

Deciding When to Switch Tasks in Time-Critical Multitasking

Yelena Kushleyeva (yk45@drexel.edu)

Dario D. Salvucci (salvucci@cs.drexel.edu)

Frank J. Lee (fjl@cs.drexel.edu)

Department of Computer Science, Drexel University
3141 Chestnut St., Philadelphia, PA 19104

Abstract

While cognitive modeling has begun to make good progress in accounting for human multitasking behavior, current models typically focus on externally-driven task switching in laboratory-task settings. In contrast, many real-world complex tasks, particularly time-critical tasks, involve internally-driven multitasking in which the person her/himself decides when to switch between tasks. In this paper we propose an adaptation of the ACT-R cognitive architecture that incorporates a notion of elapsed time for the current goal and uses time to determine when to switch away from the current task. We demonstrate the usefulness of this mechanism in an application to a dynamic, time-critical dual search task, showing how an ACT-R model can account for various aspects of human subjects' switching behavior.

Introduction

In the push to generalize to increasingly complex real-world tasks, cognitive modeling and cognitive architectures have begun to address several important challenges, including direct interaction with realistic environments and complex integration of lower-level performance with higher-level planning and decision making. One of the most important challenges today involves making the leap from single-task laboratory experiments to real-world, time-critical tasks in which a person performs tasks together — in other words, addressing the fundamental yet ill-understood skill of human multitasking in time-critical tasks.

Recent modeling work on multitasking has focused primarily on situations where external cues are provided as a signal initiating switching between tasks. Such situations provide an excellent framework for analyzing the dynamics of the reaction time, and thus the distribution of attentional resources, at the time when switching between tasks takes place — for instance, the analysis of switching-time costs (STC) and the psychological refractory period (PRP) has been done for various types of tasks, discrete and continuous, successive and concurrent (Kieras et al., 2000). However, it has been argued (e.g., Burgess et al, 2000) that relying on external signals for switching between tasks is less common in real-life multitasking situations. Instead, many real-life scenarios, particularly in time-critical environments, involve *internally-driven* multitasking in which the person decides when to switch between tasks (e.g., the tasks of driving and dialing a cell phone). While the environment may indirectly define certain parameters of task switching, reliance on direct environmental feedback as

a signal to switch would imply processing that is too slow for time-critical situations (Rasmussen, 1983). In addition, environmental feedback is not even always available in real-life situations; it may also be significantly delayed and so only useful for adjusting future, but not current, behaviors. For this reason, the ultimate distribution of task switch points under time-critical conditions must be signaled by processes internal to human cognition.

The modeling of internally-driven multitasking proves to be an interesting challenge for current cognitive architectures such as ACT-R (Anderson et al., in press), EPIC (Meyer & Kieras, 1997), and Soar (Newell, 1990). One critical component of internally-driven multitasking in time-critical contexts is the ability to implicitly reason about time and the temporal aspects of each task. Intuitively this can be understood in terms of the increasing “pressure” over time that a person working on one task may feel to switch to another task. Possibly, this psychological phenomenon is what allows people to effectively balance cognitive resources across tasks. However, the current major architectures do not incorporate time-sensitive mechanisms that affect behavior on the scale of single-task execution.

To address this issue, we turn to the ACT-R cognitive architecture (Anderson et al., in press) and show that relatively minor changes to mechanisms already built into ACT-R allow us to implicitly reason about time and capture task-switching behavior in time-critical tasks. We evaluate the proposed architectural changes in the context of a simple experimental environment whose task demands are based on our analysis of the available neuropsychological tests (see Burgess, 2000). In developing an ACT-R model that incorporates the proposed changes, we found that the model nicely captured several aspects of people's temporal sensitivity in the task and their resulting behavior of when to switch tasks.

Deciding When to Switch Tasks in the ACT-R Architecture

As a step toward modeling internally-driven task switching in time-critical environments, we first examine some of the neuropsychological bases of real-world multitasking and explore the mapping of these concepts onto the ACT-R architecture. Given this insight, we then propose a small but powerful change to the current architecture that incorporates time into the task-switching decision and thus account for empirical phenomena such as those described in the subsequent sections of the paper.

ACT-R and Neuropsychological Bases

Analysis of the demands that real-life multitasking situations place on the individual allows us to identify three major skill sets whose involvement is crucial for satisfactory multitasking performance (Burgess et al., 2000):

- The ability to create and schedule future intentions.
- The facility to remember/maintain those intentions, as well as prioritize them.
- The ability to switch from carrying out one intention to another when the appropriate moment in time is finally reached.

In identifying the brain structures that are expected to be most significantly involved in producing these types of skills, we note two important points related to modeling and multitasking. First, both creation and maintenance of intentions, or goals, through prospective cognition rely heavily on the activity of the Dorsolateral Prefrontal Cortex (DLPFC), which immediately points to the possibility that this aspect of multitasking behavior should be taken care of by the existing ACT-R “goal buffer,” since the activity of this buffer is thought to account for neural processing in DLPFC (Anderson et al., in press). In this paper, we focus on the problem of when to switch away from the current goal in the buffer; in a complementary paper, we focus on the problem of choosing what goal to execute next (Salvucci et al., 2004).

Second, the ability to initiate switching between tasks and to do so at appropriate points in time is closely related to the notion of interval timing and prospective memory, and thus, anatomically, to basal ganglia (Nenadic et al, 2003; Diedrichsen et al, 2003). This connection to basal ganglia leads us to suggest that, in ACT-R, execution of prospective goals should be managed on the level the production system, which is thought to account for the activity of the basal ganglia (Anderson et al., in press). More specifically, we claim that such management centers on the production-level conflict resolution in ACT-R, which chooses the next production instantiation based on “expected gain” (described below). It is this last aspect of human multitasking and its possible ACT-R implementation in terms of conflict resolution that is our focus in this paper.

Procedural Memory and Conflict Resolution

The contents of procedural memory in ACT-R are composed of condition-action production rules. The “condition” part of each such rule is utilized in selection of appropriate productions to be activated; this process is accomplished through matching condition parts of the rules to the current states of the system buffers, including the goal buffer, declarative memory, and the buffers of perceptual modules. However, any given combination of buffer states may potentially match the conditions of multiple production rules; yet, the serial nature of ACT-R’s cognitive processing allows only one production to be activated within one cognitive cycle. For this reason, the conflict-resolution mechanism in ACT-R selects a single production rule out of the pool of rules whose condition parts match the current state of the system.

Conflict resolution attempts to select a production rule that maximizes the *expected gain* of the rule — namely, the gain that model expects to receive from executing the next production. Expected gain is defined as the value $PG-C$. Here, P is the probability that selecting a particular production will eventually lead to successful completion of the goal objective, G is the value of the goal, and C is the expected cost of completing the goal if a given production is chosen on this particular step (Anderson et al., in press). While the conflict resolution process attempts to select a production with the highest expected utility, ACT-R also includes logistic noise on expected gain to introduce some level of randomness to the rule selection process.

Analysis of conflict resolution within the ACT-R procedural memory shows that all parameters involved in conflict resolution are encoded through temporally static variables. The parameter G , representing the subjective value of successfully accomplishing a given goal and expressed as the amount of time the model (or person) is willing to spend on that particular goal, is currently kept static at 20 seconds for all models and all goals within those models. Parameters C (the cost of accomplishing the current goal given that a particular production is chosen) and P (expected probability of successfully achieving the objective of the goal if the given production is chosen) can be modified through the ACT-R utility learning mechanisms (Anderson et al., in press). However, such modifications take place only at such points that are explicitly marked as clear successes or failures in execution of the current goal, and thus such dynamics of the parameters C and P take place at a higher timescale than what is needed to make conflict resolution reflective of the passage of time within a single instance of a goal execution. Thus, the system of conflict resolution in ACT-R can currently be regarded as temporally static if looked at on the level of the execution of a single goal.

This aspect of the ACT-R architecture can be questioned in terms of psychologically plausibility, given that people are clearly capable of adjusting their behavioral strategies with constantly varying levels of time pressure (Maule et al, 2000). Thus, we propose slight modifications to the existing formulation of the conflict resolution parameters that allow us to resolve this issue and also enable us to rely on the same essential conflict resolution mechanism for deciding when to switch away from the current task, particularly in the context of time-critical multitasking.

Proposed Architectural Modifications

We propose to redefine the parameter G as a representation of the subjective value of continuing to execute the current goal. Since utilities throughout the ACT-R architecture are expressed in terms of time, we attempt to maintain consistency with this idea and define G more specifically as the amount of time the model is willing to spend on continuing attempts to achieve the current goal.

To implement this idea, we propose that the initial value of parameter G be set at goal creation — that is, when a production rule creates and sets the new declarative goal chunk in the goal buffer. The initial value of G represents

the duration of time that the model is willing to spend on achieving the goal in its entirety, from the start and to the moment the goal is abandoned in favor of another one. From the moment of goal creation, the value of parameter G linearly decreases with the passage of time, ultimately decreasing to 0 as the desired time expires. Such a redefinition of G removes the previously static temporal nature of conflict resolution: because production utilities are calculated as $PG-C$, as the time goes and the value of G decreases, the system will go from favoring productions that have high probability P of leading to a successful state, to favoring productions that promise low future execution costs C .

The resulting dynamics of the conflict resolution process are very interesting even simply because they allow models to change their strategies for performing tasks with the passage of time. This may affect models in both single and multiple task situations. But more importantly for us here, this new mechanism enables us to rely on conflict resolution for the initiation of switching between tasks. We can now introduce one extra production rule for every goal within a model, such that this new production, when activated, causes the goal to give up control. We further assign this production very low values of parameters P and C — which is intuitively reasonable, since switching away from a goal cannot lead to a successful state *within* that goal (thus low value of P), and similarly, switching away from a goal means that the cost of executing that goal immediately goes down to nearly zero (thus low value of C).

Low values of P and C will ensure that this newly introduced production rule will be activated at or near the end of the time period that the model intended to spend on execution of the corresponding goal: because P is low, this production will not likely be activated at the beginning of goal execution, since the relatively high initial value of G will cause the conflict resolution system to favor productions with high P (and thus high product PG); However, as the value of G goes down with time, conflict resolution becomes increasingly more biased towards selecting low-cost production rules, and as G approaches zero, a production with nearly-zero cost will likely to be selected even if its P value is nearly zero as well.

In addition to ensuring that switch times be near their intended values, reliance on conflict resolution mechanism allows us to utilize the built-in logistic distribution of the expected gain noise (Anderson et al., in press) to guide our predictions about the level of precision we may expect in the distribution of switch times.

Design of the Validation Task

The two major factors guiding us in designing the experimental task were the literature on neuropsychological tests available for evaluation of human multitasking skills and the need for such level of task simplicity that would facilitate integration of the task with computational models.

Guiding Principles

There is a strong understanding in psychological, particularly the neuropsychological, literature that skills

involved in human multitasking compose a subset quite distinct from the rest of the skills enabling executive functioning. For instance, it has long been shown that certain types of brain damage nearly exclusively affect the ability to carry out previously planned intentions, thus greatly impairing the person's ability to effectively function in the everyday life, while having very limited impact on performance on traditional standardized tests of general intelligence, and even on the tests of executive functioning (Shallice & Burgess, 1995).

This notion has led to development of specialized tests that evaluate the ability to execute prior intentions, such that the level of performance on those tests would reliably correlate with the person's ability to carry out basic everyday tasks involving multitasking, such as shopping, cooking, etc. Some of the tests thus developed were the Six Element Test (Shallice & Burgess, 1991) and the Greenwich Test (Burgess, 2000). Because both tests involve performing a range of fine motor movements that are difficult to model, we were not able to employ these tests directly; instead, we analyzed the procedures of the two tests to derive guiding principles for designing our own experimental procedure:

- Encourage subtask scheduling on the scale of seconds to minutes, rather than a sub-second scale, as this appears to be optimal for evaluation of the prospective memory aspect of human multitasking, and generally better represents the demands of real-life multitasking environments.
- Provide such rules and instructions that would create a continuing conflict between the need to stay with the current task for prolonged intervals, in order to maximize the given performance measure, and the need to switch between tasks, to prevent negative consequences on performance.
- Provide feedback on progress both along the time dimension and in terms of accuracy in order to allow people to tune their performance by possibly modifying future planning strategies; however, ensure that feedback information is only available after the corresponding action and only upon specific request by the subject, and thus cannot be utilized as a performance-guiding forward cue.
- In general, exclude any environmental cues that could be interpreted as external signals for initiation and/or termination of subtasks, thus ensuring that execution of previously scheduled subtasks is governed exclusively by internal signaling.
- Encourage task prioritization; enable the experimenter to manipulate task priorities throughout the course of the experiment.

Experimental Task

These guiding principles led us to the an experimental task and procedure which can describe briefly as follows. In the experiment, the subject was simultaneously presented with two similar tasks, of which only one could be active at any given point in time. The instructions were given to work on both tasks within the allocated time. Each task required the subject to repeatedly perform visual search on a sequence of



Figure 1: Screenshot of the experimental window.

letters with the goal of determining whether or not a specified target letter is present in the sequence, as shown in Figure 1. The subject was asked to respond to each search query by pressing one of the two designated keys, one key ('y') meaning that the target is present, and the other key ('n') meaning that target is absent from the sequence. In order to motivate the subject to answer accurately, a correctness score was calculated and displayed to the subject. Each correct response earned 1 point in the scoring, while an incorrect response deducted 10 points. Subjects were encouraged to answer as fast as they could, since that would allow them to reach a higher score within the limited time interval.

In order to ensure that the subject regularly switches between the two presented tasks and does not commit to working on only one of them, each task had a timer associated with it. Each timer was reset whenever the subject answered search queries of the task corresponding to this particular timer. However, whenever a task was abandoned in favor of the other task, the timer of the first task was able to run uninterrupted at constant speed up until the point where it reached its endpoint, where the timer would "expire," thus causing a very significant decrease (25 points) in the subject's performance score. The subject was instructed to treat the timers with high priority and to be sure to switch between tasks with a frequency high enough to avoid either of the timers ever reaching their expiration points.

The tasks were further complicated by the fact that the timer of the currently inactive task was hidden until that task was activated again. Thus, the subject had to work in the situation where he or she had to switch from one task to another without being able to rely on any external cues to guide in determining the appropriate time to do so.

In order to prevent the subject from adopting the strategy of switching between tasks after every single query, we manipulated the difficulty of the two tasks: the length of the letter sequence in one task was kept at 20, and in the other task it was either 5 or 9. The subjects were explicitly encouraged to spend as much time as possible on the easier task, with the notion that such strategy would allow them to reach a higher performance score within the time allocated for the experiment.

Participants and Procedure

In the experiment, which included 10 university students (2 women, 8 men), each participant was first instructed as to the basic structure of the task. Then, the participant performed two experimental blocks in each of two conditions: one in which the timer expired in 30 seconds, and another in which the timer expired in 15 seconds. In each block, the participant was allowed to practice the task for 5 minutes, and then performed 15 minutes of actual experimental trials. For simplicity, in this paper we present only the results for the 30-second timer condition; the results for the 15-second timer condition are reasonably similar and not discussed herein.

ACT-R Model

Given the experimental task described above, we set out to develop an ACT-R model that incorporates the proposed architectural mechanism, thus validating that the mechanism can indeed generate interesting emergent predictions in task-switching behavior.

Description

We first developed an isolated model that adequately performed the required search-and-answer procedure in a single-task context. The model adhered to the following action pattern:

- encode the target letter;
- sequentially search the letter sequence (the mechanism of simultaneous encoding and attention shifts was used to achieve a search speed high enough to match the one observed in people);
- if the target is encountered during search, respond by pressing the 'y' key;
- if end of the letter sequence is reached, respond by pressing 'n'.

We do realize that this model oversimplifies the real human interaction with the experimental environment. For instance, it does not incorporate the action redundancy that was observed in the general task procedure and especially in the search strategy (evidence for this comes in part from the fact that people were able to perform search significantly faster under increased time pressure, and this speedup was only




Figure 2: Comparison of experimental results and model predictions. Experimental data were obtained using ten subjects, each performing the experimental task for 15 minutes; to match these conditions, the model was also run for ten 15-minute intervals.

poorly correlated with the increase in error rate). However, its problems notwithstanding, this model provides a close enough approximation of human performance to allow us to focus our study on the initiation of task switches.

Enabling the model to perform two independent tasks within a given time interval by regularly switching between the tasks required only those changes described earlier: we had to introduce one extra production into each task with the intention that this new production, when activated, would cause the goal to cede control to the other task goal.

To fit this model to the behavior of the experiment participants, we varied one parameter, the value of ACT-R's expected gain noise, which was set to a value of 2.0 (instead of the default value of 0.5). We explain this by the fact that the conflict resolution system in our model was relying on perfect system time, while in people the corresponding processing system is likely to receive temporal information that is already noisy. And most importantly, the very idea of having the value of the current objective as a dynamic construct may by itself call for reevaluation of expected gain noise parameters. In addition, we preset the C values of the production rules that create a new goal to include the cost of this new goal; this additional cost was derived from the average time needed by the model to complete these goals. The final model results were compiled from 10 simulations of 15 minutes each (the same amount of data as for human subjects).

Model Predictions

The most important result that we obtained from experiments with the model was that the conflict resolution system relying on the redefined dynamic goal value parameter G is capable of producing a probability distribution for the times of switching between tasks that is very close to the one obtained for human subjects, as shown

in Figure 2, with $R=0.94$ (calculated only from 2-30 s). The shape of the distribution is an emergent prediction of the time-sensitive conflict resolution; the exact size of this distribution is a function of the expected gain noise.

In addition to achieving the desired shape of the probability distribution, we were also able to account for the fact that the majority of our human subjects were more likely to switch away from a task whenever continuing with that task required them to perform a discrete time-consuming subtask — that is, both the model and human subjects were more likely to switch before a 5-letter search than before a 9-letter search, as shown in Figure 3.




Figure 3: Probability of switching given that the next trial is a 5-letter or 9-letter search.

The model predicts this effect because the C values for subgoal-launching productions include the cost of the subgoal itself, thereby making the 9-letter search more costly and thus more likely to lead to a switch. Interestingly, although we are not modeling learning in the current model, we believe that such definition of the C parameters suggests a simple and intuitive way of learning conflict resolution parameters and nicely accounts for the tendency to avoid time-consuming tasks, for production rules.

Discussion

Overall the new mechanism of a time-sensitive conflict resolution seems to account nicely for the decision of when to switch tasks, particularly during time-critical performance. One aspect of multitasking not accounted for in this work, however, is *what* to do next — that is, once you switch away from the current task, deciding what task (or goal) should follow. In a companion paper (Salvucci et al., 2004), we propose a mechanism that generalizes the goal buffer into a goal set, scheduling goals according to “urgency” and thus help to decide what to do next. Also, Altmann and Trafton (2002) have proposed an activation-based model for retrieving the next goal from declarative memory. Such work on deciding what to do next is complementary to our work here and could eventually be integrated to form a more cohesive account of task scheduling and multitasking.

Another interesting aspect of the new mechanism emerged from the treatment of time as associated with individual goals. The current mechanism focuses on the time elapsed for an individual goal, but time sensitivity does not currently cross goal boundaries. Another approach would allow the G value of the current goal to influence the G value for subgoals. The older ACT-R goal stack (Anderson & Lebiere, 1998) discounted subgoals’ G values when pushed onto the stack; however, this (now defunct) mechanism had some aspects of this propagation but was somewhat rigid in implementation. While we believe that our mechanism has the potential to generalize to such phenomena, more in-depth studies could clarify how the propagation of time sensitivity across goals might be incorporated into the architecture.

Acknowledgments

This work was supported by ONR grant N00014-03-1-0036 to the second author. We thank Alex Chavez for helpful ideas and suggestions.

References

- Altmann, E. M., & Trafton, J. G. (2002). Memory for goals: An activation-based model. *Cognitive Science*, 26, 39-83.
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (in press). An integrated theory of the mind. *Psychological Review*.
- Anderson, J.R., & Lebiere, C. (1998). *The atomic components of thought*. Hillsdale, NJ: Erlbaum.
- Burgess, P.W. (2000). Strategy application disorder: the role of the frontal lobes in human multitasking. *Psychological Research*, 63, 279-288.
- Burgess, P.W., Veitch, E., De Lacy Costello, A., Shallice, T. (2000). The cognitive and neuroanatomical correlates of human multitasking. *Neuropsychologia*, 38, 848-863.
- Diedrichsen, J., Ivry, R.B & Pressing J. (2003). Cerebellar and basal ganglia contributions to interval timing. In Meck, W.H. (Ed.), *Functional and Neural Mechanisms of Interval Timing*. 457-483. CRC Press.
- Kieras, D.E., Meyer, D.E., Ballas, J.A., & Lauber, E.J. (2000). Modern computational perspectives on executive mental processes and cognitive control: Where to from here? In S. Monsell & J. Driver (Eds.), *Control of Cognitive Processes: Attention and Performance XVIII* 681-712. Cambridge, MA: MIT Press
- Maule, A., Hockey, G. R. J., Bdzola L (2000). Effects of time pressure on decision making under uncertainty: changes in affective state and information processing strategy. *Acta Psychologica*, 104, 283 – 301.
- Nenadic, I., Gaser, C., Volz, H., Rammsayer, T., Hager, F., Sauer, H. (2003). Processing of temporal information and the basal ganglia: new evidence from fMRI. *Experimental Brain Research*, 148, 238-246.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Rasmussen, J. (1983) Skills, rules and knowledge; signals, signs and symbols, and other distinctions in human performance models. *IEEE Transactions: Systems, Man and Cybernetics*, SMC-13, 257-267.
- Salvucci, D. D., Kushleyeva, Y., & Lee, F. J. (2004). Toward an ACT-R general executive for human multitasking. To appear in Proceedings of the Sixth International Conference on Cognitive Modeling.
- Shallice T., Burgess P.W. (1995). Supervisory control of action and thought selection. In: Baddeley A.D., Weiskrantz L. (Eds.) *Attention: selection, awareness and control: a tribute to Donald Broadbent*. Oxford: Clarendon Press.
- Shallice T., Burgess P.W. (1991). Deficits in strategy application following frontal lobe damage in man. *Brain*, 114, 727-741.