THE QUARTERLY JOURNAL OF EXPERIMENTAL PSYCHOLOGY, 2013
Vol. 66, No. 1, 160–178, http://dx.doi.org/10.1080/17470218.2012.702117

Routledge
Taylor & Francis Group

Paying the price works: Increasing goal-state access cost improves problem solving and mitigates the effect of interruption

Phillip L. Morgan[1] and John Patrick[2]

[1]School of Sport, Health and Applied Social Sciences, University of Wales, Newport, Caerleon, Newport, UK
[2]School of Psychology, Cardiff University, Cardiff, UK

The aim of this paper was to investigate whether it was possible to induce more internal planning in the four-disk Tower of Hanoi (ToH) in order not only to produce more efficient problem solving but also to make it more resistant to the negative effect of interruption. The theoretical frameworks of soft constraints and the memory for goals model underpinned Experiments 1 and 2. In both experiments, three goal-state access cost conditions were used: high (mouse movements and 2.5-s delay), medium (mouse movements) and low (goal state always available). In Experiment 1, more memory-based planning was induced by the high cost condition, which resulted in fewer moves to solution and the gradual development of an efficient subgoaling strategy, resulting in more perfect solutions. In Experiment 2, the same condition protected performance against a 10-s interruption irrespective of the interrupting task (blank screen, mental arithmetic, or three-disk ToH). The more memory-based planning strategy, induced by high access cost, presumably strengthened participants' goals during planning and problem solving, making them less susceptible to decay and interference from interruption. These novel results are discussed in the context of other recent studies.

Keywords: Planning; Problem solving; Interruption; Access cost; Soft constraints; Memory for goals.

Planning is intrinsic to problem-solving efficiency even though it rarely occurs spontaneously (e.g., Atwood & Polson, 1976; Delaney, Ericsson, & Knowles, 2004), and people are reluctant to use internal as opposed to external memory resources for this purpose (e.g., Hayes-Roth & Hayes-Roth, 1979). The aim of this paper is to investigate whether it is possible to induce more internal planning that produces not only more efficient problem solving but also makes it more resistant to forgetting following interruption. The rationale for inducing a shift in the extent of memory-based planning comes from the theory of soft constraints (Gray & Fu, 2004; Gray, Sims, Fu, & Schoelles, 2006). This theory states that a more internal memory-based strategy can be effected by

Correspondence should be addressed to John Patrick, School of Psychology, Cardiff University, Cardiff, CF10 3AT, UK. E-mail: Patrickj@cardifff.ac.uk

This work was part funded by a grant from the UK Data and Information Fusion Defence Technology Centre awarded to the second author. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the UK Ministry of Defence or the UK Government. We are grateful to Sophia King, Olivier DeCondappa, Rebekah Chadwick, Tanya Patrick, and Leyanne Tiley for running participants and/or assisting with data coding. Some data from Experiment 2 appeared in the Proceedings of the 2010 Annual Meetings of the Cognitive Science Society and the Human Factors and Ergonomics Society (HFES).

the imposition of an appropriate access cost on the availability of important task-related information, which in the present study is the goal state of the Tower of Hanoi (ToH). Recently, a study by Waldron, Patrick, and Duggan (2011) found that it was possible to increase internal planning by this means but this did not result in improved problem-solving efficiency. The present study not only reexamines this issue using the more strategic ToH problem-solving task but also proposes that more internal planning will strengthen goal representation and therefore increase the resistance of problem solving to interruption, as suggested by the memory for goals model (Altmann & Trafton, 2002, 2007). To date, mitigation of the effect of interruption using a similar strategy has only been examined and found to be successful in the context of copying a visuospatial pattern in the Blocks World Task (Morgan, Patrick, Waldron, King, & Patrick, 2009). Therefore, the second experiment in this paper investigates whether this finding can be extended to a different type of task—namely, problem solving.

Exploiting the theory of soft constraints

In the tradition of rational analysis (e.g., Anderson, 1990; Anderson & Milson, 1989; Chater & Oaksford, 1999; Oaksford & Chater, 1998), the theory of soft constraints (Gray & Fu, 2004; Gray et al., 2006) focuses on interactive behaviour and proposes that low-level task strategies made up of cognitive, perceptual, and motor elements are selected on the basis of minimizing local time costs. The idea is that milliseconds matter (Gray & Boehm-Davis, 2000), and adjustments to strategy occur at the one-third- to 3-s level of task performance with more effective strategies replacing those that are less effective in minimizing performance time at a local rather than a global level (Gray et al., 2006). Such strategies are affected not only by the hard constraints of the task environment (e.g., having to click an onscreen button to access a computer function) that are fixed and determine what interactive behaviour is or is not possible, but also by soft constraints (e.g., the decision of how frequently to access the function), which are determined by strategy selection. The person has no control over the hard constraints imposed by the task but does have control over the soft constraints by choosing how to tackle the task and the nature of the strategy to be employed. The theory states that strategy is flexible and will adapt rationally to small changes at the millisecond level in how information is accessed within the task environment (Gray & Fu, 2004; Gray et al., 2006). When information is permanently available within the task environment, a strategy will prevail that relies on using this environment as an external memory resource. This has the advantage of reducing demands on internal memory. In contrast, if there is a small time delay associated with accessing information within the task environment, people will react adaptively and switch to a more internalized strategy that entails encoding more information in memory. This has the advantage of minimizing the need to access information and pay the associated time cost. Therefore, manipulation of the cost of accessing information can be exploited to affect the extent that an internal memory-based strategy is selected (Fu & Gray, 2000; Gray et al., 2006; Morgan et al., 2009; Waldron, Patrick, Morgan, & King, 2007).

There is extensive evidence in support of the theory of soft constraints from tasks requiring a significant memory resource (e.g., Fu & Gray, 2000; Gray & Fu, 2004; Gray et al., 2006; Morgan et al., 2009; Waldron et al., 2007). For example, in the Blocks World Task the aim is to recreate a target pattern of coloured blocks in a workspace window. Increasing the cost of accessing the target pattern in this visuospatial copying task by requiring a mouse movement and by imposing a time delay of only a few seconds induced a reliable shift to a more memory-based strategy (e.g., Fu & Gray, 2000; Gray et al., 2006; Morgan et al., 2009; Waldron et al., 2007). Similar effects of increased information access cost have also been demonstrated for a VCR procedural programming task (e.g., Gray & Fu, 2001). However, less is known about the effect of increased access cost on problem solving and whether this will induce greater planning that, in turn, will increase problem-solving efficiency.

The effect of increased access cost on planning in problem solving

The present study uses the theoretical framework of the theory of soft constraints (Gray et al., 2006) and manipulation of the key factor in this theory —that is, increased time to access information (here the goal state)—to examine whether more memory-based planning can be induced in problem solving.

This differs from the study by O'Hara and Payne (1998) that was carried out under the more general theoretical framework of rational analysis (e.g., Anderson, 1990; Anderson & Milson, 1989) and focused on the effect of increasing the cost of making a move rather than that of accessing information. Indeed the cost of making a move involved both a significant motor component of typing a command string and an associated time to do this, although this time was not separately analysed. Despite these differences with the present study, it is important to note that O'Hara and Payne (1998) demonstrated an increase in planning and a reduction in number of moves to solution when the cost of implementing a move in the eight-puzzle involved typing a string of characters (high cost) compared to a single key press (low cost). Evidence of more planning under high implementation cost also came from longer inter-move latencies and concurrent verbalizations. This effect was replicated in a follow-up study, where O'Hara and Payne (1999) operationalized high implementation costs by increasing the time to undo a move in the slide-jump puzzle or delaying the next move in the eight-puzzle. In both cases higher implementation costs associated with executing moves led to more planning and shorter solution paths.

The only problem-solving study we are aware of that has investigated the effect of varying information access cost, as opposed to implementation cost, on planning is by Waldron et al. (2011). In this study, the access cost in viewing the goal state was manipulated, in contrast to the cost of move making in the O'Hara and Payne studies discussed above. Waldron et al. (2011) argued that manipulation of goal-state access cost is likely to have a more direct effect on planning than implementation costs. Unless the goal state has been memorized, effective planning can only occur when the current and goal states are both simultaneously visible in such transformation problem-solving tasks. Thus, manipulating goal-state access cost enables planning to be more directly affected at the perceptual input side of the problem-solving cycle as opposed to imposing an implementation cost that occurs at the motor output side of the problem-solving cycle (i.e., executing a planned action that requires a sequence of key presses to implement). The study by Waldron et al. (2011) found that the time to make a mouse movement coupled with a 2.5-second delay to view the goal state (the high access cost condition) led to more "planning before action" as opposed to "planning during action". Planning before action occurs when a sequence of moves is evaluated, encoded, and then executed from memory compared to planning during action that is more display based involving the execution of single moves in an opportunistic manner. However, although participants in the high goal-state access condition engaged in more planning before action, there was no improvement in problem-solving efficiency as indicated by the number of moves to solution in either of the two problem-solving tasks. Perhaps this lack of improvement in problem-solving efficiency in Experiment 1 was not surprising as the participants had available many straightforward and easy pathways for moving blocks to match any given goal state. Hence there was little incentive or advantage in trying to learn and use an efficient strategy that could display the benefit of planning. However, this was not true of Experiment 2 because it used a task very similar to the "eight-puzzle" for which some strategies are more efficient, and planning therefore can reduce moves to solution (Ericcson, 1974; O'Hara & Payne, 1998). Nevertheless, Experiment 2 of Waldron et al. (2011) failed to find an effect of increased goal-state access cost on number of moves to solution. Possible explanations are that encoding and retrieving the goal state in this task imposed memorial demands that were too high for advantage to be taken of a planned and more efficient strategy or that more trials were needed for such a strategy to be fully developed.

Manipulating planning in the ToH

In the present study, in order to examine whether planning and efficiency can be improved in a problem-solving task by manipulating goal-state access cost, the four-disk ToH was used because the memorial demands of encoding and retrieving the target pattern are relatively modest in comparison with those associated with that of the eight puzzle. Also, when a 15-move solution is achieved using an efficient subgoaling strategy (e.g., Anzai & Simon, 1979; VanLehn, 1991), the ToH can be characterized as constituting a series of three subgoals that concern relocating each of the largest-out-of-place (LOOP) disks to their goal destination peg(s) in the order largest through to smallest. Planning would be expected to facilitate both the development and execution of such a strategy, resulting in more efficient problem solving. Indeed, this was confirmed by Davies (2003) who enforced an initial planning period and found that this reduced both time and moves to solution. However, Davies did caution that planning may not be helpful for more complex versions of this task, although his means of inducing planning differed from the manipulation of goal-state access cost used in the present study.

A potential difficulty of using the ToH is whether the goal-state access cost will be sufficient to induce a more internal memory-based planning strategy. In Waldron et al. (2011) Experiment 1, the task was to transform a 4 × 4 grid of 10 coloured blocks and 6 empty spaces into a goal pattern, and, in Experiment 2, the array was reduced to a 3 × 3 grid but planning was more complex given that there was only one empty square (similar to the eight-puzzle). In both these tasks, remembering the goal pattern of blocks and any plans of how to move blocks to achieve this goal-state imposed memory demands necessitating a number of visits to view the goal state and therefore payment of several access costs. The goal state was accessed over 4 and 12 times on average per problem by participants in the high access cost conditions of Experiments 1 and 2, respectively, in the study by Waldron et al. (2011), representing extra total time costs of approximately 13 and 37-s in these two high goal-state access cost conditions. In contrast, in the ToH used in the present study, the goal-state configuration of four disks should be much easier to memorize, and although this should make planning without the goal state easier than in the eight-puzzle, we do not know whether the cumulative time cost to access the goal state will be sufficient to induce more internal memory-based planning in the high access condition and consequently more efficient solutions.

Summary

Our aim is to establish whether increased goal-state access costs, after the theory of soft constraints, can be deployed to induce more memory-based planning and consequently improve problem-solving efficiency through the development of an efficient subgoaling strategy without any associated decrements in time and accuracy found in studies using the Blocks World Task (e.g., Fu & Gray, 2000; Gray et al., 2006; Morgan et al., 2009). This may be facilitated by use of the ToH task in which the goal state should be easier to remember than the more complex goal-state patterns used by Waldron et al. (2011), and this will make more resources available for planning. This was the aim of Experiment 1, and it was predicted that increased goal-state access costs would increase both initial planning and problem-solving efficiency, the latter being indicated by the number of moves to solution and the frequency of use of an efficient subgoaling strategy. The aim of Experiment 2 was to investigate whether the use of more memory-based planning, induced by increased access cost, would enable problem solving to be more resilient to the negative effects of interruption. This is another novel aspect of the current study, and we predicted that increased goal-state access cost will support the ability to maintain more efficient planned move sequences, thus leading to shorter solutions following interruption.

EXPERIMENT 1

Experiment 1 compared the effect of a high goal access cost condition (a mouse movement time to and from the goal-state window plus an access cost of 2.5-s lockout time to view the goal state) against a medium (a mouse movement to view the goal state) and a low goal access cost condition (goal state always available) on memory-based planning and problem-solving efficiency in the four-disk ToH. First, we calculate whether actual access times varied between the high and medium access cost conditions and, if they did, whether there is any difference in the number of goal-state visits between these conditions. We predict that there will be fewer visits to the goal state in the high than in the medium access cost condition that is a reflection of more memory-based planning. The extent that participants engaged in initial internal planning will be assessed by the time spent before making the first move (excluding any access costs), which has been found to be an important and sensitive measure in previous studies (e.g., Davies, 2003; Waldron et al., 2011). Finally, we examine the prediction that problem-solving efficiency, indicated by moves to solution and use of an efficient subgoaling strategy, is best in the high goal access cost condition because participants in this condition should engage in more internal planning in trying to avoid paying extra time costs associated with accessing the goal state. A corollary of this is that we predict that the increase in planning in the high goal access cost condition will facilitate the development of an efficient subgoaling strategy across trials.

Method

Participants

Forty-two Cardiff University students participated in the study for course credit. Colour-blind individuals were not used. Also students who were knowledgeable about or had performed the ToH or a similar task (as indicated by questionnaire response) were screened out. Thirty-six participants remained and were randomly assigned to one of three conditions. Of these participants, 29 were women, and 7 were men with ages ranging between 18 and 38 years ($M = 20.00$, $SD = 3.76$).

Materials

The experiment was programmed in Microsoft Visual Basic 6.0 and run from a Pentium Dual Core PC linked to a 22″ LCD monitor. The program recorded mouse movements, clicks, and key presses. Twenty ToH problems were used, each containing four coloured disks varying in size, located across three pegs, and each with different start- and goal-state disk configurations. The start or current state was presented in a workspace window at the bottom of the screen, and the goal state was presented in a same-sized window at the top of the same screen. Subgoals involved relocating each of the largest-out-of-place (LOOP) disks to their goal destination peg(s) in the order largest through to smallest. If solved error-free using an efficient subgoaling strategy, Subgoals 1, 2, and 3 consisted of Moves 1–8, 9–12, and 13–14, respectively, with Move 15 involving the movement of the smallest disk to its goal destination peg to complete the problem. Disk movement involved a mouse-controlled drag-and-drop facility within the workspace window whereby a disk was dragged from a source peg and dropped onto a destination peg. If a larger disk was dragged to a peg in which the top (or only) disk was smaller, the larger disk would be automatically returned to its source peg with an "illegal move" message. Each ToH problem had to be solved correctly before moving to the next problem. The program automatically detected when the current disk configuration in the workspace window matched that of the goal state and then instructed participants to click a button (located in the middle of the screen) when ready to proceed to the next problem. No feedback was given during or at the end of a trial concerning any aspect of performance.

Design and procedure
Goal-state access cost was manipulated between participants and had three levels: low (goal state permanently visible); medium (goal state covered with a grey mask that disappeared immediately when the mouse cursor was moved into the goal-state area and reappeared when the mouse cursor was moved out of the goal state area); and high access cost (as in medium but with an extra 2.5-s lockout cost for the mask to disappear from the goal state when the mouse cursor was in the goal-state area). Consequently, the extra time cost for viewing the goal state in the medium access cost condition concerned the time to place the mouse cursor in the goal-state area (0.23-s) and to return it into the workspace window containing the permanently visible current state (0.23-s). These average mouse cursor movement time costs were estimated using a variation of Fitts' law (MacKenzie, 1992), detailed in the Appendix. In total, one visit to and from the goal state in the medium access cost condition would come at a time cost of 0.46-s. The high cost condition also involved the extra time to move the mouse cursor to the goal-state window (0.23-s), the 2.5-s lockout cost associated with waiting for the mask to disappear from that window measured from when the mouse cursor entered the goal-state area, and the time to move the mouse cursor back into the workspace window (0.23-s), which permanently displayed the current state of the problem. In total, one visit to and from the goal state in the high access cost condition would come at a time cost of 2.96-s. It was predicted that planning would increase in this condition, indicated by fewer visits during problem solving to the goal state and by a longer latency to make the first move (time from the start of a trial to the execution of the first move, excluding any access cost time). Measures taken of problem-solving efficiency were: the number of moves to complete problems; the number of problems solved in the minimum 15 moves using an efficient subgoaling strategy; and problem completion time.

Participants were tested individually and were given written instructions on how to operate in their assigned condition and how to perform the ToH task. Participants were not instructed with respect to any specific performance criterion (e.g., speed, accuracy, number of moves). One 4-disk 15-move ToH problem in the corresponding access cost condition was given as practice.

Results and discussion

Given that the purpose of this experiment was to investigate the effect of goal-state access cost on planning and efficiency in problem solving, we initially needed to verify that the total access time costs, involving the actual total lockout and movement time costs per problem did indeed vary between the medium and high access cost conditions. This was calculated by the number of visits per problem to view the goal state multiplied by the time required to access the goal state on one visit in the medium access cost condition (0.46-s) and the high access cost condition (2.96-s). (For further details see the Method and Appendix.) The average total time per problem that participants spent accessing the goal state was indeed longer in the high than in the medium access cost condition ($M = 4.34$-s, $SD = 1.03$; $M = 1.47$-s, $SD = 0.55$, respectively), $t(22) = 8.53$, $p < .001$, $d = 3.64$. Therefore the access cost manipulation resulted in a significant difference in actual total time required to access the goal state during each problem, which Gray et al. (2006) suggested is the critical variable affecting the extent that a memory-based strategy is deployed.

Effects of goal access cost on memory and planning
The next question is whether these differences in access cost were sufficient to induce a more memory-based planning strategy in the high than in the medium and low access cost conditions. Our first indicator is the number of goal-state visits made by participants in the high and medium access cost conditions (such data were not available in the low condition where the goal state was always unmasked and visible). A more memory-based planning strategy would be reflected in a reduced need to visit the goal state. Consistent with this prediction, participants in the medium

access cost condition made more than twice as many goal-state visits per problem than those in the high access cost condition (Table 1), and this difference was significant, $t(22) = 4.82$, $p < .001$, $d = 2.1$.

A second indicator, consistent with more memory-based planning, is the time spent before making the first problem-solving move, which would correspond to making the first move in Subgoal 1 if an efficient subgoaling strategy was used. The results indicate that participants spent several extra seconds prior to executing the first move in the high access cost condition (Table 1). An analysis of variance (ANOVA) was carried out with goal access costs removed, and a difference was found between access cost conditions, $F(2, 33) = 3.61$, $MSE = 11.15$, $p < .05$, $f = 0.47$, although only the difference between the high and the low goal access cost conditions was significant ($p < .05$).

Therefore, from the evidence concerning fewer goal-state visits and an increase in time to make the first move, we infer that more memory-based planning occurred in the high access cost condition.[1] We consider next whether this shift in degree of strategy use can be linked to improvements in problem-solving efficiency.

Effects of goal-state access cost on problem-solving efficiency

It was proposed that increasing the cost of viewing the goal state would encourage greater problem-solving efficiency indicated by a reduction in the number of moves required to complete problems as well as an increase in the number of problems solved using an efficient subgoaling strategy in the minimum of 15 moves. In addition, we predicted that greater learning and execution of an optimal 15-move solution would develop across trials in the high than across those in the low and medium access cost conditions because of the increase in memory-based planning.

On average, participants in the high access cost condition made approximately 3 fewer moves per problem and solved more than twice as many problems in the minimum 15 moves using an efficient subgoaling strategy than participants in the medium and low access cost conditions (Table 1). ANOVAs found significant main effects of goal access cost for both number of moves to solution, $F(2, 33) = 6.49$, $MSE = 10.37$, $p < .01$, $f = 0.82$, and problems solved in 15 moves, $F(2, 33) = 6.82$, $MSE = 21.17$, $p < .01$, $f = 0.64$. Participants in the high access cost condition demonstrated improved problem-solving efficiency on both of these measures compared to those in both the medium and low goal access cost conditions ($ps < .05$), and there were no significant differences between those in the medium and low conditions ($ps > .05$).

A further analysis examined whether high goal access cost encouraged faster learning of an optimal strategy across trials than the medium and low access cost conditions. The average frequency of optimal solutions was calculated for four blocks of trials (i.e., 1–5, 6–10, 11–15, 16–20) for the three access cost conditions (Figure 1). There was a trend for an increase in perfect solutions across trial blocks in the high but not in the medium and low goal-state access cost conditions. A 3 (access cost condition) × 4 (blocks) mixed ANOVA confirmed not only a significant main effect of access cost, $F(2, 33) = 6.33$, $MSE = 0.05$, $p < .01$, $f = 0.63$, with high being superior to medium and low (ps $< .05$), but also a significant interaction with trial block, $F(6, 99) = 2.92$, $MSE = 0.03$, $p < .025$, $f = 0.42$. Simple main effect analyses confirmed that there was a significant improvement across trial blocks in the high goal-state access cost condition ($p < .001$) but not in the medium or low conditions ($ps > .05$). More specifically, Bonferroni comparisons found a significant improvement under high access cost in optimal performance from Trial Block 1 (1–5) to Trial Blocks 3 (11–

[1] A further analysis of move latency was carried out at what would have corresponded to the implementation of the first move at Subgoals 2 and 3 (i.e., after moving the largest-out-of-place disk to its goal peg and the second largest-out-of-place disk to its goal peg, respectively). There was no significant main effect of goal-state access cost, $F(2, 33) < 1$, and no significant interaction between goal-state access cost and subgoal, $F(2, 33) < 1$.

Table 1. *Effect of goal-state access cost on planning and problem-solving efficiency, Experiment 1*

	Access cost					
	High		*Medium*		*Low*	
Planning and problem solving measures	M	SD	M	SD	M	SD
Planning						
Number of visits to view the goal state	1.47	0.35	3.18	1.18		
Time spent before making first move (excluding 2.5-s lockout time in high and extra mouse movement time to and from the goal-state window in high and medium cost conditions) (s)	11.21	3.14	9.12	3.82	7.66	3.54
Problem-solving efficiency						
Number of moves to complete problems (min. 15)	18.21	1.93	21.68	3.82	22.74	3.58
Number of problems solved in the minimum 15 moves (max. 18)	9	4.99	4	4.75	2.33	4.01
Time to complete a problem (including 2.5-s lockout time in the high cost condition and extra mouse movement time to and from the goal window in high and medium conditions) (s)	49.48	6.76	55.74	18.26	53.3	13.21
Time to complete a problem (excluding 2.5-s lockout time in the high cost condition and extra mouse movement time to and from the goal window in high and medium conditions) (s)	45.14	6.87	54.27	18.14	53.3	13.21

15) and 4 (16–20), *ps* < .05, and no improvement for the low and medium access cost conditions across any trial block.

Past studies in which goal-state access cost has been increased have found that performance benefits are accompanied by longer task completion times (e.g., Fu & Gray, 2000; Morgan et al., 2009). However, in the present study, improved problem-solving efficiency under high access cost did not come at any extra cost in terms of time to complete problems (Table 1). This was confirmed by ANOVAs of completion times with access costs included, $F(2, 33) < 1$, and excluded, $F(2, 33) = 1.64$, $MSE = 183.51$, $p > .05$, $f = 0.32$. Even though such null results must be interpreted with caution, it is important to note that the average completion times on both these measures for participants in the high access cost condition have a nonsignificant trend to be *lower* than those in the medium and low access cost conditions (Table 1).

The overall results from Experiment 1 for participants in the high access cost condition were fewer visits to the goal state, an increase in time to make the first move, improved problem-solving efficiency in the number of moves to solution, and accelerated learning of an optimal subgoaling strategy. Together, these results provide evidence of a more cognitively intensive and successful memory-based planning strategy in the high goal-state access cost condition that gradually became more efficient across trials, partly as a consequence of more planning before the first move. Our results therefore bear on the traditional question concerning strategy acquisition that has been of concern to problem-solving researchers (e.g., Anzai & Simon, 1979; VanLehn, 1991). In our Experiment 1 with participants naïve to the ToH, relatively few perfect solutions were executed in the medium goal state access cost condition (22%) and even fewer in the low condition (13%). In contrast, participants in the high goal-state access cost condition not only solved 50% of all ToH problems error free but were the only group of participants to demonstrate a clear acquisition of an efficient subgoaling strategy over trials (Figure 1).

The novel finding in comparison with the results of Waldron et al. (2011) is a reduction in moves to solution by participants in this condition coupled with an increase in not only the use of an efficient subgoaling strategy but also its progressive deployment across trial blocks. Also this advantage in problem-solving efficiency did not come with the

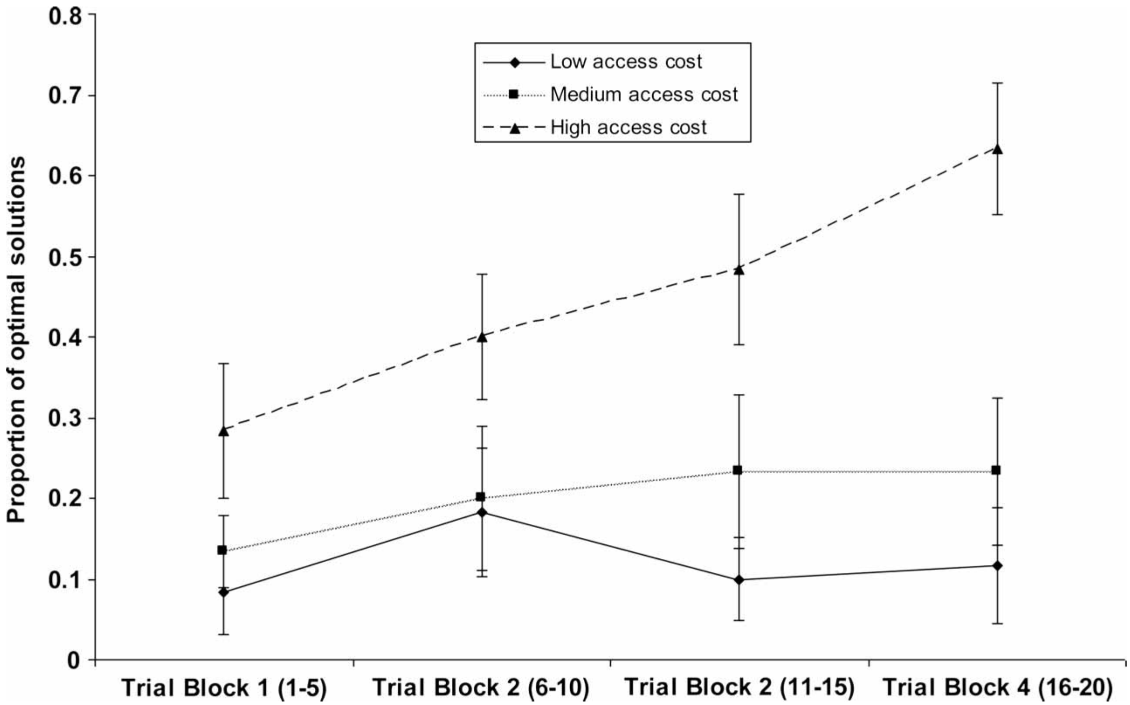

Figure 1. *The effect of goal-state access cost on the proportion of optimal solutions across trial blocks (Experiment 1). Error bars represent standard error.*

disadvantage of a reduction in the speed of completion, as found in studies using the Blocks World copying task (e.g., Fu & Gray, 2000; Morgan et al., 2009; Waldron et al., 2007). Manipulation of time to access the goal state had a powerful effect on performance of the ToH as imposition of only a few extra seconds associated with accessing the goal state, which was only visited on one or two occasions per problem, was sufficient to induce a more memory-based planning strategy that used more internal rather than external memory resources. Consequently these results affirm the potency of increasing goal-state access cost on moves to solution and highlight the utility of the theoretical framework of the theory of soft constraints (e.g., Gray et al., 2006) in the context of problem solving. The result concerning the change in the degree of memory-based planning coupled with the development and use of an optimal subgoaling strategy under high access cost is consistent with the recent exhortation by Fu (2011) to researchers in this area to identify the changes in strategy and processing as a consequence of dynamic shifts in the internal and external environments.

EXPERIMENT 2

The aim of Experiment 2 was to assess whether the same manipulation of increased goal-state access cost will improve resistance to the negative effects of interruption during problem solving as might be predicted from the memory for goals model (Altmann & Trafton, 2002, 2007). Generally interruptions have been found to produce universally negative effects in different contexts (Hodgetts & Jones, 2006b; McFarlane, 2002; Monk, Trafton, & Boehm-Davis, 2008; Trafton & Monk, 2008). However a study by Morgan et al. (2009) found that an increase in goal-state access cost can be used as a method for mitigating the effects of interruption (particularly forgetting what one was doing or intended to do prior to interruption) using the visuospatial Blocks World copying task. Experiment 2 seeks to establish whether this finding can be extended to another category of task—namely, problem solving and also whether the improved problem-solving efficiency, found in Experiment 1, can be maintained post interruption in the high access cost condition.

The memory for goals model (Altmann & Trafton, 2002) is particularly useful in explaining interruption effects. When a problem-solving task is interrupted, goals have to be suspended and usually resumed at some point in the future. Contrary to the once popular view that goal memory is infallible (e.g., Anderson, Kushmerick, & Lebiere, 1993; Miller, Galanter, & Pribram, 1960), goals committed to memory run the risk of being forgotten when they are suspended (e.g., Altmann & Trafton, 2002; Anderson & Douglass, 2001). According to the memory for goals model, there are two key processes that are essential for retrieving and resuming a suspended goal from memory. First, a goal can only govern behaviour if its memorial representation is more active than an interference level, determined by the activation strength of other goals. Activation is determined by the number of times a goal has been strengthened by encoding and rehearsal and how much it has decayed since it was last processed (e.g., Peterson & Peterson, 1959). A suspended goal that has been insufficiently strengthened or not used for some time risks being forgotten quite rapidly. Second, the suspended goal must be associatively primed with a salient mental or physical reminder cue that has to be available both immediately prior to goal suspension and when the goal is to be retrieved from memory. In their ToH example, Altmann and Trafton (2002) posit that the disk corresponding to the to-be-suspended goal could function as an ideal physical priming cue. More recently Patsenko and Altmann (2010) have suggested from eye movement recordings during the ToH not only that the disks themselves may not necessarily always serve as retrieval priming cues but that performance can be explained by constructs concerning perception and attention rather than more centrally controlled plans and goals. However, unlike the present study, participants in the Patesenko and Altmann study were instructed in and practised the goal-recursion strategy prior to solving the problems.

Our expectation from the memory for goals model is that a person's goal or subgoal(s) in performing the ToH will be strengthened by increased encoding brought about by the extra memory-based planning induced in the high access cost condition, and this, in turn, will mitigate the effects of interruption. In the high access cost condition, the problem solver is more actively engaged in seeking out goal-state information in order to plan the next action or sequence of actions, and the necessary strengthening and priming processes for the suspended goal to survive interruption are therefore more likely to have occurred. Indeed, from the findings of Experiment 1, increased access cost resulted in better memory-based planning and more efficient problem solving. Therefore, in the high access cost condition, we predict that more memory-based planning will not only provide greater protection from the negative effects of forgetting as a result of interruption but will also support more efficient problem solving. This is because either planned moves prior to interruption will be better remembered and/or the goal state will be better recalled and used for planning post interruption. This should be manifested in a greater ability for participants in the high access cost condition to resume more interrupted problems without the need to immediately revisit the goal state together with more moves made before the need to revisit the goal state (if at all). In addition, we examine whether any greater resistance to the negative effects of interruption manifested by participants in this condition will still be accompanied by superior problem-solving efficiency, in terms of number of moves to solution, as found in Experiment 1.

Method

Participants

Sixty-five Cardiff University students participated in the study for course credit. Colour-blind individuals were excluded. As in Experiment 1, all students who were knowledgeable about or had performed the ToH or a similar task (as indicated by questionnaire response) were screened out. Fifty-four participants remained and were randomly assigned to one of three conditions. Of these participants, 48 were women, and 6 were men, with an age range of 18 to 37 years (M = 20.56, SD = 2.95).

Materials

A total of 18 four-disk problems were constructed, each with a different start- and goal-state configuration, and all problems could be solved with a minimum of 15 moves. Half of the problems were not interrupted, and the other half were interrupted on one occasion only. Many interrupting tasks have been used in the literature (e.g., McFarlane, 2002; Trafton & Monk, 2008), although it is notoriously difficult to calibrate their potential for disrupting the primary task. Some interrupting tasks have been selected because of either their high processing demands (e.g., mental arithmetic, Hodgetts & Jones, 2006b; n-back task, Monk et al., 2008) or their apparent similarity to the primary task (e.g., Edwards & Gronlund, 1998; Gillie & Broadbent, 1989). This provided our rationale in selecting two interrupting tasks—a mini-version of the ToH and mental arithmetic. The mini-ToH involved three-disk versions that could be solved in a minimum of 7 moves and involved the same rules and procedure as those of the primary ToH task, although no goal-state access cost was imposed. The mental arithmetic interrupting task involved solving a series of self-paced double digit mental arithmetic addition problems, each requiring an answer that could range between 20 (i.e., 10 + 10 = ?) and 98 (i.e., 49 + 49 = ?). Answers were entered using a numeric keypad, with the option of deleting and replacing numbers prior to pressing enter to register the answer. Participants were given feedback in the form of a tick to indicate that they had answered the sum correctly or a cross if it was incorrect. An incorrectly answered sum had to be reattempted so that a new sum would not appear until the correct answer was entered. Each interrupting task covered the whole screen for 10 seconds, after which the suspended primary task was reinstated.

Design and procedure

The same three goal-state access cost conditions were used as those in Experiment 1: low, medium, and high. Interruptions occurred once during each of nine problems. They were programmed to occur after making the first move associated with the first, second, or third subgoal in the primary ToH task. Three different interruptions were used (a control blank screen, a mini-ToH, and mental arithmetic) at each of the subgoal positions with the nature of the interruption and its position counterbalanced within each goal access cost condition.

Participants were tested individually and were given instructions on how to operate in their assigned condition and how to perform the task. One 4-disk 15-move ToH problem in the corresponding access cost condition was given as practice followed by one example of each possible interrupting task for 10 seconds each.

Results and discussion

First, we investigated whether there were any performance differences among the three goal-state access cost conditions on either of the interrupting tasks in order to examine whether differential rehearsal opportunities existed between participants in different access cost conditions. There were no such differences in the number of sums correct or the average number of mini ToH problems solved, both $F(2, 51) < 1$.

Subsequent analyses investigated whether increased goal-state access cost improved problem-solving efficiency post interruption and whether there was evidence that this was due to more memory-based planning.

Effects of goal-state access cost on problem-solving efficiency following interruption

First, we examined whether the increased problem-solving efficiency, manifested by the high access cost condition in Experiment 1, was sufficiently resilient to be maintained post interruption. This is indicated by comparing the number of moves required to complete ToH problems after interruption among all three access cost conditions. Participants in the high goal-state access condition solved problems in the fewest moves after interruption (Table 2). A 3 (goal-state access cost condition) × 3 (interruption type) ANOVA confirmed a main effect of access cost condition, $F(2, 51) = 3.35$, $MSE = 2.86$, $p < .05$, $f = 0.36$. Participants in the high goal-state access cost condition completed interrupted problems using significantly fewer moves than those in the low access cost condition ($p < .05$), but not significantly fewer moves than those in the medium access cost condition ($p > .05$). There was no difference between the medium and low goal-state access cost conditions ($p > .05$), and there was no effect of interruption task ($p > .05$) and no interaction (ps $> .05$).

To complement the last measure, we also examined whether this improved efficiency in problem solving after interruption came with an overhead of increased time for problem completion (Table 2). There was no effect of access cost condition on the time taken to complete a problem following interruption both with access time costs included, $F(2, 51) < 1$, and with costs excluded, $F(2, 51) = 2.46$, $MSE = 29.81$, $p = .096$. There was, however, a main effect of interruption type both with access costs included, $F(2, 102) = 7.77$, $MSE = 52.77$, $p < .001$, $f = 0.39$, and with costs excluded, $F(2, 102) = 7.56$, $MSE = 53.94$, $p < .001$, $f = 0.38$, due to more time taken to complete the interrupted problem-solving task following a mini ToH than a blank screen interruption (ps $< .01$). There were no interactions.

Evidence of memory-based planning post interruption

One measure, indicative of the extent that participants could continue to execute moves after interruption from memory, is on how many trials they were able to resume without revisiting the goal state. A second complementary measure concerns the number of moves that participants executed following interruption without revisiting the goal state (Table 2). (Note that the low goal-state access cost condition could not be considered on either measure because in this condition the goal state was permanently uncovered.) Participants in the high access cost condition not only resumed more interrupted trials without first viewing the goal state, $F(1, 34) = 64.01$, $MSE = 0.3$, $p < .001$, $f = 1.37$, but also executed more moves post interruption before revisiting the goal state than participants in the medium access condition (Table 2), $F(1, 34) = 62.85$, $MSE = 2.85$,

Table 2. *Effect of goal-state access cost on performance following interruption (Experiment 2)*

	High		*Medium*		*Low*	
Post interruption measures	M	SD	M	SD	M	SD
Problem-solving efficiency post interruption						
Number of moves to complete a problem following interruption	9.87	0.39	10.2	0.35	11.27	0.44
Time to complete a problem following interruption (including goal-state access costs) (s)	20.48	6.25	21.21	7.97	22.87	5.67
Time to complete a problem following interruption (excluding goal-state access costs) (s)	19.23	6.17	19.54	4.38	22.87	5.67
Memory-based planning post interruption						
Number of trials resumed after interruption without first revisiting the goal state (max. 9)	6.3	1.86	1.89	1.41		
Number of moves executed following interruption without first revisiting the goal state	5.33	2.24	0.88	0.81		
Number of goal-state visits following interruption in order to complete task	0.46	0.26	1.79	0.42		

$p < .001$, $f = 1.36$. This evidence indicates greater resistance to interruption in the high access cost condition because planned moves prior to interruption could be better remembered and/or the goal state could be recalled and used for planning post interruption.

These effects of goal-state access cost did not depend on the nature of the interrupting task because there was no interaction between this factor for either interrupted trials continued without first viewing the goal state, $F(2, 68) < 1$, or number of moves post interruption before visiting the goal state, $F(2, 68) = 1.19$, $p = .31$. However, ANOVAs indicated that there were main effects of the nature of the interrupting task on both postinterruption measures: number of interrupted trials continued without first viewing the goal state, which was a maximum of three (blank screen, $M = 1.86$, $SD = 1.07$; mini-ToH, $M = 0.79$, $SD = 0.98$; mental arithmetic, $M = 1.44$, $SD = 1.16$), $F(2, 68) = 23.31$, $MSE = 0.45$, $p < .001$, $f = 0.83$; and, the number of moves post interruption before visiting the goal state (blank screen, $M = 4.31$, $SD = 3.49$; mini-ToH, $M = 1.88$, $SD = 2.67$; mental arithmetic, $M = 3.13$, $SD = 3.88$), $F(2, 68) = 9.88$, $MSE = 5.36$, $p < .001$, $f = 77$. Performance was poorest following a mini-ToH than a blank-screen interruption on both measures (ps $< .01$).

Further evidence of more memory-based planning as a consequence of high access cost comes from the average number of goal-state visits per problem after interruption. More memory-based planning should lead to fewer visits to view the goal state. Participants in the medium access cost condition visited the goal state four times more frequently than those in the high cost condition (Table 2), and this difference was significant, $t(34) = 11.6$, $p < .001$, $d = 3.98$. Indeed the mean number of goal-state visits of participants in the high access cost condition indicates that the goal state was only visited post interruption approximately on every other trial. Therefore one can infer that the more memory-based planning strategy developed prior to interruption supported improved postinterruption performance for participants in the high access cost condition

In summary, the results from Experiment 2 extend those from Experiment 1 because participants in the high goal access cost condition again engaged in more memory-based planning that enabled them to better resume problem solving after interruption. This was indicated by more trials on which moves were made without immediately viewing the target goal state and more moves executed before it was necessary to revisit this goal state. In addition, participants using more memory-based planning in the high goal access cost condition were able to maintain their problem-solving efficiency post interruption with fewer moves to solution than participants in the low access cost condition, without the need for extra completion time. These findings were independent of the nature of the interruption.

GENERAL DISCUSSION

The aim of this paper was to extend our understanding of the theory of soft constraints in the context of the effect of increasing access cost to the goal state on planning and problem-solving efficiency in the ToH. We selected the ToH as we knew that performance can benefit from planning (e.g., Davies, 2003) and the development of efficient problem-solving strategies (e.g., Anderson, 1993; Anzai & Simon, 1979; VanLehn, 1991). Also we wished to investigate whether accessing the simple goal state of the ToH, comprising the spatial arrangement of four disks, would still involve sufficient additional access time to induce the switch to a more memory-based planning strategy as observed in the study by Waldron et al. (2011) and predicted by the theory of soft constraints (Gray et al., 2006). Even though the extra time cost in the high access cost condition involved 2.96-s for each occasion that the goal state was inspected, this additional time cost might have been too infrequently paid to effect more memory-based planning given the relatively low memory load associated with encoding the position among the four disks. Indeed, our results in Experiment 1 found that participants in the high access cost condition only incurred, on average, approximately three extra seconds access time cost per problem compared to the medium access cost condition. Nevertheless, this small amount of extra time was sufficient to produce a dramatic shift to a more memory-based planning strategy, which, in turn, improved problem-solving efficiency in terms of number of moves to solution. Participants in the high access cost condition made just over three fewer moves to solution and more than twice as many perfect solutions than those in the medium access cost condition.

It is also of interest to note that the imposition of high goal-state access cost in Experiment 1 had a progressively positive effect across trials on the acquisition of an efficient subgoaling strategy involving perfect 15-move solutions. The beneficial effect of imposing a high goal-state access cost on this strategy was evident over the second half of the trials (i.e., Trials 11–20). It should be noted that participants in our experiments were untrained in the use of an efficient strategy, unlike many ToH studies in which researchers are interested in the effect of various manipulations on the execution of an efficient strategy in which participants are instructed (e.g., Anderson & Douglass, 2001; Anderson et al., 1993; Anderson & Lebiere, 1998). Therefore, high goal-state access cost not only promotes memory-based planning but also results in the gradual acquisition of a more efficient subgoaling strategy that is manifested without any training.

These findings provide further corroboratory evidence for the theory of soft constraints (Gray & Fu, 2004; Gray et al., 2006) but not with other views that assume that cognition operates to conserve cognitive effort (e.g., Ballard, Hayhoe, & Pelz, 1995; Cary & Carlson, 2001; Wilson, 2002). The more memory-based planning strategy induced in the high access cost condition involves more encoding of goal state and internal planning that result in the development of a more efficient problem-solving strategy. The only other study to have examined the effect of increased goal-state access cost on planning in a problem-solving task was by Waldron et al. (2011). Contrary to the findings of that study that used an eight-puzzle-like problem, the high access cost condition in our Experiment 1 induced more efficient problem solving in terms of moves to solution and more error-free solutions. It is also important to note that this improvement in efficiency did not come at any extra cost in terms of time to complete problems, even when access costs were included in the analysis. Arguably, the development of a good strategy in the ToH is not only easier than developing one in the eight puzzle but also supports more of a cascading effect of subsequent moves with respect to the achievement of Subgoals 1 and 2.

The results of Experiment 1 also demonstrate that a switch to a more memory-based planning strategy comes about from a relatively small extra access price paid per problem by the high access cost condition in comparison with the access costs paid by participants in the same condition in Experiments 1 and 2 of Waldron et al. (2011). Participants in the present high access cost condition of Experiment 1 spent a little less than an extra three seconds accessing the goal state per problem (out of an average completion time of 49.48-s for each problem). In the Waldron et al. (2011) study, participants spent approximately 13 and 37 extra seconds accessing the goal state (out of an average completion time of 85-s and 332-s) in Experiments 1 and 2, respectively. These few extra seconds access cost had a large effect on strategy in our Experiment 1, presumably because there are more benefits from a little planning in the ToH than in the eight-like puzzle used in Experiment 2 of the Waldron et al. (2011) study. Of course it is unlikely that the same amount of time will have the same value and ramifications for participants across problem-solving tasks that differ in their strategies, solution times, memory demands, and so on.

Having established that a relatively small increase in access cost results in more planning and efficient problem solving, the second aim of this study was to investigate whether this more active cognitive strategy would mitigate the negative effects of interruption and still maintain efficiency following interruption. These effects are predicted by the memory for goals model (Altmann & Trafton, 2002, 2007) as discussed above. The only study to date that has investigated this found that increased goal access cost had a protective effect against interruption during performance of the Blocks World Task (Morgan et al., 2009). A more memory-based strategy protected against forgetting moves following interruption, even with different types of interrupting task and particularly when interruptions were neither very infrequent nor very frequent. The results from the present Experiment 2 indicate that a more memory-based planning strategy can also protect against interruption in a problem-solving task. Participants in the high goal-state access cost condition resumed more interrupted trials from memory than those in the medium access cost condition and executed more moves after interruption before revisiting the goal state (or reaching solution). Also, despite the increased time caused by the interruption itself, participants in this condition were able to maintain their greater problem-solving efficiency post interruption by completing interrupted trials in fewer moves than those in the low access cost condition. Also given that the access cost condition did not interact with the type of interruption, these conclusions are independent of the nature of the interrupting tasks used in the present study. These findings are consistent with the memory for goals model (Atmann & Trafton, 2002, 2007) as the enhanced memory-based planning in the high access cost condition would have strengthened the representation(s) of goal(s) and therefore made them more resistant to decay and forgetting, and therefore interruption.

Besides theoretical considerations, some practical ramifications follow from Experiment 2. In the literature, three main methods have been proposed for reducing the negative effects of interruption: coordinating interruptions with performance of the primary task (McFarlane, 2002; McFarlane & Latorella, 2002); using reminder cues (e.g., Cutrell, Czerwinski, & Horvitz, 2001; Czerwinski, Cutrell, & Horvitz, 2000; Franke, Daniels, & McFarlane, 2002; McDaniel, Einstein, Graham, & Rall, 2004); and manipulating the time to encode goals prior to interruption by inserting an interruption lag (e.g., Altmann & Trafton, 2002, 2004; Hodgetts & Jones, 2006a; Trafton, Altmann, Brock, & Mintz, 2003). Manipulating goal-state access cost in order to induce more encoding was a fourth method for mitigating the effects of interruption proposed by Morgan et al. (2009). The results of the present study confirm the efficacy of this method for a problem-solving task without the negative side-effect of increased completion time found by Morgan et al. (2009) using the Blocks World Task. Avoiding such a negative side-effect of increased completion time is potentially important for the practical application of this method together with the finding of increased problem-solving efficiency in comparison with when the goal state is always available.

Some limitations of the present study need acknowledgement. First, our assumption has been that people dislike and are motivated to avoid experiencing the increased delay in accessing the goal state in the high condition, and to achieve this a more memory-based strategy is adopted. However, it is possible that the lockout time itself could act as an opportunity and provocation for extra thinking/planning (E. M. Altmann, May 2012, personal communication) although the nature of any planning before seeing the goal state would necessarily be limited (even though the current state is visible during lockout time). In order to clarify whether planning was occurring during the lockout and even mouse movement time to the target, a future study could increase cognitive load during this delay time via a secondary task. Second, the effects that we have observed in both experiments are confined to the particular values of various parameters used. This study has not systematically increased the value of access time costs to identify whether there is a progressive shift to a more memory-based planning strategy or whether there is some threshold effect. Using the Blocks World Task, Gray et al. (2006) found strong linear relationships between increasing lockout time and various indicators of an increasingly memory-based strategy. Similarly, interruption duration was not varied in Experiment 2 to identify values beyond which manipulation of goal-state access cost might fail to provide any mitigation to the negative effect of interruption. This would be predicted by the memory for goals model that incorporates a decay function for the representation of a goal that has not been reactivated over a period of time. Altmann and Trafton (2007) plotted the time-course of recovery after interruptions varying in duration from 30 to 45 seconds. Also, Monk et al. (2008), using a VCR programming task, found that resumption lag increased with interruption duration, and this effect was amplified by more demanding interrupting tasks. Future research should examining how goal-state access cost is affected by increasing interruption duration together with the extent that findings generalize across tasks or whether they are task or task-type specific.

In conclusion, the present study corroborates the prediction of the theory of soft constraints (Gray & Fu, 2004; Gray et al., 2006) that increased access time would result in more memory-based planning. The results provide novel evidence concerning the effect of increased goal-state access cost, resulting in more memory-based planning and the development of a more efficient problem strategy, and how this can be capitalized upon in order to mitigate the negative effect of interruption. It is important to note that only a few seconds time delay in accessing the goal state per problem was necessary to effect this change to more memory-based planning, which resulted in fewer moves to problem solution. Furthermore, this change to a more memory-based planning strategy meant that the representation of goals was more activated during performance, and therefore some protection against interruption was afforded as suggested by the memory for goals model (Altmann & Trafton, 2002, 2007). Furthermore, there was evidence that the increased problem-solving efficiency in the high access cost condition survived interruption because postinterruption solution was achieved in fewer moves. Finally, manipulation of goal-access cost in the ToH had these positive effects without any attendant increase in completion time, which makes it a more feasible practical solution to be explored in interface design for mitigating the negative effect of interruption in problem solving.

Original manuscript received 17 June 2011
Accepted revision received 28 May 2012
First published online 30 August 2012

REFERENCES

Altmann, E. M., & Trafton, G. J. (2002). Memory for goals: An activation-based model. *Cognitive Science*, *26,* 39–83.

Altmann, E. M., & Trafton, G. J. (2004). Task interruption: Resumption lag and the role of cues. *Proceedings of the 26th Annual Meeting of the Cognitive Science Society* (pp. 42–47). Hillsdale, NJ: Lawrence Erlbaum Associates.

Altmann, E. M., & Trafton, J. G. (2007). Timecourse of recovery from task interruption: Data and a model. *Psychonomic Bulletin & Review*, *14,* 1079–1084.

Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Anderson, J. R. (1993). Problem solving and learning. *American Psychologist*, *48*(1), 33–44.

Anderson, J. R., & Douglass, S. (2001). Tower of Hanoi: Evidence for the cost of goal retrieval. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *27,* 1331–1346.

Anderson, J. R., Kushmerick, N., & Lebiere, C. (1993). The Tower of Hanoi and goal structures. In J.R. Anderson (Ed.), *Rules of the mind* (pp. 121–142). Hillsdale, NJ: Lawrence Erlbaum Associates.

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Lawrence Erlbaum Associates.

Anderson, J. R., & Milson, R. (1989). Human memory: An adaptive perspective. *Psychological Review*, *96,* 703–719.

Anzai, Y., & Simon, H. (1979). The theory of learning by doing. *Psychological Review*, *86,* 124–140.

Atwood, M. E., & Polson, P. P. (1976). A process model for water jug problems. *Cognitive Psychology*, *8,* 191–216.

Ballard, D. H., Hayhoe, M. M., & Pelz, J. B. (1995). Memory representations in natural tasks. *Journal of Cognitive Neuroscience*, *7*(1), 66–80.

Card, S. K., English, W. K., & Burr, B. J. (1978). Evaluation of mouse, rate-controlled isometric joystick, step keys and text keys for text selection on a CRT. *Ergonomics*, *21,* 601–613.

Cary, M., & Carlson, R. A. (2001). Distributing working memory resources during problem solving. *Journal of*

Experimental Psychology: Learning, Memory & Cognition, *27,* 836–848.

Chater, N., & Oaksford, M. (1999). Ten years of the rational analysis of cognition. *Trends in Cognitive Sciences*, *3*(2), 57–65.

Cutrell, E., Czerwinski, M., & Horvitz, E. (2001). Notification, disruption and memory: Effects of messaging interruptions on memory and performance. In M. Hirose (Ed.), *Human–computer interaction—INTERACT 2001* (pp. 263–269). Tokyo, Japan: IOS Press, IFIP.

Czerwinski, M., Cutrell, E., & Horvitz, E. (2000). Instant messaging and interruption: Influence of task type on performance. In C. Paris, N. Ozkan, S. Howard, & S. Lu (Eds.), *OZCHI 2000 conference proceedings* (pp. 356–361). Sydney, Australia.

Davies, S. P. (2003). Initial and concurrent planning in solutions to well-structured problems. *The Quarterly Journal of Experimental Psychology*, *56A,* 1147–1164.

Delaney, P. F., Ericsson, A. K., & Knowles, M. A. (2004). Immediate and sustained effects of planning in a problem-solving task. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *30,* 1219–1234.

Edwards, M. B., & Gronlund, S. D. (1998). Task interruption and its effects on memory. *Memory*, *6,* 665–687.

Ericcson, K. A. (1974). *Problem-solving behavior with the 8-puzzle II: Distribution of latencies* (Report No. 432). Stockholm, Sweden: University of Stockholm, Department of Psychology.

Franke, J., Daniels, J., & McFarlane, D. (2002). Recovering context after interruption. In W. Gray & C. Schunn (Eds.), *Proceedings of the 24th Annual Meeting of the Cognitive Science Society* (pp. 310–315). Hillsdale, NJ: Lawrence Erlbaum Associates.

Fu, W.-T. (2011). A dynamic context model of interactive behavior. *Cognitive Science*, *35,* 874–904.

Fu, W.-T., & Gray, W. D. (2000). Memory versus perceptual-motor tradeoffs in a blocks world task. *Proceedings of the 22nd Annual Conference of the Cognitive Science Society* (pp. 154–159). Hillsdale, NJ: Lawrence Erlbaum Associates.

Gillie, T., & Broadbent, D. (1989). What makes interruptions disruptive? A study of length, similarity, and complexity. *Psychological Research*, *50,* 243–250.

Gray, W. D., & Boehm-Davis, D. A. (2000). Milliseconds matter: An introduction to microstrategies and their use in describing and predicting interactive behavior. *Journal of Experimental Psychology: Applied*, *6,* 322–335.

Gray, W. D., & Fu, W.-T. (2001). Ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head: Implications of rational analysis for interface design. *CHI Letters*, *3*(1), 112–119.

Gray, W. D., & Fu, W.-T. (2004). Soft constraints in interactive behaviour: The case of ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. *Cognitive Science*, *28,* 359–383.

Gray, W. D., Sims, C. R., Fu, W.-T., & Schoelles, M. J. (2006). The soft constraints hypothesis. A rational analysis approach to resource allocation for interactive behaviour. *Psychological Review*, *113,* 461–482.

Hayes-Roth, B., & Hayes-Roth, F. (1979). A cognitive model of planning. *Cognitive Science*, *3,* 275–310.

Hodgetts, H. M., & Jones, D. M. (2006a). Contextual cues aid recovery from interruption: The role of associative activation. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *32*(5), 1120–1132.

Hodgetts, H. M., & Jones, D. M. (2006b). Interruption of the Tower of London task: Support for a goal-activation approach. *Journal of Experimental Psychology: General*, *135*(1), 103–115.

MacKenzie, I. S. (1992). Fitts' law as a research and design tool in human–computer interaction. *Human-Computer Interaction*, *7*(1), 91–139.

McDaniel, M. A., Einstein, G. O., Graham, T., & Rall, E. (2004). Delaying execution of intentions: Overcoming the cost of interruptions. *Applied Cognitive Psychology*, *18*(5), 533–547.

McFarlane, D. C. (2002). Comparison of four primary methods for coordinating the interruption of people in human–computer interaction. *Human-Computer Interaction*, *17,* 63–139.

McFarlane, D. C., & Latorella, K. A. (2002). The scope and importance of human interruption in human–computer interaction design. *Human-Computer Interaction*, *17,* 1–61.

Miller, G. A., Galanter, E., & Pribram, K. H. (1960). *Plans and the structure of behavior*. New York, NY: Holt.

Monk, C. A., Trafton, J. G., & Boehm-Davis, D. A. (2008). The effect of interruption duration and demand on resuming suspended goals. *Journal of Experimental Psychology: Applied*, *14,* 299–313.

Morgan, P. L., Patrick, J., Waldron, S. M., King, S. L., & Patrick, T. (2009). Improving memory after interruption: Exploiting soft constraints and manipulating information access cost. *Journal of Experimental Psychology: Applied*, *15*(4), 291–306.

Oaksford, M., & Chater, N. (1998). *Rationality in an uncertain world*. Hove, UK: Psychology Press.

O'Hara, K. P., & Payne, S. J. (1998). The effects of operator implementation cost on planfulness of problem solving and learning. *Cognitive Psychology*, *35,* 34–70.

O'Hara, K. P., & Payne, S. J. (1999). Planning and the user interface: The effects of lockout time and error recovery cost. *International Journal of Human-Computer Studies*, *50,* 41–59.

Patsenko, E. G., & Altmann, E. M. (2010). How planful is routine behaviour? A selective-attention model of performance in the Tower of Hanoi. *Journal of Experimental Psychology: General*, *139*(1), 95–116.

Peterson, L. R., & Peterson, M. J. (1959). Short-term retention of individual verbal items. *Journal of Experimental Psychology*, *58,* 193–198.

Trafton, J. G., Altmann, E. M., Brock, D. P., & Mintz, F. E. (2003). Preparing to resume an interrupted task: Effects of prospective goal encoding and retrospective rehearsal. *International Journal of Human-Computer Studies*, *58,* 583–603.

Trafton, J. G., & Monk, C. M. (2008). Task interruptions. In D. A. Boehm-Davis (Ed.), *Reviews of human factors and ergonomics* (Vol. 3, pp. 111–126). Santa Monica, CA: Human Factors & Ergonomics Society.

VanLehn, K. (1991). Rule acquisition events in the discovery of problem solving strategies. *Cognitive Science*, *15,* 1–47.

Waldron, S. M., Patrick, J., & Duggan, G. B. (2011). The influence of goal-state access cost on planning during problem solving. *Quarterly Journal of Experimental Psychology*, *64*(3), 485–503.

Waldron, S. M., Patrick, J., Morgan, P. L., & King, S. L. (2007). Influencing cognitive strategy by manipulating information access costs. *The Computer Journal*, *50*(6), 694–702.

Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review*, *9,* 625–636.

APPENDIX

Estimation of mouse cursor movement time

Mouse movement times between the goal-state window and current-state windows were calculated using a variation of Fitts' law (MacKenzie, 1992). This is a universally accepted approximation of human movement in human–computer interaction and was also used by Gray et al. (2006) in calculating movement time for the Blocks World Task (BWT). The Fitts' law equation used was $MT = a + b \log_2(A/W + 1)$, where MT is movement time, A is amplitude (or movement distance) and W is the tolerance or width of the target area. We used the Adaptive Control of Thought - Rational (ACT-R) parameters for Fitts' law ($a = 0.05$; $b = 0.10$), derived by Card, English, and Burr (1978), that are cited and used by Gray et al. (2006) as providing a good fit to moving a mouse cursor around a computer screen,. The estimated time to make a mouse movement between the goal and workspace window in either direction was 231-ms.