ACKNOWLEDGMENTS

The writing of this article was sponsored in part by a research grant, a graduate scholarship, an E. W. R. Steacie Memorial Fellowship from the Natural Sciences and Engineering Research Council of Canada, by the Jerome Clarke Hunsaker Distinguished Visiting Professorship at MIT, and by Honeywell International. We thank Tom Stoffregen, Pascal Van Lieshout, Susan McCahan, Bill Mace, Chris Pagano, and an anonymous reviewer for their contributions.

REFERENCES

Bernstein, N. A. (1996). On dexterity and its development. In M. L. Latash & M. T. Turvey (Eds.), *Dexterity and its development* (pp. 1–244). Mahwah, NJ: Lawrence Erlbaum Associates, Inc. (Original work published 1991)

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

Hajdukiewicz, J. R. (2000). Adapting to change in complex work environments. In G. Szwillus & T. Turner (Eds.), *Proceedings of the conference of human factors in computing systems (CHI 2000;* pp. 89–90). New York: ACM.

Hajdukiewicz, J. R. (2001). *Adapting to change in dynamic worlds: A study of higher-level control and key success factors in a process control microworld.* Unpublished dissertation, University of Toronto.

Hajdukiewicz, J. R., & Vicente, K. J. (2002). Designing for adaptation to novelty and change: The role of functional information and emergent features. *Human Factors, 44,* 592–610.

Irmer, C., & Reason, J. T. (1991). Early learning in simulated forest fire-fighting. In simulations, evaluations, and models: Proceedings of the Fourth MOHAWC Workshop (pp. 1–20). Roskilde, Denmark: Riso National Laboratory.

Kelso, J. A. S. (1995). *Dynamic patterns: The self-organization of brain and behavior.* Cambridge, MA: MIT Press.

Kelso, J. A. S., Tuller, B., Vatikiotis-Bateson, E., & Fowler, C. A. (1984). Functionally specific articulatory cooperation following jaw perturbations during speech: Evidence for coordinative structures. *Journal of Experimental Psychology: Human Perception and Performance, 10,* 812–832.

Loftus, G. R. (1993). A picture is worth a thousand p values: On the irrelevance of hypothesis testing in the microcomputer age. *Behavior Research Methods, Instruments, & Computers, 25,* 250–256.

Pask, G., & Scott, B. C. (1972). Learning styles and individual competence. *International Journal of Man-Machine Studies, 4,* 217–253.

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

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

Schmidt, R. C., Christianson, N., Carello, C., & Baron, R. (1994). Effects of social and physical variables on between-person visual coordination. *Ecological Psychology, 6,* 159–183.

Torenvliet, G. L., Jamieson, G. A., & Vicente, K. J. (2000). Making the most of ecological interface design: The role of individual differences. *Applied Ergonomics, 31,* 395–408.

Turvey, M. T., & Carello, C. (1996). Dynamics of Bernstein's level of synergies. In M. L. Latash & M. T. Turvey (Eds.), *Dexterity and its development* (pp. 339–376). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Vicente, K. J. (1999). *Cognitive work analysis: Toward safe, productive, and healthy computer-based work.* Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Vicente, K. J., & Rasmussen, J. (1990). The ecology of human-machine systems II: Mediating "direct perception" in complex work domains. *Ecological Psychology, 2,* 207–250.

Vicente, K. J., & Torenvliet, G. L. (2001). The earth is spherical ($p < .05$): Alternative methods of statistical inference. *Theoretical Issues in Ergonomics Science, 1,* 248–271.

Yu, X., Lau, E., Vicente, K. J., & Carter, M. W. (2002). Toward theory-driven, quantitative performance measurement in ergonomics science: The abstraction hierarchy as a framework for data analysis. *Theoretical Issues in Ergonomics Science, 3,* 124–142.

APPENDIX
Assessing Level of Control With Across-Level State Patterns

This Appendix describes the methods used to assess level of control for the three types of local perturbations: output valves, heaters, and input valves. Unlike most of the other measures collected, these were not based on automated data collection, but instead required visual inspection of the data and then classification of those data based on the qualitative criteria described here.

The frame of reference used to assess level of control with across-level state patterns is shown in Figure A1 (top row). In general, a lower level variable (i.e., component setting) is compared with a functionally linked higher level variable (i.e., flow rate or mass/energy variable) based on the abstraction hierarchy shown in Figure 4. Specific patterns in this setting-flow space infer level of control based on the local perturbation at the component when compared with a normal (unperturbed) condition.

Three representations are shown in the rows of Figure A1. The top row shows across-level patterns when mapping component setting against its functionally linked higher level variable. The second row shows the temporal patterns of these two variables with respect to their target regions for a normal (unperturbed) trial condition. The third row shows the temporal patterns of these two variables for a trial with a local perturbation at the component.

Three general patterns of action with respect to function that were discovered in this study are shown in the columns of Figure A1. For the first pattern (column A), the participant adjusted the component to the target setting and waited until the actual flow rate reached the target region. For the second pattern (column B), the participant exploited the component dynamics by first adjusting the component to its maximum position to increase the rate of change of flow rate and then readjusting the component to the target setting once the flow rate was in the target region. For the third pattern (column C), referring to any other pattern, the participant made multiple component adjustments before the target region was attained.

Next, we discuss each type of perturbation and the patterns associated with level of control—higher level, lower level, and other types of control.

Output Valve Perturbation

The function of the output valve is to provide enough mass flow to meet the associated demand in the target region (D1g or D2g). Using Figure A1, *setting* refers to

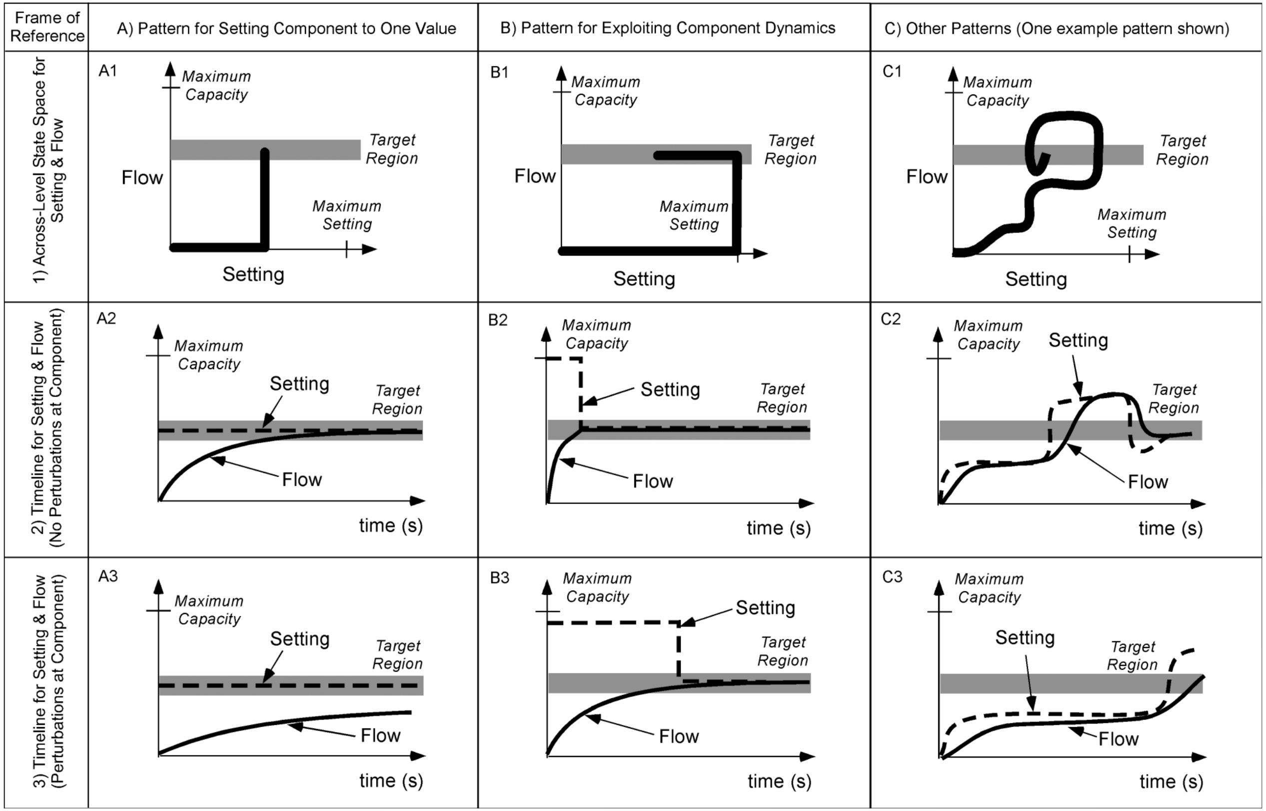

FIGURE A1 Across-level patterns used to assess level of control based on component perturbations. Adapted from Hajdukiewicz (2001).

the output valve setting for VO2 (or Physical Function) on the horizontal axis and its associated higher level variable (mass outflow) on the vertical axis. The outflow demand for the reservoir is shown as the target region. For normal trials, the dynamics of the output valve are fairly responsive, and participants are expected to set the output valve to the required demand (refer to cells A1 and A2 in Figure A1). However, if the dynamics of this component become very slow, the control strategy would have to be adjusted (e.g., column B in Figure A1—exploiting component dynamics) to complete the trial in a shorter period of time.

Three different types of patterns are shown in the columns of Figure A1, corresponding to three types of control when a perturbation is encountered at the output valve (see bottom row): higher level, lower level, and other. With the pattern in column B the valve is initially set near its maximum value and then brought back to the desired value, thus exploiting the dynamics of the valve by increasing its rate of change of flow. In this way, a different control strategy is used in operating the valve to maintain the higher level variables on a similar temporal path as during normal (unperturbed) trials; note the similarities in temporal trajectories between cells A2 to B3 in Figure A1 with the switch in control strategy at the onset of a perturbation. This consistency in maintaining the higher level variables by using the output valve in a way that exploits its dynamics is evidence for higher level control for an output valve (e.g., VO2) perturbation. With the pattern in column A there is consistency at a low level because the output value is always set only once to the target region and left there in the same way as during normal trials. There is no consistency in maintaining the same temporal trajectory for the higher level variables because the local dynamics of the output valve have changed; note the deviations in temporal trajectories between cells A2 to A3 in Figure A1 with no switch in control strategy. This is evidence for lower level control during the output valve perturbation. Any other pattern during the output valve perturbation (i.e., cell C3 in Figure A1) is evidence for other types of control.

Heater Perturbation

The function of the heater is to increase the temperature in the reservoir to the target region. The patterns for level of control for the heater perturbations are shown in Figure A2. Using Figure A1, the horizontal axis represents the heater setting, and the vertical axis represents the associated higher level variable of heat transfer rate and, indirectly, temperature. The temperature target region could be determined analytically for heat transfer rate based on a steady state condition.

With the pattern in column B of Figure A1 participants exploited the dynamics of the heater by first setting it near its maximum value. Once the higher level variables are in the region mapping to the target, the heater setting is adjusted downward to maintain that higher level variable (shown as a horizontal band in the goal region); this provides evidence for higher level control taking into account that a perturbation has occurred at the heater. With the pattern in column A, the heater

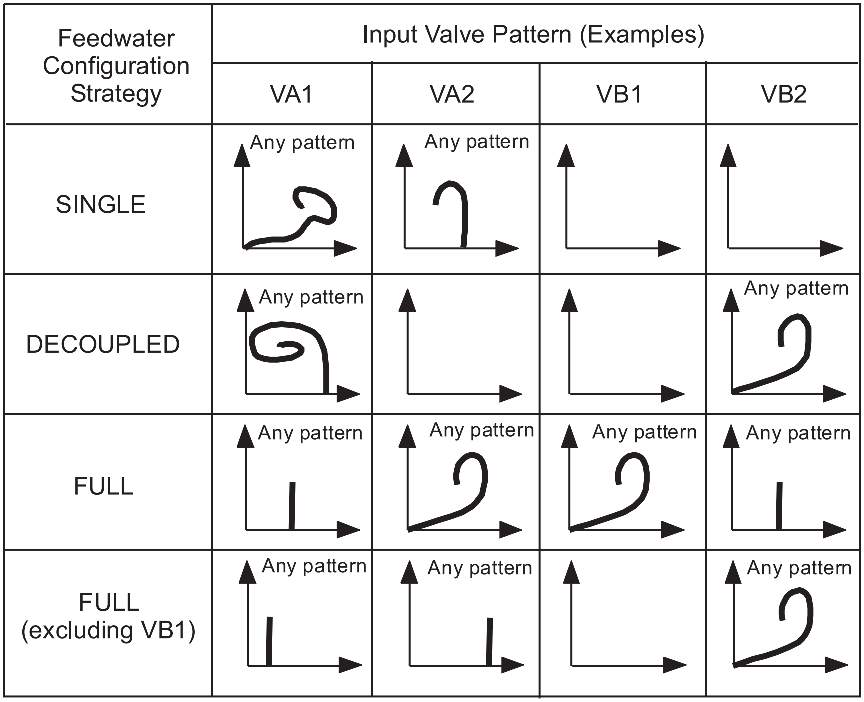

FIGURE A2 Patterns across components indicating feedwater configuration strategy (Hajdukiewicz, 2001).

is roughly set to one final value from the beginning of the trial. This was evidence for lower level control with respect to this perturbation. Any other pattern during this local perturbation (column C) was evidence for other types of control.

Input Valve Perturbation

The main function of the input valve is to provide mass flow to the reservoirs. In general, these valves can be configured in three ways: single, decoupled, and full (Vicente, 1999). If one of the input valves was responding very slowly, its effect on the outcome of the trial depended on which strategy was adopted. For example, if the slowly reacting valve was never used, the outcome was not impacted by the perturbation. However, if the participant used the affected valve, certain patterns in the use of the valve itself and changes in the configuration of other valves could reveal insights into the level of control. In particular, two alternative ways of compensating for this change could be used to complete the trial in a shorter period of time: (a) adjusting settings to offset the perturbation effect or (b) using other components that map onto the same higher level functions.

For the first alternative, the affected input valve was still used. Using Figure A1, *setting* refers to the valve setting for the affected input valve VA1 or VB2 (Physical Function) on the horizontal axis and its associated higher level variable (mass inflow rate contributed by the input valve) on the vertical axis. When the dynamics of this component become very slow, different patterns (represented by columns in Figure A1) were indicative of particular levels of control similar to those discussed for the output valve perturbation.

In Figure A1, if the input valve was used in a way that exploited its dynamics (i.e., pattern in column B) when the input valve perturbation occurred, this would be indicative of higher level control. If the valve was set at roughly one value (pattern in column A), this would be indicative of lower level control. Any other pattern (pattern in column C) is indicative of other types of control.

The second alternative is the use of other components to achieve the same higher level functions, as described. For example, if the input valve perturbation occurred at VA1, VA1 could be shut off and VB1 turned on to support the mass flow into Reservoir 1. This switch in feedwater configuration was indicative of higher level control because alternative structural paths were invoked to support the higher level functions.

A method to qualitatively assess feedwater configuration strategy is to plot the trajectories of four state diagrams for the input valves VA1, VA2, VB1, and VB2 using the same frame of reference as in Figure A1 (i.e., valve setting on the horizontal axis and mass inflow from the valve on the vertical axis). Figure A2 shows different patterns corresponding to the configuration strategies. In the top pattern of Figure A2, participants used only two valves (VA1 and VA2) from the same feedwater stream (FWS A) to feed both reservoirs; this was evidence for a single feedwater configuration strategy. In the next pattern following, participants used only two valves, but from different feedwater streams (FWS A and FWS B); this was evidence for a decoupled feedwater configuration strategy. In the third pattern, participants used all valves; this was evidence for a full feedwater configuration strategy. In the last pattern, participants used three valves; this was evidence for a full feedwater configuration strategy, excluding VB1.

This method was used to identify feedwater configuration strategies for each participant for each trial in the perturbation phase. Feedwater strategies for the perturbation trials (i.e., perturbation strategies) were then compared with reference strategies invoked in normal trials with similar target conditions as the perturbation trials. Strategies that were different and minimally used the perturbed valve were indicative of higher level control. Strategies that were the same were indicative of lower level control or other types of control.

A detailed map showing the possibilities for feedwater configuration strategy switching is shown in Figure A3. The patterns associated with input valve configuration strategy transitions are ways of implying level of control. For a reference strategy that is decoupled, higher level control is implied by a switch to a full, decoupled, or single strategy, with minimal use of the slowly responding valve. Lower

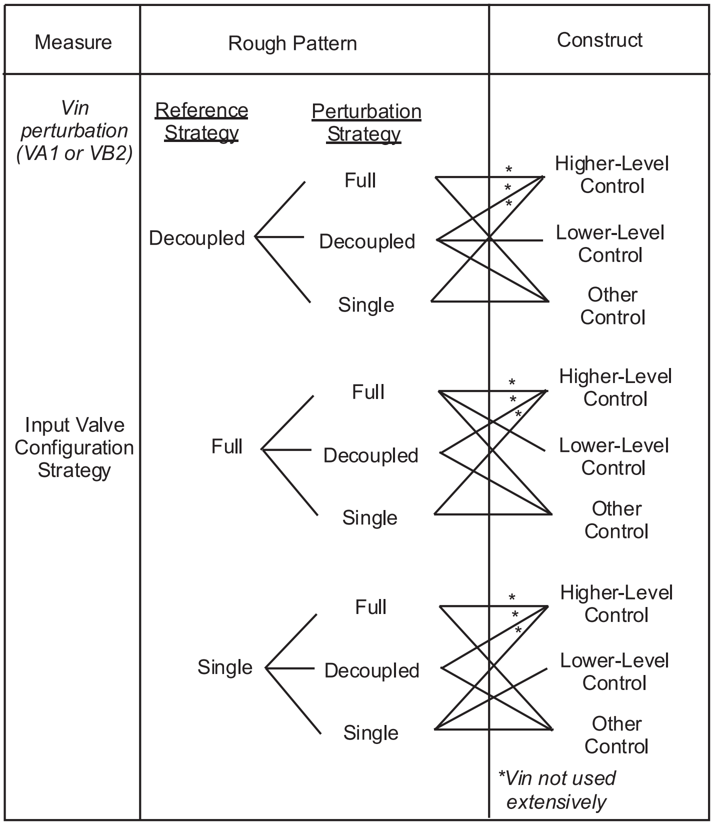

FIGURE A3 Patterns in topological feedwater configuration strategy transitions for VA1 and VB2 perturbations. The asterisk indicates that the slowly reacting valve is minimally used (Hajdukiewicz, 2001).

level control is implied by the continued use of the decoupled strategy (same as reference). All other uses of the components imply other types of control. For a reference strategy that is full, higher level control is implied by a switch to a full, decoupled, or single strategy, with minimal use of the affected valve. Lower level control is implied by the continued use of the full strategy (same as reference). All other uses of the components imply other types of control. For a reference strategy that is single, higher level control is implied by a switch to a full, decoupled, or single strategy, with minimal use of the slowly responding valve. Lower level control is implied by the continued use of the single strategy (same as reference). All other uses of the components imply other types of control.