The Temporal Dimension of Dynamic Function Allocation

Michael Hildebrandt and Michael Harrison

The Dependability Interdisciplinary Research Collaboration

Department of Computer Science
University of York, York YO10 5DD
United Kingdom

Tel: +44-1904-433376, Fax: +44-1904-432767
{Michael.Hildebrandt, Michael.Harrison}@cs.york.ac.uk

Abstract

Current Dynamic Function Allocation methods are designed to switch adaptively between levels of automation on the human-automation resource dimension in order to avoid excessive workload levels or provide backup if parts of the system fail. These methods usually consider functions in isolation, and assume that all functions should be serviced as soon as possible. In other words, they are concerned with *who* should service a function, not *when* (or if) the function should be serviced. This paper introduces a temporal perspective on function allocation by discussing how functions can be scheduled on a joint human/automation timeline. This includes the options to postpone, drop or swap functions. Central to the Dynamic Function Scheduling approach is the concept of value-based function scheduling and strategy selection. Finally, psychological constraints which could affect the operator's temporal reasoning and decision making are discussed.

Introduction

Technological advances over the last decades have not led to a complete replacement of the human operator by automated systems. While automation makes an essential contribution to safe, dependable and efficient operations, the operator's ability to perform adaptively in unexpected situations is unlikely to be matched by automation in the foreseeable future. As a consequence, joint human-automation operation has become the dominant paradigm for designing modern work situations, especially in high-risk domains. Such models go beyond the dichotomy of complete automation vs. complete human operation by defining levels of automation between these two extremes, where a function is serviced by an interaction between human and automation. Dynamic Function Allocation refers to a mode of operation where several levels of automation are provided and a decision procedure (also potentially automated) is used to switch between them.

Dynamic Function Allocation has a number of advantages over static Function Allocation methods, where only one level of automation is selected at the design stage. Early Function Allocation approaches (e.g. Fitts, 1951) were based on general 'Men are better at – Machines are better at' lists. Consideration of the agents' strengths and weaknesses is part of any design process. The problem of these early approaches, as Hancock and Scallen (1998) noted, “does not lie in ignoring environmental context as an influence; it is in assuming environmental context is constant and predictable […]. Variations that are intrinsic to performance, such as changes associated with learning, fatigue, stress, and anxiety, remain largely unincorporated into the design picture.”

Recent Function Allocation methods (Harrison, Johnson and Wright, 2002; Parasuraman, Sheridan and Wickens, 2000) have provided a much finer level of analysis by decomposing functions into their information acquisition, decision making, implementation and monitoring component, and considering the automation options and implications for each of these components. While these approaches provide useful solutions in a variety of work domains, they do not address the problem of adaptation to dynamically changing operational parameters. As no alternative levels of automation are provided to switch between at runtime, the solution selected at the design stage will have to fit the entire operational context.

Structure of the paper

The following two section will briefly describe how Dynamic Function Allocation provides adaptation through re-distribution of functions along the resource dimension. The next section will discuss how current Dynamic Function Allocation methods are limited in their allocation options because either they consider functions in isolation, or at the very best in a narrow, usually retrospective, temporal window around the decision point. In the Dynamic Function Scheduling approach, temporal organisation is relevant both for managing concurrent functions and for selecting among strategies for servicing a function. Two examples will illustrate the options to postpone, drop or swap functions. Contrary to the present paradigm, where levels of automation are assumed to provide equal quality of solution, we argue that Dynamic Function Scheduling and Allocation decisions often involve value-based trade-offs, including speed/quality compromises. Frameworks for representing the temporal properties of functions and agents are presented. Finally, we will discuss psychological limitations of temporal planning and control, and the role of temporal awareness.

Dynamic Function Allocation

Classifications of levels of automation with various degrees of human and machine involvement in function implementation began to emerge in the late 1970s (Sheridan and Verplank, 1978). This has led to a proliferation of research on design methods and analyses of Dynamic Function Allocation (for a review, see Scerbo, 1996). Despite the variety, three broad modes of use of Dynamic Function Allocation can be identified.

Re-allocation as an emergency or backup solution

Adaptive switches between levels of automation can be used as a reaction to unexpected and intolerable workload increases. Workload peaks may be due to a sudden increase in task arrival rates, decreases in the operator’s service rate, or critical incidents like system failures. Here automation acts as a backup processor to support or substitute the operator (e.g. Chu and Rouse, 1979; Rencken and Durrant-Whyte, 1993). Most of the basic research on performance benefits of adaptive automation, the effects on operator workload levels, the problem of loss of situation awareness, and issues of complacency (Parasuraman, Molloy and Singh, 1993) and trust (Moray, Inagaki and Itoh, 2000) in automation has been conducted in this area.

Allocation as operation strategy

In these applications, switches between levels of automation are part of a deliberate, pro-active strategy. For example, there are clear rules and conditions for engaging and disengaging the autopilot during cruise and takeoff/landing. The operator initiates the strategy in the context of normal operations, although the autopilot might also be useful as an emergency backup solution. Loss of situation awareness and automation surprises are less likely to be a problem here than in the backup case unless an automation failure forces the operator to take control unexpectedly.

Allocation in adaptive problem solving

Automation can support higher-level control tasks such as information gathering, analysis and decision making at different levels depending on operational parameters and cognitive strategy. In time-critical situations, a high degree of data aggregation and decision support may be desirable (provided that the automation is reliable), whereas for other problems, access to raw data may be required to assess a situation. Operation in these conditions often involves close interaction between the operator and automation, and changes of levels of automation are usually initiated by the operator. It is here that Sheridan's (1981) ten levels of allocation apply most easily.

Triggers of re-allocation

Besides providing different levels of automation, at least as important to Dynamic Function Allocation is the decision procedure used to switch between the levels. Three broad classes of decision models can be distinguished: critical event models, measurement models, and operator-driven models. Critical event models are used most commonly for emergency- or backup-type automation. Re-allocation is triggered by detection of an unforeseen or irregular system state, assuming that the workload produced by the incident exceeds the operator's capabilities. While it is widely agreed that in most domains the ultimate authority for switching between levels of automation should lie with the human operator, it has been argued that some rapidly developing incidents should trigger automatic fault management or shutdown procedures without allowing for a temporal window in which the human operator can override the decision (Moray, Inagaki and Itoh, 2000). It is therefore reasonable to see the decision procedure itself as a function which can be allocated dynamically between human and automation.

Measurement models are also used to manage excessive workload. Instead of detecting critical system states, these models infer workload changes either indirectly from changes in system performance or directly from physiological measurements of the operator's stress level. Both critical event and measurement models are reactive, or retrospective, in nature. The trigger for re-allocation is based on the circumstances leading up to the decision point. Differences lie in the sampling period used for detecting the re-allocation trigger; whether an instant as in the case of critical event models or a period as in measurement models. Both assume that immediate action is necessary to reduce workload, and that workload reduction is to be achieved by shifting the current allocation of function along the human-automation resource dimension.

Operator-driven allocation shifts apply mainly to the domains of automation as operation strategy and automation in adaptive problem solving. The need to re-allocate may be related to workload increases, but re-allocation in this case is usually a pro-active decision, taken in relation to a specific work strategy on the operator's planning horizon.

From Dynamic Function Allocation to Dynamic Function Scheduling

Current Dynamic Function Allocation methods consider functions in isolation, and consider the causes and effects of the allocation decision in a narrow temporal window around the decision point. The re-allocation decision is mostly reactive, based on sampling into the past (from the perspective of the decision point). The aim of the re-allocation decision is to reduce current workload. The workload distribution will remain at the set level until another shift occurs. More specific effects on the agents' future work schedule are not normally considered. The term 'snapshot allocation' for this type of allocation emphasises that the decision is dynamic on the resource dimension, but not on the temporal dimension.

Current approaches also assume that functions should be serviced as soon as possible. For many situations, particularly emergencies, this is a reasonable assumption. However, even in these critical situations there might be the need to prioritise function servicing. This can be achieved either by allocating the function to another agent for immediate servicing (as in current Dynamic Function Allocation), or by re-allocating it on the agents' timeline (Fig. 1). As a consequence, relaxing the immediate-processing constraint opens a new dimension for designing function servicing: Dynamic Function Scheduling.

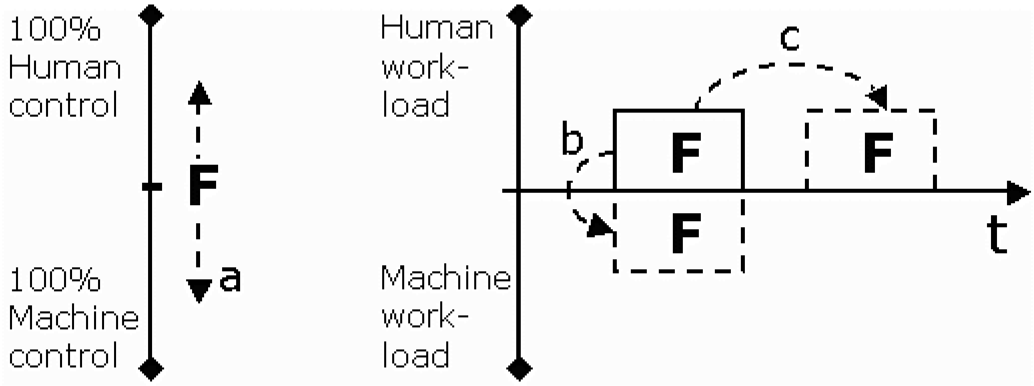

Figure 1. Dynamic Function Allocation (left) allocates on the resource dimension (a). Dynamic Function Scheduling (right) allocates on the resource (b) and/or the temporal dimension (c).

Managing functions by considering scheduling opportunities and constraints has received little attention compared to the body of literature on distributing tasks along the resource dimension. This is not very surprising, considering that time itself is a largely under-researched topic in human factors as well as in many other disciplines (for exceptions, see for instance Decortis, De Keyser, Cacciabue and Volta, 1991; De

Keyser, 1995; Grosjean and Terrier, 1999; Hollnagel, 1991, 2000, 2001; Smith, Hill, Long and Whitefield, 1997). When time is considered, it is mainly aspects like time pressure (Svenson and Maule, 1993) or the duration and control of elementary movements and cognitive operations that are being discussed (Card, Moran and Newell, 1980). The active aspects of time and work, such as temporal reasoning, decision making and synchronisation, are less well understood, not least because the theoretical and methodological foundations are incomplete (see section *The Psychology of Time* for some interesting contributions). By the crucial shift in perspective from reactive to proactive temporal behaviour, and from control to planning, time is no longer seen as a uniform background variable, but as a multi-faceted, subjective concept.

Scheduling and planning are an integral part of most work situations, and indeed of most aspect of life in general, though these decisions are often taken implicitly or rely on habitual actions schemas and strategies. Explicit scheduling decisions, especially in human multi-agent systems, can become very complex by involving multi-factorial decision processes, uncertainty about future states, rapidly changing contexts and unexpected events. This could be one of the reason why in high-consequence domains like aviation, designers often try to provide as much proceduralisation as possible, thereby reducing uncertainty about ordering and temporal planning. On the other hand, operational diversity, including the discretion to schedule adaptively and to construct novel action sequences, can be important for the robustness of a system. Research into Dynamic Function Scheduling promises to provide analysis tools for understanding the temporal structure of these processes and might lead to the development of support tools for taking scheduling decisions online. At the very least, it can provide design heuristics by defining the limitations of operator temporal awareness and making transparent the conditions where temporal errors are likely to occur. The following two examples illustrate that scheduling, sequencing and temporal coordination is relevant both in everyday life and in high-risk emergency situations.

Scenario 1: Supermarket checkout

Consider for example packing assistance at supermarket checkouts. In a typical scenario, the prices are scanned by the cashier, while the customer packs the products. The input to the scanning function is a queue of products waiting to be processed. The cashier scans as quickly as possible and can usually perform this function more quickly than the customer can pack the items. Products build up waiting to be packed. As the next customer is not served before the current one has paid, a backlog of items will reduce the overall service rate of the checkout system. Therefore, the cashier can decide to postpone scanning to assist the customer with packing. Only when the backlog is cleared will the cashier resume the primary function.

This example involves both Dynamic Function Scheduling and Allocation. By postponing the scanning, the cashier re-organises his/her schedule when the value of packing becomes greater than that of scanning. The decision is dynamic because it depends on the customer's service rate. If the customer packs quickly enough, the cashier will not have to re-schedule (for certain customers or situations it might be appropriate to allocate packing to another member of staff by asking the customer in advance if he/she requires assistance). From a function-centred perspective, there is a partial re-allocation of workload from customer to cashier, so the level of allocation for the packing function has shifted along the resource dimension from purely customer-operated to customer-and-cashier operated. A combination of scheduling and allocation is characteristic of many work systems where adaptation to endogenous or exogenous parameters is required.

Scenario 2: Hydraulics fault

The second example has no resource allocation component although various ways of automating the functions could be imagined. The scenario (taken from Fields and Merriam, 1998) involves fault management of an aircraft hydraulics system. A leak can occur in either of two reservoirs or the servos operating the rudder and the aileron (Fig. 2). The problem is signalled to the operator by a drop in the hydraulics pressure readings.

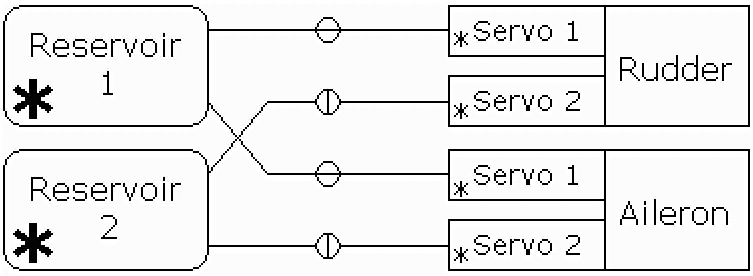

Figure 2. Simplified aircraft hydraulics system. Leaks (*) can occur in any of the servos or reservoirs .

Two strategies for servicing the fault can be identified. The first strategy analyses the fault by going through a checklist procedure of selectively setting valves connecting the reservoirs with the servos. If the leak is in one of the servos then only this servo can be disconnected and both reservoirs can still be used. If the leak is in one of the reservoirs, the reservoir has to be disconnected and the other reservoir has to be used. The alternative strategy does not go through an initial diagnostic process, but switches to the other reservoir as soon as the leak is noticed. This immediately fixes the problem, as no matter where the problem was located (the reservoir itself or the servos), all potentially faulty components are now completely disconnected. However, if the problem lies in one of the servos, then an intact reservoir has been disconnected and is unavailable.

Considering 'diagnosing hydraulics problem' and 'fixing hydraulics problem' as two separate but related functions, the scenario gives a choice between two strategies for servicing the problem: diagnose first, fix second, or fix first, diagnose second (or later, or not at all). One or the other strategy can be used depending on the operational parameters. If there are other high-value functions waiting to be serviced, detailed diagnosis can be postponed until sufficient time is available. If the reservoir is needed during the rest of the mission (as the other reservoir might be failing later), then it is important to diagnose if the reservoir is intact after all, and the problem is due to a leaking servo. Depending on temporal parameters, even the option to ignore the problem completely can be considered: when the aircraft is close to landing, servicing the function could be dropped altogether if the leak will not have a critical effect before touchdown.

Value-based scheduling

Value-based function scheduling

A number of factors have to be taken into account in a function scheduling decision: How important is the function? How important are other concurrent functions? What are their deadlines? Which strategies (including automation) are available for servicing the functions? What are the temporal properties of the different strategies? What quality of solution do they provide? How do they fit into the overall schedule? What costs and benefits does a particular decision (strategy selection and scheduling) involve? Clearly, these considerations are not purely temporal, but include a notion of quality or value. Interestingly, value-based scheduling has received considerable attention in Flexible Scheduling, an area in Real-Time Systems research concerned with management of tasks competing for scarce processing resources (Krishna, 1997; Liu, 2000). In such systems, tasks are prioritised according to their 'value', i.e. the contribution a task makes to the overall objectives of the system. Values in these models are fixed and are usually set at the design stage by expert judgement. For example, in civil aviation the highest priority is given to safety-related functions, followed by passenger comfort and operations efficiency. Compared to the situations Function Allocation is usually concerned with, the operational environment of Real-Time Systems is often far more clearly defined. Though the scheduling algorithms can become complex, the decision trade-offs are normally very transparent so as to be expressible in computational terms. However, though Flexible Scheduling algorithms will not be applicable to a wider work design context without modification, its basic concepts provide a useful framework for discussing Dynamic Function Scheduling. Some scheduling patterns will be described in the next section.

Value-based strategy selection

The value of a function is critical for prioritising and scheduling concurrent functions. However, value and, by implication, time, is also relevant in deciding how and by whom a function should be serviced. It is therefore surprising that value-based strategy selection is not discussed in Dynamic Function Allocation. Most current Dynamic Function Allocation methods assume or require that all available levels of automation should provide equal quality of solution, i.e. equal value. However, in many situations, as we have seen in the hydraulics example, function allocation is part of a process of strategy selection, where strategies involve different levels of automated operations. Strategies may differ in their timing properties (e.g. execution speed, sequential rigidity), but, more crucially, in the quality of solution they provide. If automation is used for backup or emergency purposes, it may not produce the same quality a human operator would, but this lower level of quality will be preferable to a complete breakdown of operations. Expert systems may not match the abilities of a human expert, but can be of great benefit if the expert is unavailable and a decision has to be taken by a less experienced operator. The decision procedure will have to consider compromises such as cost/benefit and speed/accuracy trade-offs to select a strategy that provides the best solution (as indicated by the strategy's value) under the given operational parameters. It should be noted that value-based strategy selection assumes that a number of strategies are available in the form of fairly proceduralised schemas or scripts, in which there is some durational flexibility, but the order of actions is largely fixed. In some situations, however, no strategy or set of strategies may be available, or non of the strategies satisfies the operational constraints (e.g. deadlines). Under these circumstances, the operator has to engage in a process of goal-directed action planning and resource allocation to compose a novel strategy that meets at least some of the system's core objectives.

We thereby obtain 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. The other 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. Both 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 relate the concept of value in Flexible Scheduling to the notions of value and utility in the psychological literature on judgement and decision making.

Patterns of Dynamic Function Scheduling

Flexible Scheduling provides a number of options for managing tasks dynamically. We will briefly describe the three most relevant alternatives, namely postponing, dropping and swapping, and their relation to the temporal allocation of functions. It is important to note that the applicability of these management options depends critically on the temporal properties of the functions (see discussion below).

Postponing of functions

The most obvious option for temporal allocation is postponement. A good example of postponement can be found in the medical domain. During high-workload periods, staff will sometimes postpone the writing of patient records if higher-value functions are waiting to be serviced. Discretion to postpone depends on time available and time required to service the function or task, i.e. its urgency. An automated system that is aware of the time required and available could display a measure of urgency (and possibly value) to the operator, or compare levels of urgency for different service strategies requiring different amounts of execution time to aid the decision process. Postponing of functions is risky, as other tasks might arrive unexpectedly at the period allocated for servicing the function, when no temporal buffer will be available to accommodate all functions. As a consequence, the temporal window for function servicing should be considered at design stage, and assistive technology to aid online postponement decisions should be provided.

Dropping of functions
If a function cannot be executed by a hard deadline, its value is effectively zero. As a consequence, the function could be dropped so that the processing effort can be invested in other tasks that *can* be executed in time. Many mundane examples of dropping a function can be imagined; if a submission deadline is close, the author may decide not to have the paper proof-read by a colleague (hoping that at least the spelling errors will have been spotted by the automatic spell checker). This kind of satisficing is very common in everyday work, but should receive much more explicit attention. Considerations of the option to drop or postpone a function depend critically on the temporal structure of the function, as not all functions have hard deadlines, or produce all-or-nothing results if the deadline is or is not met.

Dropping can also be considered in the context of concurrent functions. If an excessive number of functions occur, and there is no option to re-allocate them either on the resource or temporal dimension, a value-based decision for dropping some of the functions will have to be made. The criticality of the function, its causal relation to other functions, and the reduction in workload resulting from dropping it are important parameters to be taken into account. Lower-value functions could be dropped in order to free resources. An automated decision support system might be helpful in carrying out these multi-factorial considerations. Alternatively, if the decision to drop a function is taken by the human operator, a mechanism for communicating it to the automation may be necessary.

Swapping of functions
Functions are swapped if they are serviced in reverse order of arrival (in event-driven systems), or if the operator can choose the order in which functions are executed (in operator-driven interaction; see hydraulics example above). This option is not available if functions are sequentially rigid, i.e. when there is a precedence relationship between them (e.g. lowering the landing gear and landing). Swapping is most likely to occur if functions queue up and have to be prioritised in order to be processed sequentially. Automation could again provide analyses of the values and urgencies of functions, and the allocation and scheduling options.

While it is feasible to design scheduling options for related functions such as ‘diagnose-fix’ or ‘fix-diagnose’ in the hydraulics example, the combinations of unrelated functions are usually too unpredictable to be considered at the design stage. Swapping of unrelated functions will be based on re-ordering the functions by measures of urgency and value.

Temporal properties of functions: formal models

Real-Time Systems and temporal logics
Any scheduling decision depends critically on the temporal properties of the functions and the agents. Again, these aspect are not reflected in current Dynamic Function Allocation methods. No comprehensive category system for describing the temporal structure of a work domain is available. An inventory of this type will be a pre-

condition for any kind of temporal design. There are, however, some approaches which might contribute to a better understanding of Dynamic Function Scheduling. As discussed above, the Flexible Scheduling and Real-Time Systems literature provides concepts for reasoning about value-based scheduling as well as a variety of scheduling patterns. Another important source for constructing formal models of time is Artificial Intelligence, and, in particular, temporal logics (Allen, 1991a; Allen and Ferguson, 1994). Temporal logics provide representations and operators for manipulating states and events. Some of this literature specifically addresses temporal reasoning and planning (McDermott, 1982; Allen, 1991b). However, the problems addressed in these models are usually quite simple and, crucially, well-defined, which cannot be assumed for real-world work situations, where decisions, especially for planning on a wider temporal horizon, are subject to considerable uncertainty. While these models seem to accommodate fixed ‘values’ quite naturally, they are less well suited for representing dynamically changing values and non-linear value functions. Both Flexible Scheduling and temporal logics operate on a Newtonian notion of time, where time is a constraint of the environment to which the system’s behaviour has to be adjusted. These computational approaches do not capture the richness of psychological time, which inevitably is necessary for understanding Dynamic Function Scheduling decisions. It also seems difficult to represent within these approaches the critical distinction between functions and strategies

Triggers

Some relevant temporal properties of functions should be mentioned briefly. Obviously, the processing effort required for servicing a function will affect processing time. But more structural differences in temporal properties can be identified; some functions, such as monitoring a patient’s vital signs during surgery, are continuous, i.e. occur throughout the course of operations (although workload may differ over time). Functions may also occur periodically. As periodic events are predictable, processing resources can be reserved for these functions on the timeline. A special form of predictability occurs when functions are deliberately designed to shift between levels of automation, forcing the operator to take manual control at regular intervals. This cycling between automated and human control has been suggested as a means for avoiding the problems of out-of-the-loop performance and loss of situation awareness (Parasuraman, 1993). Other functions, such as fault servicing, are invoked sporadically. The schedule should provide sufficient resources above those reserved for continuous or periodic functions, or provide rules for prioritising, to accommodate these unexpected functions. Mean function arrival rates and standard deviations can provide some measure of distribution for periodic and sporadic tasks.

Deadlines

Functions not only differ in their triggering conditions, but also in terms of their termination. There are many examples of hard deadlines, such as resuscitation of a patient before brain damage is caused. If the deadline is ‘soft’, the system *should* meet it, but missing it will not have a catastrophic effect; the value of such functions does not drop to 0 as the deadline is missed, but decreases as a function of time when the deadline has been passed. For other functions, such as resolving a production line fault, function servicing should be finished as soon as possible, but there are no explicit deadlines.

Temporal aspects of operator behaviour: The psychology of time

Psychological aspects of time 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 on this fine-grained level of behaviour, many temporal phenomena on a wider temporal horizon are still elusive. There is, unfortunately, no unified literature on the psychology of time (see for instance Block, 1990; De Keyser, Ydevalle and Vandierendonck, 1998; Fraser, 1978; Friedman, 1990; Macar, Pouthas and Friedman, 1992; Michon and Jackson, 1985). The majority of research falls into two categories: time perception is concerned with how accurately people can judge, memorise and produce durations. The second area, temporal memory, asks if and how temporal markers are stored and retrieved in episodic memory. While these two problems affect the operator's ability to schedule functions, the following research domains seem to be more directly relevant.

Control

The term 'control' refers to the operator's problem of deciding what to do next, and how to do it, given the set of current operational constraints. Control can be lost and regained, and covers all decisions taken towards achieving the system's objectives or handling unforeseen events. Control decisions can differ in the temporal horizon they take into account, both in terms of the past (diagnostic reasoning) and future (predictive reasoning). There is usually a correlation between the temporal horizon and the quality of the control decisions: 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 necessarily 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. As the operator's control mode, i.e. the temporal horizon, is critical for both function scheduling and strategy selection, system design should ensure the necessary conditions for maintaining a high level of control, or provide the appropriate assistance for regaining it.

Planning

Where control is concerned with the *next* action, planning is concerned with assembling the sequence of actions necessary to achieve a certain goal. Most psychological models of planning are based on an Artificial Intelligence approach (Newell and Simon, 1972; Hoc 1988). They are usually used for planning towards a single goal, and not, as in Dynamic Function Scheduling, for managing multiple functions. As noted above, dynamically changing values are not easily represented. Also, as Hill, Long, Smith and Whitefield (1995) note, these models "view plans as complete and fully elaborated behaviour sequences which ensure task goal achievement." They argue for a more realistic perspective where "plans need not be complete and fully-elaborated, but rather they may be partial […] and/or general." As with control, planning is strongly related to the operator's temporal horizon. 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 interesting control rules and planning heuristics for plan maintenance and revision, interruption handling, task switching and sharing, and prioritisation.

Temporal awareness

Most behaviour requires close temporal coordination and synchronisation, both in terms of sequence and duration. Usually these temporal contingencies are represented implicitly, bound up into cognitive schemas or motor scripts for skilled action (Michon, 1990). However, sometimes explicit reasoning about durations and deadlines is necessary. In these cases, time is used as information. Temporal awareness refers to the operator's knowledge of process durations, deadlines and arrival rates, and the ability to perform more productively and safely by temporal planning and anticipation of temporal landmarks. In one of the few experimental studies on this topic, Grosjean and Terrier (1999) concluded that "temporal awareness proves to be a good indicator of performance, both in terms of errors committed and multiple-goal optimisation."

Temporal errors

As Hollnagel (1991) notes, "Surprisingly, few of the existing action and error taxonomies include the aspect of time, but rather describe and classify human error on an atemporal (static) basis […] In many domains it is, however, necessary to include time in a much more conspicuous way, as, perhaps, one of the principal 'mechanisms' or 'error areas' of human action […] This is particularly true with respect to planning and scheduling." Temporal errors relate both to the order and duration of events and actions, to the perception and memory of time, and to the implicit and explicit use of temporal information. Classifications can be found in Decortis, De Keyser, Cacciabue and Volta (1991), De Keyser (1995) and Hollnagel (1991).

Decision making under time stress

A large interest in judgement and decision making under time pressure was triggered by Wright's (1974) study on the 'harassed decision maker' (see Svenson and Maule, 1993, for a recent collection). In their review, Edland and Svenson (1993) identified, among others, the following effects of time pressure: "…increased use of many pieces of information but in a more shallow way […] the accuracy of human judgement decreases […] the use of noncompensatory decision rules becomes more frequent than compensatory rules requiring value tradeoffs […] time pressure leads to a tendency of locking in on a strategy and to decrease competence of finding alternative strategies in problem solving." As scheduling decisions are often taken under conditions of high workload and time pressure, it is essential to support or simplify the decision process to avoid the risks of decision biases and errors.

Concluding remarks

This paper is not intended to suggest that temporal processes and temporal behaviour in work situations can or should be determined and designed in every detail – we are not proposing a return to the tradition of time-and-motion studies. Indeed, Dynamic Function Scheduling can be seen as the exact opposite of a procrustean approach to work, as it emphasises the operator's active role in defining and controlling the work schedule. Where a Taylorist approach decomposes work into elementary operations, Dynamic Function Scheduling intends to provide an integrated perspective on a wider temporal horizon and to make scheduling options and trade-offs explicit.

As scheduling is an ubiquitous and sometimes challenging element of everyday life and work, it is surprising that our understanding of the temporal organisation of interactive systems is incomplete. We have introduced a perspective to function allocation that identifies the temporal properties of re-allocation triggers, deadlines, agents and functions itself. This allows us to reason about patterns of adaptive temporal function re-allocation. Dynamic Function Scheduling could complement existing Dynamic Function Allocation methods in managing workload peaks during emergencies as well as normal operations. More basic research is necessary to describe, understand and design the temporal course of work processes. Richer notions of value, urgency and utility will be needed. To integrate temporal design into the system engineering life cycle, a method for making assumptions about task arrival rates and agents' service rates should be developed (queueing theory is a natural candidate, see Walden and Rouse, 1978; Chu and Rouse, 1979). This will provide the basis for taking explicit decisions about the sequential rigidity and deadlines of the system, and potentially for providing automatic scheduling assistance.

Dynamic Function Scheduling will introduce risks as well as opportunities. The longer the temporal window under consideration, the more uncertain predictions about future states become. The temporal horizon for designing and managing the agents' timeline will be restricted by the dynamics of the system and by cognitive limitations of the operator. It will be useful to apply Hollnagel's (2000, 2001) notion of control modes to determine the temporal window the operator is likely to be in control of. As sequencing is an important component of the Dynamic Function Scheduling approach, consideration has to be given to the operator's task of decomposing functions into an operation strategy, and assessing the semantic and temporal parameters of the different options. This process of 'task analysis at runtime' is likely to be more demanding than non-temporal allocation decisions in current Dynamic Function Allocation methods. As functions are broken down into tasks and sequenced by the operator or automation at runtime, there is again the need to schedule, now on the level of tasks instead of functions (for instance, some task sequences may be pre-emptable, others may not allow pre-empting).

At this stage, our considerations are mainly conceptual – they make few assumptions about technological feasibility, psychological constraints, or the pragmatic aspects of providing a design method for Dynamic Function Scheduling. However, we believe this broad approach is necessary so that subsequent research and development can build on a comprehensive set of scheduling options. Once a temporal perspective on work processes has been established, suitable analytical, descriptive and normative models can be produced. Scheduling notions from the Real-Time Systems- and temporal logics literature provide a valuable frame of reference, but will have to be adapted to be useful in a complex human-automation work context. Extensive empirical research will be necessary to understand and assist the temporal course of joint human-automation task servicing. A diverse literature on the psychology of time will have to be integrated and applied to the scheduling problem. Finally, from a systems-engineering point of view, the interaction between the resource dimension and the temporal dimension, between Dynamic Function Allocation and Dynamic Function Scheduling, will have to be investigated. Can a joint decision procedure for allocation and scheduling be developed?

Acknowledgements

This work was supported in part by the UK EPSRC DIRC project (www.dirc.org.uk), Grant GR/N13999.

References

Allen, J.F. (1991a). Time and time again: The many ways to represent time. *International Journal of Intelligent Systems, 6,* 341-356.

Allen, J.F. (1991b). Planning as temporal reasoning. *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning* (pp. 3-14). San Mateo, CA: Morgan Kaufmann Publishers.

Allen, J.F. & Ferguson, G. (1994). Actions and events in interval temporal logic. *Journal of Logic and Computation, 4,* 531--579.

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

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

Chu, Y. & Rouse, W.B. (1979). Adaptive allocation of decisionmaking responsibility between human and computer in multitask situations. *IEEE Transactions on Systems, Man, and Cybernetics, 12,* 769-778.

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.

Edland, A. & Svenson, O. (1993). Judgement and descision making under time pressure. In O. Svenson & A.J. Maule (Eds.). *Time Pressure and Stress in Human Judgement and Decision Making* (pp. 27-40). New York: Plenum Press.

Fields, R.E & Merriam, N.A. (1998). Inference and information resources: A design case study . In P. Markopoulos and P. Johnson (Eds.), *Proceedings of the 5th Eurographics Workshop on Design, Specification, and Verification of Interactive Systems* (pp. 41-56). Wien: Springer.

Fitts, P.M. (Ed.) (1951). *Human engineering for an effective air navigation and traffic control system*. Washington, DC: National Research Council.

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.

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.

Hill, B., Long, J., Smith, W. & Whitefield, A. (1995). A model of medical reception – the planning and control of multiple task work. *Applied Cognitive Psychology, 9*, 81-114.

Hoc, J-M. (1988). *Cognitive Psychology of Planning*. London: Academic Press.

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

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.

Krishna, C. M. & Shin, K.G. (1997). *Real-time systems*. New York: McGraw-Hill.

Liu, J.W. (2000). *Real-time systems*. Englewood Cliffs, NJ: Prentice Hall.

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

McDermott, D. (1982). A temporal logic for reasoning about processes and plans. *Cognitive Science, 6*, 101-155.

Michon, J.A. (1990). Implicit and explicit representations of time. In R.A Block (Ed.). *Cognitive Models of Psychological Time* (pp. 37-58). Hillsdale, NJ: Lawrence Erlbaum Associates.

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.

Newell, A. & Simon, H.A. (1972). *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.

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.

Parasuraman, R., Sheridan, T.B. & Wickens, C.D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics, 30*, 286-297.

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.

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

Sheridan, T.B. & Verplank, W.L. (1978). *Human and computer control of undersea teleoperators* (Technical Report). Man-Machine Systems Laboratory, Department of Mechanical Engineering, Massachusetts Institute if Technology, Cambridge, Massachusetts.

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.

Svenson, O. & Maule, A.J. (Eds.). (1993). *Time Pressure and Stress in Human Judgement and Decision Making*. New York: Plenum Press.

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.

Wright, P. (1974). The harassed decision maker: Time pressure, distractions, and the use of evidence. *Journal of Applied Psychology, 59*, 551-561.