Time-Related Trade-Offs in Dynamic Function Scheduling

Michael Hildebrandt and Michael Harrison
Department of Computer Science
University of York, York YO10 5DD, UK
Tel: +44-1904-433376, Fax: +44-1904-432767
{Michael.Hildebrandt, Michael.Harrison}@cs.york.ac.uk

Abstract: A possible route to managing workload peaks facilitated by advances in technology is to use Dynamic Function Allocation, in other words to design work so that it is possible to switch adaptively between levels of automation. In the main, current approaches to Dynamic Function Allocation assume that functions are to be serviced as soon as possible, and in order of arrival. These methods utilise online allocation decisions along the human-automation resource dimension. Dynamic Function Scheduling takes a different approach and considers the organisation of functions along a joint human-automation timeline using scheduling mechanisms developed for real-time embedded systems. This paper highlights the limitations of Dynamic Function Allocation as currently considered and argues for the introduction of a temporal dimension to work design. Time-related trade-offs faced by the system designer (e.g. flexibility vs. simplicity) and the operator (e.g. value-based scheduling decisions) are discussed.

Introduction

Time is an ubiquitous and often inconspicuous property of physical and psychological processes. At almost every level of granularity, temporal structures can be identified. "Time is nature's way of keeping everything from happening at once", as Woody Allen put it. However, although processes necessarily unfold in time, this is not of itself a property of primary scientific interest. Indeed, many disciplines adopt a Newtonian view and treat time as a background variable that "flows equably, without relation to anything external." This is true in both psychology and computer science, though notable exceptions can be found. Computer science, despite the strong influence of non-temporal logics and computational theory, is also concerned with designing systems that can adapt reliably to the temporal contingencies and requirements of the environment. This focus has resulted in useful models for scheduling concurrent tasks under conditions of scarce processing resources. In human factors engineering, queuing models (e.g. Walden and Rouse, 1978) have been used to address similar problems. In psychology, time perception was an issue for many early researchers such as Wilhelm Wundt and William James. Interest in time subsided when psychology adopted the information processing paradigm and state transition models from Artificial Intelligence, where temporality is reduced to pure sequence. Recent years have seen a revival in the psychology of time, and research is now going beyond the traditional interest in the psychophysics of time to cognitive models of temporal memory, temporal perspective and time as information. Few attempts have been made at unifying the diverse notions of time across different disciplines. Exceptions are Fraser's (1978) model of 'temporalities' and, with a more socio-psychological focus, Doob's (1971) 'taxonomy of time'.

Time and work: Psychological aspects of time in human factors are often reduced to problems of reaction times and the duration of elementary actions and cognitive operations. While time is fairly well understood and modelled at this fine-grained level of behaviour (e.g. the 'Keystroke-Level Model', Card, Moran and Newell, 1980), many temporal phenomena on a wider temporal horizon are still elusive. Advances in the cognitive psychology of time (see for instance Block, 1990; Friedman, 1990; Macar, Pouthas and Friedman, 1992; Michon and Jackson, 1985; Prabhu, Drury and Sharit, 1997; Roeckelein, 2000) have triggered a new interest in temporal issues in human factors (for instance, Decortis, De Keyser, Cacciabue and Volta, 1991; De Keyser, 1995; De Keyser, Ydevalle and Vandierendonck, 1998; Grosjean and Terrier, 1999; Hollnagel, 1991, 2001; Svenson and Maule, 1993). These studies are concerned with temporal awareness and anticipation, temporal planning and control, temporal errors, and decision making under time stress. Despite this progress in human factors, the work design and automation literature has so far given little consideration to temporal organisation.

It is important to emphasise that this line of research is not following a Taylorist agenda – we are not proposing a return to time-and-motion studies. On the contrary, where Taylorism sees the operator as a mainly reactive, event-driven agent who has to adapt to the rhythm of the system, our interest is in the operator's active shaping of the joint human-automation timeline. Instead of breaking work down into elementary, disconnected units, this approach aims at understanding behavioural integration on the operator's temporal horizon.

Structure of the paper: The next section introduces Dynamic Function Allocation, a work design concept, and discusses some unresolved issues and limitations of this approach, relating both to system design and operation. To provide a temporal perspective on work design, an outline of the Dynamic Function Scheduling approach (Hildebrandt and Harrison, 2002) is presented. Time-related trade-offs in system design and operations are discussed.

Dynamic Function Allocation
One of the defining features of modern work is its dynamism. Processes unfold rapidly and sometimes in unexpected ways, resource constraints have to be accommodated online, actions have to be synchronised and coordinated, information needs to be updated and distributed, plans have to be revised and adapted. Automation, introduced to help the human operator handle this complexity, can produce new problems by removing the operator from the control loop and leaving him/her unaware of the state of the system in case of a failure. To address the problems of all-or-nothing automation and static Function Allocation methods, where a level of automation is selected at the design stage, Dynamic Function Allocation (sometimes also called 'Adaptive Automation') provides systems with multiple levels of automation, and decision rules to switch between them at runtime (see Scerbo, 1996, for an overview). Empirical evaluations, mostly based on microworld simulations of production line tasks, air traffic control or aviation scenarios, suggest significant improvements in situation awareness, handling of faults and workload peaks, and overall productivity (e.g. Endsley and Kaber, 1999; Moray, Inagaki and Itoh, 2000; Parasuraman, 1993; Walden and Rouse, 1978; Rencken and Durrant-Whyte, 1993; Tattersall and Morgan, 1997).

The Dynamic Function Allocation literature is diverse. Studies differ in the problems they address (mainly workload and situation awareness), the control over level-of-automation switches (human-initiated, automation-initiated, or comparisons of distinct blocks of trials under different automation levels), the levels of automation provided (full automation vs. full human control or automation scale), and the decision rule used to switch between them (human-initiated, critical event logics, workload- or model-based logics). Despite the multitude of empirical basic research, these approaches have not yet been translated into a unified, mature design method (see Hancock and Scallen, 1998, for some recommendations). As few Adaptive Automation systems are available outside the aviation domain (e.g. Morrison, 1993), the long-term benefits and problems of this approach are as yet difficult to assess. The following two sub-sections discuss a number of unresolved issues relating both to the design and operations of Adaptive Automation systems.

Design considerations: To be more adaptive than all-or-nothing automation approaches, Dynamic Function Allocation provides a number of different levels of automation for a given system. For

instance, Sheridan's (1981) widely cited automation scale, which applies most readily to information processing and problem solving purposes, comprises 10 distinct levels (for a more recent scale, see Endsley & Kaber, 1999). Implementing this diversity is likely to be a major challenge. Not only must the designer develop and test a variety of different solutions for the same function, but also provide a sensitive and reliable decision logic, which might involve workload and context measures. The costs and benefits of this development effort are not currently discussed, and it is unclear how easily the current scales can be adapted to a variety of application domains.

Current research in this area tends to assess the effects of Adaptive Automation for single, isolated functions. In these studies, the relevant aspect of the automation decision is the effect on workload and situation awareness, and not the potential, more specific implications for the servicing of other functions. Even when multi-task paradigms are used, the functions are often not strongly causally related. However, as functions in modern socio-technical systems are usually highly inter-connected, the effects of a mode change in one function might have significant implications for a whole network of other functions. The requirements of the specific problem or problem-solving strategy might be a much stronger constraint on the automation decision than workload reduction and maintaining situation awareness (see next section). Before Dynamic Function Allocation can develop into a mature work design method, it has to be able to take account of the inter-dependencies of functions and the contexts in which they might occur (see Harrison, Johnson and Wright, 2002, for an example of such an approach in static Function Allocation).

The only option for workload balancing in Dynamic Function Allocation is automation – 'Dynamic' here refers to a decision on the resource axis, not on the timeline. In so far as the decision is based on performance data or critical events, the method has a temporal element, but it often takes into account only a narrow, retrospective temporal window around the decision point. As the specific effects of the allocation decision for the future timeline are not usually considered, this approach can be characterised as 'snapshot allocation'. However, understanding workload in practice will need to allow considerations of the operator's pro-active, future oriented behaviour.

Operator considerations: Among the primary concerns for Dynamic Function Allocation methods is the loss of situation awareness. As this phenomenon can occur under long periods of automation, Parasuraman (1993) suggested that automation levels should switch periodically, even without being triggered by critical workloads, to keep the operator in the control loop. However, this approach could be problematic if the manual operation cycles of various different functions are not well synchronised, creating the risk of task interference. Confusion can also be caused if automation levels switch too quickly and frequently as a result of insufficient inertia in the decision logic, or if the decisions are intransparent to the operator.

Another critical issue in Dynamic Function Allocation is complacency or over-reliance on automation (Parasuraman, Molloy and Singh, 1993). Especially if the automation is fairly reliable (but still not perfect), operators could be lulled into a false sense of security and thereby neglect their supervisory duties. The authors suggest that vigilance could be encouraged by "simulat[ing] a variable-reliability system by including (at variable intervals) artificial failures that would require an operator response". On the other hand, some studies (Harris, Hancock and Arthur, 1993; Tattersall and Morgan, 1997) have documented human failure to engage automation even when available. More specifically, Harris et al. report that fatigued participants failed to use automation, even though they might benefit most from automatic support. Unfortunately, as with most studies in this field, these papers report summary results and not individual strategies, making it difficult to generate explanations for these results.

A problem of Dynamic Function Allocation, especially if allocation shifts are to be triggered by the human operator, is the added processing demand induced by the decision process (note that a similar problem occurs in real-time systems, where there is a trade-off between more sophisticated and effective scheduling algorithms and the processing time required to execute them). This problem becomes aggravated the more adaptivity a system provides, as more levels of automation have to be considered. Thus, a compromise has to be found between flexibility and simplicity. The computational complexity can be reduced when automation levels for different functions are not seen as independent of each other, but instead as bound up into automation configurations, with each configuration

appropriate for a certain operation scenario. This perspective, seeing automation in the context of strategy choice, is not strongly developed in current approaches.

Most current Dynamic Function Allocation concepts assume or require that all available levels of automation provide equal quality of solution, so that re-allocation decision can be based purely on the required workload reduction. While this assumption is feasible for some isolated automation scenarios, under a more naturalistic perspective, function servicing strategies often involve satisficing decisions and trade-offs. For instance, in a medical context, expert systems could be used by more junior staff as part of a backup strategy if advice from senior staff is unavailable. Similarly, unavailability of automatic medical equipment such as blood gas monitors or ventilators might require higher manual involvement, even though the quality of this treatment may be lower. In fault analysis, different levels of data integration (e.g. high integration with decision support or access to raw data) will be chosen according to the cognitive strategy of the operator, not necessarily for the workload reduction they provide.

Dynamic Function Scheduling

Dynamic Function Scheduling (Hildebrandt and Harrison, 2002) brings a temporal perspective to workload-related problems in high-consequence systems, and also aims at understanding and designing a broader range of scheduling and satisficing phenomena in normal operations. It considers allocation along the joint human-automation timeline as a strategy in multi-task servicing (Fig. 1). In this sense it goes further than the automation option considered in Dynamic Function Allocation. In addition to asking *who* should perform a function, it asks *when* and *if* a function should be performed, taking into account the agents' current and predicted workload, available resources, service rates, and the configuration of other functions on the joint timeline. Scheduling options include postponing, swapping and dropping of functions. For instance, Hildebrandt and Harrison (2002) discuss a fault servicing scenario for an aviation hydraulics system and identify conditions where different scheduling strategies are appropriate (diagnose fault first, then fix it; switch to redundant circuit, then diagnose fault; drop function, i.e. ignore problem, if leak will not become critical before touch-down). Arguing that scheduling is an ubiquitous problem, the authors also discuss a supermarket checkout scenario, where both function allocation and scheduling can be observed: if a customer cannot pack the items quickly enough, the cashier will often switch from his/her primary function of scanning the items to assisting the customer in packing in order to optimise overall throughput. From an allocation perspective, part of the packing function has been re-distributed to the cashier. From a scheduling perspective, the operator has postponed the primary function (scanning) to increase performance of the joint packing function (note that this decision could be context dependent: if the cashier is fatigued, the delay in packing could provide a welcome break). A combination of scheduling and allocation is characteristic of most multi-agent systems.
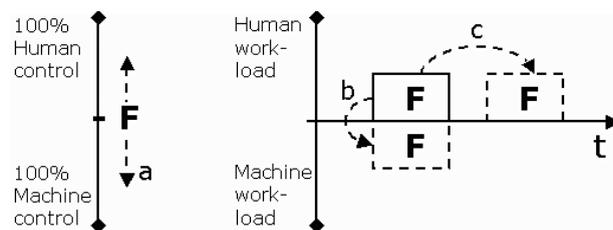

Figure 1. Conceptual differences: Dynamic Function Allocation (left) allocates on the resource dimension (a). Dynamic Function Scheduling (right) allocates on the resource (b) and/or the temporal dimension (c).

Value-based function scheduling / strategy selection: Dynamic Function Scheduling considers both temporal and quality-related aspects ('value') of a function, and considers the trade-offs involved in trying to accommodate concurrent functions in a given time frame. To address some of the limitations of current Dynamic Function Allocation, the approach distinguishes between functions and the strategies available for servicing a function. This results in two different notions of value: one is a measure of the contribution a function makes to the overall system objectives and is used in planning, i.e. to prioritise and order concurrent functions by comparing their values (for example, in aviation the highest priority is

given to safety-related functions, followed by passenger comfort and economy). In the above example, the value of assisting in packing becomes greater than the value of continuing scanning when items pile up. The other notion of value is a measure of the quality of solution a particular strategy (possibly involving a certain level of automation) provides in servicing a certain function. It is used to select among the different strategies available for servicing a function. Seeing the hydraulics example as a case of strategy selection (though it also involves scheduling), the decision to 'diagnose first, fix second', 'fix first, diagnose later' or 'drop function' will depend on the utility of obtaining a closer diagnosis, the time required for the diagnosis, the time available for fixing the problem, current workload, and the stage of the mission. Though these computations can, in theory, become very complex, most expert operators will have developed efficient heuristics and decision rules to assess the dynamics of the problem and resolve speed-quality trade-offs in strategy selection (e.g. Amalberti and Deblon, 1992).

Both notions of value are closely related; a lower-value, but faster, strategy may have to be selected if there is insufficient time (or resources) for executing the higher-value, but slower, strategy by the function's deadline. The quality of the selected strategy will, in turn, affect the value of the function itself. To reason about such relations, it is useful to introduce the notion of urgency, which can be obtained by relating the time required and the time available for servicing a function or executing a strategy. The urgency approaches 1 as the function gets closer to its deadline. If the ratio exceeds 1, the function cannot be serviced in time, and might have to be dropped.

A further dimension is added by assuming that values change over time. The value of servicing a function may be lower when the function is far from its deadline than when it is very close to it. Similarly, a strategy that requires a shorter execution time than an alternative strategy will have a higher relative value when the deadline is close than when the deadline is still a long time away. When applied to actual work situations the concept of value will have to be extended to represent dynamic changes over time and to allow for linear or non-linear value functions. It will also be necessary to integrate the notions of value and utility in the psychological literature on judgment and decision making.

System design trade-offs: For the designer, the main challenge related to Dynamic Function Scheduling is in deciding on the sequential flexibility or rigidity of the functions in the system. The order in which functions should be serviced can be constrained by their physical and logical nature (e.g. lowering the landing gear and landing), or by requirements and limitations of the human operator (e.g. biases in temporal reasoning or tendency to omit actions and confuse sequence in high workload situations). In many high-consequence domains such as aviation and power plant control, there is a need to provide rigid sequentialisation in the form of checklist procedures to avoid omissions and to ensure correct ordering. In other situations, procedural diversity might be necessary to operate in a dynamic environment. Flexibility is also necessary if the operator has to find solutions to unforeseen failures. Thus a compromise has to be found between the risk and the diversity provided by flexible temporal organisation. In terms of the hydraulics example mentioned above, this would involve analysing the benefits of the different strategies (diagnose-fix, fix-diagnose, drop) in different scenarios, considering the operator's decision effort for matching a strategy to a situation, and possibly considering a redesign of the function (using automation) and the physical system.

The designer should also be aware of the overall temporal properties of the system. This includes the assessment of the expected function arrival rates, temporal properties of the functions (e.g. continuous, periodic, sporadic), service rates for human and automation, and the ability of the combined system to accommodate unexpected events on the timeline. This temporal inventory of the domain will be the basis for designing levels of redundancy and a function distribution policy that can achieve the required performance within acceptable workload levels.

Operator trade-offs: The operator's value-based function scheduling and strategy selection often involves online satisficing decisions and speed-quality trade-offs (see discussion above). This can take the form of more shallow processing (e.g. in problem solving and decision making, Payne and Bettman, 1988), use of an alternative processing strategy (Sperandio, 1978), or 'buying time' by slowing down the process itself (e.g. a production line). While these decision trade-offs strongly depend on the semantics of the specific function, temporal reasoning itself involves costs and benefits. Higher levels of temporal reasoning and awareness (Grosjean and Terrier, 1999) might support problem solving and situation

awareness (provided that functions are sufficiently predictable), but will require close familiarity with the system and absorb attentional resources.

A similar trade-off exists between *control* (of the immediate system state) and *planning* (assembling a goal-directed action sequence or strategy). A more elaborate plan will simplify control decisions. With a rough or incomplete plan, control decisions will require more online reasoning. Either strategy might be appropriate depending on characteristics such as predictability, time pressure and operator capabilities. For instance, Amalberti and Deblon (1992) report that expert fighter pilots plan flight routes in more detail and consider more problem scenarios than less experienced pilots.

There is often a correlation between the quality of planning and control decisions and the operator's temporal horizon: if causes and effects are only assessed for the short term, or not at all, decisions tend to be erratic and based on arbitrary situational cues. Reasoning about a wider temporal window will take more potentially relevant factors into account. Hollnagel's (2000) Contextual Control Model captures these differences in the quality of control by the notion of control modes (scrambled, opportunistic, tactical, strategic). Hollnagel (2001) explicitly discusses the role of time in losing and regaining control.

Few studies have addressed temporal issues in planning directly. Smith, Hill, Long and Whitefield (1997) modelled planning and control of multiple task work in secretarial office administration and identified a number of control rules and planning heuristics for plan maintenance and revision, interruption handling, task switching and sharing, and prioritisation.

Conclusion

This paper introduced a temporal dimension to function allocation and discussed some of the trade-offs of temporal work organisation, both for the system designer and operator. To overcome the limitations of current Dynamic Function Allocation concepts, allocation along the joint human-automation timeline should be considered in addition to allocation on the human-automation resource dimension. If Dynamic Function Allocation is to be applied to a wider set of problems, automation decisions should be seen in the context of value-based strategy selection, allowing for speed-quality trade-offs. Dynamic Function Scheduling is a conceptual framework that has the potential to analyse a wide range of scheduling and planning behaviour and provide guidance for the designer in assessing the risks and benefits of temporal flexibility in a system. Future work, using both microworld experimentation and case studies, should address problems of temporal reasoning, awareness, and temporal planning and control.

Acknowledgements

Referenecs

Amalberti, R. & Deblon, F. (1992). Cognitive modelling of fighter aircraft process control: a step towards an intelligent on-board assistance system. *International Journal of Man-Machine Studies, 36*, 639-671.

Block, R.A. (Ed.). (1990). *Cognitive Models of Psychological Time*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Card, S.K., Moran, T.P. & Newell, A. (1980). The keystroke-level model for user performance with interactive systems. *Communications of the ACM, 23*, 396-410.

Decortis, V., de Keyser, V., Cacciabue, P.C. & Volta, G. (1991). The temporal dimension of man-machine interaction. In G.R.S. Weir & J.L. Alty (Eds.), *Human-Computer Interaction and Complex Systems*. London: Academic Press.

De Keyser, V. (1995). Time in ergonomics research. *Ergonomics, 38*, 1639-1660.

De Keyser, V., Ydevalle, G. & Vandierendonck, A. (Eds.). (1998). *Time and the Dynamic Control of Behavior*. Göttingen: Hogrefe.

Doob, L.W. (1971). *Patterning of Time*. New Haven, CT: Yale University Press.

Endsley, M.R. & Kaber, D.B. (1999). Level of automation effects on performance, situation awareness and workload in a dynamic control task. *Ergonomics, 42*, 462-492.

Fraser, J.T. (1978). *Time as Conflict*. Basel: Birkhäuser.

Friedman, W. (1990). *About Time: Inventing the Fourth Dimension*. Cambridge: MIT Press.
Grosjean, V. & Terrier, P. (1999). Temporal awareness: Pivotal in performance? *Ergonomics, 42*, 1443-1456.

Hancock, P.A. & Scallen, S.F. (1998). Allocating functions in human-machine systems. In R.R. Hoffman, M.F. Sherrick & J.S. Warm (Eds.), *Viewing psychology as a whole: The integrative science of William M. Dember* (pp. 509-539). Washington, DC: American Psychological Association.

Harris, W.C., Hancock, P.A. & Arthur, E.J. (1993). The effect of taskload projection on automation use, performance, and workload. *Proceeding of the 7th International Symposium on Aviation Psychology*.

Harrison, M.D., Johnson, P.D. & Wright, P.C. (2002). Automating functions in multi-agent control systems: supporting the decision process. *Proceedings of the Tenth Safety-critical Systems Symposium* (pp. 93-106). London: Springer.

Hildebrandt, M. & Harrison, M. (2002). The temporal dimension of Dynamic Function Allocation. Paper to be presented at *Eleventh European Conference on Cognitive Ergonomics*.
Hollnagel, E. (1991). The phenotype of erroneous actions: Implications for HCI design. In G.R.S. Weir & J.L. Alty (Eds.), *Human-Computer Interaction and Complex Systems*. London: Academic Press.

Hollnagel, E. (2000). Modeling the orderliness of human action. In N. Sarter & R. Amalberti (Eds.), *Cognitive engineering in the aviation domain*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Hollnagel, E. (2001). Time and control in joint human-machine systems. *Proceedings of the 2nd IEE People In Control Conference* (pp. 246-253).

Macar, F., Pouthas, V. & Friedman, W.J. (Eds.). (1992). *Time, Action, Cognition*. Dordrecht, The Netherlands: Kluwer.

Michon, J.A. & Jackson, J.L. (Eds.). (1985). *Time, Mind and Behavior*. Berlin: Springer.

Moray, N., Inagaki, T. & Itoh, M. (2000). Adaptive automation, trust, and self-confidence in fault management of time-critical tasks. *Journal of Experimental Psychology: Applied, 6*, 44-58.

Morrison, J.G. (1993). *The adaptive function allocation for intelligent cockpits program: Interim research and guidelines for the application of adaptive automation* (Technical report). Warminster, PA: Naval Air Warfare Center, Aircraft Division.

Parasuraman, R. (1993). Effects of adaptive function allocation on human performance. *Proceeding of the FAA/NASA Conference on Artificial Intelligence and Human Factors in Air-Traffic Control and Aviation Maintenance*. Daytona Beach, FL: Embry-Riddle Aeronautical University.

Parasuraman, R., Molloy, R. & Singh, I.L. (1993). Performance consequences of automation-induced 'complacency'. *International Journal of Aviation Psychology, 3*, 1-23.

Payne, J.W. & Bettman, J.R. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory and Cognition, 14*, 534-552.

Prabhu, P., Drury, C. & Sharit, J. (1997). Using temporal information in time-constrained tasks. *Proceedings of the Human Factors Society 41st Annual Meeting*.

Rencken, W.D. & Durrant-Whyte, H.F. (1993). A quantitative model for adaptive task allocation in human-computer interfaces. *IEEE Transactions on Systems, Man, and Cybernetics, 23*, 1072-1090.

Roeckelein, J.E. (2000). *The Concept of Time in Psychology: A Resource Book and Annotated Bibliography*. Westport, CT: Greenwood.

Scerbo, M.W. (1996). Theoretical perspectives on adaptive automation. In R. Parasuraman & M.

Mouloua (Eds.), *Automation and Human Performance: Theory and Applications* (pp. 38-63). Hillsdale, NJ: Lawrence Erlbaum Associates.

Sheridan, T.B. (1981). Understanding human error and aiding diagnostic behavior in nuclear power plants. In J. Rasmussen & W.B. Rouse (Eds.), *Human detection and diagnosis of system failures* (pp. 19-36). New York: Plenum Press.

Smith, W., Hill, B., Long, J. & Whitefield, A. (1997). A design-oriented framework for modelling the planning and control of multiple task work in secretarial office administration. *Behaviour & Information Technology, 16*, 161-183.

Sperandio, J-C. (1978). The regulation of working methods as a function of work-load among air traffic controllers. *Ergonomics, 21*, 195-202.

Tattersall, A.J. & Morgan, C.A. (1991). The function and effectiveness of dynamic task allocation. In D. Harris (Ed.), *Engineering Psychology and Cognitive Ergonomics, vol. 2* (pp. 247-255). Aldershot: Ashgate.

Walden, R.S. & Rouse, W.B. (1978). A queueing model of pilot decision making in a multitask flight management situation. *IEEE Transactions on Systems, Man, and Cybernetics, 8*, 867-875.