

Brief Paper

Ironies of Automation*

LISANNE BAINBRIDGE†

Key Words—Control engineering computer applications; man-machine systems; on-line operation; process control; system failure and recovery.

Abstract—This paper discusses the ways in which automation of industrial processes may expand rather than eliminate problems with the human operator. Some comments will be made on methods of alleviating these problems within the 'classic' approach of leaving the operator with responsibility for abnormal conditions, and on the potential for continued use of the human operator for on-line decision-making within human-computer collaboration.

Irony: combination of circumstances, the result of which is the direct opposite of what might be expected.

Paradox: seemingly absurd though perhaps really well-founded statement.

The classic aim of automation is to replace human manual control, planning and problem solving by automatic devices and computers. However, as Bibby and colleagues (1975) point out: "even highly automated systems, such as electric power networks, need human beings for supervision, adjustment, maintenance, expansion and improvement. Therefore one can draw the paradoxical conclusion that automated systems still are man-machine systems, for which both technical and human factors are important." This paper suggests that the increased interest in human factors among engineers reflects the irony that the more advanced a control system is, so the more crucial may be the contribution of the human operator.

This paper is particularly concerned with control in process industries, although examples will be drawn from flight-deck automation. In process plants the different modes of operation may be automated to different extents, for example normal operation and shut-down may be automatic while start-up and abnormal conditions are manual. The problems of the use of automatic or manual control are a function of the predictability of process behaviour, whatever the mode of operation. The first two sections of this paper discuss automatic on-line control where a human operator is expected to take-over in abnormal conditions, the last section introduces some aspects of human-computer collaboration in on-line control.

1. Introduction

The important ironies of the classic approach to automation lie in the expectations of the system designers, and in the nature of the tasks left for the human operators to carry out.

The designer's view of the human operator may be that the operator is unreliable and inefficient, so should be eliminated from the system. There are two ironies of this attitude. One is that

designer errors can be a major source of operating problems. Unfortunately people who have collected data on this are reluctant to publish them, as the actual figures are difficult to interpret. (Some types of error may be reported more readily than others, and there may be disagreement about their origin.) The second irony is that the designer who tries to eliminate the operator still leaves the operator to do the tasks which the designer cannot think how to automate. It is this approach which causes the problems to be discussed here, as it means that the operator can be left with an arbitrary collection of tasks, and little thought may have been given to providing support for them.

1.1. *Tasks after automation.* There are two general categories of task left for an operator in an automated system. He may be expected to monitor that the automatic system is operating correctly, and if it is not he may be expected to call a more experienced operator or to take-over himself. We will discuss the ironies of manual take-over first, as the points made also have implications for monitoring. To take over and stabilize the process requires manual control skills, to diagnose the fault as a basis for shut down or recovery requires cognitive skills.

1.1.1. *Manual control skills.* Several studies (Edwards and Lees, 1974) have shown the difference between inexperienced and experienced process operators making a step change. The experienced operator makes the minimum number of actions, and the process output moves smoothly and quickly to the new level, while with an inexperienced operator it oscillates round the target value. Unfortunately, physical skills deteriorate when they are not used, particularly the refinements of gain and timing. This means that a formerly experienced operator who has been monitoring an automated process may now be an inexperienced one. If he takes over he may set the process into oscillation. He may have to wait for feedback, rather than controlling by open-loop, and it will be difficult for him to interpret whether the feedback shows that there is something wrong with the system or more simply that he has misjudged his control action. He will need to make actions to counteract his ineffective control, which will add to his work load. When manual take-over is needed there is likely to be something wrong with the process, so that unusual actions will be needed to control it, and one can argue that the operator needs to be more rather than less skilled, and less rather than more loaded, than average.

1.1.2. *Cognitive skills.*

Long-term knowledge: An operator who finds out how to control the plant for himself, without explicit training, uses a set of propositions about possible process behaviour, from which he generates strategies to try (e.g. Bainbridge, 1981). Similarly an operator will only be able to generate successful new strategies for unusual situations if he has an adequate knowledge of the process. There are two problems with this for 'machine-minding' operators. One is that efficient retrieval of knowledge from long-term memory depends on frequency of use (consider any subject which you passed an examination in at school and have not thought about since). The other is that this type of knowledge develops only through use and feedback about its effectiveness. People given this knowledge in theoretical classroom instruction without appropriate practical exercises will probably not understand much of it, as it will not be within a framework which

* Received 16 December 1982; revised 23 May 1983. The original version of this paper was presented at the IFAC/IFIP/IFORS/IEA Conference on Analysis, Design, and Evaluation of Man-Machine Systems which was held in Baden-Baden, F.R.G. during September 1982. The published proceedings of this IFAC meeting may be ordered from Pergamon Press Ltd, Headington Hill Hall, Oxford OX3 0BW, U.K. This paper was recommended for publication in revised form by editor A. Sage.

† Department of Psychology, University College London, London WC1E 6BT, U.K.

makes it meaningful, and they will not remember much of it as it will not be associated with retrieval strategies which are integrated with the rest of the task. There is some concern that the present generation of automated systems, which are monitored by former manual operators, are riding on their skills, which later generations of operators cannot be expected to have.

Working storage: The other important aspect of cognitive skills in on-line decision making is that decisions are made within the context of the operator's knowledge of the current state of the process. This is a more complex form of running memory than the notion of a limited capacity short-term store used for items such as telephone numbers. The operator has in his head (Bainbridge, 1975) not raw data about the process state, but results of making predictions and decisions about the process which will be useful in future situations, including his future actions. This information takes time to build up. Manual operators may come into the control room quarter to half an hour before they are due to take over control, so they can get this feel for what the process is doing. The implication of this for manual take-over from automatically controlled plant is that the operator who has to do something quickly can only do so on the basis of minimum information, he will not be able to make decisions based on wide knowledge of the plant state until he has had time to check and think about it.

1.1.3 Monitoring. It may seem that the operator who is expected solely to monitor that the automatics are acting correctly, and to call the supervisor if they are not, has a relatively simple task which does not raise the above complexities. One complexity which it does raise of course is that the supervisor too will not be able to take-over if he has not been reviewing his relevant knowledge, or practising a crucial manual skill. Another problem arises when one asks whether monitoring can be done by an unskilled operator.

We know from many 'vigilance' studies (Mackworth, 1950) that it is impossible for even a highly motivated human being to maintain effective visual attention towards a source of information on which very little happens, for more than about half an hour. This means that it is humanly impossible to carry out the basic function of monitoring for unlikely abnormalities, which therefore has to be done by an automatic alarm system connected to sound signals. (Manual operators will notice abnormal behaviour of variables which they look at as part of their control task, but may be equally poor at noticing changes on others.) This raises the question of who notices when the alarm system is not working properly. Again, the operator will not monitor the automatics effectively if they have been operating acceptably for a long period. A classic method of enforcing operator attention to a steady-state system is to require him to make a log. Unfortunately people can write down numbers without noticing what they are.

A more serious irony is that the automatic control system has been put in because it can do the job better than the operator, but yet the operator is being asked to monitor that it is working effectively. There are two types of problem with this. In complex modes of operation the monitor needs to know what the correct behaviour of the process should be, for example in batch processes where the variables have to follow a particular trajectory in time. Such knowledge requires either special training or special displays.

The second problem is that if the decisions can be fully specified then a computer can make them more quickly, taking into account more dimensions and using more accurately specified criteria than a human operator can. There is therefore no way in which the human operator can check in real-time that the computer is following its rules correctly. One can therefore only expect the operator to monitor the computer's decisions at some meta-level, to decide whether the computer's decisions are 'acceptable'. If the computer is being used to make the decisions because human judgement and intuitive reasoning are not adequate in this context, then which of the decisions is to be accepted? The human monitor has been given an impossible task.

1.2. Operator attitudes. I know of one automated plant where the management had to be present during the night shift, or the operators switched the process to 'manual'. This raises general issues about the importance of skill to the individual. One result of skill is that the operator knows he can take-over adequately if

required. Otherwise the job is one of the worst types, it is very boring but very responsible, yet there is no opportunity to acquire or maintain the qualities required to handle the responsibility. The level of skill that a worker has is also a major aspect of his status, both within and outside the working community. If the job is 'deskilled' by being reduced to monitoring, this is difficult for the individuals involved to come to terms with. It also leads to the ironies of incongruous pay differentials, when the deskilled workers insist on a high pay level as the remaining symbol of a status which is no longer justified by the job content.

Ekkers and colleagues (1979) have published a preliminary study of the correlations between control system characteristics and the operators' subjective health and feeling of achievement. To greatly simplify: high coherence of process information, high process complexity and high process controllability (whether manual or by adequate automatics) were all associated with low levels of stress and workload and good health, and the inverse, while fast process dynamics and a high frequency of actions which cannot be made directly on the interface were associated with high stress and workload and poor health. High process controllability, good interface ergonomics and a rich pattern of activities were all associated with high feeling of achievement. Many studies show that high levels of stress lead to errors, while poor health and low job satisfaction lead to the high indirect costs of absenteeism, etc. (e.g. Mobley and colleagues, 1979).

2. Approaches to solutions

One might state these problems as a paradox, that by automating the process the human operator is given a task which is only possible for someone who is in on-line control. This section will discuss some possible solutions to problems of maintaining the efficiency and skills of the operator if he is expected to monitor and take over control; the next section will introduce recent proposals for keeping the human operator on-line with computer support.

Solving these problems involves very multi-dimensional decision making; suggestions for discussion will be made here. The recommendations in any particular case will depend on such factors as process size and complexity, the rate of process change, the speed and frequency of process or automatic control failure, the variability of the product and the environment, the simplicity and cost of shut down, and the qualities of the operator.

2.1. Monitoring. In any situation where a low probability event must be noticed quickly then the operator must be given artificial assistance, if necessary even alarms on alarms. In a process with a large number of loops there is no way in which the human operator can get quickly to the correct part of the plant without alarms, preferably also some form of alarm analysis. Unfortunately a proliferation of flashing red lights will confuse rather than help. There are major problems and ironies in the design of large alarm systems for the human operator (Rasmussen and Rouse, 1981).

Displays can help the operator to monitor automatic control performance, by showing the target values. This is simple for single tolerance bands, but becomes more complex if tolerances change throughout batch processing. One possible solution is to show the currently appropriate tolerances on a VDU by software generation. This does not actually get round the problems, but only raises the same ones in a different form. The operator will not watch the VDU if there is a very low probability of the computer control failing. If the computer can generate the required values then it should also be able to do the monitoring and alarms. And how does the operator monitor that the computer is working correctly, or take over if it obviously is not? Major problems may be raised for an operator who is highly practised at using computer generated displays if these are no longer available in an emergency. One ironic but sensible suggestion is that direct wired displays should be used for the main process information, and software displays for quantitative detail (Jervis and Pope, 1977).

'Catastrophic' breaks to failure are relatively easy to identify. Unfortunately automatic control can 'camouflage' system failure by controlling against the variable changes, so that trends do not become apparent until they are beyond control. This implies that the automatics should also monitor unusual variable movement. 'Graceful degradation' of performance is quoted in 'Fitts List's of

man-computer qualities as an advantage of man over machine. This is not an aspect of human performance to be aimed for in computers, as it can raise problems with monitoring for failure (e.g. Wiener and Curry, 1980); automatic systems should fail obviously.

If the human operator must monitor the details of computer decision making then, ironically, it is necessary for the computer to make these decisions using methods and criteria, and at a rate, which the operator can follow, even when this may not be the most efficient method technically. If this is not done then when the operator does not believe or agree with the computer he will be unable to trace back through the system's decision sequence to see how far he does agree.

One method of overcoming vigilance problems which is frequently suggested is to increase the signal rate artificially. It would be a mistake, however, to increase artificially the rate of computer failure as the operator will then not trust the system. Ephrath (1980) has reported a study in which system performance was worse with computer aiding, because the operator made the decisions anyway, and checking the computer added to his workload.

2.2. Working storage. If the human operator is not involved in on-line control he will not have detailed knowledge of the current state of the system. One can ask what limitations this places on the possibility for effective manual take-over, whether for stabilization or shut-down of the process, or for fault diagnosis.

The straightforward solution when shut-down is simple and low-cost is to shut down automatically. The problems arise with processes which, because of complexity, cost or other factors (e.g. an aircraft in the air) must be stabilized rather than shut-down. Should this be done manually or automatically? Manual shut-down is usable if the process dynamics can be left for several minutes while the operator works out what is happening. For very fast failures, within a few seconds (e.g. pressurized water nuclear reactor rather than an aircraft), when there is no warning from prior changes so that on-line working storage would also be useless, then reliable automatic response is necessary, whatever the investment needed, and if this is not possible then the process should not be built if the costs of failure are unacceptable.

With less fast failures it may be possible to 'buy time' with overlearned manual responses. This requires frequent practice on a high fidelity simulator, and a sufficient understanding of system failures to be sure that all categories of failure are covered. If response to failure requires a larger number of separate actions than can be made in the time available then some must be made automatically and the remainder by a highly practised operator.

2.3. Long-term knowledge. Points in the previous section make it clear that it can be important to maintain manual skills. One possibility is to allow the operator to use hands-on control for a short period in each shift. If this suggestion is laughable then simulator practice must be provided. A simulator adequate to teach the basic behaviour of the process can be very primitive. Accurate fast reactions can only be learned on a high fidelity simulator, so if such reactions are necessary then this is a necessary cost.

Similar points can be made about the cognitive skills of scheduling and diagnosis. Simple pictorial representations are adequate for training some types of fault detection (Duncan and Shepherd, 1975), but only if faults can be identified from the steady-state appearance of the control panel, and waiting for the steady-state is acceptable. If fault detection involves identifying changes over time then dynamic simulators are needed for training (Marshall and Shepherd, 1981). Simple recognition training is also not sufficient to develop skills for dealing with unknown faults or for choosing corrective actions (Duncan, 1981).

There are problems with the use of any simulator to train for extreme situations. Unknown faults cannot be simulated, and system behaviour may not be known for faults which can be predicted but have not been experienced. This means that training must be concerned with general strategies rather than specific responses, for example simulations can be used to give experience with low probability events, which may be known to the trainer but not to the trainee. No one can be taught about unknown properties of the system, but they can be taught to

practise solving problems within the known information. It is inadequate to expect the operator to react to unfamiliar events solely by consulting operating procedures. These cannot cover all the possibilities, so the operator is expected to monitor them and fill in the gaps. However, it is ironic to train operators in following instructions and then put them in the system to provide intelligence.

Of course, if there are frequent alarms throughout the day then the operator will have a large amount of experience of controlling and thinking about the process as part of his normal work. Perhaps the final irony is that it is the most successful automated systems, with rare need for manual intervention, which may need the greatest investment in human operator training.

3. Human-computer collaboration

By taking away the easy parts of his task, automation can make the difficult parts of the human operator's task more difficult. Several writers (Wiener and Curry, 1980; Rouse, 1981) point out that the 'Fitts list' approach to automation, assigning to man and machine the tasks they are best at, is no longer sufficient. It does not consider the integration of man and computer, nor how to maintain the effectiveness of the human operator by supporting his skills and motivation. There will always be a substantial human involvement with automated systems, because criteria other than efficiency are involved, e.g. when the cost of automating some modes of operation is not justified by the value of the product, or because the public will not accept high-risk systems with no human component. This suggests that methods of human-computer collaboration need to be more fully developed. Dellner (1981) lists the possible levels of human intervention in automated decision making. This paper will discuss the possibilities for computer intervention in human decision making. These include instructing or advising the operator, mitigating his errors, providing sophisticated displays, and assisting him when task loads are high. Rouse (1981) calls these 'covert' human-computer interaction.

3.1. Instructions and advice. Using the computer to give instructions is inappropriate if the operator is simply acting as a transducer, as the computer could equally well activate a more reliable one. Thompson (1981) lists four types of advice, about: underlying causes, relative importance, alternative actions available, and how to implement actions. When following advice the operator's reactions will be slower, and less integrated than if he can generate the sequence of activity himself; and he is getting no practice in being 'intelligent'. There are also problems with the efficient display of procedural information.

3.2. Mitigating human error. Machine possibilities for counteracting human error range from simple hardware interlocks to complex on-line computation. Except where specific sequences of operations must be followed it is more appropriate to place such 'checks' on the effects of actions, as this does not make assumptions about the strategy used to reach this effect. Under manual control human operators often obtain enough feedback about the results of their actions within a few seconds to correct their own errors (Ruffell-Smith, 1979), but Wiener and Curry (1980) give examples of humans making the same types of errors in setting up and monitoring automatic equipment, when they do not get adequate feedback. This should perhaps be designed in. Kreifeldt and McCarthy (1981) give advice about displays to help operators who have been interrupted in mid-sequence. Rouse (1981) suggests computer monitoring of human eye movements to check that instrument scanning is appropriate, for example to prevent tunnel vision.

3.3. Software generated displays. The increasing availability of soft displays on VDUs raises fascinating possibilities for designing displays compatible with the specific knowledge and cognitive processes being used in a task. This has led to such rich veins of creative speculation that it seems rather mean to point out that there are difficulties in practice.

One possibility is to display only data relevant to a particular mode of operation, such as start-up, routine operations, or maintenance. Care is needed however, as it is possible for an interface which is ideal for normal conditions to camouflage the development of abnormal ones (Edwards, 1981).

Goodstein (1981) has discussed process displays which are compatible with different types of operator skill, using a classification of three levels of behaviour suggested by Rasmussen (1979), i.e. skill based, rule based and knowledge based. The use of different types of skill is partly a function of the operator's experience though the types probably do not fall on a simple continuum. Chafin (1981) has discussed how interface design recommendations depend on whether the operator is naive/novice/competent/expert. However, he was concerned with human access to computer data bases when not under time pressure. Man-machine interaction under time pressure raises special problems. The change between knowledge-based thinking and 'reflex' reaction is not solely a function of practice, but also depends on the uncertainty of the environment, so that the same task elements may be done using different types of skill at different times. It could therefore confuse rather than help the operator to give him a display which is solely a function of his overall skill level. Non-time-stressed operators, if they find they have the wrong type of display, might themselves request a different level of information. This would add to the work load of someone making decisions which are paced by a dynamic system. Rouse (1981) has therefore suggested that the computer might identify which type of skill the operator is using, and change the displays (he does not say how this might be done). We do not know how confused operators would be by display changes which were not under their own control. Ephraim and Young (1981) have commented that it takes time for an operator to shift between activity modes, e.g. from monitoring to controlling, even when these are under his control, and one assumes that the same problems would arise with changes in display mode. Certainly a great deal of care would be needed to make sure that the different displays were compatible. Rasmussen and Lind's recent paper (1981) was about the different levels of abstraction at which the operator might be thinking about the process, which would define the knowledge base to be displayed. Again, although operators evidently do think at different levels of complexity and abstraction at different times, it is not clear that they would be able to use, or choose, many different displays under time stress.

Some points were made above about the problems of operators who have learned to work with computer generated displays, when these displays are no longer available in abnormal conditions. Recent research on human memory (Craik, 1979) suggests that the more processing for meaning that some data has received the more effectively it is remembered. This makes one wonder how much the operator will learn about the structure of the process if information about it is presented so successfully that he does not have to think about it to take it in. It certainly would be ironic if we find that the most compatible display is not the best display to give to the operator after all! (As usual with display choice decisions this would depend on the task to be done. A highly compatible display always supports rapid reactions. These points speculate whether they also support acquisition of the knowledge and thinking skills needed in abnormal conditions.)

A few practical points can be suggested. There should be at least one source of information permanently available for each type of information which cannot be mapped simply onto others, e.g. about layout of plant in space as opposed to its functional topology. Operators should not have to page between displays to obtain information about abnormal states in parts of the process other than the one they are currently thinking about, nor between displays giving information needed within one decision process. Research on sophisticated displays should concentrate on the problems of ensuring compatibility between them, rather than finding which independent display is best for one particular function without considering its relation to information for other functions. To end on a more optimistic note, software displays offer some interesting possibilities for enriching the operator's task by allowing him to design his own interface.

3.4. Relieving human work-load. A computer can be used to reduce human work-load either by simplifying the operator's decisions, as above, or by taking over some of the decision making. The studies which have been done on this show that it is a complex issue. Ephraim and Young (1981) found that overall control performance was better with manual control of a single loop, but was also better with an autopilot in the complex

environment of a cockpit simulator. This suggests that aiding is best used at higher work loads. However, the effect of the type of aiding depends on the type of work-load. Johannsen and Rouse (1981) found that pilots reported less depth of planning under autopilot in abnormal environmental conditions, presumably because the autopilot was dealing with the conditions, but more planning under emergency aircraft conditions, where they suggest that the autopilot frees the pilot from on-line control so he can think about other things. Chu and Rouse (1979) studied a situation with both computer aiding and autopilot. They arranged for the computer to take over decision making when the operator had a queue of one other task item to be dealt with and he was controlling manually, or after a queue of three items if the autopilot was controlling. The study by Enstrom and Rouse (1977) makes it clear why Rouse (1981) comments that more sophisticated on-line methods of adapting computer aiding to human work-load will only be possible if the work-load computations can be done in real time. (It would be rash to claim it as an irony that the aim of aiding human limited capacity has pushed computing to the limit of its capacity, as technology has a way of catching up with such remarks.) Enstrom and Rouse also make the important point that the human being must know which tasks the computer is dealing with and how. Otherwise the same problems arise as in human teams in which there is no clear allocation of responsibility. Sinaiko (1972) makes a comment which emphasizes the importance of the human operator's perception of the computer's abilities: "when loads were light, the man appeared willing to let the computer carry most of the assignment responsibility; when loads were heavy, the men much more often stepped in [and] over-rode the computer". Evidently, quite apart from technical considerations, the design of computer aiding is a multi-dimensional problem.

4. Conclusion

The ingenious suggestions reviewed in the last section show that humans working without time-pressure can be impressive problem solvers. The difficulty remains that they are less effective when under time pressure. I hope this paper has made clear both the irony that one is not by automating necessarily removing the difficulties, and also the possibility that resolving them will require even greater technological ingenuity than does classic automation.

References

- Bainbridge, L. (1975). The representation of working storage and its use in the organisation of behaviour. In W. T. Singleton and P. Spurgeon (Eds.), *Measurement of Human Resources*. Taylor and Francis, London, pp. 165-183.
- Bainbridge, L. (1981). Mathematical equations or processing routines? In J. Rasmussen and W. B. Rouse (Eds.), *op. cit.*, pp. 259-286.
- Bibby, K. S., F. Margulies, J. E. Rijnsdorp and R. M. J. Withers (1975). Man's role in control systems. *Proc. 6th IFAC Congress*, Boston.
- Chafin, R. L. (1981). A model for the control mode man-computer interface. *Proc. 17th Ann. Conf. on Manual Control*, UCLA. JPL Publication 81-95, pp. 669-682.
- Chu, Y. and W. B. Rouse (1979). Adaptive allocation of decision making responsibility between human and computer in multi-task situations. *IEEE Trans. Syst., Man & Cybern.*, SMC-9, 769.
- Craik, F. M. (1979). Human memory. *Ann. Rev. Psychol.*, 30, 63.
- Dellner, W. J. (1981). The user's role in automated fault detection and system recovery. In J. Rasmussen and W. B. Rouse (Eds.), *op. cit.*, pp. 487-499.
- Duncan, K. D. (1981). Training for fault diagnosis in industrial process plant. In J. Rasmussen and W. B. Rouse (Eds.), *op. cit.*, pp. 553-573.
- Duncan, K. D. and A. Shepherd (1975). A simulator and training technique for diagnosing plant failures from control panels. *Ergonomics*, 18, 627.
- Edwards, E. (1981). Current research needs in manual control. *Proc. 1st European Ann. Conf. on Human Decision Making and Manual Control*, Delft University, pp. 228-232.
- Edwards, E. and F. P. Lees (Eds.) (1974). *The Human Operator in Process Control*. Taylor and Francis, London.

- Ekkers, C. L., C. K. Pasmooij, A. A. F. Brouwers and A. J. Janusch (1979). Human control tasks: A comparative study in different man-machine systems. In J. E. Rijnsdorp (Ed.), *Case Studies in Automation Related to Humanization of Work*. Pergamon Press, Oxford, pp. 23-29.
- Enstrom, K. O. and W. B. Rouse (1977). Real-time determination of how a human has allocated his attention between control and monitoring tasks. *IEEE Trans. Syst., Man & Cybern.*, **SMC-7**, 153.
- Ephrath, A. R. (1980). Verbal presentation. *NATO Symposium on Human Detection and Diagnosis of System Failures*, Roskilde, Denmark.
- Ephrath, A. R. and L. R. Young (1981). Monitoring vs. man-in-the-loop detection of aircraft control failures. In J. Rasmussen and W. B. Rouse (Eds.), *op. cit.*, pp. 143-154.
- Goodstein, L. P. (1981). Discriminative display support for process operators. In J. Rasmussen and W. B. Rouse (Eds.), *op. cit.*, pp. 433-449.
- Jervis, M. W. and R. H. Pope (1977). Trends in operator-process communication development. Central Electricity Generating Board, E/REP/054/77.
- Johannsen, G. and W. B. Rouse (1981). Problem solving behaviour of pilots in abnormal and emergency situations. *Proc. 1st European Ann. Conf. on Human Decision Making and Manual Control*, Delft University, pp. 142-150.
- Kreifeldt, J. G. and M. E. McCarthy (1981). Interruption as a test of the user-computer interface. *Proc. 17th Ann. Conf. on Manual Control*, UCLA. JPL Publication 81-95, pp. 655-667.
- Mackworth, N. H. (1950). Researches on the measurement of human performance. Reprinted in H. W. Sinaiko (Ed.), *Selected Papers on Human Factors in the Design and Use of Control Systems* (1961). Dover Publications, New York, pp. 174-331.
- Marshall, E. C. and A. Shepherd (1981). A fault-finding training programme for continuous plant operators. In J. Rasmussen and W. B. Rouse (Eds.), *op. cit.*, pp. 575-588.
- Mobley, W. H., R. W. Griffith, H. H. Hand and B. M. Meglino (1979). Review and conceptual analysis of the employee turnover process. *Psychol. Bull.*, **86**, 493.
- Rasmussen, J. (1979). On the structure of knowledge—a morphology of mental models in a man-machine system context. Riso National Laboratory, Denmark, RISO-M-2192.
- Rasmussen, J. and M. Lind (1981). Coping with complexity. *Proc. 1st European Ann. Conf. on Human Decision Making and Manual Control*, Delft University, pp. 70-91.
- Rasmussen, J. and W. B. Rouse (Eds.) (1981). *Human Detection and Diagnosis of System Failures*. Plenum Press, New York.
- Rouse, W. B. (1981). Human-computer interaction in the control of dynamic systems. *Computing Surveys*, **13**, 71.
- Ruffell-Smith, P. (1979). A simulator study of the interaction of pilot workload with errors, vigilance, and decisions, NASA TM-78482.
- Sinaiko, H. W. (1972). Human intervention and full automation in control systems. *Appl. Ergonomics*, **3**, 3.
- Thompson, D. A. (1981). Commercial air crew detection of system failures: state of the art and future trends. In J. Rasmussen and W. B. Rouse (Eds.), *op. cit.*, pp. 37-48.
- Wiener, E. L. and R. E. Curry (1980). Flight-deck automation: promises and problems. *Ergonomics*, **23**, 995.