Now, where was I? A cognitive experimental analysis of the influence of interruption on goal-directed behaviour

Phillip L. Morgan

Thesis submitted for the degree of
Doctor of Philosophy

School of Psychology
Cardiff University

July 2005

DECLARATION

This work has not previously been accepted in substance for any degree and is not being currently submitted in candidature for any degree.

Signed .....

Date .....18/11/2005.....

STATEMENT 1

This thesis is the result of my own investigations, expect where otherwise stated. Other sources are acknowledged by footnotes giving explicit references. A bibliography is appended.

Signed .....

Date .....18/11/2005.....

STATEMENT 1

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

Signed .....

Date .....18/11/2005.....

To Joanne,

Together we can achieve anything

Acknowledgements

I would like to thank Professor Dylan M. Jones for providing me with the opportunity and confidence to undertake the PhD and for his endearing encouragement throughout the project. I am gratefully indebted to him for his knowledge and expertise, especially during the preparation of this thesis.

Special thanks to those who provided unconditional help and support throughout my research: Robert Houghton (for preparing the initial program), Helen Hodgetts (the other half of 'The Interruption People'), Andrew Howes, and Bill Macken. I am also grateful to the clerical and technical staff of the School of Psychology at Cardiff University, and to all who participated in the experiments.

I wish to thank my friends and colleagues for being there for the good and the occasional not so good times. First, to my officemates and lasting friends for putting up with my moans and groans: Giles Greene (friend and partner in crimes against the professional game of snooker), Fiona ('have you finished yet?') Phelps, and Gareth (Mr 'G') Linsmith. Second, to a bevy of stars: Alastair McKenzie-Kerr, Alex Johnston, Anna Sagerson, Anthony Brown, Ben Wellens, David Sanderson, Geoff Duggan, John ("gee Ernie") Marsh, Jo Haddon (mines a double JD), Melissa Allman, Nick Perham, Rob Hughes, and Steve Chambers (the VB connoisseur). Third, to Ray and all the 'Ladz' down at the snooker lounge. Last but by no means least, to past and present members of 64 Park Place and those from the School of Psychology that made it all worthwhile.

I would like to thank my parents for their continued encouragement. The scale and timely demands of a work of this size have been tolerated by some lifelong friends: Gareth Roberts for immeasurable friendship, Matthew Hall and Stuart Bailey.

Thanks are also due to the Economic and Social Research Council for funding me from 2001 – 2004. I would also like to thank Professor Hadyn Ellis for providing a life changing moment of clarity at a point of blinded uncertainty.

Finally, I wish to thank my fiancée Joanne Thomas for her constant love, support and patience throughout this journey. Without her, this work would not have been possible and I would not be as strong a person as I am today ('we did it!').

Summary

Task interruption is a pervasive applied problem despite a dearth of experimental work and the absence of a developed theoretical framework. Using a novel experimental approach (interrupting problem solving in the Tower of Hanoi task), and theoretical guidance from ACT-R-based models of goal suspension and resumption (Altmann & Trafton, 2002; Anderson & Douglass, 2001), nine experiments were conducted to assess how goal-directed behaviour is affected by interruption.

A cost of interruption was exhibited mainly by extended times to resume an interrupted goal compared to an uninterrupted goal. The first empirical series established performance impairments in the form of long resumption latencies for promptly suspended goals and decrements in move accuracy, especially when interruption fell before or during a complex goal-sequence, with performance impaired further by secondary tasks that were similar to primary tasks. The second empirical series revealed that participants opportunistically encode promptly suspended goals for retrieval, a process supported by the associative activation provided by a salient colour priming cue and impaired when such a cue had changed colour and/or location. With a brief time lag before secondary task initiation, participants were able to encode a suspended goal more efficiently, reflected in faster resumption latencies even when secondary tasks were similar and when interruption fell within a complex goal sequence.

The findings suggest that suspended goals do not reside in a heightened level of activation such that retrieval is definite (e.g., Goschke & Kuhl, 1993); neither is retrieval always abandoned at longer retention intervals (as suggested by Anderson & Lebiere, 2001). Instead, goals decay as a power function of the time since they were last processed and suffer retroactive interference from other goals, but can be reactivated if appropriately rehearsed and associated with salient retrieval cues (in support of Altmann & Trafton, 2002). In contrast to Altmann and Trafton, participants exhibit retrieval-like behaviour even when interruption is un-signalled, with efficiency augmented by experience of problem solving in the task domain *and* experience of being interrupted. The current experiments provide a novel insight into interruption management behaviours, particularly that humans are able and willing to adapt strategies to support faster and more efficient transitions back into the primary task.

Table of Contents

Figures

Tables

Chapter 1

GENERAL INTRODUCTION

1.1. Establishing the area

1.1.1. Overview

If a task is interrupted, performance is usually impaired. Effects include, the time taken to resume the suspended activity (e.g., Altmann & Trafton, 2004), frequency of mistakes (e.g., Speier, Vessey & Valacich, 2003), and the time taken to complete the interrupted task (e.g., Gillie & Broadbent, 1989). Differences in performance deficits are often attributed to factors such as, secondary task complexity (e.g., Gillie & Broadbent, 1989), interruption position within a primary task (e.g., Monk, Boehm-Davis & Trafton, 2002), and similarity between tasks (e.g., Czerwinski, Chrisman & Rudisill, 1991a; Edwards & Gronlund, 1989). Empirical research has largely focused on identifying the consequences of being interrupted, with the chief motivation of trying to identify methods that may alleviate such effects (e.g., McFarlane, 2002). For instance, if the time taken to resume a primary task is reduced, if the user controls interruption timing (e.g., McFarlane, 2002), the interrupted task is sufficiently practiced (e.g., Edwards & Gronlund, 1998), or retrieval aids are made available (e.g., Lahlou, Kirsh, Rebotier, Reeves & Remy, 2002). As informative as this research appears, little effort has been invested in trying to understand the cognitive processes affected by task interruption.

There is a dearth of theoretical understanding concerning the cognitive processes affected by task interruption. The most pervasive view is that dealing with an interruption breaks the cognitive focus of performing an ongoing task (Miyata & Norman, 1986). One extension of this view is that the cognitive resources dedicated to primary task performance may be disrupted during different stages of interruption coordination (e.g., the transitory period between the primary and secondary tasks, McFarlane & Latorella, 2002), leading to varying performance deficits. For instance, some suggest that a suspended goal cannot be retrieved unless it is encoded immediately

prior to interruption, highlighting the consequence of prompt interruption with little opportunity to prepare (e.g., Altmann & Trafton, 2002). Others propose that at long retention intervals, humans routinely choose not to encode suspended goals, and instead, offload mnemonic demands to the task environment (e.g., Anderson & Douglass, 2001). Both views assume that being interrupted affects the ability to maintain a cognitive representation of a suspended goal.

A variety of theoretically grounded experimental methods may be used to inform a conceptual understanding of the cause of interruption effects. Firstly, performance deficits are reliably found when having to manage numerous goals in uninterrupted goal-directed behaviour (e.g., Anderson & Douglass, 2001). It has recently been suggested that the time taken to resume a suspended goal is dependent upon how it was encoded prior to being suspended (Anderson & Douglass, 2001). Interruption may disrupt the cognitive representation of an encoded goal, resulting in forgetting of its associated action. Secondly, forgetting to prospectively perform a future action is often assumed to reflect inefficient association of a suspended intention with a reminder cue during encoding (e.g., Einstein, McDaniel, Williford, Pagan, & Dismukes, 2003; Trafton et al., 2003). Failing to notice a cue to perform a future intention (e.g., seeing a telephone) may result in forgetting to perform its associated action at the appropriate time (e.g., making a telephone call). Thirdly, performance deficits are found when switching between different tasks, and are often attributed to problems with alternating, redirecting, or reconfiguring processing resources (see Monsell, 2003, for a review). Finally, an interrupting secondary task may enforce attentional alternation between tasks (e.g., Adams, Tenney & Pew, 1995), which might result in a loss of task situational awareness. All methodologies seem to be tightly related to the event of being interrupted, and I will review each of them in depth throughout the current chapter.

The experiments presented in the current thesis were designed to investigate the affects of being interrupted, using a primary task with a well-established paradigm for which there is a wealth of conceptual understanding (the Tower of Hanoi (ToH), e.g., Anderson & Douglass, 2001). In a recent paper, the costs associated with suspending and resuming goals were modelled using performance data from the ToH task (the goal-activation model, Altmann & Trafton, 2002). The authors subsequently contemplated the

value of the parameters set in the model for predicting the processes affected by task interruption. Motivated by the goal-activation model, the ToH was used as an exploratory task in all experiments presented in this thesis. Since the ToH is a well-structured problem solving task, it allows for an acceptable level of control over the positioning of interruptions; an important design factor affecting the interpretation of results from a number of previous studies.

The empirical chapters within this thesis were designed to investigate specific factors that contribute to performance attenuation following an interruption episode. These were: Interruption position and processing complexity; task similarity; preparation for interruption; and the use of retrieval aids. Before considering their contribution to interruption effects, singular or combined, it is worth taking a step back to build a more detailed understanding of how an interruption is co-ordinated into an ongoing task.

1.1.2. Characteristics of interruption

1.1.2.1. The stages of interruption

The general stages involved in a task interruption are well established (e.g., McFarlane, 2002; McFarlane & Latorella, 2002). An interruption is initially detected when the *interrupt alert* (e.g., an e-mail notification) captures attention. Sometime afterwards, a *secondary task* is initiated forcing redirection of processing resources from the primary task. The secondary task is usually performed through to completion, whereby the primary task is fully reinstated and requires *resumption*.

An interruption should not be thought of as an event peripheral to the primary task. A recent diagrammatic example of an alerted interruption is illustrated in Figure 1.1 (Altmann & Trafton, 2004), introducing periods that may be critical for recovering the representation of a suspended goal. Firstly, the *interruption lag* is, "...a brief transitional interval immediately preceding an interruption, during which the operator knows of the pending interruption but is not yet engaged in it" (Altmann & Trafton, 2004). This has formed an independent measure in recent experimental attempts to alleviate the negative affects of task interruption (e.g., Altmann & Trafton, 2004; Miller, 2002; Trafton et al., 2003). It should however be acknowledged, that such a preparatory lag may not always

be available, especially when the interruption is unexpected and requires prompt attendance. For instance, Covey (1980, pp. 150-152) stressed that an interruption often, "requires immediate attention", and, "insists on action". Secondly, the *resumption lag* is the time taken to execute the first action following primary task reinstatement. It is during this period that the suspended activity might either be retrieved from memory or re-established from the external environment.

1.1.2.2. What does an interruption interrupt?

In embarking upon a study of task interruption, it is important to consider the characteristics of the task that is being interrupted. A *task* is generally considered to be a set of actions carried out to achieve an overall goal (e.g., Newell & Simon, 1972). Take for example the goal-directed task of making a photocopy of this page. To service the photocopy or the *master goal*, the task may have to be first decomposed into a number of smaller manageable *subgoals*, such as the series of actions associated with setting up the photocopier (e.g., setting the copy size, choosing the contrast and so on). Such subgoals can be further decomposed to the individual actions required to achieve each subordinate goal – the level at which a task will be performed irrespective of complex cognitive processing.

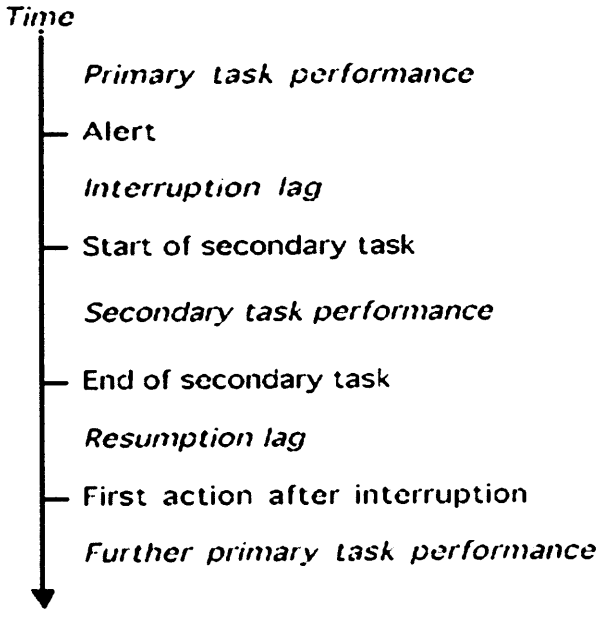

Figure 1.1. The time course of an interruption (taken from Altmann & Trafton, 2004)

The effects of interruption on task performance have been considered by analysing various post-interruption performance measures, resulting in subsequent controversy regarding exactly how performance is impaired. For example, the components of the secondary task may affect the efficiency of performing the remaining actions within the primary task (e.g., Gillie & Broadbent, 1989).

Researchers have since stressed that effects may in large be confined to the specific goal structure interrupted (i.e., the subgoal containing the goal) (e.g., Monk et al., 2002). Consequently, the time taken to resume the primary task may consist only of the time taken to retrieve the suspended goal (e.g., Altmann & Trafton, 2002; Speier et al., 2003) and not the time taken to reconstruct the whole task. It may be the efficiency of this period that determines the extent of residual performance outcomes within the primary task (e.g., mistakes, and completion time).

In the next section, I will review the pioneering research that cleared the way for studying the effects of being interrupted. I hope to highlight the limitations of generalising such results to the current study of task interruption.

1.2. Interruption theory

1.2.1 Models and frameworks of interruption

1.2.1.1 The interruption management stage model (IMSM, Latorella, 1996b)

The IMSM (Latorella, 1996; McFarlane & Latorella, 2002) is a framework that identifies the information-processing stages affected by being interrupted. It considers task, operator, and work environment characteristics that may contribute to the deleterious effects of interruption on performance at each of four interruption processing stages. The stages are: *Detection* of an interruption annunciation; *interpretation* of interruption meaning; *integration* of a secondary task; and *primary task resumption.* Each of these stages are likely to be affected by interruption resulting in performance deficits. *Diversion* from the primary task will occur if an interruption annunciation signal is salient enough to exceed a sensory detection threshold. Such diversion is proposed to come at little cost to the efficiency of subsequent performance on the primary task,

because of our ability to maintain attentional focus when distracted (e.g., Hockey, 1970). *Distraction* to the cognitive continuity of primary task processing will occur if the meaning of the interrupting event is derived, affecting subsequent performance even if the secondary task is not initiated. *Disturbance* to the cognitive focus of the primary task will occur if the secondary task is integrated and performed. *Disruption* is experienced when trying to reinitiate the primary task following cessation of the secondary task. The level of disruption will reflect the difficulty in maintaining a representation of the primary task during performance of the secondary task.

Five methods of interruption implementation are also considered by the IMSM, each independently contributing to the way interruptions are co-ordinated into ongoing activities. *Oblivious dismissal* occurs if the interruption alert is not detected. *Unintentional dismissal* occurs if the interruption alert is detected, but its meaning is not derived and the secondary task is not performed. *Intentional dismissal* occurs if the annunciation is interpreted but the operator makes a conscious decision not to take up the interruption. *Preemptive integration* occurs if the primary task is suspended and the secondary task is taken up immediately. Finally, *intentional integration* occurs if the operator attempts to synchronise the performance of the primary and secondary tasks. Typically, preemptive integration results in the largest overall performance deficits, especially during primary task resumption (Latorella, 1996), and may reflect the level of disruption caused by ubiquitous interruption (e.g., Adams et al., 1995).

The IMSM provides a valuable framework for studying task interruption effects and potential management strategies. In its current form, the IMSM is somewhat limited in providing an in-depth explanation of the cognitive processes affected by interruption (e.g., McFarlane, 2002; Latorella & McFarlane, 2002). Additionally, it seems to rely upon the assumption that all tasks are serial (i.e., all actions occur in a coherent sequential order), and in doing so, underestimates the affects of interrupting complex goal structures (e.g., where backtracking to previously processed actions may be required). Experiments within this thesis aim to provide a more in-depth coverage of the cognitive processes effected by task interruption. Such in-depth analyses included the effects of interruption on goal structures that are not always processed sequentially (e.g., where performance of a goal is dependent upon the outcome of its predecessors).

1.2.1.2. The definition and taxonomy of human interruption (McFarlane, 1997, 1998, 2002)

The deleterious effects often caused by having to unexpectedly perform a secondary task are described in a comprehensive definition and taxonomy of human interruption (McFarlane, 1997, 1999, 2002). The taxonomy incorporates a literature driven review aimed at informing human-computer interaction (HCI) interface design, but generates recommendations using little theoretical support. McFarlane identifies eight dimensions that, when manipulated, result in different post-interruption performance outcomes. These are the: (1) source of the interruption; (2) individual characteristics of the person receiving interruption; (3) method of co-ordination; (4) meaning of the interruption; (5) method of expression; (6) channel of conveyance; (6) human activity changed by interruption; and (8) effect of interruption. The method of co-ordination has been empirically supported as a major factor in determining the performance efficiency effected by interruption, and has four levels (e.g., McFarlane, 2002).

1. The *immediate* method is where an interruption forces an instantaneous attentional shift to the secondary task, and usually results in the highest overall performance deficits (e.g., the longest time taken to resume the primary task).
2. The *negotiated* method is where there is a fixed period of time between interruption annunciation and secondary task initiation. This presumably allows preparation to occur before switching to the secondary task, and usually results in the lowest overall performance deficits (e.g., the shortest time taken to resume the primary task).
3. The *mediated* method is where switching to the secondary task is governed by an external entity (i.e., another person or technological device). This method has the advantage of delaying interruption induced distraction and diversion during complex task performance, but is limited to the number of contexts in which it can be applied (e.g., non-safety critical).
4. The *scheduled* method is where the switch to the secondary task is predetermined, allowing the user conscious awareness as to when the primary task is to be suspended. Similar to the mediated method, scheduling

interruption is only suitable to certain work situations (e.g., a nurse carrying out hourly patient observations).

The taxonomy of human interruption should not be considered as a theoretical framework for understanding the process of task interruption. The literature driven review details the effects of being interrupted, but fails to consider their cause at any sort of conceptual level. Therefore, it was used within this thesis to guide techniques that may alleviate the deleterious effects of task interruption.

1.2.1.3. A framework of interruption and decision-making (Speier, Vessey & Valacich, 2003)

Using decision making as a source for investigation, it has been proposed that task interruption can be understood by considering factors contained within only three dimensions (Speier, Vessey & Valacich, 2003). The first dimension contains the *cognitive processing characteristics* of interruption, including, frequency, duration, content, complexity, and timing. They are uncontrollable by the recipient and are likely to affect performance efficiency after secondary task completion. The second dimension contains the *social characteristics of the interruption*. These consist of, the *form of the interruption* (e.g., human or machine communication); the *person or object generating the interruption*; and the *social expectancies existing due to organisational or regional culture* (e.g., when it is appropriate to communicate with colleagues). The third dimension contains *processing mechanisms,* which can be, sequential; preemptive (i.e., process interruptions as they occur); or simultaneous (i.e., interleave both tasks at once).

Two theoretical sources were used to inform the investigation of the cause of interruption effects (Speier et al., 1999, 2003). Firstly, Distraction Conflict Theory (DCT, Baron, 1986) claims that being distracted can facilitate subsequent performance on simple tasks, but can impair performance on more complex tasks. Being distracted might increases stress, and cause a narrowing of attention to primary task information that is only relevant to short-term performance outcomes. For simple tasks, irrelevant task information can be dismissed from attention, freeing up processing resources that can otherwise be dedicated to making more accurate and faster decisions. By contrast,

processing resources are available in smaller quantities when performing complex tasks, which may encourage uncritical choice of which material to keep in active focus (e.g., Baron, 1986). As a result, performance deficits may ensue in the form of less accurate and slower decisions. Being interrupted might cause an automatic increase in primary task complexity, because of the appreciable processing resources required to keep to-be-interrupted material in active consciousness. Only the simplest of primary tasks should therefore be facilitated by task interruption (Speier et al., 2003).

Secondly, Cognitive Fit Theory (CFT, Vessey, 1991) proposes that a secondary task that demands similar processing resources to a primary task should affect subsequent performance in the primary task. Speier et al. (2003) compared symbolic tasks (e.g., pattern matching) and spatial tasks (e.g., navigation) when considering the effects of task interruption on performance. They suggested that when task demands are low, people are able to use analytic processes to establish cues to direct them to suspended goals. This is particularly useful for symbolic tasks that are difficult to decompose into easily distinguishable cues. In contrast, when task demands are high, people may have to replace the desire to achieve performance accuracy with the need to perform as well as possible in a limited time. Thus, they may have to over-rely upon perceptual processes (e.g., cue extraction) which takes less time but usually results in reduced performance efficiency (e.g. Anderson & Douglass, 2001). This is particularly useful for spatial tasks that can be decomposed into cues to guide future memory performance. In the case of interruption, it is proposed that tasks that are normally suited to analytic processing may have to be processed using perceptual cues. Thus, interruption may negatively affect performance efficiency on symbolic primary tasks but not on those that are spatial.

From conducting only one experiment, it was concluded that experience of dealing with task interruption influences the way in which information is processed (Speier et al., 2003), which may have implications for the experiments conducted within this thesis. Both DCT and CFT accounts were only partially supported. For instance, decision accuracy and speed was increased for an interrupted simple spatial task. However, only decision speed was increased for an interrupted simple symbolic task and not accuracy.

The presented framework may have implications for understanding the costs associated with resuming an interrupted task. It provides an exploratory insight into interrupting tasks of varying complexity, and is the first account to suggest that resumption accuracy is sought following an interruption and its effectiveness may be dependent upon a timely decision. In this case, it could be argued that task interruption might facilitate general post-task resumption performance (e.g., fewer mistakes, shorter completion times), but not the time taken to resume the task following an interruption episode. However, there was limited control over the way interruptions were operationalised in the Speier et al. (2003) experiment. That is, interruptions were embedded in simple tasks (those requiring few actions) and complex tasks (those requiring many actions), but their position was not controlled. Additionally, although post-interruption resumption time costs were identified in the Speier et al. (2003) experiment, an in-depth analysis of their cause (such as interruption timing, content and duration) was not considered. The experiments presented in this thesis provide a fine-grained control of such factors, to further inform theoretical understanding of such processes.

1.2.1.4. A cognitive framework of multiple task management (Adams, Tenney & Pew, 1995)

A conceptual understanding of interruption may be informed by the *cognitive framework of multiple task management* proposed by Adams et al. (1995) in their review of the role of situation awareness in complex task settings. The framework provided an insight into how people cope with the demands of multitasking (i.e., interleaving the performance of two or more tasks), and how interruptions are integrated into performance. It was developed for informing aviation based technological design, and incorporates theories of perception and cognition (Neisser, 1976), and text comprehension (Sanford & Garrod, 1981). When working on complex tasks, it is assumed that attention is directed towards highly relevant information only, and all else is peripheral (Neisser, 1976). The theory proposed by Sanford & Garrod (1981) suggests that three cognitive processes control such an event governed by higher-order knowledge-schemas (i.e., chunk like declarative memory structures). The processes are:

Explicit focus, which contains the material in current attention and WM; *Implicit focus,* which contains all other task relevant information not in current attentional focus; and *long-term memory (LTM)*, which contains the inactive declarative memory chunks.

The framework proposed by Adams et al. (1995) provides testable constraints that may guide the understanding of interruption effects. Generally, it is proposed that performance efficiency is dependent upon the processing demands of primary and secondary tasks and the compatibility of the processing resources between these tasks. For instance, if a simple task requiring few attentional resources (e.g., filing journal articles) is interrupted by having to perform a simple secondary task sharing similar processing resources (e.g., removing all articles reviewing situation awareness), performance attenuation will be minimal. The explanation for this is that, both tasks require limited attentional resources, share similar processing resources and therefore both are compatible with a single knowledge-schema. Disruption to the primary task is minimal because its explicit contents do not need to be updated or removed to accommodate the processing demands associated with the secondary task. However, if the secondary task required incompatible processing resources (e.g., typing a text message on a mobile telephone), primary task information in explicit focus may be displaced. Reactivating this information should require a time consuming process of accessing knowledge contained in implicit focus (e.g., filing all types of PhD related material). For complex primary and interrupting tasks, the incompatibility of processing resources will result in the displacement of explicit and implicit information, and post-interruption performance will be dependent upon accessing LTM knowledge (e.g., general knowledge of how to file). Retrieval from LTM is assumed to be the most time consuming due to the necessity to activate 'inactive' knowledge-structures, and is prone to mistakes.

In sum, Adams et al. (1995) provide a useful framework for understanding differential interruption effect sizes. It achieves this by considering the structure of tasks and how they may be processed during complex cognitive activity. The framework will serve to support experimental hypotheses within this thesis, particularly when considering the contributory influence of task processing demands during an interruption episode.

I now turn to a review of related conceptual domains, each with a wealth of theoretical support, which may provide supportive and/or alternative explanations for interruption effects; particularly the costs associated with the suspension and resumption of goals.

1.3. Related conceptual domains

1.3.1. The Zeigarnik effect and intention superiority

Almost 80 years ago, it was proposed that memory retrieval performance was better for an interrupted task rather than an uninterrupted task (Lewin, 1926) Ovsiankina, 1928; Zeigarnik, 1927). Suspending an intention to perform task was assumed to result in a mental state referred to as *goal tension* (e.g., Lewin, 1926). Goal tension was described to reflect a suspended intention residing in a higher level of conscious awareness compared to intentions that had been completed.

In Zeigarnik's (1927) classic paper, a 'heightened activation' account of suspended goals was empirically supported through a series of experiments. In her first experiment she observed participants whilst they completed a series of 22 tasks under the instruction to perform as quickly and efficiently as possible. Half of the participants were instructed to suspend one randomly chosen task within the series to work on an interrupting secondary task (e.g., threading beads). The suspended task was never resumed, and instead, participants returned to the next task within the series. A free recall test implemented after completion of all tasks, revealed that details of the interrupted task were remembered better than details of the uninterrupted tasks; a phenomenon since referred to as the *Zeigarnik effect*. This effect was replicated when: Interruption fell within most of the 22 tasks; using different tasks; using children as participants; *and* if interrupted tasks were completed (Zeigarnik, 1927). Taken together, the results suggest that the mental representation formed of an interrupted task is in some way more enduring than for one that has been completed. However, the results do not in any way inform theory of how the process of task interruption affects performance, in part because Zeigarnik's manipulations did not require participants to remember what they were doing immediately before or following interruption.

Although the Zeigarnik effect has been difficult to replicate in subsequent task interruption studies (e.g., Van Bergen, 1968), related research has identified similar effects. In particular, a small number of studies have produced evidence of superior memory for certain forms of intentions (e.g., Goschke & Kuhl, 1993; Marsh et al., 1998). This, the *intention superiority effect* (ISE), has been documented when participants are asked to recall information regarding uncompleted intentions that were to be executed by themselves; an effect that is abolished if the uncompleted intention is observed rather than performed. For instance, Marsh et al. (1995) reported Zeigarnik like effects when demonstrating that memory for suspended intentions was better than that for completed intentions. Since the ISE is usually person specific, it has been suggested that uncompleted intentions reside in a heightened level of activation in memory (e.g., Goschke & Kuhl, 1993; Marsh et al., 1995). It has recently been suggested that increased activation of uncompleted intentions may account for the ISE (Anderson & Lebiere, 1998), whereas rapid inhibition may explain poor memory for completed intentions. The Zeigarnik effect may be nothing more than a demonstration of intention superiority: A phenomenon that is difficult to replicate when participants have to resume an activity immediately following an interruption.

The theoretical basis for the idea of heightened activation of suspended intentions may be limited for explaining the processes affected by task interruption. If a suspended intention resides in a higher state of activation compared to other items, this would implore convincing evidence regarding a limited working-memory (WM) activation supply (e.g., Lovett et al., 1999). A heightened activation account could also imply that suspended intentions may be stored in a different representational form to other memories, again an unlikely event given evident cognitive limitations (e.g., Byrne & Bovair, 1997). Reber and Mycielska (1982) describe four common errors unexplainable by an ISE account. These are: Performing the same action twice; choosing the wrong item to act upon; mixing actions (e.g., taking your keys out of you pocket when walking up to the wrong car); and actions out of sequence (e.g., getting in the bath with your socks still on). The most convincing evidence against the heightened activation claim comes from the goal-directed memory literature. For instance, efficient retrieval of a goal seems to be based upon the amount of activation that goal receives before it is

suspended (e.g., Altmann & Trafton, 2002, 2004). If a suspended goal is not encoded, then it is likely that it will be forgotten and will have to be reconstructed (e.g., Anderson & Douglass, 2001).

1.3.2. Prospective memory

Prospective memory (PM) typically involves remembering to perform a future intention (e.g., remembering to buy milk when passing the supermarket), and is relevant to the study of interruptions. The cognitive processes involved in PM include, encoding, storage, activation/inhibition, retrieval, and updating (e.g., Mäntylä, 1996). PM is usually studied in a different way to task interruption, although there are many similarities between the two research areas. The first thing to note is that PM researchers generally refer to the suspended item as an intention (i.e., the course of action one wishes to follow), whereas interruption researchers prefer to class it as a goal (i.e., the end product of a plan). In fact, classic PM research was motivated to understand PM processes (i.e., remembering when to do something) in isolation of those that are retrospective (i.e., remembering what it is you wanted to do)(e.g., Meacham & Leiman, 1982). However, the types of tasks utilised in PM studies (e.g., remembering to press a keyboard key at a certain time), inherently contain a retrospective component (i.e., which key to press or *the action*). Thus it is unsurprising that it has recently been suggested that an interrupted task becomes a prospective memory task (Dodhia & Dismukes, 2003). Secondly, the length of the retention interval between intention formation and execution is generally longer in PM studies. Retention interval or *interruption length* varies considerably in interruption studies from a few seconds (e.g., Hodgetts & Jones, 2003) to minutes (e.g., Gillie & Broadbent, 1989). Thirdly, knowing when to perform a suspended intention is usually cued by reinstatement of a primary task following an interruption episode; a factor not always operationalised in PM studies.

Many factors that improve post-interruption performance also improve prospective remembering. These include, cognitive demands of the tasks between encoding and retrieval (e.g., Marsh, Hancock & Hicks, 2002), length of the storage retention interval (e.g., Einstein, Holland, McDaniel & Guynn, 1992), and the availability

and salience/distinctiveness of retrieval cues (e.g., McDaniel & Einstein, 1993; Guynn, McDaniel & Einstein, 1998).

PM has been studied using a range of paradigms. These include: *Event-based*, where the intention is performed in response to the presentation of a specific stimulus; *time-based* where intention is performed at specific time intervals (e.g., every 5 minutes); and *delayed execute,* where participants are cued to perform an intention but only after completing their current task. The delayed execute paradigm has a significant overlap with a typical interruption paradigm (e.g., Einstein, McDaniel, Williford, Pagan & Dismukes, 2003). Both involve temporarily (e.g., seconds) suspending the performance of an intention to complete an interleaved task. Einstein et al. (2003) found that performance of a delayed intention was mitigated compared to a non-suspended intention. This deficit was enlarged if participants were unexpectedly interrupted (for 15 s) during the delayed retention interval. Interestingly, McDaniel, Einstein, Graham and Rall (2004) have recently demonstrated that delayed execute performance deficits are almost completely alleviated if a simple mnemonic cue (i.e., a blue dot positioned on the visual interface) is made available during the retention interval, even when interrupted.

When considering the theoretical basis for PM effects, a number of explanations have been proposed, each of which may inform the study of interruptions. Mäntylä (1996) suggests that the product of three interacting processes will determine the retrieval outcome of a suspended intention. These are defined as component processes and are: *Trace-dependent* (the activation level of the intention); *cue-dependent* (the perceptual attributes of cueing stimuli); and *capacity-dependent* (the available processing resources to mange the retrieval of the intention). It is believed that the strength of cue-dependency as well as trace-dependency can be increased before an intention is suspended (e.g., during planning and encoding) (e.g., Goschke & Kuhl, 1993). Evidence for this comes from the finding that interrupting the encoding phase of a PM task abates the chance of retrieving the future intention (e.g., Mäntylä & Sgaramella, 1997). So, improving PM performance may be reliant upon the strength of the association between a prospective intention and its retrieval cue.

Of direct relevance to the study of interruptions is the 'Noticing + Search' model (Einstein & McDaniel, 1996). The model estimates that the familiarity associated with

an intention retrieval context (e.g., being in a shopping centre) may elicit search strategies for an appropriate retrieval cue (e.g., a shop that sells milk). Upon identifying a cue (e.g., a grocery shop), a 'directed search' process may occur (e.g., deciding whether or not the shop chosen will have any milk left to sell that day) in order to determine the utility of performing the associated action. It is proposed that an unexpected interruption episode should similarly produce a feeling of uncertainty when returning to the interrupted task. So, a suspended goal may be re-established, but that does not guarantee that its associated action will be immediately executed.

In sum, PM research provides a number of explanations for the often-negative effects associated with being interrupted. In many ways, established PM theory is similar to that of goal-directed memory (i.e., proposes the importance of activation, cue strength and capacity limitations). However, PM research has investigated how the event of having to suspend an intention might affect its future performance. In particular, it accords to an earlier point that I stressed: The processing a goal receives prior to interruption annunciation (e.g., during planning) might affect its retrieval efficiency. Additionally, some PM studies have contemplated the costs associated with performing an intention, even after it has been retrieved. This could suggest that restarting a task following an interruption episode might involve more than re-establishing the suspended goal, for example, evaluating its utility to subsequent performance. Experiments presented in this thesis aim to investigate this claim.

1.3.3. Task Switching

A typical task switching (TS) experiment involves switching between two short tasks, which is similar to the requirement to alternate between tasks when interrupted. Having to return to a suspended task is what distinguishes task interruption from the general paradigm used to study TS. The general paradigm used in TS research boasts experimental control, evidenced by effect reliability. Mostly, there are larger switch costs associated with switching between different tasks (e.g., AB) than when alternating between the same task (e.g., AA); a phenomena that has been replicated many times (e.g., Allport, Styles & Hsieh, 1994; Jersild, 1927; Rogers & Monsell, 1995). Unsurprisingly, post-interruption resumption deficits are mostly larger than the costs incurred from

switching between tasks (e.g., Altmann & Trafton, 2004). Nevertheless, TS research affords a fine-grained conceptual understanding of the cognitive processes involved in alternating between tasks, which may inform interruption theory.

Two accounts provide conceptual explanations for such effects. The first account, 'task-set reconfiguration (TSR)', advocates an active transformation of attentional resources causing the diversion of attentional resources from the first task to the second task (e.g., Monsell, 2003; Rogers & Monsell, 1995). This takes more time for 'switch trials' where attention has to be reconfigured for processing of the second task to occur. The second account, 'task-set inertia (TSI)', stresses that both passive priming of a possible task switch, and inhibition of the old task need occur to incorporate the new task (e.g., Allport et al., 1994). Using an 'explicit task cueing paradigm', where participants are presented with a passive priming cue from the second task before switching, has been shown to result in reduced switch costs (e.g., Sudevan & Taylor, 1987; Meiran, 1996). However, the time taken to perceive, encode and act upon a cue may contribute to the switch costs exhibited (e.g., Logan & Bundesen, 2003; Mayr & Kliegl, 2003; Monsell, 2003). Inhibition may serve to constrict retrospective interference experienced from older goals, resulting in larger costs in switch trials. Recently, it has been proposed that both accounts may inform a unified theory of the processes involved in TS (e.g., Monsell, 2003). Both accounts also share an overlap with factors already considered to contribute to interruption affects (e.g., limited activation, interference and decay).

The TS literature has also identified a number of methods that reduce the size of switch costs; similar to those suggested for alleviating the negative effects of interruption. These include, allowing a preparation period before alternation (e.g., Gopher, Armony & Greenspan, 2000; Meiran, 1996), practice (e.g., Nieuwenhuis & Monsell, 2002; Rogers & Monsell, 1995), reducing task difficulty (e.g., Rubenstein, Meyer & Evans, 2001), and allowing control over when to switch (e.g., Arrington & Logan, 2004). Given the successfulness of these methods at attenuating performance deficits for TS and task interruption, they may also inform a theoretical understanding of the cognitive processes affected by interruption. In particular, advanced preparation may be a key determinant factor in reducing interruption effects. An important question however, is *when* this preparation should occur to yield the highest payoff to post-interruption performance?

Experiments presented in this thesis will compare preparation at all stages of the pre-interruption episode in an attempt to address this question.

1.3.4. Long-term working-memory (LTWM, Ericsson & Kintsch, 1995)

Performing the same task multiple times often leads to more efficient performance of that task (e.g., Schneider & Shiffrin, 1997, Schneider & Detweiler, 1988). The LTWM model (Ericsson & Kintsch, 1995) offers one account for this phenomenon by assuming that memory is a unitary storage mechanism where WM is an active subset of LTM. When performing a task, WM is used for the temporary storage and retrieval of items directly relevant to advancing short-term performance (e.g., Baddeley & Hitch, 1974). This short-term WM is a time limited, capacity dependent store, that is responsible for manipulating information, switching attention between tasks, and may also be responsible for selecting relevant information and inhibiting irrelevant information (e.g., Baddeley, 1986).

The model provides a number of constraints that may highlight some of the cognitive factors involved in the suspension and resumption of goals. During encoding (e.g., before suspension), information incidentally enters LTWM during skilled performance. The speed and efficiency of this process increases with: Practice; if new information is being mapped onto pre-existing LTM knowledge schemas; and if the selectivity of material is not restricted by high task demands. Resumption via retrieval is largely based upon the strength of associations. Specifically, information is encoded into retrieval structures in LTM and these can be activated by cues or pointers available in STM (e.g., Chase & Ericsson, 1982).

In the case of task interruption, experience of performing a task may attenuate the need to rely upon a restricted short-term WM store for remembering suspended goals. If the goal is part of a LTWM knowledge structure, identification of a retrieval cue may be sufficient to retrieve that goal. Additionally, if LTWM operates via the strength of connections between short-term WM and LTM, practising a task should lead to faster storage and retrieval of its component goals.

In sum, task expertise together with an environment rich in mnemonic cues (both before and following interruption) might result in the dissipation of performance deficits

caused by being interrupted. If one is sufficiently practiced at a task, encoding may sufficiently strengthen the association between a suspended goal and its LTWM cue. However, such encoding may not have to occur immediately before being interrupted to effectively prime a LTWM representation, although its associated cue must be available upon primary task re-instatement. For efficient retrieval, this cue should not be masked by competing stimuli. Experiments presented in this thesis will explore the relationship between encoding and the availability of retrieval cues for both minimally practiced and highly practiced tasks.

1.3.5. Summary of related conceptual domains

It is proposed that the conceptual domains reviewed within this section may inform the development of interruption theory, at least when considering goal-directed behaviour. The mechanisms underlying the intention superiority effect may not be the same as those used to recover an interrupted goal, although the event of being interrupted may induce a state of tension to resume an uncompleted goal. The PM literature highlights limitations in the ability to store and retrieve suspended intentions. Additionally, it stresses that in situations of uncertainty (like being interrupted), humans are prone to evaluate the consequences of performing a suspended intention, even after it has been successfully re-established (e.g., Speier et al., 2003). The TS literature supports the premise that preparing to resume an interrupted task may reduce post-interruption performance deficits, but suggests that there will always be a cost of switching between tasks that require different processing resources. Specifically, switching attention to another task may require the activation of new knowledge structures as well as inhibition of those supporting the task that is switched from (e.g., Monsell, 2003). Finally, although people can acquire expertise at performing a task through practice, it may be that this experience has little utility when dealing with unexpected interruption episodes.

1.4. Developing a conceptual approach for studying task interruption

1.4.1. Goal-directed behaviour

Established theoretical constructs regarding goal-directed memory may be useful for informing a conceptual understanding of the process of being interrupted. An investigation of goal-directed behaviour led to the development of a recent cognitive model that predicts the processes and constraints involved in suspending and resuming goals (the *goal-activation model*, Altmann & Trafton, 2002). Before reviewing the goal-activation model, it is important to understand the empirically grounded theory that motivated its development – the Adaptive Control of Thought (ACT, Anderson, 1976) and the more recent Adaptive Control of Thought-Rational (ACT-R, e.g., Anderson & Lebiere, 1998).

1.4.1.1. ACT-R (e.g., Anderson & Lebiere, 1998)

ACT-R is a cognitive architecture that attempts to capture the nature of human knowledge, particularly how it is acquired and deployed. It predicts that human behaviour is guided by the current goal (e.g., writing this paragraph), which is usually that within attentional focus. How this goal is processed depends upon the way in which it is modified by three connected memory systems: The *goal stack* (reviewed in the next section), which manages and encodes the hierarchical structure of current goals guiding behaviour; *procedural memory,* containing production rules that act upon these goals; and *declarative memory,* containing memory chunks relevant to performing goals (i.e., the things we 'know', Ryle, 1949). ACT-R functions through these systems interacting during discrete processing cycles.

The flow of behaviour in ACT-R is dependent upon the firing of 'production rules'. These are condition-action parings, specifically declarative knowledge facts that serve to modify the goal[1]. An example of a production rule is:

[1] The deployment of a production rule is dependent upon its cognitive fit with a declarative memory chunk, but ACT-R can also learn new rules at a subsymbolic level. The subsymbolic level involves the processing of continuously varying qualitative perceptual quantities, which through statistical regularity can be transformed to declarative knowledge.

```
IF the goal is to classify a shape [the goal]
   and the shape has four equal sides [the condition]
THEN classify the shape as a square [the action]
```

In the given example, if declarative memory contained an active representation of a square, then the action resulting from the firing of the production rule should occur within 50 ms – 1 s, thus transforming the goal. Transformations generally involve; *goal pushing*, where a goal is modified and stored for future use because it cannot be executed at the present time; and *goal popping*, where a goal can be executed, and is removed from the stack. The selection of the best production rule to modify a specific goal is generally dependent upon the process of *conflict resolution*. If a goal is only associated with one production rule (e.g., 8 + 4 = '?') then it has no need for conflict resolution (i.e., '?' = 12). However, if it is associated with more than one production rule (e.g., 'A' '?' 4 = 12), then the system has to decide which is the best production to achieve the goal (e.g., 'A' = 8, '?' = +, or, 'A' = 3, '?' = x, and so on). This will affect the time taken to process a goal, and thereby introduces the idea of varying retrieval latencies.

If a goal is processed to a level that it is encoded in memory, ACT-R assumes that it can be retrieved. Speed of performance and speed of retrieval are thus highly related. Retrieval latency consists of, the sum of time taken to retrieve a chunk, the time taken to select and implement an appropriate production rule, and the activation strength of the goal to be retrieved. The retrieval process is affected by the connection between declarative and procedural memory and is controlled by the activation level of a declarative chunk. The formal activation of a chunk (**Ai**) is quantifiably formulated in equation 1:

$$\mathbf{Ai} = \mathbf{Bi} + \sum \mathbf{Wj}\ \mathbf{Sji}$$

(Equation 1)

Bi refers to the *base-level activation* of the chunk, and is an activation value determined by how recently and frequently that chunk has been accessed. **Wj** refers to the *source activation* of a chunk, and results from the amount of attention given to the slot values of the goal (e.g., '8', +, '4' and '=' are slot values in the calculation 8 + 4 =

12). The received opinion is that the source activation for any given goal is limited, as activation is in short supply and has to be carefully distributed among all goals (e.g., Lovett, Reder & Lebiere, 1999). Each slot value will receive the same level of source activation, and further processing of one element may have the consequence of reducing the active strength of others. S_{ji} refers to strength of association and is how often chunk i was needed when j was an element of the goal (e.g., the past experience of calculating the sum 12 when 8 was a slot value). ACT-R does additionally suggest that retrieval can fail, usually due to the insufficient activation strength of one or more elements contained within Equation 1. However, it has been suggested that retrieval failure using one production rule usually leads to another retrieval attempt using a different production rule, sometimes resulting in partial retrieval (e.g., Siegler, 1988). The more production rules accessed the longer the retrieval latency.

1.4.1.1.1. ACT-R and the goal stack

The goal stack is of interest to the study of task interruptions because it provides a hypothetical construct for considering how goals may be represented in memory. Anderson and colleagues have posited that goals are processed and executed in a 'last-in first-out' (LIFO) order, thus satisfying the main cognitive constraint of ACT/ACT-R, that the most recently processed goal will govern behaviour (e.g., Anderson, 1993, Anderson, Kushmerick & Lebiere, 1993, Anderson & Lebiere, 1998). Simply, if uninhibited satisfaction of this goal is possible (i.e., the outside world does not constrain its execution), then it is executed and 'popped', and the next goal is pushed to the top of the stack to govern behaviour. Inhibition may occur because of having to execute other goals first. In such a case the current goal is modified to become a subgoal and is pushed further down the stack for future use.

ACT-R, and earlier versions of the theory (e.g., ACT, Anderson, 1993) have used the ToH task to describe and empirically test the processes involved in 'stack-like' goal memory. The ToH task involves the movement of a series of discs varying in size across three pegs to get them from a start configuration to a goal configuration. Larger discs can never be moved on top of smaller discs, thus enforcing the encoding, suspension, and resumption of goals at different task points. ACT-R dictates that the ToH has algorithmic

characteristics, meaning that problems can be solved by considering each goal in a sequential fashion, and by using the perceptual cues within the environment as declarative knowledge (so therefore declarative chunk retrieval is not always necessary). Although the ToH task will be described in detail later on in this chapter, the production rules that are applicable to its performance need to be reviewed to understand the processes involved in goal suspension and resumption. ACT-R denotes that production rule firing will change the representation of a goal, even if it is not popped from the stack or modified. There are six ways a goal can be acted upon: *No change*; *goal elaboration*; *side effect*; *return result*; *pop unchanged*; and *pop changed.* Goal elaboration underlies subgoaling. If performance of a goal within the ToH is interrupted, its representation is modified (elaborated on), but it is not popped/executed. Following secondary task completion, the goal will be on the top of its stack and will be retrieved in its modified form.

The concept of a goal stack, although plausible, does not have universal support from studies of goal suspension and retrieval. In the next section, I will review some of these studies, and discuss the limitations of the goal stack in explaining the processes affected by task interruption.

1.4.1.1.2. The assumption of perfect memory for goals

Controversy surrounds the perfect goal memory assumption. Stemming from the goal stack premise that goal memory operates using a hierarchical and non-effortful storage and retrieval mechanism (i.e., uninhibited LIFO), many have assumed that such memory must be infallible (e.g., Anderson et al., 1993; Goschke & Kuhl, 1993). However, empirical evidence has clearly shown that this is not the case (e.g., Altmann & Trafton, 2002; VanLehn, Ball & Kowalski, 1991).

A number of experimental manipulations have identified costs associated with the hierarchical storage and retrieval of goals. For instance, forgetting to execute a subordinate goal (e.g., removing the original from the photocopier) because a superordinate goal has been achieved (e.g., removing the copy) is said to reflect a common human cognitive processing deficit known as *post-completion error* (e.g., Byrne & Bovair, 1997). Such error cannot be explained by the goal stack, which posits that all

goals – even superordinate – have to popped from the stack to satisfy a 'heightened activation' goal memory process. Additionally, many advocate the view that the activation supply made available to process goal-sets is limited, and must be carefully distributed to achieve efficient performance across all goals (e.g., Altmann & Trafton, 2002; Anderson & Lebiere, 1998; Lovett et al., 1999). Consequently, not all goals will have the same level of activation, and some may be forgotten, or confused with those that have similar representations (e.g., Altmann & Trafton, 2002). This leads to the suggestion that different goals will have different activation levels reflected by their importance in advancing goal-directed behaviour; something the goal stack does not support.

Recently however, the goal stack concept is undergoing serious reconsideration. Anderson and Lebiere in their 1998 book, *The Atomic Components of Thought* (p. 40), confessed to, "...not being particularly sanguine about this feature", and suggest that "...perhaps future research in this domain will expose problems with ACT-R's assumption of perfect goal memory". Experiments within this thesis will test this premise. Evidence of post-interruption resumption latencies longer than those proposed by ACT-R's stack retrieval mechanisms would further denunciate the idea of infallible goal memory, encouraging structural modification.

1.4.1.1.3. A cost of delaying goal stack performance

ACT-R assumes that a single task is performed using a single hierarchically organised goal stack, and does not satisfactorily explain how other tasks may interrupt or be interleaved with this process. It has been simply posited that task performance can be postponed in response to an in urgent interrupt (e.g., somebody shouting the word 'fire' during performance of a primary task), where 'high-priority' production rules will supersede those being deployed to perform the primary task. The processes governing the deployment of these production rules is not elaborated upon; neither is the cost of suspending one goal stack to incorporate another. Additionally, managing multiple tasks is thought to occur by interleaving their goals into a one stack governed by a single overall goal (e.g., 'perform all tasks'). Given cognitive limitations, a single stack may not support such a complex hierarchy of goals, whereby ACT-R assumes that tasks may

instead be performed sequentially. However, Anderson & Lebiere (1998, p. 40), do not attempt to flesh out the affects of such unpredictable goal suspension on the cognitive management of goals, and instead state that, "Cognitive psychology has tended not to be in the business of creating such emergency interruptions and studying the cognition that results". They go on to suggest that, "...we cannot say that ACT-R's mechanism is the right mechanism for modelling such interrupt handling because there is no data with which to assess it." This thesis will attempt to show that task interruptions do not have to be perceived as being urgent to disrupt goal memory performance. It will also produce a range of behavioural data from which future modelling can be conducted.

1.4.1.2. An ACT-R model of goal suspension and retrieval (Anderson & Douglass, 2001)

The goal stack assumption, that little cognitive effort is required to retrieve goals is subject to theoretical and experimental scrutiny (e.g., Altmann & Trafton, 2002; Anderson & Douglass, 2001). ACT-R theory proposes that the ToH task environment can be used a declarative retrieval source for expert users, thus reducing the need to use declarative memory (e.g., Anderson & Lebiere, 1998). This view was endorsed by Anderson & Douglass (2001). They taught participants how to solve ToH problems by using a 'sophisticated perceptual strategy', whereby the use of hierarchical declarative memory was minimised. Specifically, participants had to get the largest-out-of-place disc (LOOP) disc to its goal position, by moving all other discs to the peg that was neither the source nor destination of the goal move, and to repeat this process until the problem was solved. Participants had to solve problems by posting (i.e., stacking) all goals they could not execute immediately to a goal stack that was hidden when moving disks. This was to encourage subgoaling behaviour to generate data for the time taken to encode and retrieve goals and goal structures. To avoid practice effects by repeating the same ToH start and goal states over and over again (e.g., Ruiz, 1987), a range of arbitrary problems were used in the experimental phase. Move latencies and eye-tracking data suggested that participants tended to abandon the process of encoding goals that were to be performed in the distant future; instead they opted to sequentially post goals relevant to the performance of current subgoals. However, goal retrieval times were high, especially for those positioned at the beginning of complex goal structures (e.g., within subgoals). Eye-

tracking data also demonstrated evidence of goal forgetting, where 4% of the overall experimental time was spent re-planning problems using environmental cues.

Data generated from the experiment was modelled using an adapted version of ACT-R's declarative memory, with the goal stack omitted (Anderson & Douglass, 2001). An effort driven forgetting factor was incorporated, although the model was designed to dismiss any cost associated with goal rehearsal[2]. This too, revealed costs associated with storing goals for future use. The probability of accurate goal retrieval was estimated at: 75% for a goal stored for 13 s; 50% for a goal stored for 18 s; and 25% for a goal stored for 25 s. In all, the data provided strong evidence for the costs associated with retrieving goals, even when participants are familiar with the task *may* not have to rely upon declarative memory retrieval processes.

Conversely, I propose that being interrupted may encourage strategic decisions to dedicate more resources to goal encoding, given the unpredictable chance of having to suspend goals (e.g., Altmann & Trafton, 2002). Moreover, the strategy utilised to reduce encoding demands may be restricted to situations where people are confident that the stimulus array, both prior to and following interruption, can support goal reconstruction without relying upon declarative memory. Generally, assuming that people routinely offload complex goal management to the task environment, irrespective of utility, is suspect in contexts other than problem solving (e.g., Wilson, 2002; Payne, Richardson & Howes, 2000). It is this premise that leads on to the next review of a model that predicts the costs associated with both the suspension *and* resumption of goals; the goal-activation model (Altmann & Trafton, 2002).

1.4.1.3. The goal-activation model (Altmann & Trafton, 2002)

The goal-activation model is a framework that models goal-directed behaviour in problems that require the suspension and resumption of goals. Moving away from the concept of a goal stack, Altmann and Trafton (2002) suggest that goal memory is both effortful and vulnerable to errors. It also challenges the view of Anderson & Douglass

[2] This approach was adopted because the authors assumed that there was little cost associated with goal encoding. However, the data used was generated by participants who were very familiar with the ToH task, through training and previous exposure (Anderson et al., 1993). A recent model using the same data *has* identified costs of goal encoding (Altmann & Trafton, 2002)

(2001) in stressing that there are costs associated with goal encoding as well as goal retrieval. Using data generated from the ToH task (Anderson et al., 1993), Altmann and Trafton (2002) modelled the costs associated with suspending and resuming goals using adapted quantitative predictions from the ACT-R cognitive architecture (e.g., Anderson & Lebiere, 1998).

1.4.1.3.1 Activation, decay and interference

An activation-based approach is used to explain the short-term processing and retention of goals. The likelihood of retrieving a suspended encoded goal is co-dependent upon the efficiency of both higher-order cognitive processing (neural activation) and the use of environmental cues (associative activation). The time spent encoding a goal is likely to affect its performance efficiency during a retrieval episode (e.g., Altmann & Trafton, 1999, 2002). Memory for a suspended goal will decay as a temporal function of the time elapsed since it was last processed. This decay is functional in that it allows future goals to govern behaviour if the current goal is not executed (e.g., Altmann & Gray, 1999). The goal-activation model also predicts that a retrieval attempt of a suspended goal will be affected by retroactive interference from recently processed goals (e.g., Altmann & Gray, 1999; Altmann & Schunn, 2002). All of these processes are controlled by a limited activation supply (e.g., Lovett et al., 1999), governed by a similar sampling process to conflict resolution in ACT-R. Specifically, this is a retrieval cycle that occurs every 50 ms – 1 s that is sensitive to retrieving only the goal that resides at the highest activation level.

In their model, Altmann and Trafton (2002) define three main constraints for the encoding, hierarchical organisation and retrieval/resumption of a suspended goal. These are: The interference level; strengthening; and priming, each of which is illustrated in Figure 1.2 and described in the following subsections.

1.4.1.3.2. The interference level

The form of the goal-activation model views human working-memory as a collection of gradually decaying items that when trying to retrieve a goal that is specific to advance behaviour create a 'mental clutter'. The interference threshold is a construct

proposed to produce some order to this clutter. Formally its level is determined by the mean activation of the most active distractor goal. The goal-activation model harmonises with ACT-R in agreeing that only one goal can govern behaviour during a temporal episode, and will be that which exceeds the interference threshold during a retrieval cycle.

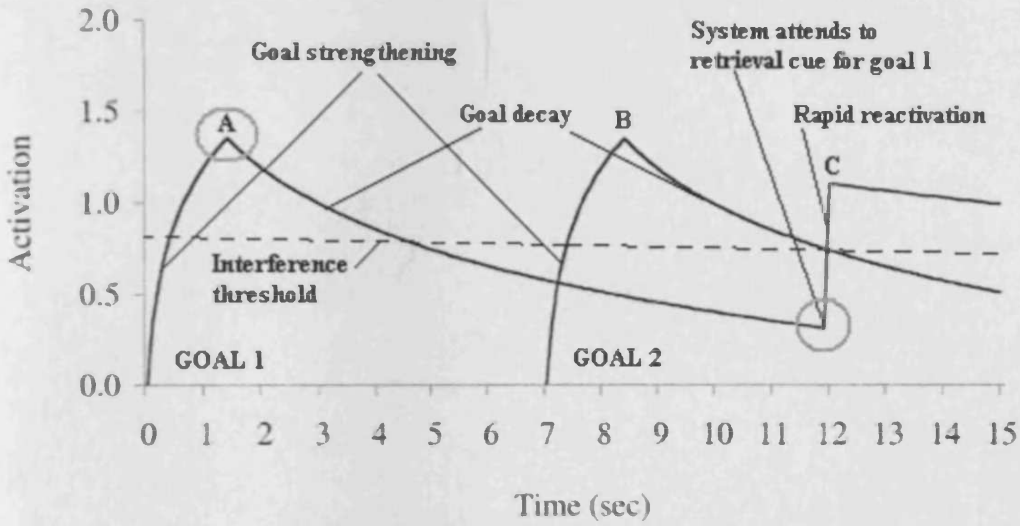

Figure 1.2. Strengthening, decay and associative priming in the goal-activation model (taken from Altmann & Trafton, 2002). Unfilled circles represent points where a suspended goal might be associated with a retrieval cue at encoding and attempted retrieval.

1.4.1.3.3. Strengthening

For a goal to govern behaviour it has to be activated via a process of strengthening for it to exceed the interference threshold. When a goal is first processed, its activation level rapidly increases to achieve this status within approximately 1.5 s – 2 s. For example, Figure 1.2 shows that 'goal 1' is strengthened to exceed the interference threshold, thereafter governing behaviour if it is sampled during a retrieval cycle. If goal 1 receives no further strengthening, its representation will gradually decay. Strengthening of 'goal 2' will take the same time before it governs behaviour ('B'); meanwhile goal 1 will continue to decay. Strengthening, as a process is responsible for, initially activating a goal, maintaining its activation (to avoid decay), and rapidly

increasing the activation of a decaying goal (the value of which will become apparent in the next section).

The amount of strengthening a goal receives determines whether or not it will be sampled during a retrieval cycle. The actual activation level is based upon an adaptation of the fundamental ACT-R base-level activation equation (e.g., Anderson & Lebiere, 1998). This is formulated in equation 2:

$$m = ln\left(\frac{n}{\sqrt{T}}\right)$$

(Equation 2)

where m is the overall activation of a specific goal, n is the number of times the goal has been sampled (throughout a lifetime) and T is the duration of the goal's lifetime since it was first processed. Thus, the activation level of a goal is determined by, the number of times it has been processed, the temporal proximity of these processing cycles, and the overall processing time.

There may be a functional role for strengthening a goal and the rate in which it decays between its suspension and resumption (e.g., Altmann, 2002, Altmann & Gray, 2002). Strengthening a goal before interruption through rehearsal might increase its base-level activation (e.g., Altmann & Trafton, 1999), making it more resilient to retrospective interference from other competing goals.

1.4.1.3.4. Priming

The goal-activation model suggests that the efficiency of resuming an interrupted goal through retrieval will vary given processing constraints imposed by, a limited activation supply, goal decay, and retroactive interference from older intervening goals. The fluctuating interference level links all such constraints. Firstly, over-strengthening of any secondary task goal (thereby increasing the mean activation of the most active distractor) will cause the interference level to rise, consequently making it harder to reactivate the suspended goal. Secondly, the representation of a goal is likely to decay if not regularly activated during performance of a secondary task (which is not always

possible due to secondary task demands). Finally, goals sharing similar representational attributes may have similar mnemonic representations, increasing the risk of retroactive interference leading to errors at retrieval.

The goal-activation model mechanism for retrieving a suspended goal is priming, which operates through goal association with a retrieval cue that has to be available before the goal is suspended and when the primary task is reinstated. Within the goal-activation model, *associative activation* – a source of goal strengthening available from the current context (i.e., priming cues) – is in limited supply, and is equally distributed amongst all goals (e.g., Anderson et al., 1996). When expertise in the task domain is low, associative activation must be carefully distributed amongst a limited number of goals, and when task expertise is higher, distribution can occur at the goal chunk level (with each chunk containing more than one goal). Implications are that for expert task users, an associative cue can be a declarative mental representation of the goal chunk, but for non-experts, a cue has to be a particular environmental stimulus: A perceptual attribute that, "...primes that goal but no (or few) others" (Altmann & Trafton, 2002, p. 52).

In its current form, the goal-activation model is only in hybrid form, asserting *ad hoc* predictions regarding the costs associated with suspending and retrieving goals (Altmann & Trafton, 2002). Given the fact that it was developed using existing theory and only modelled on one data set, it is unsurprising that it may have certain limitations. The current thesis set out to address these limitations and to, wherever possible, steer the form of the goal-activation model.

1.4.1.4. Summary of goal-directed behaviour

It may be useful at this stage to recollect the key implications that goal-directed behaviour research offers to a conceptual understanding of task interruption.

1. The goal-stack is a restricted concept in that it underestimates the limitations of human memory, assuming minimal performance attenuation following an interruption. The accepted view is that the suspension and resumption of goals can be a fallible process, with performance likely to be effected by being interrupted (e.g., Altmann & Trafton, 2002).

2. Anderson & Douglass (2001) suggest that with the expectation of long retention intervals (as is usually a consequence of interruption) humans choose strategically abandon fallible goal encoding and storage processes, and opt to reconstruct goals from the perceptual task environment. By contrast, Altmann & Trafton (2002) comparatively modelled retrieval and reconstruction strategies, and found that the former was far more predictive of human goal-directed behaviour in the ToH task. The fact that retrieval occurred at all suggests that goal encoding, even over long retention intervals, might be a strategy used to maintain the representation of a suspended goal.
3. The retrieval functions provided by both ACT-R theory and the goal-activation model may not necessarily reflect the time taken to resume an interrupted task. ACT-R is firm in suggesting that retrieval of a suspended goal takes in the region of 50 ms – 1 s. The goal-activation model is more conservative in suggesting that: Following successful encoding of a goal (through strengthening and associative priming), retrieval can be almost instantaneous. However, recent research has shown that both views may largely underestimate the time to re-establish a goal following a task interruption (e.g., Hodgetts & Jones, 2003).

Given the firm theoretical grounding within the goal-directed behaviour literature, particularly with regard to the suspension and resumption of goals, the paradigm served as a basis for experimental investigation within the current thesis. As such, performance costs associated with interruption could be explored with strict conceptual guidance lacking in many previous empirical studies.

1.5. Experimental evidence of task interruption

I now turn to a review of the existing research that has investigated the behavioural affects of being interrupted. As stated from the outset of this chapter, the rationale for much of the existing interruption literature was to identify methods that may alleviate interruption affects, whilst largely neglecting theoretical development. Therefore, I aim

to review the main areas surrounded by theoretical controversy. I initially turn to a review of experimental evidence documenting the deleterious affects of task interruption in applied settings.

1.5.1. Applied studies

Given the ecological importance of understanding the effects of being interrupted (e.g., Eyrolle & Cellier, 2000) it is unsurprising that there is a wide range of non-laboratory based studies; even though their theoretical implications are generally limited. Task interruption in applied settings results in various performance deficits. These include, resumption lags (e.g., Zijlstra, Roe, Leonara & Krediet, 1999), increased task completion times (e.g., Eyrolle & Cellier, 2000; Paquiot, Eyrolle & Cellier, 1986), accuracy decrements (Bainbridge, 1984), and negative effects on worker well-being (e.g., Johanasson & Aronsson, 1984; Roe, Zijlstra & Leonova, 1996). Performance deficits have been documented within a wide range of applied settings including, the flightdeck (e.g., Dismukes et al., 1998) the office (e.g., Jackson, Dawson, & Wilson, 2001, 2002, 2003; Paquiot et al., 1986; Tétard, 2000), police radio dispatcher centres (Kirmeyer, 1988), nuclear power plants (Griffon-Fouco & Ghertman, 1984), and engineering sites (e.g., DeMarco & Lister, 1998; Perlow, 1999).

The negative effects of being interrupted in the workplace are widely documented when the outcomes effect the efficiency of making decisions and performing safety-critical actions. Below are two examples of such settings where effective interrupting technologies are essential: The modern office environment and the flightdeck.

1.5.1.1. The office

The frequency of interrupting events appears to be increasing in the modern office environment, an affect partially attributable to the surge in communicative interruption technologies (e.g., Loukopoulos et al., 2001; Dismukes, 2003). For instance, it was estimated that the most popular Instant Messenger (IM) applications had a combined demand of over 70 million users in 1999 (*Media Metrix, December, 1999*); a figure rising to 201 million for just 1 application (*MSN*) in a year (*Microsoft PressPass, July, 2000*). Previous research has shown that approximately 10 minutes of every working hour is

spent dealing with interruptions (O'Conaill & Frohlich, 1995). In 1998, it was been demonstrated in the U.S. that knowledge-workers were sent an average of 204 Emails per day (Pitney Bowes, 1998); although more recently it is reported that the average office worker is interrupted up to 73 times per working day (Cubesmart Inc., 2002). Although recommendations are in place informing people to finish important tasks before switching to an interrupting task, many studies have shown that people prefer to perform the switch immediately (e.g., Jackson et al., 2002, 2003; Tétard, 2000).

Disruptive interruptions within the work environment were documented well before the surge in Internet based communication technologies; the most common being telephone calls and 'pop-in visitors' (e.g., Dahms, 1988; O'Connail & Frohlich, 1995). These are often unscheduled and intrude upon primary task performance output (e.g., Coates, 1990; Vernon, 1990). Knowledge-workers for example, allow such interruptions to take precedence over other work activities (e.g., Jones & McLeod, 1986). Alarmingly, recovering a pre-interruption performance level following a telephone call can take up to 15 minutes (DeMarco & Lister, 1987).

1.5.1.2. The Flightdeck

The modern flightdeck has undergone considerable changes with semi-autonomous and multi-tasking technologies enforcing operators into more of a supervisory role (e.g., monitoring system changes) rather than one that is manual (e.g., manually maintaining altitude). Pilots have to manage multiple activities and frequently update their situational awareness because of the need to integrate new tasks. Despite the numerous benefits of modern technology within the flightdeck, HCI problems have been documented. Such problems include: Trying to effectively manage alarms (e.g., Stanton, Booth & Stammers, 1992); integrating multi-modal tasks (e.g., Latorella, 1998; Williams, 1995); and dealing with communications from Air Traffic Control (ATC) and colleagues (e.g., Barnes & Monan, 1990; Loukopoulos et al., 2001). Generally, performance deficits are believed to reflect lapses in attention and/or memory due to the demands imposed by the secondary task, although theoretical support for this is limited (e.g., McFarlane, 2002).

Studies have confirmed that flightdeck interruptions frequently occur from a number of sources, resulting in a range of effects (e.g., Damos & Tabachnick, 2001; Loukopoulos et al., 2001). For instance, Dismukes et al. (1998) describe how flightdeck accidents are partially attributable to being interrupted and/or distracted. Their main focus was on the cognitive continuity of performing checklist procedures. The performance in such tasks was found to be effected adversely by, communication, head down work (e.g., monitoring the flight management system), searching for air traffic, and responding to abnormal situations (e.g., weather changes). Similarly, Loukopoulos et al. (2001, 2003) using a flight simulator test-bed, detail the effects interruptions have on overlooking procedural steps such as, failing to check fuel gauges before takeoff, improper setting of pressurisation gauges before takeoff, and forgetting to set landing gear flaps during runway clearance.

Various reports have identified that interruption is at least partially involved in a number of aviation incidents. For instance, the National Transportation Safety Board (NTSB) reviewed incidents in U.S. air carriers and found that almost half were attributable to, 'lapses of attention associated with interruptions, distractions, or preoccupation with one task to the exclusion of another task' (Dismukes et al., 1998). Similarly interruption was cited as a major contributory factor in approximately 175/2500 incidents reviewed by the Aviation Safety Reporting System (ASRS) before 1979 (Monan, 1979). Interruptions to checklist procedures have been shown to be involved 58% of in-flight incidents (Turner & Huntley, 1991). More recently, an ASRS survey reported that 21 of 107 tasks were neglected during a critical in-flight period, due to preoccupation with another task (Dismukes et al., 1998).

Interruption is a factor that has also been attributable to a number of aviation accidents, some of which have been fatal. For instance, in December 1972, the crew of flight L-1011 was preoccupied with a malfunctioning landing gear light and failed to notice that the autopilot was not switched on. Simultaneously, they were entering a rapid descent, which resulting in a fatal crash killing 94/163 passengers and 5/13 crewmembers (NTSB, 1973). In 1988, another NTSB report revealed that shortly after takeoff, the crew of Northwest Airlines' flight 255 was unavoidably interrupted by ATC communications before verifying the status of the landing flaps. Only minutes after, an emergency

occurred, and due to the status of the flaps, the aircraft crashed killing all 148 passengers and 6 crewmembers. Both examples reflect the most extreme consequences of failing to recover from unavoidable interruptions: The cost of human life.

Within the flightdeck literature, it is evident that the driving motivation for research is to identify how best to recover from being interrupted rather than investigating the cognitive processes that may cause the negative affects in the first place. Informative guidelines have been provided by Smith and Mosier (1986), which recommended: Minimising the disruptiveness of the interruption annunciation stimulus (p. 364); using distinct controls for different interruption methods (p. 277); and to use indicators of the status of the suspended task (p. 280). One general rule of thumb is to explicitly mark where a task was interrupted using, for example, pointing at the interrupted task (e.g., Degani & Weiner, 1990). Although such methods might have some worth in reducing the disruptive affects of interruption, theoretical grounding is weak, and in some cases, lacking in experimental support (e.g., McFarlane & Latorella, 2002).

1.5.2. Laboratory studies

The laboratory setting provides a means to study interruption effects whilst controlling for extraneous variables; nevertheless advancements in theoretical understanding regarding their effects on cognitive performance has been limited. This may be in part due to the wide range of variables studied within laboratory based task interruption experiments, with little emphasis on focusing on the understanding of the affect of individual factors. These include factors that may cause negative interruption effects: Interruption complexity (e.g., Gillie & Broadbent, 1989); interruption position within a primary task (e.g., Adamczyk & Bailey, 2004; Monk et al., 2002; 2004); similarity between tasks (e.g., Czerwinski et al., 1991a, b; Gillie & Broadbent, 1989; Edwards & Gronlund, 1998); interruption length (e.g., Gillie & Broadbent, 1989; Lahlou et al., 2000); task modality (Arroyo & Selker, 2003; Latorella, 1998); and psycho-physiological measures (e.g., Kirmeyer, 1988; Zijlstra et al., 1999). Other factors have been investigated as methods of alleviating negative interruption effects: Warning (e.g., Miller, 2002); the availability of retrieval cues (e.g., Czerwinski et al., 2000; Franke,

Daniels & McFarlane, 2002) and rehearsal (e.g., Gillie & Broadbent, 1989). Even when studying individual factors, the theories used to explore reasons for effects vary considerably. For instance larger performance deficits are usually found when interrupting complex tasks/subtasks compared to those which are easier to perform. This effect has been attributed to: Attentional resource allocation limitations (e.g., Adamczyk & Bailey, 2004; Speier et al., 1997, 2003); workload demands (e.g., Bailey, Konstan & Carlis, 2000b) and the activation levels and priming cue availability of suspended goals (e.g., Altmann & Trafton, 2002, 2004; Trafton et al., 2003). To my knowledge, no existing theory has provided concrete evidence for supporting any one or all of these theories. Additionally, primary tasks used across laboratory studies have varied substantially and include: Calculators (Kreifeldt & McCarthy, 1981); arithmetic tasks (e.g., Detweiler et al., 1994); shopping based retrieval tasks (Gillie & Broadbent, 1989; Edwards & Gronlund, 1998); computer based productivity packages (e.g., Adamczyk & Bailey, 2004; Czerwinski et al., 2000a, b); driving tasks (e.g., Monk et al., 2002, 2004); and resource allocation tasks (e.g., Trafton et al., 2003). Finally, there has been a wide range of dependent measures used across laboratory experiments, including: Resumption latency; primary task completion time (sometimes including resumption latency); overall 'time-on-task' (ToT); mistakes/errors; and psycho-physiological measures.

Throughout subsections 1.5.2.1. - 1.5.2.5, I will review factors that may be largely responsible for the negative effects caused by being interrupted and methods used to mitigate these effects; each of which will be experimentally investigated throughout empirical chapters. Each will be reviewed independently with the aim of using discussed theory to explain their potential effects on cognitive performance. Summaries will be provided at the end of each section, forming the foundations of the main hypotheses echoed throughout empirical chapters.

1.5.2.1. Interruption position and processing complexity

The point in which an interruption occurs within a primary task seems to influence the effects it has on subsequent performance within that task. Firstly, the relative size of interruption effects may be related to the processing complexity of the goal or goal-set that is interrupted (e.g., Cellier & Eyrolle, 1992; Navon & Gopher,

1979). For instance, a goal whose execution is reliant upon the status of other goals may be more effected by interruption than a goal that can be independently executed (e.g., Bailey et al., 2001a, b).

The processing complexity associated with the performance of an interrupting secondary task may also effect primary task retrieval efficiency. For a single task, the greater the number of component parts to process, the harder it will be to maintain a high representational strength of any one component (e.g., Wood, 1986). This account is supported by those who maintain that activation is a limited resource, which has to be carefully distributed across components relevant for task completion, for efficient performance to ensue (e.g., Altmann & Trafton, 2002; Lovett et al., 1999). So, in the case of integrating a secondary task, the representational strength of primary task items may be abated by the 'limited processing and memory capacity' of STM (Miyata & Norman, 1986, p. 268). It has also been suggested that the nature of the processing required to perform primary and secondary tasks (e.g., visuo-spatial versus verbal tasks) may contribute to subsequent performance deficits in the primary task (e.g., Cellier & Eyrolle; Hess & Detweiler, 1994). Some believe that tasks requiring similar processing resources for their execution will result in small performance deficits (e.g., Adams et al., 1995), whereas others suggest the opposite (e.g., Gillie & Broadbent, 1989; Edwards & Gronlund, 1998).

Current theoretical views vary considerably in their propositions regarding how the position and complexity of a secondary task might effect primary task performance. In one view, both factors are predicted to have little or no effect on subsequent performance efficiency within the primary task. For instance, Anderson and Douglass (2001) – using an ACT-R model incorporating a retrieval factor unaffected by processing demands – would predict no effect of interruption position. Their model also predicts no cost of secondary task complexity, reflected by a parameter that excludes goal interference as a factor effecting retrieval. This is a general assumption of ACT-R related models that tend to avoid elaborating upon the cognitive affects of suspending one task to perform another (e.g., Anderson & Lebiere, 1998). In a contrasting view interruption position and secondary task complexity are predicted to have marked effects on subsequent performance efficiency. For instance, the form of the goal-activation model

seems to suggest a cost of interruption position in terms of how much cognitive effort is required to associatively prime a suspended goal with a retrieval cue (Altmann & Trafton, 2002). The rate of decay during secondary task performance may also negatively effect the representational strength of a suspended goal. The goal-activation model does however explicitly predict a negative effect of secondary task complexity on primary task performance. This is as conceptualised as an effect of interference; the number of items between goal suspension and attempted retrieval. The more items encoded during this period, the more retroactive interference suffered by representation of the suspended goal. However, the goal-activation model suggests that if a suspended goal is strongly associated with an available priming cue (both prior to and following interruption), the complexity associated with performing the secondary task may have little effect on retrieval performance.

The negative effects on performance of interruption position as a function of processing complexity are also endorsed in related research. Performance deficits are found in the PM literature: When an intervening task requires complex processing (e.g., Einstein et al., 2003; Smith, 2003); when the number of intervening items are increased (e.g., Nakajima & Sato, 1989; Posner & Konick, 1986); and when priming cues are unavailable (e.g., Einstein et al., 2003). Similarly, the TS literature reveals larger switch costs: When tasks are not well practiced (e.g., Garcia-Ogueta, 1993); if switching between complex arithmetic tasks such as multiplying and dividing (e.g., Rubenstein, Meyer & Evans, 2001), and when inadequate switch preparation time is given (e.g., De Jong & Monsell, 2002; Rogers & Monsell, 1995). However, neither the PM or TS literature provides an adequate explanation of what happens to the representation of an interrupted goal that has received substantial encoding prior to suspension.

The literature supports that interruption position as a function of the processing complexity required to complete a primary task, largely contributes to subsequent primary task performance deficits. This effect is even found in tasks that can vary in terms of their structure but allow for sophisticated measures of cognitive demands during performance (e.g., Iqbal & Bailey, 2005; Iqbal, Adamcyzk, Zheng & Bailey, 2005). For instance, post-interruption task completion times are longer following complex multi-component office based secondary tasks (e.g., searching and sorting information)

compared to those that are simpler (e.g., reading and registering information) (e.g., Bailey et al., 2000a). This effect has been supported by Bailey et al. (2000b) with complex tasks (e.g., adding and counting) taking longer to complete than simple tasks (e.g., reading and selecting). Similarly, longer completion times are found if interruptions are positioned within a primary task containing no fixed sequence of actions compared to those that are practiced in a fixed sequence (e.g., Edwards & Gronlund, 1998). With greater control over interruption timing, Hess and Detweiler (1994) found larger primary task completion times following interruptions positioned at points requiring high processing demands (i.e., during performance of a goal sequence) compared to those requiring lower processing demands (i.e., during performance of a single goal). Most recently, Monk et al. (2004) found that resumption lags were longer for interruptions positioned in the middle of subgoals compared to those at the beginning and end (see also Botvinick & Bylsma, 2005, in press). In all, results seem to suggest a cost associated with maintaining a representation of a suspended goal. This may reflect the difficulty in keeping the goal active when it is part of a complex task suggesting that the goal is forgotten or the task in which it was embedded needs to be re-established (e.g., Adams et al., 1995).

The literature also suggests that the processing demands associated with performing the secondary task may also contribute to the negative effects caused by being interrupted. Longer completion times were found for a shopping based retrieval primary task interrupted by a task requiring the coding of letters into numbers before performing a calculation compared to standard arithmetic tasks using only numbers (Gillie & Broadbent, 1989). Similarly, looking up a telephone number during a text editing task resulted in shorter time to resume a text editing task than after having to perform an additional short editing task (Zijlstra et al., 1999). It has been suggested that such effects occur when the processing demands of primary and interrupting tasks are low due to decreased workload or because they are highly practised. Such an account would support the findings of a few studies that suggest that routine interruptions might actually improve subsequent primary task performance (e.g., O'Connail & Frohlich, 1995; Zijlstra et al., 1999). Although it should be acknowledged that one study actually found no difference between performance on a text-editing task following arguably simple 'conversational' interruptions and more complex logical reasoning tasks.

Few of the reported studies are supported by existing theory, and thus experiments presented within this thesis attempt to address this issue. Specifically, I aim to determine the level of disruption caused by different attributes of interruption position within a primary task (e.g., overall task complexity, goal/subgoal complexity, workload and so on), as well as the influence of the processing demands imposed by performing the secondary task (e.g., complexity and similarity). Thus, I now turn to a review of an attribute that has created much controversy regarding the disruptive nature of secondary interrupting tasks; their similarity to primary tasks.

1.5.2.2. Similarity between primary and secondary tasks

The relationship between two tasks, in terms of their processing similarity, may effect the performance on one or both tasks. In the case of being interrupted, studies have provided mixed evidence regarding the consequences of dealing with similar versus dissimilar interruptions. Interrupting a primary task with a secondary task requiring similar processing resources generally causes sizeable performance deficits (e.g., Czerwinski et al., 1991a; Gillie & Broadbent, 1989), although some studies have revealed weaker effects (e.g., Edwards & Gronlund, 1998; Latorella, 1996). Nevertheless, the general theoretical perspective regarding the effects of task similarity on primary task performance following completion of a secondary task is largely mixed.

On one view, performing a secondary task requiring similar processing resources to a suspended primary task could be thought to result in less disruption to post-interruption performance than one that is dissimilar. ACT-R models (e.g., Anderson & Lebiere, 1998) support this view, and generally predict a beneficial effect of task similarity on memory retrieval performance when processing two or more tasks. This is derived from the premise that activation within WM is in limited supply, and has to be spread amongst all memory items. Thus, processing new items – as for a dissimilar secondary task – is likely to overload WM capacity; consequently increasing the likelihood that an older suspended item(s) will be displaced. This account is similar to that of Adams et al. (1995) who suggest that similar secondary task items (those not directly interfering with the representation of the suspended item) will be supported by knowledge-structures already active in implicit memory. Thus returning to a primary

task following completion of a similar secondary task will not require as much cognitive task-set reconfiguration as is predicted to be required following a dissimilar. Finally, the TS literature provides extensive support for smaller time switch costs between two tasks that are similar compared to two tasks that are dissimilar (e.g., Arrington, Altmann & Carr, 2003). So, there is theoretical support for the beneficial effects of task similarity, although much evidence is generated from paradigms where tasks are performed in alternation rather than when they are interleaved (as is the case when being interrupted).

In a contrasting view, a secondary task requiring similar processing resources is proposed to interfere with the representation of a suspended goal (e.g., Edwards & Gronlund, 1998). The goal-activation model (Altmann & Trafton, 2002) supports this contention, whilst assuming similar WM storage restrictions to ACT-R. The goal-activation model would predict that the strength of a suspended goal's representation would be threatened by any prospectively encoded similar goals. This might occur for two reasons. Firstly, two previously processed goals may reside in close representational proximity of one another, given a limited WM activation supply. Thus, the difficulty in retrieving one of the goals via reactivation may be more difficult because of the increased chances of retrieving its similar counterpart. Secondly, if a suspended goal is successfully associated with a priming cue upon its suspension, encountering a similar cue in the secondary task may result in premature reactivation of the goal. Consequently, the cue may not have the same associative strength when returning to primary task, due to the fact that its purpose will seem to have been satisfied. The Search of Associative Memory (SAM) model seems to offer a similar account (Gillund & Shiffrin, 1984; Raaijmakers & Shiffrin, 1980). The model conceptualises memory as containing items that are linked to one another based upon associative strengths of connections. Two items that share similar attributes are expected to occupy the same representational space, thus making retrieval of either much more difficult than if they were not highly associated. Additionally, the SAM predicts that the more items competing for the same representational space, the higher the competition will be during attempted retrieval. So conversely, there is also much support for the negative effects of task similarity, especially when considering the costs of storing memory representations that can potentially compete with one another during retrieval or recognition attempts.

The interruption literature, although largely lacking in theoretical grounding, reveals larger primary task performance deficits following similar secondary tasks. Using their shopping based retrieval task, Gillie and Broadbent (1989) found longer primary task completion times using a secondary task which they argued was highly similar, and involved verbalisation of a list of items followed by a surprise free recall test. However, their interruption manipulations have since been scrutinised for two reasons. Firstly, Latorella (1996) poses that the similarity overlap between the primary and secondary task (i.e., both required similar processing resources) was insufficient to conclude *just* an 'interruption-similarity effect'. She suggests that tasks have to be closely matched on three factors to determine the level of similarity overlap: Processing resources (e.g., visuo-spatial versus verbal processing); the form of the information (e.g., searching versus sorting); and the semantic content (e.g., word relatedness). Indeed, related research has found only small performance deficits for tasks sharing similar processing resource requirements, with larger effects for those also sharing similar information form (e.g., Hess & Detweiler, 1994; Detwieler et al., 1994), and semantic similarity (e.g., Czerwinski et al., 1991a). Secondly, Gillie and Broadbent, also found that secondary task complexity was a strong predictor of primary task performance deficits; a potential confound of their similar secondary task (e.g., Latorella, 1996; McFarlane, 2002). This factor was partially used to explain the null effects of interruption similarity in Latorella's (1996) comprehensive study of task interruption. She argued that reducing secondary task complexity (e.g., allowing the task to be performed using rote memory), negates the typical effects of similarity.

The levels of similarity overlap between tasks as well as their complexity have been regarded as important factors for predicting performance deficits following an interruption episode. Using the same primary task used by Gillie and Broadbent (1989), Edwards and Gronlund (1998) investigated the effects of interrupting the retrieval of, a series of sequentially learned items (the fixed group) and compared this to a series of items with no such order information (the permuted group). Secondary tasks involved performing actions that were either: Completely similar (organising to-be-retrieved items within the shopping task); partially similar (unscrambling anagrams of words related to the to-be-retrieved items); or dissimilar (mental arithmetic problems) to those in the

primary task. Both completely similar and partially similar secondary tasks resulted in longer primary task completion times, but only for the permuted group. Thus, reducing the complexity associated with performing the primary task (i.e., by increasing associations between items) appeared to abolish the negative effects of performing similar secondary tasks. Finally, interruption-similarity effects can be attenuated by simply informing participants that they will have to recall what they were doing before being interrupted (Czerwinski, Chrisman & Schumacher, 1991).

In sum, the literature seems to suggest that task similarity can affect subsequent performance when returning to an interrupted primary task. However, this effect may be mediated by the difficulty in maintaining a representation of an interrupted item in the case of a similar interruption. If task demands are low (primary, secondary or both), then typical similarity effects are weakened or in some cases completely abated. I must again stress however, that the theoretical support for such effects is limited, and thus a number of experiments presented within this thesis will address this issue further. It is expected that increased task similarity will result in higher post-interruption performance deficits, but the size of these decrements will be affected by the complexity associated with performing the suspended goal and/or the secondary task.

1.5.2.3. Rehearsal and the interruption lag

Rehearsal is a pervasive strategy for maintaining information in the short-term (e.g., Atkinson & Shiffrin, 1968) and is commonly regarded as a prerequisite to encode the representation of an interrupted goal (e.g., Altmann & Trafton, 2004; Trafton et al., 2003). That rehearsal may mitigate the negative effects caused by interruption is not a new phenomenon (e.g., Gillie & Broadbent, 1989; Storch, 1992), although interpretation of what is being done at rehearsal should sometimes be met with caution (e.g., McFarlane, 2002). The experimental method used by Czerwinski et al. (1991a), where 30-seconds rehearsal time before interruption did not attenuate performance deficits, would certainly conform to such a claim. However, in a follow up study, where participants were explicitly told that they would be tested on interrupted information, allowing rehearsal led to an improvement in performance, although this effect did not quite reach significance (Czerwinski et al., 1991b). Perhaps only through encouragement

of strategy deployment might rehearsal to lead to even moderately faster resumption times (e.g., Detweiler et al., 1994). When faced with the opportunity to rehearse, people might instead adapt to task constraints to minimise metal load at the point of interruption, by relying on externalisation to support the recovery of interrupted state information (e.g., Anderson & Douglass, 2001). For instance, participants can adapt to preparation periods by choosing to use memory aids instead (e.g., pointing to interrupted information, Miller, 2002).

There is controversy regarding the performance utility of being able to prepare for an inevitable interruption episode. In one view, pre-interruption preparation is regarded as essential for improving post-interruption resumption performance (Altmann & Trafton, 2002). For instance, the goal-activation model poses that an interruption lag is required to strengthen the representation of a suspended goal as well as to successfully associate it with an appropriate retrieval cue. In another view, the utility of preparation is regarded as extremely low, given that people may actively choose not to encode a goal that cannot be executed for some time (Anderson & Douglass, 2001).

The interruption lag may seem a plausible solution to reducing post-interruption resumption costs, but experimental evidence for this effect is weak. Using a resource allocation primary task, Trafton et al. (2003) found that an 8-second interruption lag (cued by a flashing visual stimulus) resulted in an overall reduction in the length of resumption lags compared to an immediate interruption condition. They rationalised that resumption performance was better in the interruption lag group because the suspended goal had been strengthened to a greater extent. Additionally, verbal protocols revealed that participants preferred to use this time to prospectively encode a future intention rather than to retrospectively rehearse the current state. However, upon closer inspection of the data, the interruption lag was considerably longer than the reduction in the length of resumption lags. This was reflected across three testing sessions with resumption lag reductions of approximately, 3.5 seconds (session 1), 1.1 seconds (session 2) and 0.4 seconds (session 3). Resumption lags in session 3 did not significantly differ between group, suggesting that practice within the primary task as well as at dealing with interruptions, negates any utility of an interruption lag. Recently, Altmann & Trafton (2004) found reduction in resumption lag length for interruption lags up to 4 seconds in

duration, with only marginal reductions for those between 6 and 8 seconds. For instance, an 8-second interruption lag reduced resumption lag length by approximately 1 second, thus incurring a cost of 7 seconds to overall task performance. This effect was partially attributed to the complexity of the primary task used, in terms of the difficulty in locating suitable retrieval cues. The evidence accumulated from both studies has prompted the suggestion that interruption resumption costs may have been underestimated in the goal-activation model (Altmann & Trafton, 2002).

In similar studies, the cost-benefit ratio of an interruption lag to primary task resumption performance has also been scrutinised. In one study, Detweiler et al. (1994) found that memory performance for a mathematical equation was slightly improved for a complex secondary task preceded by a 3-second interruption lag. This effect was not found for a simple secondary task, perhaps because rehearsal could occur during performance of an easier task that requires fewer processing resources for completion. Using the Tower of London as a primary task, Hodgetts & Jones (2003) revealed a 6.05 second resumption lag for a promptly initiated interruption requiring a verbal reasoning response. This was reduced to only 3.77 seconds if the same secondary task was preceded by a 3-second interruption lag. In fact, one study revealed larger performance impairments due to the inclusion of an interruption lag compared to when not including an interruption lag (Miller, 2002).

Even within related areas of research, the utility of preparation is not clear-cut. In the TS literature, the benefits of preparing to switch between tasks certainly outweigh the costs of no preparation (e.g., Garcia-Ogueta, 1993). Indeed, a number of TS studies have supported that switch costs can be reduced if preparation time – a *response cue interval* – is allowed before the switch occurs (e.g., Meiran, 1996; Meiran, Chorev & Sapir, 2000). However, even with preparation, a SC (although reduced) still ensues, and with too much preparation, typical SCs can in fact be increased (e.g., Koch, 2001). Conceptually, it seems plausible that preparation allows a TSR type process to occur before switching (e.g., Rogers & Monsell, 1995), but this comes at an additional cost to overall performance (i.e., the time taken to prepare). In the PM literature, manipulating the time allocated to encode a prospective intention has delivered similar results to manipulating the length of the interruption lag. For example, Einstein et al (2003), found no

improvement on prospective remembering if participants were given 6 seconds to encode an intention rather than 2 seconds. Like Anderson & Douglass (2001), this may reflect a strategy of rationalising the benefits of deploying resources to encode a goal to be performed in the future, when environmental cues alone may support its re-establishment.

In sum, it seems that the interruption lag may not be such a good 'window of opportunity' for reducing interrupted task resumption lags as was originally suggested by Altmann & Trafton (2002, p. 65). Experimental work presented in this thesis aims to build upon the claim by Altmann & Trafton (2002) that their original model may have underestimated the cost of resuming an interrupted task. I propose that an in-depth investigation of the goal structure being interrupted may reveal that there are benefits of an interruption lag, but in restricted contexts (e.g., interruption of a single goal structure). Additionally, if the general cost-benefit trade-off of an interruption lag to resumption speed is negative, then what other methods can be incorporated to alleviate interruption induced performance deficits? Experimental work within this thesis explores alternative methods (e.g., enforced planning, and training) that may provide answers to such a question.

1.5.2.4. Retrieval aids

In everyday life, we rely on the use of retrieval aids to remind of things that we otherwise may forgot if using memory only. Typical examples include, a shopping list, the use of a diary, writing on hands, and 'post-it' notes. Indeed, Anderson and Douglass (2001) provide evidence that external memory aids may replace the need to rely on a limited capacity STM store. However, such memory aids may be effective when sufficient time is allocated to prepare for a future task, but may not be viable when time restrictions are high; such as in the case of unanticipated interruption. In such cases, it has been suggested that retrieval aids can still be used, but these may be more subtle attributes of stimuli (e.g., the position of a mouse cursor).

Certainly, it has been conceptualised that implicit priming may lead to better post-interruption performance (e.g., Altmann & Trafton, 2002), although supporting experimental evidence is extremely limited. For example, van Dantzich, Robbins, Horvitz and Czerwinski (2002) developed an application called 'Scope' that filters the

availability of important salient information during multiple interruption intrusions. They suggested that limiting the amount of information available before switching to a secondary task might increase the likelihood that people will choose to use that information following interruption. Nevertheless, further research is needed to investigate the utility to of implicitly priming a suspended goal before interruption.

In the case of interruption, it may be that explicit priming is the only effective way of ensuring that a suspended goal has been associated with a priming cue. This is a key construct of the goal-activation model. Specifically, the model suggests that a suspended goal must be associated with a priming cue (externally or internally generated) immediately before switching to an interrupting task, and that the same cue has to be available at resumption. To avoid retrieval problems, this cue has to be associated the suspended goal only, and none (or very few) others. The cue, it is argued (Altmann & Trafton, 2002), will provide a rapid reactivation boost to a suspended goal that has decayed below the interference threshold. However, given the dearth of experimentation concerning the role of retrieval aids before goal suspension, it might be useful to review experiments that have manipulated cues at goal resumption.

One line of research has shown that post-interruption retrieval aids support faster resumption and more efficient primary task reorientation when made. In one study, Kreifeldt and McCarthy (1981) demonstrated that one calculator design that displayed pre-interruption actions supported better reorientation performance than one that did not. Similarly, Field (1987) found that participants were able to resume a primary task faster if they were only shown the specific section of information they had been working on previously rather than the whole screen. An experiment by Harrison, Ishii, Vicente and Buxton (1995) revealed that when performing two or more tasks simultaneously, the availability of visual markers supported the re-orientation from one task to another. Being able to review goals executed before interruption has also been shown to support resumption performance (e.g., Cypher, 1986). Providing participants with a log of pre-interruption activities, for example, has been shown to increase reorientation performance. Recently, Franke, Daniels & McFarlane (2002) provided evidence for the effectiveness of a 'read back' application for marines which provided either a log of pre-interruption activities or just the suspended item only. Finally, in a recent PM study,

McDaniel et al. (2004) eliminated the cost of interrupting the execution of a delayed intention by simply presenting a red dot on the screen following interruption completion. So, retrieval aids appear to effectively produce speedier reorientation when returning to an interrupted task.

Another line of research has shown that retrieval aids may not be effective at improving resumption performance following an interruption episode. For instance, Czerwinski et al. (2000b) demonstrated that a mouse cursor covering interrupted information only led to faster resumption when the associated action involved a single non-semantic response. When response involved multiple semantic responses, the mouse cursor had no effect on attenuating resumption deficits. This effect was subsequently replicated, although a more thorough data analysis revealed that participants choose to use the marker early on in experimental trials, but this also had no effect on resumption performance (Czerwinski et al., 2001). Thus, retrieval aids might not support the reactivation of complex information sequences. This may be because of insufficient encoding of the information in the first place as well as a rationalisation process of deciding cost-benefit ratio of using such cues.

In sum, retrieval aids can support more efficient performance when alternating between tasks, but their support as effective memory aids in interruption research is limited (e.g., Czerwinski et al., 2000b, 2001). This effect seems to be restricted to contexts where the interruption is imposing upon the execution of single rather than multiple actions. Before drawing specific conclusions as to how this process occurs (e.g., memory retrieval or task reconstruction), there is a need for further investigation. Empirical work presented within this thesis will assess the effectiveness of retrieval aids (e.g., availability, salience), both before and following an interruption episode.

1.5.2.5. Practice and expertise

Practising a task can lead its automation (e.g., tying ones shoelaces). Automated processes require fewer attentional resources because sufficient task knowledge is already stored in memory (e.g., most of us can tie our shoelaces without looking at what we are doing). This phenomenon can be explained by considering the Automatic Attention Response (AAR, e.g., Schneider & Shiffrin, 1977). The AAR assumes that mental tags

are attached to familiar high priority stimuli, which can lead to significant reductions in response and retrieval times (e.g., Schneider & Detweiler, 1988). Additionally, training within a task has been shown to lead to associative learning between a memory representation in LTM and a reliable priming cue (e.g., Fisk & Rogers, 1991).

An important question is whether or not task practice will interact with interruption recovery performance? For multi-tasking at least, performance is improved with practice. People have been shown to concurrently: Read and write (e.g., Hirst, Spelke, Reaves, Caharack & Neisser, 1980); shadow speech and play the piano (e.g., Allport, Antonis & Reynolds, 1972); and type descriptions of visual material whilst reciting nursery rhymes (e.g., Shaffer, 1975).

ACT-R's mechanism for selecting the best production to service a goal after experience in the task domain is *expected gain* (e.g., Anderson & Lebiere, 1998). Expected gain (*E*) is a measure of how quickly a production (e.g., a strategy to encode a goal) will be used based upon its history of predicting the outcome of a goal (e.g., retrieval). As such, with experience of a particular production leading to a successful goal outcome, certain preparation strategies will be used more often than others based upon expected gain. The equation for expected gain in ACT-R is:

$$\mathbf{E = PG - C \pm noise}$$

(Equation 3)

P is probability, and represents how often a production has successfully[3] led to an outcome in the past (e.g., encoding has led to retrieval) and is a determined by:

Number of eventual successes/(Number of eventual successes + Number of eventual failures)

(Equation 4)

[3] ACT-R allows for a retrieval attempt to be made more than once before attempted goal retrieval is abandoned (see Section 3.2.5 of Chapter 3). As such, a success might evolve from more than one production firing causing the eventual retrieval of a goal, and a failure is not only when a single production firing has not led to the retrieval of the goal.

The higher the value of P (e.g., the more retrieval successes than failures), perhaps from experience of preparation strategy leading to eventual goal retrieval, the more likely the production will be chosen in the future. **G** is a global variable that represents how much time the system is willing to dedicate to the processing of the goal (e.g., maintaining an active representation of the goal), and has a default value of 20 s before production firing will be abandoned. However, a recent investigation of problem solving using the ToH task predicts that this value may be more in the region of 20 s – 50 s to accommodate the time need to maintain goals for long time periods (Gunzelmann & Anderson, 2001). **C** is the anticipated time cost of achieving the goal using the production, and is based upon history of how long the production has taken to retrieve the goal in the past. Finally, **noise** – slight variations in the time taken for a production to fire because of unexplained variance (e.g., extraneous influences on strategy control) – provides some degree of variation to the expected gain (Gray et al., 2004). Too little noise might result in settling on a strategy that best suits a high P and low C early on in problem solving, whereas too much noise might result in no specific strategy ever being settled upon.

In the case of task interruption research, Edwards and Gronlund (1998) documented that a secondary task requiring similar processing resources to that of a primary task were disruptive to performance; but this effect was abolished if the primary task had already been sequentially learned. In contrast, some have shown that this effect only occurs if participants are experienced at dealing with immediate interruption episodes in a specific primary task (Hess & Detweiler, 1994; Trafton et al., 2003). Recently, the constraints proposed by the LTWM model (Ericsson & Kintsch, 1995) were empirically tested using an interruption paradigm in a HCI context (Oulasvirta & Saariluoma, 2004). The authors posited that, "Task interruption provides a practical HCI problem in which LTWM is involved" (p. 54). Based on the premise that LTWM cannot be easily disrupted by having to perform a secondary task – given the availability of a mnemonic pointer – the authors supported that expertise may lead to an attenuation of interruption effects. However, disrupting the encoding of a to-be-interrupted goal significantly reduces the effectiveness of this safeguarding measure. Specifically, performance deficits were caused by, abrupt interruption episodes, high task processing

demands during encoding, and if there was high semantic relatedness between tasks. Consequentially, unhindered encoding of a to-be-suspended goal may be necessary to bind its short-term WM representation to a LTWM representation. If this encoding is disallowed, partially suppressed or affected by cue selection competition, LTWM representations are difficult to access following interruption.

In sum, task expertise together with an environment rich in mnemonic cues (both before and following interruption) might result in the dissipation of performance deficits caused by being interrupted. If one is sufficiently practiced at a task, encoding may sufficiently strengthen the association between a suspended goal and its LTWM cue. However, such encoding may not have to occur immediately before being interrupted to effectively prime a LTWM representation (which is in contrast to the goal-activation model), although its associated cue must be available upon primary task re-instatement (which is in agreement with the goal-activation model). For efficient retrieval, this cue should not be masked by competing stimuli. Experiments presented in this thesis will explore the relationship between encoding and the availability of retrieval cues for both minimally practiced and highly practiced tasks. It is expected that practice at a task may lead more efficient performance, but practice at dealing with interruption may have a larger impact on reducing the time taken to resume an interrupted task. Finally, anticipating when an interruption might occur from experience may facilitate the amount of attention allocated to remembering a suspended item as well as the use of mnemonic cues (e.g., Nagata, 2003).

1.6. A paradigm for an experimental analysis of task interruption

1.6.1. The Tower of Hanoi (ToH) task

The goal-activation model effectively demonstrated the plausibility of utilising the ToH task for investigating the effects of interruption on goal-directed behaviour (Altmann & Trafton, 2002). Given its wealth of empirical support as a task that encourages the encoding, hierarchical organisation and deployment of a series of planned goals it was reasoned that it would form an excellent platform for studying the effects of interruption in goal-directed behaviour.

Figure 1.3 illustrates an example of a three-peg four-disk ToH problem. In this traditional 'pyramid' example, there are four discs, varying in size located on one peg of a possible three. The overall or *master* goal is to move all the discs to the third peg whilst maintaining the pyramid stack shape. However, when solving the problem, only one disc can be moved at a time (that is not blocked by another), and a larger disc can never be placed on top of one that is smaller.

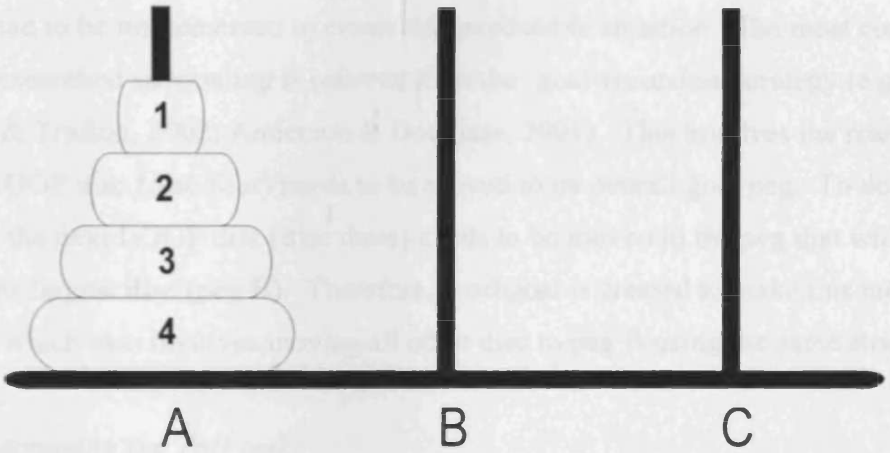

Figure 1.3. An example of a four-disc Tower of Hanoi problem. Discs one - four must be moved to peg C within fifteen-moves to achieve error-free completion

The ToH (and similar variants, e.g., the ToL) is an example of a well-defined problem solving task where operators (i.e., making moves) are applied in a logical way to achieve the goal-state. The ToH is a well-defined problem solving task because it presents a 'problem space' in which participants have to deploy operators in a specific sequence to achieve a goal-state whilst aiming not to make any move mistakes or straying off course from an error-free solution sequence (e.g., Hayes, 1978). A problem space refers to the possible move sequences given the problem context and the information-processing capacity of the solver (e.g., Newell & Simon, 1972).

Novices tend to initially approach ToH problems using a problem-solving strategy often referred to as 'hillclimbing' or 'selective search' whereby moves are executed in an exploratory fashion (e.g., Anzai & Simon, 1979), until experience leads to development and deployment of more effective strategies. With limited practice, participants learn to

encode, store and deploy a number of deliberately planned goals across a solution-sequence by setting achievable subgoals (e.g., Altmann & Trafton, 2002; Anderson et al., 1993; Anderson & Douglass, 2001; Davies, 2000, 2003; Simon, 1975). Anzai & Simon (1979) provided a detail protocol of strategy development within the ToH task, and demonstrated that a participant used four strategies (including goal search) across only four ToH problems. Each of the last three involved the realisation that the LOOP disc needed to be moved to its ultimate goal peg, and therefore some kind of unblocking strategy had to be implemented to create this executable situation. The most common form of researched subgoaling is referred to as the 'goal-recursion' strategy (e.g., Altmann & Trafton, 2002; Anderson & Douglass, 2001). This involves the realisation that the LOOP disc (disc four) needs to be moved to its overall goal peg. To do this however, the next LOOP disc (disc three) needs to be moved to the peg that will not contain the largest disc (peg B). Therefore, a subgoal is created to make this move possible, which also involves moving all other disc to peg B using the same strategy.

1.6.2. Interrupting the ToH task

The ToH has recently been reviewed as a task that may support the study of task interruption (Altmann & Trafton, 2002). Uninterrupted performance data unearths differences between move latencies, likely to reflect task demand variations at different positions within the solution-sequence (e.g., Anderson & Douglass, 2001; Anderson & Lebiere, 1998). For example, Anderson et al. (1993) identified three points within perfectly solved fifteen-move ToH problems where average inter-move latencies are exceptionally high. These are before making the first (9.7 s), ninth (6.6 s) and thirteenth (5.4 s) moves. These latencies are believed to reduce in a linear fashion because each subgoal reflects an opportunity to plan a new subgoal; each containing fewer moves to process. This may reflect behavioural adaptation to the difficulty in planning move-sequences in complex problems (e.g., Anderson & Douglass, 2001; Davies, 2000, 2003), and/or a tendency to stray from an initial plan (e.g., Gilhooly, Phillips, Wynn, Logie & Della Sala, 1999). Different task states require deployment of different levels of processing resources. Thus, interruptions are likely to affect performance within the ToH

task *per se*, but this level of disruption caused may interact with demands imposed by the task at the point of interruption (e.g., Altmann & Trafton, 2002).

The task importantly allows methodological flexibility for interruption research purposes. Firstly, problems can be decomposed into controllable subgoals that require different levels of processing for their execution. Therefore, experimental control over interruption timing relative to performance on the interrupted task can be achieved (e.g., Altmann & Trafton, 2002). Secondly, the task can be manipulated in terms of processing complexity. This can be achieved by varying the number of moves required for error-free completion, usually by manipulating the number of discs to be moved within the stimulus-array (e.g., Anderson & Douglass, 2002; Anzai & Simon, 1979). Thirdly, adequate instruction and practice can cause participants to adopt different strategies when solving problems, or to develop existing strategies during a test phase (e.g., Anzai & Simon, 1979; Fum & Del-Missier, 2001). Performance differences are found from implementing even the simplest instruction such as informing participants of the number of moves required to solve problems error-free (e.g., Davies, 2003). Fourthly, the task can be manipulated in terms of visual presentation. Attributes such as colour, size, shape and position can be varied which may be potentially interesting for exploratory purposes. Finally, the task can be manipulated in terms of the amount of planning participants are allowed before attempted to solve problems. This can be achieved by varying the time allocated before solving problems and by administering guidelines to ensure utilisation of this time (e.g., Anderson & Douglas, 2002).

Complex problem solving behaviour seems inherently viable to opportunistic planning (e.g., Byrne, 1977; Hayes-Roth & Hayes-Roth, 1979; Patalano & Seifert, 1997). However, research has shown reliable beneficial effects to performance from forming an initial plan to well-structured problem solving tasks (e.g., Anzai & Simon, 1979; Davies, 2000, 2003; VanLehn, 1991). It has been suggested however, that such benefits may only be evidenced when utilising moderately complex problem solving tasks (e.g., Phillips, Wynn, McPherson & Gilhooly, 2001; Davies, 2003; Ward & Allport, 1997). Thus it is also likely that the level of pre-problem execution planning and opportunistic online planning, will affect the activation level of an interrupted goal.

Some argue that little or no higher-order memory resources are required to perform the ToH task (e.g., Anderson & Douglass, 2001) especially at non-subgoal creation points within the solution-sequence. Although others maintain that the task requires memory at more than just subgoaling opportunities (e.g., Altmann & Trafton, 1999; Davies, 2003) Interruption might however have the consequence of making the ToH task one that requires memory (even if only during the event of interruption) when one asks 'what was I doing?' (Altmann & Trafton, 2002).

Given the flexibility of the ToH task and the theoretical support for its use as a tool for investigating the effects of interruptions, it will form the primary task in all following empirical sections.

1.6.3. Implications for interrupting the ToH task

In their model, Altmann and Trafton posit that an interruption will affect ToH primary task performance in a number of ways. For novice users, the retrieval of a suspended goal should be highly dependent on how effectively it is encoded during an *interruption lag* – a transitory period between the interruption alert and secondary task initiation. This brief opportunity might be utilised to associate the to-be-suspended goal with a cue located within the primary task's stimulus array. Successfully locating this cue following primary task reinstatement will result in a process of 'rapid reactivation' of the suspended goal, thus bringing it above the interference threshold; resulting in likely sampling during a retrieval cycle. Implications for an unexpected task interruption requiring prompt initiation is that a suspended goal may not be adequately associated with a priming cue, and thus it *cannot* be subsequently retrieved from memory; forcing reconstruction and/or re-planning. For expert users, associative priming is believed to be an automatic process because of declarative knowledge of the cues within the solution-sequence; largely the chunking of such associations in memory. A production rule linked to a suspended goal should fire upon attendance of a previously encoded priming cue, serving to modify its current state – through retrieval and consequent updating of the goal's base-level activation (e.g., Anderson & Lebiere, 1998).

In their goal-activation model, Altmann and Trafton (2002) make firm behavioural assertions regarding the necessity of associative priming and cue selection in

the ToH task. First, if a goal is not encoded through associated priming, the goal will not be retrieved. Secondly, they propose that the likely cues within the ToH task will be the discs themselves. Specifically, "...the presence or absence, or relative availability, of disks as cues should affect performance by affecting the ability to retrieve goals" (Altmann & Trafton, 2002, p. 53).

Chapter 2

EMPIRICAL SERIES 1

Recovering from un-signalled interruption: A closer look at the point of interruption and secondary task demands

---

2.1. Overview

With conceptual guidance from recent ACT-R based models of goal-directed behaviour (e.g., Altmann & Trafton, 2002; Anderson & Douglass, 2001), the current experiments explored the impact of un-signalled interruption on performance within 4-disc ToH problems. Interruption caused significantly longer resumption latencies than comparable move latencies. Resumption latencies were longer and redundant moves increased in number the earlier interruption fell within solution-sequences (Experiment 1). Resumption latencies were shorter for ToH problems that required fewer moves for completion, with no additional impact of increasing the complexity of secondary tasks (Experiment 2). Resumption latency was however longer when interruption fell within a ToH subtask than when it fell at the end of a subtask, and was inflated further when secondary tasks were similar to primary tasks (Experiment 3). In both Experiments 2 and 3, redundant moves were fewer in number following an interruption. Results suggest that promptly suspended goals are not forgotten. Strategic encoding occurs prior to interruption leading to more timely resumption, and is augmented with experience in the task domain.

2.2. Introduction

2.2.1. General Introduction

Interrupting the execution phase of a primary task often leads to performance degradation in the form of longer latencies to execute the first action within that task and/or subsequent mistakes (see McFarlane, 2002, for a comprehensive review). Such effects are often loosely attributed to breaking the cognitive focus of a primary task, although factors pertaining to secondary tasks such as similarity to the primary task often receive attribution of cause (e.g., Burmistrov & Leonova, 2001, Gillie & Broadbent, 1989; Lahlou et al., 2000). In fact, most of the existing literature lacks grounded theoretical support regarding the processes involved in resuming a task that is suspended because of interruption. However, recent ACT-R based models, that view goal suspension and resumption as a costly process (Altmann & Trafton, 1999; 2002; Anderson & Douglass, 2001), have cleared the way for a more conceptually-driven investigation of the impact of being interrupted (e.g., Altmann & Trafton, 2004; Hodgetts & Jones, 2003; Trafton et al., 2003). Within this chapter, I will use these models as a conceptual framework for investigating task interruption.

The ToH task has been used widely to study goal-directed behaviour (e.g., Altmann & Trafton, 1999, 2002; Anderson & Douglass, 2001; Anzai & Simon, 1979; VanLehn, 1991). Due to its strong conceptual base, the ToH will be used as the primary task within the current experiments to explore its susceptibility to interruption induced performance deficits, and to investigate the effects of suspending and resuming different phases of its execution.

Within the current experimental series, interruption was not signalled, and various attributes of secondary tasks were manipulated. This allowed an investigation of the involvement of internal processes (e.g., encoding, storage and retrieval) and external factors (e.g., priming cues) that might support the restoration of interrupted state information, especially for goal structures that vary in their difficulty to execute. Overlap in processing and content similarity between primary and secondary tasks was also manipulated to address the controversy regarding similarity-based interruption effects introduced in Chapter 1 (e.g., Gillie & Broadbent, 1989; Latorella, 1996). Additionally,

the number of component actions required to perform a secondary task was manipulated to speak to another controversy concerning the effects of secondary task complexity (e.g., Gillie & Broadbent, 1989; Zijlstra et al., 1999). From an interrupt-driven goal-directed behaviour perspective, this allowed for an exploration of the consequences on recovery performance of manipulating the retention interval between goal suspension and resumption (e.g., Altmann & Trafton, 2002; Anderson & Douglass, 2001).

Three broad experimental hypotheses are proposed each relating to the strategies that might support the resumption of goal structures that are promptly suspended because of interruption. Firstly, reactivation (and possibly retrieval) of a suspended goal might depend upon the strength of its representation both before and following interruption (e.g., Adams et al., 1995; Altmann & Trafton, 2002). Secondly, interruption might overload cognitive processing resources (e.g., Speier et al., 1997, 2003) and/or encourage abandonment of goal storage processes (e.g., Anderson & Douglass, 2001), forcing the recovery of state information from the display. Thirdly, there might need to be a careful integration between cognition and the reprocessing of interrupted state information (i.e., associative cues) in order to re-establish a suspended goal, perhaps through reconstruction of a suspended goal and/or reconfiguration of the problem space.

2.2.2. Task interruption and goal-directed memory

The archetypal view of memory for goals relies upon the goal stack concept (e.g., Miller et al., 1960), which has dominated much of the subsequent goal-directed behaviour literature (e.g., Anderson, 1983, 1993; Anderson & Lebiere, 1998; Newell, 1990). If interrupted, say by a task containing a single goal, the theory assumes that this goal would be 'pushed' onto the top of the hierarchical stack, thereby governing behaviour. Consequently, the interrupted goal would be pushed down the stack into second position, meaning it will not govern behaviour until the top goal is acted upon. Upon executing (or 'popping') the interrupting goal, the interrupted goal should immediately regain its top position on the stack, again governing behaviour. This process should be unaffected by factors that often mediate the retrieval efficiency of retrospective memory representations, such as, the inability to rehearse during retention intervals (Muter, 1980), temporal based decay (Burgess & Hitch, 1999), and activation

competition from other processed items (e.g., Keppel & Underwood, 1962; Bower, 1967).

The major assumption of the goal stack concept of a perfect goal memory has been subject to much recent criticism (e.g., Altmann & Trafton, 1999, 2002; Anderson & Douglass, 2001). For instance, there is empirical evidence of: Goal forgetting (e.g., Byrne & Bovair, 1997); differences in the number of executed actions required to achieve a goal state when trained on a task (e.g., Anderson et al., 1993; Anderson & Douglass, 2001); and goal selection errors (e.g., Anzai & Simon, 1979; VanLehn, 1991). The stack has also been regarded as conceptually inadequate for explaining performance deficits caused by interruption (Altmann & Trafton, 2002); a supposition that is receiving growing empirical support (Altmann & Trafton, 2004; Hodgetts & Jones, 2003; Monk et al., 2002, 2004, Trafton et al., 2003). Consequently, the stack in its current form seems to be ill-equipped to explain performance costs associated with interruption. It is due to these issues that I now revisit two ACT-R inspired models that were introduced in Chapter 1. They differ fundamentally from the goal stack by conceptualising goals in the same way as other declarative memory items (Altmann & Trafton, 2002, Anderson & Douglass, 2001).

Using empirical support from the ToH task, recent ACT-R based hybrid models proposed by Anderson and Douglass (2001) and Altmann and Trafton (2002) predict performance deficits because of interruption. Both assume that the retrieval of a suspended goal is largely influenced by its *base-level activation*: The limiting factors being the duration of its processing 'lifetime', the number of times it has been processed within that lifetime, and the time elapsed since it was processed last. However, the models differ in position regarding the processes that might underlie the resumption of a suspended goal. Together with an ACT-R model, Anderson and Douglass (2001) suggest that resumption of an interrupted goal that has been postponed for more than 18-seconds should largely involve a reconstruction process; not due to retrieval failure, but because of abandoning the storage processes prior to interruption. By contrast, Altmann & Trafton (2002) propose that there must be other strategies for recovering suspended goals, given that people frequently remember to resume postponed activities (see Byrne & Bovair, 1997, for a review). According to this view, a suspended goal should be

repeatedly strengthened for its activation level to exceed an interference level, and there must be a strong associative link formed between that goal and a priming cue. To retrieve the decaying representation of the suspended goal, the priming cue (e.g., a ToH disc) has to be encoded immediately before goal suspension (preferably during an encoding opportunity) and must be perceptually available after the reinstatement of the primary task.

Given that the main predictions within each model are by and large *ad hoc,* especially in the context of interruption, their parameters might be under-specified. One diagnostic issue for the model of Anderson and Douglass (2001), is the unpredictable nature of secondary task duration during an interruption, and as such, it might be slightly premature to assume that there are fixed strategies for goal suspension and resumption. Note that after interruption, the primary task is usually removed from perceptual awareness, even if only temporarily. Such perceptual starvation might encourage greater (or more efficient) deployment of cognitive resources for encoding actions (e.g., Ballard, Hayhoe & Pelz, 1995) that might otherwise be performed using less demanding perceptual-motor strategies (e.g., Zhang, 1997). Strategic adaptation to the task constraints (e.g., realising the futility of offload mnemonic demands to the visual array) might increase the efficiency of encoding and the deployment of strategies that may consequently affect the efficiency of goal re-establishment even without explicit opportunity to prepare (in contrast to Altmann & Trafton, 2002). Additionally, re-establishment of a suspended goal might not involve retrieval, reconstruction or reconfiguration, but rather a combination of all processes given task knowledge and constraints imposed by the task domain.

In the following two sections, I will consider how both models may inform an understanding of the effects of interrupting the ToH task. Specifically, I will focus upon how they cope with un-signalled interruption to goal structures that vary in their execution difficulty. Additionally, I will review the broader literature that shows differential effects caused by manipulating the point at which an interruption occurs, and the controversies regarding the effects of manipulating primary and secondary task similarity.

2.2.3. *The point of un-signalled interruption and secondary task complexity*

The previously reviewed ACT-R models propose quantitatively and qualitatively similar effects of un-signalled interruption on the performance of different stages of a goal-directed task. In their model, Anderson and Douglass (2001) suggest that an opportunity to prepare a suspended goal for retrieval will have no impact when secondary tasks are likely to be of a long duration. In their view, the time to resume a suspended goal in the case of un-signalled interruption should increase as a function of the duration of the retention interval and the difficulty associated with reconstructing the goal using the visual array. Similar effects are predicted by Altmann and Trafton (2002), but only if an interruption is not preceded by a brief period where the goal can be prepared for resumption through strengthening and associative priming. Without such a period, performance degradation will occur when attempting to re-establish the suspended goal, irrespective of the point of interruption. With opportunity to prepare, retrieval time should be the similar, as long as the goal has received sufficient associative priming with a cue that is available immediately before interruption and immediately following secondary task completion.

Empirical evidence reveals that the point of interruption within a primary task seems to affect subsequent performance within that task. Of the studies that have considered the structure of a primary task, most conclude that performance is degraded to a greater extent when interruption imposes upon a complex activity compared to one that is simpler (e.g., Cutrell et al., 2000, 2001; Speier et al., 2003). For instance, there are usually bigger performance deficits when interruption falls at a point within a task believed to require high processing demands, compared to a phase with lower processing demands (e.g., Detweiler et al., 1994; Hess & Detweiler, 1994). However, the definition of 'processing demands' at the point of interruption is often spurious. More often than not high processing demands are said to reflect being engaged within a larger task activity, where there are lots of co-dependent activities, with the opposite characteristics when processing demands are believed to be low (e.g., Bailey et al., 2001a, b; Cutrell et al., 2000, 2001). Nevertheless, some recent studies have begun to consider the point of interruption at more of a fine-grained level, investigating the impact of interruption at different points within the solution-sequences of tasks (e.g., Iqbal & Bailey, 2005; Monk

et al., 2002, 2004,). For example, Monk et al. (2002) found longer resumption latencies following interruption of information selection actions of a VCR programming task (e.g., selecting a recording time that requires the programming of a specific month, day, and hour) compared to interruption co-ordinated before a selection action. It is such fine-grained control of the intended point of interruption, I believe, will open the way for a more stringent analysis of interactions between task difficulty and interruption.

Obtaining a measure of the complexity of individual task components, to assess the impact of interruption at a task-goal level, might be achieved by using a well-structured problem solving task such as the ToH problem (e.g., Altmann & Trafton, 2002; Ormerod, 2005). The ToH task has a grounded history for which there is a great amount of empirical regularity in performance data, for measures such as inter-move latency (e.g., Anderson et al., 1993; Anderson & Lebiere, 1998; VanLehn, 1991). Furthermore, the task allows problems to be created that contain goal structures that are controlled in a number of different ways, and as so can be manipulated in different ways also. These include, problem difficulty (usually indexed by the number of moves required for error-free completion), interruption position (within or between main subgoals, smaller subtasks or two sequential moves), and the number and availability of perceptual priming cues (e.g., ToH discs, Altmann & Trafton, 2002). Thus within the current experimental series, a measure of the difficulty in executing an individual goal could be captured by implementing interruptions at points within ToH solution-sequences that call on different levels of processing demand. Non-interrupted control data, generated by problems with the same solution-sequences, and hence similar difficulty, could be compared to data generated following the implementation of an un-signalled secondary task in interrupted problems.

Guidance from the models of goal-directed behaviour sets reconstruction as the chief resumption strategy in the current experiments (e.g., Anderson & Douglass, 2001), on the grounds that the method of interruption initiation is not one that encourages associative priming of a suspended goal (e.g., Altmann & Trafton, 2002). Nevertheless, that a suspended goal might be encoded to some extent even when interruption is un-signalled should not be ruled out. Goals that are easier to re-establish (e.g., those with fewer co-dependent goals and stronger perceptually salient priming cues), might have

more chance of being encoded, perhaps leading to differences in resumption performance.

2.2.4. Task similarity

In their goal-activation model, Altmann and Trafton (2002) suggest that performing a similar secondary task might cause larger primary task resumption deficits compared to one that is dissimilar. Specifically, similar secondary task goals might require the utilisation of a limited associative activation supply; some of which may be dedicated to maintaining an associative link between memory and the task context. In their model, Anderson and Douglass (2001) are somewhat mute about how task similarity may affect the performance efficiency of a suspended goal. ACT-R theory acknowledges that a goal's activation might be affected by the amount of associative activation required to perform the secondary task (supporting the idea of a limited activation supply, e.g., Lovett et al., 1999). However, Anderson and Douglass (2001) do not regard this as a consequence of the similarity between secondary task goals and the goal that is suspended; instead they attribute the proposed effect to the projected length of the retention interval. Elsewhere, it has even been suggested that a dissimilar secondary task will require a greater redeployment of processing resources: Given similar processing resources, a similar secondary task may impinge less on the representational strength of a suspended goal (Adams et al., 1995).

Empirical evidence of the effects of task similarity on post-interruption performance also shows lack of uniform agreement (e.g., Czerwinski et al., 1991a, b; Gillie & Broadbent, 1989; Lahlou et al., 2000). Although task similarity is often regarded as a common disruptive attribute of a secondary task (e.g., Edwards & Gronlund, 1998; McFarlane, 2002), this has occasionally been shown not to be the case (e.g., Bailey et al., 2000b; Latorella, 1996). In fact, the most frequently-cited study that apparently provides evidence of an interruption-similarity effect (Gillie & Broadbent, 1989, Experiment 3), is arguably flawed in both its definition of similarity (e.g., Latorella, 1996) as well the results attributed to similarity effects being confounded by secondary task complexity (e.g., McFarlane & Latorella, 2002). Additionally, of the studies that have provided more convincing evidence of interruption-similarity effects,

some show attenuation, and even elimination, of such effects by training participants to solve primary tasks sequentially (e.g., Edwards & Gronlund, 1998) or through interruption experience (e.g., Hess & Detweiler, 1994).

The most persuasive view is one that secondary task similarity can influence the size of performance deficits following interruption, but only when experience within the task domain and of dealing with interruption is limited (e.g., Trafton et al., 2003). I aim to further investigate this supposition within the current experimental series, by assessing the impact of task similarity on performance following different levels of primary task experience, as well as different levels of experience at dealing with interruptions.

2.2.5. Controlling interruption using the Tower of Hanoi (ToH) task

Placing control over the point of interruption may be achieved using a well-structured problem solving task, and for this purpose, I propose the ToH task. The ToH exhibits clear evidence of the deployment of subgoal structures through longer inter-move latencies usually generated at points with the solution-sequence at which subgoals are created (e.g., Anzai & Simon, 1979; Anderson & Douglass, 2001; Simon, 1975; VanLehn, 1991). Large inter-move latency peaks before moves 1, 9, and 13 (for problems that require fifteen moves and are solved error-free), reflect the initial encoding of ToH subgoals. These are points within the problem space where goals may be hierarchically planned (e.g., Anderson et al., 1993, Anderson & Douglass, 2001; Davies, 2001, 2003). At these points, the allocation of processing resources required to discover an executable operator is greater than at all other positions (e.g., Altmann & Trafton, 1999, 2002; Anderson et al., 1993; Goel, Pullara & Grafman, 2001). Indeed, the time taken to make the first move within a subgoal has been shown to increase linearly with the number of subgoals to be considered (e.g., Anderson et al., 1993; Ruiz, 1987). At subgoal positions, participants seem to try to determine: Where discs currently are, where they want them to be, and how they will decompose larger goals in order to deploy an initial solution-sequence (up to three to four moves in advance, e.g., Davies, 2003).

By contrast, shorter move latencies usually occur for goals that only involve the direct movement of a single disc (e.g., Altmann & Trafton, 1999, 2002). This is the case, for instance, for even-numbered moves, where simple perceptual problem solving

heuristics such as 'Don't-undo' can be applied; lessening the burden on memory (e.g., Altmann & Trafton, 2002). Thus, these moves are likely to require fewer processing resources in uninterrupted contexts; and can be executed using only perceptual-motor processes without the necessity of potentially fallible memory processes (e.g., Larkin, 1989; Zhang, 1997). Intermediate complexity states (e.g., move three in fifteen-move problems) are those in which a smaller disc blocks the movement of one that is larger. These may require more processing than direct moves, but not as much usually deployed at the beginning of subgoals. It is argued that latencies for these moves reflect the time taken to locate a disc cue and retrieve a previously strengthened goal (e.g., Altmann & Trafton, 1999, 2002). Cognition arguably dictates stochastically whether or not to attempt retrieval of such goals (with a probability of 0.5%, Altmann & Trafton, 2002). Resultant retrieval latencies are on average longer than other moves (i.e., one to two seconds), reflecting 50% chance use of a domain-specific perceptual heuristic of 'One-follows-two': Always move the smallest disc following movement of that which is larger than it. Concurrent planning can also occur at different points within the solution-sequence, usually reflecting the overall difficulty in solving the ToH problem and the point reached within the deployment of a move sequence (e.g., Davies, 2003).

With regards to the current experimental series, I propose that certain task constraints will encourage participants to encode to-be-suspended goals to some extent before interruption. Specifically, encouraging planning at the beginning of problems should ensure that 'most of the initial steps' are encoded before problem execution (Davies, 2003, p. 1154; see also Gunzelmann & Anderson, 2003), especially after some experience within the task domain (e.g., Burgess, 1997). In the current experiments, interruption fell only at points within a ToH solution-sequence where moves were likely to have been considered prior to execution (i.e., at the beginning of main subgoals and within the first subtask). The implementation of such subgoal-based planning is regarded as a strategy quickly employed to minimise the working memory (WM) demands that would be required to of plan and store all of the steps within complex problems (e.g., Anderson & Douglass, 2001; Delany, Ericsson & Knowles, 2004; Ward & Allport, 1997). For the ToH task, this level of subgoaling sophistication has been shown to develop after exposure to only three 5-disc problems (e.g., Anzai & Simon, 1975), and

occurs at a subgoal level even when using ToH isomorphs (see Kotovsky, Hayes & Simon, 1985).

2.3. Experiment 1

2.3.1. Introduction

Twenty-four 4-disc ToH problems requiring fifteen moves for error-free completion were to be solved. Half were interrupted once by an un-signalled secondary task. Each problem was matched with an identical control problem (i.e., containing the same start-state and goal-state disc configurations) from which comparative performance measures could be taken. Secondary tasks were 3-disc ToH problems requiring seven moves for error-free completion. As such, primary and secondary task completely overlapped in terms of the three similarity dimensions proposed by Latorella (1996): Processing resources; form of information; and semantic similarity. Interruptions could occur at the beginning of each of three main subgoals (see Figure 2.1 for an example). These were, (1) before making the first move within a problem (2) before making the first following positioning of the LOOP disc in its goal-peg position, and (3) before the first move following positioning of the second LOOP disc in its goal-peg position. Four interruptions occurred at each position (one per interruption condition).
Disc-peg configurations at the start and goal states of 12 (x2) problems were different (see Figure A2 of Appendix A for examples) to control for the potential confound of practice effects that might accrue from repeating the same tower-to-tower problem (e.g., Ruiz, 1988). The difficulty in performing each problem was uniform in that each required the same number of moves for error-free completion, contained an equal number of subgoals located in the same positions of solution-sequences, and required individual moves of equal difficulty at each point within these solution-sequences (see Anderson & Douglass, 2001).

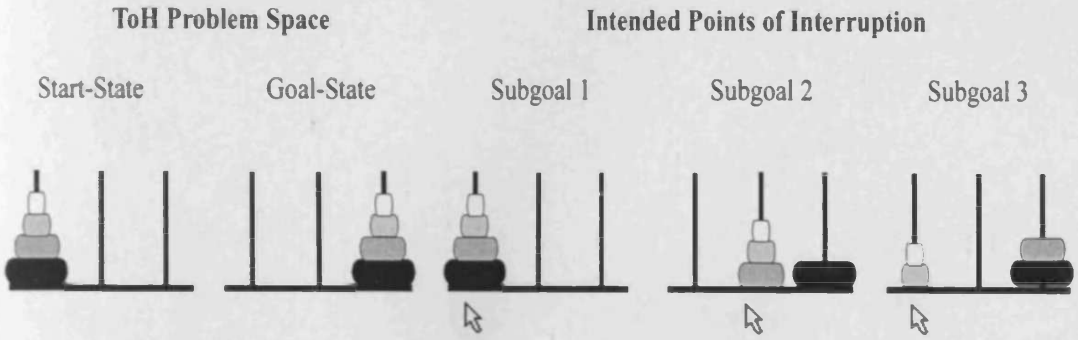

Figure 2.1. A diagrammatic representation of a ToH problem space with examples of the three intended points of interruption in Experiment 1. Interruption would only occur when discs were in the depicted configurations and when the host peg was selected using the mouse cursor (i.e., when selecting the disc that is to be moved)

Following sufficient practice (supported by 3 practice problems) participants usually approach subgoal encoding opportunities using problem solving strategies that involve encoding a series of goals (e.g., Anzai & Simon, 1979; Davies, 2003; VanLehn, 1991). Opportunistic (or on-line) planning might also be used when planning to far in advance might not seem to benefit task performance, perhaps because of exceeding working-memory limitations (e.g., Bishop, Aamodt-Leeper, Creswell, McGurk and Skuse, 2001; Gilhooly et al., 2001; Ward & Allport, 1997).

Given the exploratory nature of this experiment, a number of general hypotheses are posed. Firstly, it is expected that interruption will cause appreciable deficits to primary task performance. It is predicted that goals are processed like all other declarative memory items and will require some form of effortful re-establishment due to temporal-based decay (e.g., Altmann & Trafton, 2002; Anderson & Douglass, 2001). Secondly, given that the duration of secondary tasks is not explicitly stated, effortful goal storage strategies might be abandoned when interrupted (e.g., Anderson & Douglass, 2001). Subsequently, a suspended goal might require reconstruction using the visual array (e.g., Anderson & Douglass, 2001). If this holds true, resumption latencies are expected to be longer for subgoals that require more moves for completion (i.e.,

potentially more moves to reconsider): The deficit should be larger for subgoal 1 compared to subgoal 2, and larger again for subgoal 3.

Alternatively, if subgoal-encoding opportunities are treating as encoding opportunities (e.g., Altmann & Trafton, 2002), resumption latencies should not only be reduced but should be of comparative length irrespective of subgoal size. In this case, resumption latency should be in the region of 1 – 2 s, consisting of the time taken to, locate a cue, rapidly reactivate the representation of the goal, and execute the motor program to make the move (e.g., Altmann & Trafton, 2002). Nevertheless, retroactive interference caused by processing similar goals within secondary tasks may effect the speed and efficiency of this process (e.g., Altmann & Trafton, 2002).

2.3.2. Method

2.3.2.1. Participants

Twenty-seven undergraduate students from the School of Psychology at Cardiff University participated in fulfilment of a course requirement. All were naive to the ToH task.

2.3.2.2. Apparatus and Materials

Primary and secondary tasks were generated using *Microsoft Visual Basic 6.0* and *Microsoft Excel 2000* software. The experiment was run on a *Pentium III 800MHz PC* and tasks were presented through a monitor (measuring 42.3 cm in diameter). An example of the 4-disc ToH primary task visual display is presented in Figure A1 of Appendix A1. Clicking on a host peg and then on a destination peg with a computer mouse controlled all disc-peg moves. Movement was constrained is such a way to prevent larger discs being moved on top of smaller discs. If such an illegal operation was attempted, a command reading 'try again' would immediately appear on the screen, and disc movement would not occur. For a legal move, the trajectory of each disc was constrained to take 500 ms following the mouse click of the destination peg. The secondary task consisted of a 3-disc version of the ToH task, which upon initiation

covered the primary task visual array until its completion. Secondary tasks were controlled in exactly the same way as the primary task. For both primary and secondary tasks, the program recorded: Each legal disc-goal peg executed action (e.g., disc 1[smallest] from peg 1 to peg 3); the time taken to execute each move; the total number of moves executed per problem; and the intended point of interruption. Also, for ease of explanation, ToH discs will be referred to using the following numeric system: Smallest is 'disc 1', medium is 'disc 2', larger is 'disc 3', and giant is 'disc 4'.

2.3.2.3. Design

There were 24 experimental trials, and half were randomly selected to contain interruptions. There were three independent variables: Trial type (interruption vs. control); move execution type (resumption vs. pre-interruption vs. control); and the intended point of interruption (upon attempting to execute the first move of: Subgoal one vs. two vs. three).

There were two main dependent measures, motivated in large by the constraints imposed by the goal-activation model (Altmann & Trafton, 2002). Firstly, the intended point of interruption within this experiment allowed for comparisons between, resumption latency data for each interrupted move, pre-interruption move latency, and the latency to execute the same move in a comparative control trial. The second measure was redundant moves, which are the number of moves within primary task ToH problems executed above that required for error-free completion, following the intended point of interruption. These were, the number of moves above: Fifteen for the first subgoal; seven for the second subgoal; and three for the third subgoal.

Four additional dependent measures were investigated due to the exploratory nature of the current experiment. Firstly, because problem completion time following the intended point of interruption was likely to be confounded by the (1) the number of moves to completion in each subgoal and (2) the number of redundant moves executed within solution-sequences, an adjusted mean completion time (**x**) was taken for each condition. For each condition, this involved dividing the number of moves executed after the point of interruption (actual moves = a) by the completion time (time = **t**) and

multiplying the outcome by the number of moves required for error-free completion (expected moves = **e**). This translates to:

$$\mathbf{x = (a/t)e}$$

(Equation 5)

Given the large differences in subgoal size (i.e., the number of moves required for error-free completion), creating a potential violation of normality, only within subgoal comparisons for interruption and control trials were made for adjusted completion times. Secondly, to further explore completion times, the inter-move latency between the time taken to make the first and second moves within subgoals (referred to as 'interruption plus 1' or IP+1) was taken. It was reasoned that completion times might be inflated because of short-term rather than longer-term performance in ToH problems after the intended point of interruption. However, IP+1 was removed from further analyses for failing to exhibit any significant differences across interruption and control trials. Thirdly, the average time taken to complete secondary task ToH problems was investigated as was the number of moves executed to complete secondary tasks. Secondary task performance measures (i.e., time to complete ToH problems and number of moves to completion) were taken as a fourth measure

All dependent measures within interruption and control problems were compared on a trial-by-trial basis to investigate problem solving experience and experience of being interrupted. There were no significant differences between any condition and thus this measure received no further analysis.

Finally, three full Latin Square trial sequences were used to ensure that conditions satisfied all counterbalancing possibilities, with an equal number of participants tested in each.

2.3.2.4. Procedure

Participants were first asked to read a set of standardised instructions. These contained all relevant information regarding task constraints and operations, as well as encouragement to solve all problems in the minimum fifteen moves required for error-

free completion. They were instructed that planning time would be unlimited at the beginning of problems, but to try and begin executing a solution–sequence as soon as they could. Instructions further stated that during the performance of primary ToH tasks, a secondary ToH task could 'suddenly appear on the screen at any point within the problem', which had to be solved for the experiment to proceed. Participants performed three fifteen-move ToH practice problems; one of which was interrupted immediately after making the eighth move (irrespective of the distance between the start-state and goal-state) with a 3-disc ToH secondary task requiring seven-moves for error-free completion. The experiment took approximately 1 hour to complete.

2.3.3. Results and discussion

2.3.3.1. Move Latency

As illustrated in Figure 2.2, average resumption latencies following interruption are consistently higher than both move latencies before interruption and move latencies in comparative control problems. This pattern is consistent irrespective of subgoal position. The average pre-interruption move latencies and those in control conditions appear very similar; suggesting that similar time was taken to plan the first move in subgoals. Move latencies within the first subgoal conditions contain a large amount of variance, whereas variance is considerably reduced for the second and third subgoals.

Due to occasional violations of assumptions of linearity, and evident heteroscedasicity, all move latency data were transformed using a logarithmic function that served to normalise the linear relationship between the variability in each condition. A repeated measures 3 (move execution type) x 3 (intended point of interruption) ANOVA was conducted on all transformed data. This revealed significant main effects of move execution type, $F(2, 52) = 71.78$, $MSE = .226$, $p < .001$, intended point of interruption, $F(2, 52) = 346.24$, $MSE = .180$, $p < .001$, and a significant interaction between both variables, $F(4, 104) = 13.85$, $MSE = .145$, $p < .001$. This interaction was further explored using Bonferroni post-hoc tests, which revealed a significant simple main effect between resumption latency and pre-interruption move execution latency for

the second subgoal, F (2, 25) = 50.47, $p < .001$, and the third subgoal, F (2, 25) = 151.90, $p < .001$. Specifically, resumption

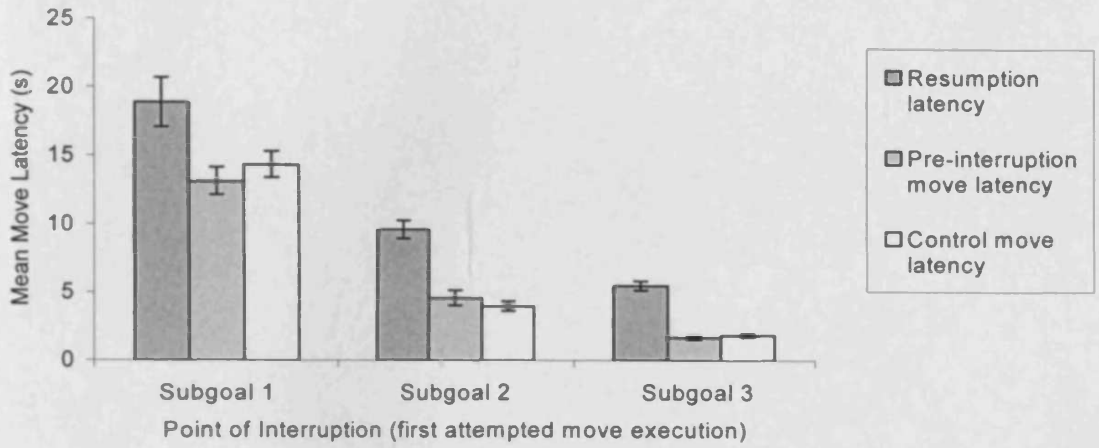

Figure 2.2. Un-transformed mean move execution latencies (±SE) for the first move of subgoals in Experiment 1.

latency was significantly longer than pre-interruption move execution latency. Another significant simple main effect was revealed between resumption latency and control move latency also for the second subgoal, F (2, 25) = 50.47, $p < .001$, and the for the third subgoal, F (2, 25) = 151.90, $p < .001$. This was again due to longer resumption latencies. There were no significant differences between resumption latencies and pre-interruption move latencies ($p = .06$) or control latencies ($p = .33$) at subgoal position one.

There were appreciable resumption deficits due to interruption. The size of deficits seem to favour an interruption recovery strategy of task reconfiguration (e.g., Anderson & Douglass, 2001), and are difficult to explain by a goal retrieval account (e.g., Altmann & Trafton, 2002). Indeed, the pattern of data seems to support the prediction within both models, that goal forgetting occurs because of a declination in the source activation of a goal. However, resumption latencies are longer than the amount of time originally invested in planning each subgoal, suggesting that recovery involves processes above and beyond that of simply re-planning the interrupted state. The reasons for such

elevated resumption latency when interruption fell before earlier subgoals will be investigated in further experiments.

2.3.3.2. Redundant Moves

Table 2.1 illustrates the mean number of redundant moves executed in primary task ToH problems following the intended point of interruption in interruption and control trials. More redundant moves were made in interrupted problems compared to control problems. As predicted, redundant moves decrease as a function of subgoal position, irrespective of the intended point of interruption.

Table 2.1

Mean redundant moves (±SE) within primary task ToH problems following the intended point of interruption in Experiment 1

	Subgoal 1		Subgoal 2		Subgoal 3	
	M	SE	M	SE	M	SE
Interruption	4.56	0.64	1.06	0.16	0.17	0.07
Control	2.88	0.42	0.57	0.13	0.07	0.06

Frequency of redundant move mean ranks were calculated for the first subgoal (interruption Mean Rank = 5.30, control Mean Rank = 4.91), the second subgoal (interruption Mean Rank = 3.63, control Mean Rank = 2.94), and the third subgoal (interruption Mean Rank = 2.28, control Mean Rank = 1.94). A Friedman test indicated that a significant difference occurred between scores for redundant moves for each of the six conditions $\chi^2(5) = 82.28, p < .001$. A series of Wilcoxon Signed Rank tests were carried out to provide *post-hoc* comparisons of mean ranks for each subgoal across trial type. These tests showed significant differences between frequency of mean ranks for redundant moves in interruption and control conditions for, the first subgoal $Z(1) = -2.97$, $p < .005$, and the second subgoal $Z(1) = -3.15, p < .005$. This difference was only marginally non-significant for the third subgoal $Z(1) = -1.91, p = .058$. The present findings suggest that more redundant moves were made in interrupted problems compared to control problems following the intended point of interruption, although *post-*

hoc tests revealed significant differences between trial type only for the first and second subgoals.

2.3.3.4. Move-error frequency

Redundant moves might not be the best indicator of error within the ToH task, especially for problems that require different numbers of moves for error-free completion following interruption. Such analyses may not wholly reflect performance error in the ToH task. Specifically, executing a single move that strays from the error-free solution-sequence of a ToH problem will cause one redundant move, but can also lead to a succession of redundant moves. Failing to 'undo' the initial stray move and continuing with a different route through the problem space will lead to differential numbers of redundant moves, depending on where in the solution-sequence that move occurred. Some stray moves (e.g., moving the LOOP disc to its non-goal peg) have greater potential consequences than others (e.g., moving disc 1 to the wrong temporary destination peg) in terms of the number of redundant moves that might follow the first. These occasions are more likely to occur for earlier subgoals in fifteen move ToH problems, because their completion inherently requires the movement of larger discs.

Perhaps then, a better measure for error data is the number of trials in which redundant moves were made (referred to as 'error frequency' hereafter), after the point of interruption. Table 2.2 presents error frequency data for all participants who committed one or more redundant moves out of a possible 648 interruption and control problems. Cochran's Q tests were conducted on these data due to the binary nature of the non-parametric coding method used (i.e., 0 = no error, 1 = error) (Morgan, Reichert & Harrison, 2002). Participants solved 402 problems error-free, $Q\ (5, N = 27) = 156.13, p < .001$. Pairwise comparisons using a Bonferroni correct ($p = 0.017$) revealed that move-error frequency was higher in interrupted problems compared to control problems for, the second subgoal, $Q(1, N = 27) = 11.77, p < .001$, and the third subgoal, $Q(1, N = 27) = 4.57, p < .05$. Move-errors were greater in interrupted problems compared to control problems for the first subgoal with the difference nearing significance, $Q(1, N = 27) = 3.6, p = 0.58$. Generally, the frequency of move-errors was higher in interrupted problems compared to control problems that were not interrupted.

The actual number of redundant moves executed following the intended point of interruption is higher for those problems that are interrupted, but only for earlier subgoals. However, these values may have been inflated because redundant moves are sometimes higher if the first stray move leads to a sequence of redundant moves, which more frequently occurs for earlier subgoals. When considering only the frequency data for occasions when a redundant move led to one or more redundant moves, values were only significantly higher in interruption conditions for the first and second subgoals. Perhaps error performance for interruption and control trials is similar for the third subgoal because there are so few moves required for error-free completion. That is, there are fewer opportunities to make redundant moves within this subgoal compared to earlier subgoals. Nevertheless, conceptual frameworks that advocate perfect goal memory, are questioned given redundant move and error data from Experiment 1 (e.g., Anderson et al., 1993; Anderson & Lebiere, 1998)

Table 2.2

Move-error frequency following the intended point of interruption in Experiment 1

	Subgoal 1	Subgoal 2	Subgoal 3
Interruption	74	55	14
Control	62	35	6

Note. Maximum move-error frequency per condition is 108 moves

2.3.3.3. Adjusted primary task completion time

Figure 2.4 illustrates the adjusted primary task completion times following the intended point of interruption for interrupted and control trials. For the first subgoal, participants seem to take more time to complete an interrupted ToH problem compared to one that is not interrupted. There are no such apparent differences between move latency for the second subgoal or the third subgoal.

For further analysis paired-samples two-tailed t-tests were conducted on adjusted completion time data comparing each subgoal across trial type. There was only a significant difference between adjusted completion times for the first subgoal, $t(26) = 2.3, p < .05$, with non-significant differences for the second and third subgoals. A similar

amount of time is spent completing ToH problems that have been interrupted at subgoal boundaries within the solution-sequence compared to problems that have not been interrupted. For the first subgoal, participants might abandon the storage of a suspended goal (e.g., Anderson & Douglass, 2001) and goals that have

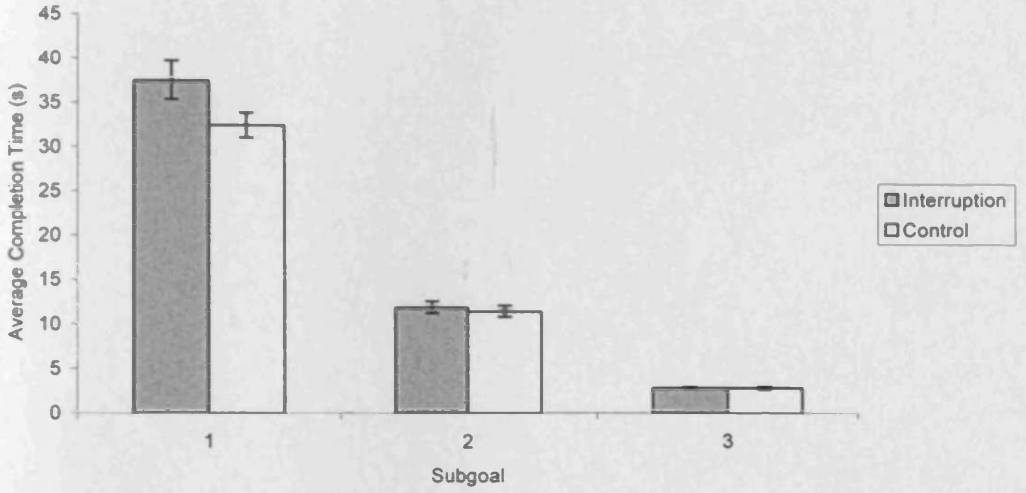

Figure 2.4. Mean adjusted primary task completion time (±SE) for interruption and control primary task problems following the intended point of interruption in Experiment 1

been processed further into the solution-sequence and/or the interference threshold is higher than that of other subgoals (e.g., Altmann & Trafton, 2002). For subgoals that require fewer moves for error-free completion, there might be either residual activation of previously strengthened goals and/or a lower interference threshold (e.g., Altmann & Trafton, 2002).

2.3.3.4. Secondary task performance

Table 2.3 illustrates the time spent performing secondary task ToH problems and the number of moves executed to achieve their goal-states. Completion time was similar for each point of interruption (M = 15.7 s). This is below the 18 s threshold proposed by Anderson and Douglass (2001) whereby they model retrieval accuracy at 50%. Participants tended to solve secondary task ToH problems close to error-free. A one-way repeated measures ANOVA confirmed non-significant differences for secondary task

completion time, whereas a Friedman test confirmed non-significant differences between moves executed within secondary tasks.

Table 2.3.

Mean secondary task completion times (±SE) and mean number of moves (±SE) executed within secondary tasks in Experiment 1

	Subgoal 1		Subgoal 2		Subgoal 3	
	M	SE	M	SE	M	SE
Completion Time	15.8	0.66	16	0.68	15.28	0.58
Number of Moves	7.34	0.15	7.48	0.13	7.59	0.19

Note. Time is in seconds and minimum number of moves is 7

2.3.4. Discussion

Experiment 1 has established successfully that interruption appreciably disrupts performance within the ToH task. These results from a goal-directed behaviour task domain stand in contrast to the few studies that find beneficial effects of interruption on performance efficiency (e.g., Speier et al., 1997; Zeigarnik, 1927; Zijlstra et al., 1999). Performance deficits within the current experiment generally took the form of long primary task resumption latencies, and a greater number of redundant moves caused by move-errors. Resumption latencies were monotonically longer for earlier subgoals, suggesting that the point of interruption affects resumption time. Interestingly, resumption latencies were longer than pre-interruption move latencies, as well as the time taken to execute the same moves in control problems. The frequency of move-errors was also higher in interrupted problems compared to control problem, which subsequently caused the high numbers of redundant moves evidenced. However, only primary task performance following interruption was affected: There was no effect of interruption on pre-interruption primary task performance or on secondary task performance.

Prima facie, the results suggest that participants reconfigure their position within a primary task following interruption, at least in the case of the first subgoal. However, this might not reflect re-establishment of the problem space through re-planning only, given longer resumption latencies compared to other move latencies that are likely to

involve planning (e.g., pre-interruption move latency). Neither can it simply reflect retrieval failure, where retrieval is estimated take no longer than 2 s, else the relevant goal is not retrieved (e.g., Altmann & Trafton, 2002, Anderson & Milson, 1989). Perhaps then, the efficiency of reconfiguring the problem space is effected by retroactive interference caused by performing a similar secondary task (e.g., Shulman, 1970). However, such interference effects are expected to in large effect the process of retrieval (e.g., Altmann & Trafton, 2002) and not reconfiguration. Specifically, performance degradation because of competition between items is often attributed to a limited capacity working-memory, where higher-order processing of information is involved (e.g., Nairne, 1990, 2002).

Participants might attempt to retrieve the representation of a suspended goal; a process met with failed conflict resolution due to decayed representation of the goal (e.g., Anderson & Milson, 1989). Reconstruction of the suspended goal using the visual array as a source of declarative memory fits resumption latency data for the second subgoal (e.g., Anderson & Douglass, 2001), but not for the first subgoal. For the first subgoal, participant have to reconfigure the problem space, such is the case during initial planning (e.g., Davies, 2003), but this might take longer because of the interference caused by being interrupted, perhaps an initial failed attempt at retrieving the goal from memory.

The re-establishment of a suspended goal seems to be effected by resumption strategy (e.g., retrieval, reconstruction and so on), the efficiency of which might be affected by other task demands. For example, the number of similar items contained within the secondary task, as well as its duration (e.g., Anderson & Douglass, 2001) might disrupt retrieval of a suspended goal, resulting in forced reconfiguration the problem space. That is, the secondary tasks used within the current experiment may utilise the associative activation that has been attributed to the suspended goal, thus decreasing the likelihood of its retrieval. Such disruption might be more amplified at earlier ToH subgoals, where goal states are more complex and perceptual heuristics less obvious (Altmann & Trafton, 2002): This might encourage participants to 'rethink' what they wanted to do before being interrupted (e.g., Speier et al., 2003). The cost of interruption, above that of the time taken to reconfigure the problem state, might reflect failed or abandoned retrieval. Finally, resumption latency appears to increase as a

function interruption position, with duration reflecting the number of moves remaining within the problem (similar to uninterrupted performance, e.g., Altmann & Trafton, 2002).

The remaining experiments within the current series will attempt to speak to this issue by, investigating whether the effect is one of primary task problem complexity (especially at the point of interruption), and by exploring the potential influence of the number component parts contained within ToH secondary tasks.

2.4. Experiment 2

2.4.1. Introduction

While Experiment 1 provided evidence of substantial performance deficits due to the point of interruption, it is somewhat clear the effect sizes were influenced by the complexity associated with performing the primary task, but not so clear whether the complexity of the secondary task also impacted upon performance. These issues are addressed within the current experiment by reducing the difficulty associated with performing interrupted moves *within* simpler ToH primary tasks, and by manipulating the amount of similar information to be processed within secondary tasks.

In the current experiment, primary task problems required only nine moves for error-free completion. ToH problems that require fewer than fifteen-moves for error-free completion are usually solved with greater accuracy (e.g., Bishop et al., 2001; Davies, 2003). All problems in Experiment 2 consisted of the first subgoal and the first move of the second subgoal of problems used in Experiment 1[4]. ToH secondary tasks varied in the amount of moves required for error-free completion (i.e., four vs. seven). Participants were explicitly informed that secondary tasks might vary slightly in duration, but that generally secondary tasks would be 'short'. This was to reduce the likelihood of

[4]This resulted in problems containing the first main subgoal and the first move of subgoal 2.The additional move (move 9 if solved error free) was included to try and reduce the potential inhibitory effects of previously encoded goal sequences (i.e., the previous trial) on the performance of future goal sequences (i.e., the current trial) (e.g., Mayr & Keele, 2000; Monsell, 1996).

abandoning goal storage because of the expectancy long retention intervals (e.g., Anderson & Douglass, 2001). Eight problems were interrupted (each with a matched control); 4 for each intended point of interruption within the first 12 problems, and 4 for each intended point of interruption within the last 12 problems. The remaining 8 problems were included as fillers, to further reduce anticipation of interruption (e.g., Anderson & Douglass, 2001).

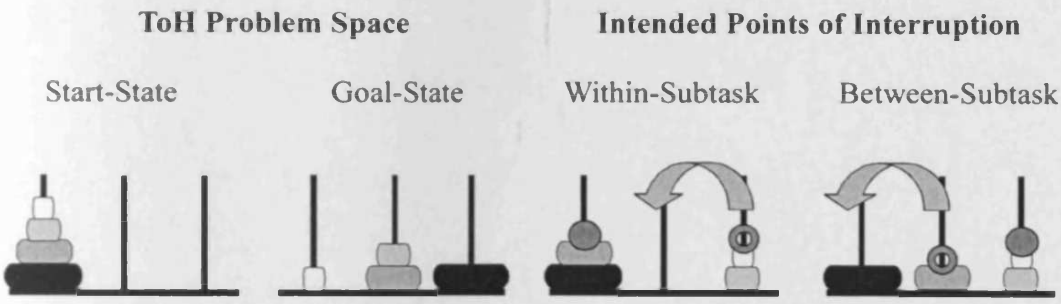

Figure 2.5. A diagrammatic representation of a ToH problem space with examples of the two intended points of interruption in Experiment 2. Interruption would only occur when discs were in the depicted configurations (i.e., after making the last move). Filled circles represent the correct disc cue. Empty circles represent incorrect disc cues that if moved would cause a move-error. The arrow represents the legal but incorrect move of 'undoing' the previously executed move

The intended point of interruption was manipulated within subtasks rather than subgoals. As introduced in Chapter 1, a subtask is a move sequence contained within a subgoal that involves: Identifying a blocked disc (not the LOOP) and its temporary destination peg (*initial* move); moving all blocking discs to the non-temporary destination peg (*middle* move(s)); and finally moving the blocked disc to its temporary destination peg (*end* move). Execution latencies for these moves are often different (i.e., the move at the beginning of a subtask has a longer latency than that at the end), although these differences are relatively small compared to the time taken to initiate ToH subgoals (e.g., Altmann & Trafton, 2002; Anderson et al., 1993). Interruptions occurred before having the opportunity to execute the last move of the first subtask (referred to as

'within-subtask' in Experiment 2) or before having the opportunity to execute the first move of the second subtask (referred to as 'between-subtask' in Experiment 2). See Figure 2.5 for an example of each point of interruption within Experiment 2.

In Experiment 1, it was noted that participants invested different amounts of time in the initial planning of problems, which might have affected the amount of goal activation strengthening (Altmann & Trafton, 2002). Here, a 10-second planning restriction was included at the beginning of each primary task problem in order to gain closer control over the amount of strengthening participants could attribute to initial solution-sequences before execution. Fifteen seconds is arguably enough time to allow participants to adequately plan a sequence of initial moves within fifteen move ToH problems (e.g., Davies, 2003), and so 10 s should be adequate for problems that require only nine moves. It is noted however that participants might not always configure a coherent plan at initial problem encoding stages, even when encouraged (e.g., Davies, 2003; Ward & Allport, 1997), and even when a plan is formed it may not always be adhered to at execution (e.g., Phillips et al., 1999).

A number of predictions are made in Experiment 2 that build upon those in Experiment 1. Firstly, given that interruption annunciation is un-signalled, the level of associative activation required to prime a goal to a ToH disc might be insufficient (e.g., Altmann & Trafton, 2002). However, within the current experiment, I am confident that participants will have enough time to allocate some associative activation to a potential priming cue before interruption, the effects of which will become apparent at resumption. In the goal-activation model it is stated that, "...at retrieval time, which comes immediately after the system moves a disc (and has to look for another disc to move), it makes sense to ask what other disc might now be moveable" (Altmann & Trafton, 2002, p. 53). Within the current experiment, disc movement was constrained so that it took 500 ms for a disc to complete a single trajectory between pegs following completion of peg clicks to move the disc. So although interruption was un-signalled, an interruption lag of sorts was inexplicably available so that associative priming might occur immediately before secondary task initiation. Secondly, reconstruction of the suspended goal or reconfiguration of the problem space could occur because of the un-signalled nature of interruption: Retrieval being arguably impossible if priming were not to occur (e.g.,

Altmann & Trafton, 2002). Resumption latencies should therefore be longer in the between-subtask condition compared to the within-subtask condition, because of the additional planning demands required in the former. However, Anderson & Douglass (2001) suggest that goal storage can occur if the retention interval is perceived to be short, thus increasing the likelihood of retrieval. In this case, resumption latency might reflect the time taken boost the goal's base-level activation, which should be supported by the availability of cues located within the stimulus array (e.g., Altmann & Trafton, 2002).

Secondary task attributes may also influence the hypothesised effects. Given the predictions of ACT-R (e.g., Anderson & Lebiere, 1998), the actual (and not expected) duration of a secondary task should negatively effect retrieval time due to activation decay. Additionally, Altmann & Trafton (2002) propose that if participants attempt to retrieve a suspended goal, the efficiency of this process will be affected by retroactive interference caused by similar secondary task goals. Moreover, given the sizeable resumption latencies found in Experiment 1, participants might use a combination of resumption strategies to resume interrupted goals. Again, larger resumption latencies than would normally be expected might accrue if retrieval fails (e.g., Siegler, 1987, 1988), due to the additional time taken up from having to initialise and deploy processes that enable the reconfiguration of the problem space.

2.4.2. Method

2.4.2.1. Participants

Thirty-six undergraduates from the School of Psychology at Cardiff University participated in fulfilment of a course requirement. All were naive to the ToH task. Eight participants were removed (leaving a total of 28) for failing to produce ToH solution-sequences in interruption and control problems that could cause interruption initiation within 4 moves of the intended point of interruption[5].

[5] In Experiment 2 interruptions were implemented at points within a solution-sequence after the first move had been executed. Participants were removed from the study for failing to reach points within one or more interruption and control problems that would normally lead to an interruption within 2 moves above the specified point of occurrence (e.g., more than five moves in the within-subtask condition). This criteria

2.4.2.2 Apparatus and materials

The same program used in Experiment 1 was modified to contain primary task ToH problems that required nine moves for error-free completion. Additionally, a fixed planning period of 10 s was cued at the beginning of each problem. This was signalled by the command 'plan' appearing on the screen for 10 s, during which time all response buttons were inactive. The buttons were reactivated upon expiration of this period, which was indicated by the command 'solve' appearing on the screen. Secondary tasks consisted of 4 randomly selected seven-move problems from Experiment 1. Another 4 secondary task problems were randomly selected and modified so that they required only the first four moves for error-free completion.

2.4.2.3. Design

There were three independent variables, each with two levels: Trial type (same as in Experiment 1); intended point of interruption (*within-subtask* one vs. *between*-subtasks one and two); and secondary task complexity (*low complexity* referring to four-move ToH problems vs. *high complexity* referring to seven-move ToH problems). Dependent variables were, move latency, redundant moves, move-error frequency, the time spent performing secondary tasks, and the number of moves to solution in secondary tasks. Error-free performance in primary tasks following the intended point of interruption required six moves for within-subtask conditions and five moves for between-subtasks. Any moves executed above these values were regarded as redundant, with any instance of redundancy providing a measure of error. Average completion time and IP+1 failed to exhibit significant differences across any condition and were removed from further analysis. All dependent measures were also further analysed to assess the potential influence of problem solving experience and experience of being interrupted: Specifically, performance during the first 12 trials compared to that of the last 12 trials. Move latency was the only variable to exhibit any significant differences across experimental half and as such will be considered in the results and discussion section.

---

was set to allow for a maximum of one instance where an 'undo' heuristic was applied or when the same disc was moved more than once consecutively before reaching its temporary destination peg.

Pre-execution planning time was constrained to 10 s at the beginning of each problem, with encouragement via instructions to begin executing following this period. Four full Latin Square trial sequences were used to ensure that conditions satisfied all counterbalancing possibilities, with an equal number of participants tested in each.

2.4.2.4. Procedure

Participants were explicitly instructed about the 10 s pre-problem execution planning constraint and the fact that secondary tasks would be 'short'. The whole experiment took approximately 45 minutes to complete. Everything else was the same as in Experiment 1.

2.4.3. Results and Discussion

2.4.3.1. Move latency

As illustrated in Figure 2.6, interruption resulted in longer resumption latencies across all interruption problems compared to comparable move latencies in control conditions. Additionally, resumption latencies are higher in between-subtask conditions compared to within-subtask conditions, with similar (although considerably shorter) differences between control move latencies for the same moves. Resumption latency does not appear to be affected by secondary task complexity at either point of interruption. Finally, variance for between-subtask control conditions is considerably lower than that for comparable within-subtask conditions. This is an indication that most participants execute a within-subtask 'goal' move – moving thesecond LOOP disc to its goal destination peg – within a very similar temporal parameter, compared to movement of a disc to a non-goal peg, perhaps suggesting more planning in the latter case.

For further analysis a 2 (trial type) x 2 (intended point of interruption) x 2 (secondary task complexity) ANOVA was conducted, revealing significant main effects of trial type, $F(1, 27) = 288.53$, $MSE = 2.48$, $p < .001$, and intended point of interruption, $F(1, 27) = 28.04$, $MSE = 1.96$, $p < .001$. There was a non-significant main effect of secondary task complexity, and none of the variables significantly interacted.

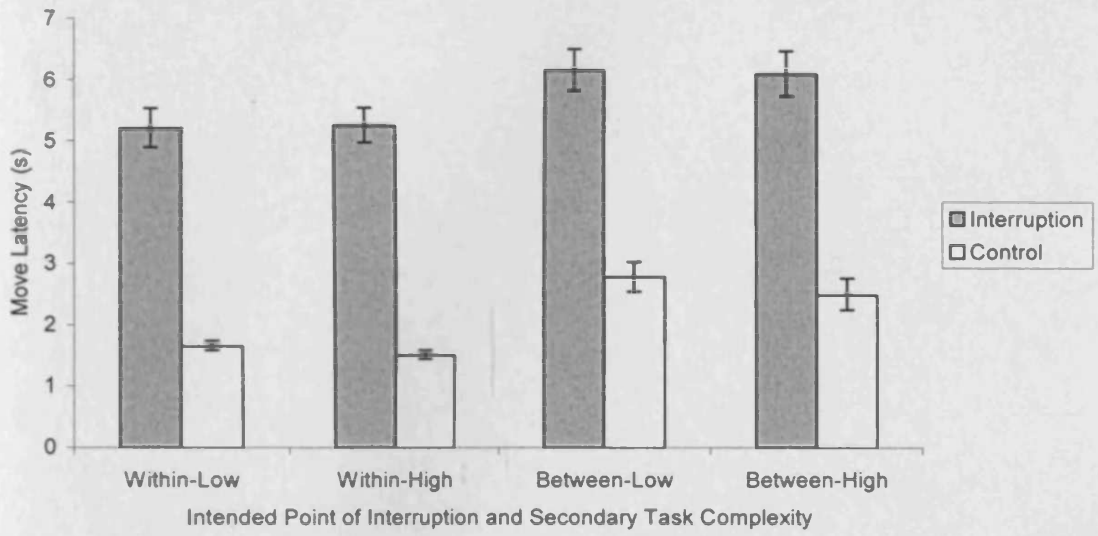

Figure 2.6. Mean move execution latencies (±SE) for each intended point of interruption, following short and long secondary tasks, within interrupted trials and control trials in Experiment 2.

Analyses show that move latencies for control problems are representative of those found within the wider non-interruption experimental literature (e.g., Anderson et al., 1993; Davies, 2003). It takes considerably longer to resume a ToH problem that has been interrupted compared a ToH problem that has not been interrupted. Resumption move latencies and control move latencies are longer for between-subtask conditions than for within-subtask conditions. Finally, resumption latency was not affected by secondary task complexity when interruption fell within a subtask or between a subtask.

2.4.3.2. Move latency and experience

Figure 2.7 illustrates mean move latencies for each intended point of interruption move for trials 1 – 12 compared to trials 13 – 24. There is no evident reduction in control move latency. Additionally, within-subtask resumption latencies are only marginally shorter across problems across experimental half, irrespective of secondary task complexity. Interestingly, this difference is much more elevated for between-subtask resumption latencies, but again this is not influenced by secondary task complexity.

Resumption latencies are however very similar for problems 13 - 24 irrespective of the point of interruption and secondary task complexity.

Control move latencies demonstrated no significant improvement across experimental half and were thus excluded from further analyses. For resumption latencies, a 2 (experimental half) x 2 (intended point of interruption) x 2 (secondary task complexity) repeated measures ANOVA was conducted. This revealed significant main effects of, experimental half, $F(1, 27) = 15.94$, $MSE = 3.84$, $p < .001$, intended point of interruption, $F(1, 27) = 9.79$, $MSE = 3.91$, $p < .005$, with no effect of secondary task complexity. Experimental half and intended point of interruption were the only variables to interact significantly, $F(1, 27) = 6.6$, $MSE = 3.39$, $p < .025$. A Bonferroni *post-hoc* analysis revealed a significant reduction in resumption latency for between-subtask conditions, $F(1, 27) = 17.07$, $p < .001$, which was not the case for within-subtask conditions.

Results suggest that experience of interruption and experience of problem solving within the task domain can reduce resumption latency in interrupted ToH problems. This effect was only significant for between-subtask interruption conditions, suggesting that participants did not benefit from more experience of problem solving in the ToH task when interruption fell before a move that should require less consideration through encoding. When interruption fell before a move that was likely to require more complex processing resumption latency might have reduced because of increased knowledge of the solution-sequence (e.g., the ability to chunk goal-sequences, Davies, 2003), or perhaps because of adaptation of more efficient goal maintenance strategies (e.g., Monk, 2004). The fact that control move latency did not reduce across experimental half suggests that just experience of problem solving in the task domain is not sufficient to cause move latencies that are faster; rather participants adapt better strategies to encode and resume an interrupted goal.

2.4.3.3. Redundant moves

Table 2.4 illustrates the mean number of redundant moves executed in primary task ToH problems following the intended points of interruption. Fewer redundant moves were made in interrupted problems compared to control problems (with an

apparent ceiling effect within all interrupted problems). Variance is however exceptionally high in end-subtask control conditions, although lower in within-subtask. Of the redundant moves executed in control problems, these ranged from 1 to 7.5 moves above that required for error-free completion, thus causing the inflated variance compared to interrupted conditions. Within interrupted problems, the majority of problems were solved error-free, more so than in control problems (see Table 2.5). Of the redundant moves made in interrupted problems, these ranged from 0.5 to 1 move above that required for error-free completion.

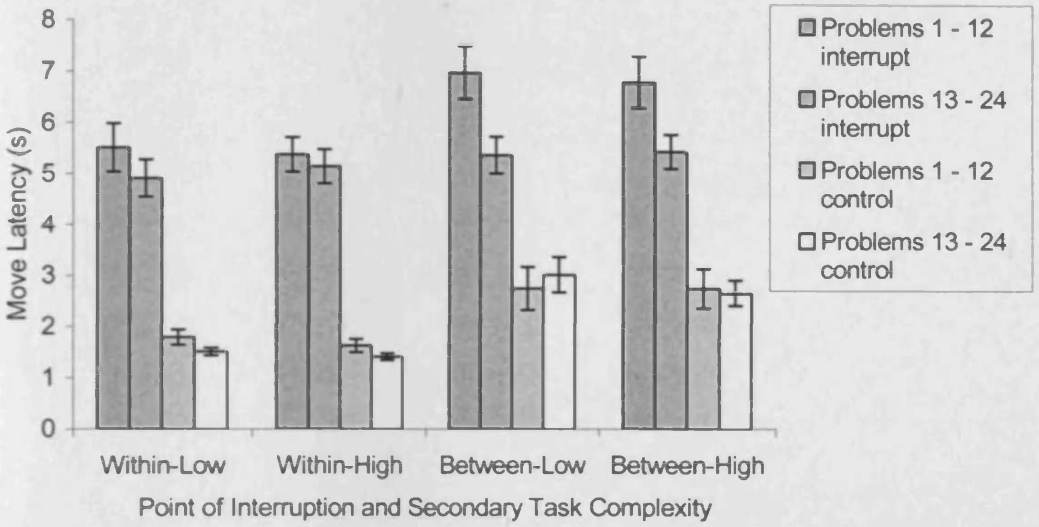

Figure 2.7. Mean move latencies (±SE) for trials 1 – 12 to trials 13 – 24 of the test phase of Experiment 2.

Table 2.4

Mean redundant moves (±SE) within primary task ToH problems following the intended point of interruption in Experiment 2

	Within-Low		Within-High		Between-Low		Between-High	
	M	SE	M	SE	M	SE	M	SE
Interruption	0.05	0.04	0.04	0.02	0.02	0.02	0.02	0.02
Control	0.27	0.09	0.6	0.2	0.79	0.31	0.66	0.30

Frequency of redundant move mean ranks were calculated for the within-short condition (interruption Mean Rank = 3.95, control Mean Rank = 4.63), the within-long condition (interruption Mean Rank = 3.96, control Mean Rank = 5.52), the between-short condition (interruption Mean Rank = 3.84, control Mean Rank = 5.13), and the between-long condition (interruption Mean Rank = 3.82, control Mean Rank = 5.16). A Friedman test indicated a significant difference between scores for redundant moves for each of the eight conditions $\chi^2(7) = 38.72, p < .001$. A series of Wilcoxon Signed Rank tests were carried out to provide *post-hoc* comparisons of mean ranks for each condition across trial type. These showed significant differences between frequency of mean ranks in interruption and control conditions for, the within-subtask low-complexity condition, $Z(1) = -1.93, p < .05$, the within-subtask high-complexity condition $Z(1) = -2.9, p < .005$, the between-subtask low-complexity condition, $Z(1) = -2.68, p < .010$, and the between-subtask high-complexity condition, $Z(1) = 2.68, p < .010$.

The present findings suggest that more redundant moves were made in control problems compared to interrupted problems following the intended point of interruption: A result that contrasts to the original hypothesis. Results impart that performance, in terms of the number of moves taken to complete ToH problems, is better following an interruption compared to when problem solving is uninterrupted. Possible causes for such an effect are discussed in the next section.

2.4.3.4. Move-error frequency

Table 2.5 presents move-error frequency data for all participants who committed one or more redundant moves out of a possible 448 interruption and control ToH problems. Participants solved 391 problems error-free, $Q(7, N = 28) = 47.31, p < .001$. Pairwise comparisons using a Bonferroni correct ($p = .0125$) revealed that more errors were made in interrupted problems compared to control problems for, the within-subtask low-complexity condition, $Q(1, N = 28) = 4.46, p < .05$, the within-subtask high-complexity condition, $Q(1, N = 28) = 10.29, p < .001$, the between-subtask low-complexity condition, $Q(1, N = 28) = 13, p < .001$, and the between-subtask high-complexity condition, $Q(1, N = 28), = 11.27, p < .001$. Thus, the hypothesis that more

move-errors would be made in interrupted trials was again unconfirmed, with analyses showing that more move-errors were actually made in control trials.

Table 2.5

Move-error frequency following the intended point of interruption in Experiment 2

	Within-Low	Within-High	Between-Low	Between-High
Interruption	2	2	1	1
Control	9	14	14	14

Note. Maximum move-error frequency for each condition is 58 moves

Participants made significantly more redundant moves and move-errors in control problems compared to interrupted problems. This contrasts to the original hypothesis, which was largely motivated by the predictions of Altmann & Trafton (2002) and results for the same dependent measures in Experiment 1. That fewer redundant moves and move-errors were executed within interrupted problems of Experiment 2 might indicate a positive effect of being interrupted (e.g., Zijlstra et al., 2003). Possibly, elevated resumption latencies improve the efficiency of the subsequent solution-sequence by reducing the occurrence of redundant moves (although this is not supported by significant correlations between the two variables). The idea here, is that when given the opportunity to break from a problem state (as is the case when interrupted), participants may spend more time processing a solution-sequence that has a greater utility in terms of reaching the goal-state in fewer moves than if this opportunity were not available. Given the redundant move and move-error data from Experiment 1, it seems that this may be a strategy applied when primary task complexity is reduced: When fewer moves are required to achieve the goal-state. A question stemming from this premise is to whether evaluation of the remaining solution-sequence occurs if the suspended goal is re-established through retrieval and/or reconstruction as well as reconfiguration of the problem space? This issue will be addressed within Chapter 3.

2.4.3.5. Secondary task performance

Table 2.6 illustrates the time spent performing secondary task ToH problems and the number of moves taken to achieve their goal-states. More time was taken to complete secondary task problems that required more moves for error-free completion. Subsequently the number of moves to completion was higher in problems that required more moves for error-free completion. Participants made fewer redundant moves in low complexity secondary tasks (M = 0.1 moves) compared to high complexity secondary tasks (M = .057 moves).

Table 2.6

Mean secondary task completion times (±SE) and mean number of moves (±SE) executed within secondary tasks in Experiment 2

	Within-Short		Between-Short		Within-Long		Between-Long	
	M	SE	M	SE	M	SE	M	SE
Completion Time	12.96	0.49	12.89	0.69	19	0.95	19.53	0.86
Number of Moves	4.06	0.36	4.31	0.46	8.29	0.57	8.2	0.49

Note. Time is in seconds

A repeated measures 2 (point of interruption) x 2 (secondary task complexity) ANOVA confirmed a significant main effect of secondary task complexity, $F(1, 27) = 116.64$, $MSE = 10.06$, $p < .001$. Frequency of secondary task move mean ranks were calculated for the within-low secondary task (Mean Rank = 1.5), the within-high secondary task (Mean Rank = 3.46), the between-low secondary task (Mean Rank = 1.5), and the between-high secondary task (Mean Rank = 3.54). A Friedman test indicated that a significant difference occurred between scores for number of moves for each of the for secondary tasks $\chi^2(3) = 77.80$, $p < .001$. A series of Wilcoxon Signed Rank tests were carried out to provide *post-hoc* comparisons of mean ranks for moves executed within each secondary task. These tests showed significant differences between frequency of

executed moves between, the within-low and within-high secondary tasks, $Z(1) = -4.87, p < .001$, and the between-low and between-high secondary tasks $Z(1) = -4.69, p < .001$.

As expected, completion times were longest for secondary tasks that required more moves for error-free completion, as were the number of moves executed within those problems. Participants also solved low complexity problems with fewer redundant moves than high complexity secondary tasks. In fact, solution-sequences were close to optimal in the case of low-complexity secondary tasks. There was however no influence of the point of interruption on either secondary task complexity or the number of moves executed. So, as predicted, high complexity secondary tasks proved to be more costly in terms of the number of mistakes that could be made and the time taken to reach the goal-state.

2.4.4. Discussion

Experiment 2 has provided further evidence for performance deficits in the form of sizeable resumption latencies when resuming an interrupted ToH problem. Both redundant moves and move-error frequency were higher in control ToH problems than after completion of a secondary task in interrupted problems. Resumption latencies were however appreciably shorter than those evidenced for the first and second subgoals in Experiment 1.

Why should resumption latencies be shorter in Experiment 2 compared to those in Experiment 1? For the primary tasks used in Experiment 2, the distance between the initial-states and goal-states – in terms of the number of moves required for error-free completion – was reduced in comparison to Experiment 1. This might have encouraged and/or allowed participants to invest more effort in planning the initial steps of a primary task solution-sequence (e.g., Anderson & Douglass, 2001; Davies, 2003). Perhaps increasing the base-level activation of goals within a to-be-interrupted goal-sequence renders that sequence less immune to the disruptive effects of interruption (see Experiment A in Appendix A2 for empirical support for such a claim). Another possibility, is that being told that secondary tasks would be relatively short, might have led to more frequent and more efficient storage of suspended goals thus increasing the chances of successfully retrieving a suspended goal.

From a goal-activation perspective, it should be noted that in Experiment 2, interruption fell at task points that typically require lower processing demands than were likely to be required at the points of interruption within Experiment 1. One, between-subtasks, usually reflects a longer move latency, because of goal recursion, whereas the other, within-subtask, can be executed using a simple perceptual-motor heuristic (i.e., movement of the second LOOP disc to its goal destination peg), with little need for higher-order processing (e.g., Altmann & Trafton, 1999). Perhaps a cue that could be used to associatively prime a suspended goal was more salient in the within-subtask condition, which participants might have been able to process even without an explicitly cued interruption lag. Such an effect is hinted towards in the goal-activation model through the assumption that for every move executed within a ToH problem, a means-ends cue 'likely to be attended both when a goal is suspended and when it is retrieved (Altmann & Trafton, 2002, p. 53). By contrast, the between-subtask condition represents a point within the ToH solution-sequence (a natural break point) where planning may need to occur again to play out the consequences of executing a move on the solution-sequence (e.g., Laird, 1984; Newell, 1990). Consequently, the move may not have been encoded either during initial problem planning (as supported by Experiment A in Appendix A2) or before the goal was suspended. That participants are faster to resume in between-subtask conditions with experience of problem solving in the task domain and experience of being interrupted, suggests more efficient utilisation of cues located within the available stimulus array (e.g., Altmann & Trafton, 2002); perhaps because of less online-planning (e.g., Davies, 2003).

Nevertheless, retrieval time based upon current ACT-R model estimates (e.g., Altmann & Trafton, 2002; Anderson & Douglass, 2001) cannot accommodate the resumption latencies obtained in Experiment 2. The current data seem to support the contention of Altmann and Trafton (2002), that an interruption lag might be required for efficient priming of a suspended goal to occur. Without an interruption lag, the opportunity to rehearse interrupted state information and to form an associative connection between a suspended goal and a priming cue is hindered.

In contrast to Experiment 1, there was an inverse relationship between being interrupted and the number of redundant moves executed within primary task ToH

problems, as well as the frequency of move-errors. This can be explained by assuming at least part of the resumption latency is spent evaluating the utility of executing a re-established move (e.g., Speier et al., 2003). Redundant move data and move-error frequency data in Experiment 2 suggest that participants are able to somehow refine a better solution-sequence, perhaps when the problem space contains fewer moves for completion (in contrast to Experiment 1).

In sum, Experiment 2 has demonstrated that when a secondary task is similar to the primary task, its complexity (as determined by the number of actions required for its completion) does not effect the time taken to resume the interrupted problem (in contrast to Gillie & Broadbent, 1989). However, might it be the case that the similarity overlap between primary and secondary tasks is the cause of resumption latencies that are far in excess of those predicted by current retrieval based models of goal suspension and resumption? Does the similar processing requirements of the secondary task, irrespective of its number of component parts (or duration), utilise associative activation that may be vital for maintaining the representation of a suspended goal? Experiment 3 set out to investigate just this, by assessing the similarity overlap between primary and secondary tasks, in terms of, processing resources, information form, and semantic similarity (Latorella, 1996). Another purpose of Experiment 3 is to further explore the effects priming cue availability within the primary task stimulus array when interrupted. In Experiment 2, it was more likely that a priming cue would have been strengthened in the within-subtask condition compared to the between-subtask condition (which might require planning with or without interruption).

2.5. Experiment 3

2.5.1. Introduction

While Experiments 1 and 2 provided substantial evidence of marked performance deficits because of being interrupted, it is not clear whether the magnitude of these costs was influenced by the similarity between primary and secondary tasks. In both

experiments, secondary tasks were ToH problems that required similar processing resources and contained similar items to the primary tasks in which they were embedded. Some argue that similarity between tasks leads to higher subsequent performance deficits than when performing secondary tasks that are dissimilar to primary tasks (e.g., Gillie & Broadbent, 1989; Czerwinski et al., 1991a, b), although others present contrasting evidence (e.g., Bailey et al., 2000b; Latorella, 1996). Incidentally, the disruptive nature of task similarity has been identified within studies of dual-task performance (see Pashler, 1998, for a review), although similarity between simple tasks (e.g., colour naming) is also shown to reliably reduce typical switch-costs found in task switching studies (see Monsell, 2003, for a review). Nevertheless, in the case of interruption, Altmann & Trafton (2002) suggest that task similarity should have a negative impact upon the retrieval efficiency of a suspended goal. If within the current experiments participants are able to use retrieval as a resumption strategy for a task that is interrupted then such a process might be impinged upon by retroactive interference generated by intervening goals.

In Experiment 3, secondary tasks were either highly similar or highly dissimilar to primary task ToH problems. The highly similar secondary tasks were the 3-disc seven-move ToH problems used in Experiment 2. The highly dissimilar secondary task was a simple digit-recall task that requires verbal working-memory resources: A task type shown not to correlate with performance in the ToH task (e.g., Handley, Capon, Copp & Harper, 2002). Seven digits were indicated one at a time, and the task of the participant was to recall the digits in the serial order in which they were presented following a cue to recall. Remembering 7 items (or chunks of items) is often regarded as being within the memory capacity of a normal adult (see Jones, 1997, for a comprehensive review), although the classic memory literature suggests this figure is closer to 5, plus or minus 2 items (Miller, 1956). The digit-recall task was differentiated from the ToH secondary task in terms of processing resources (e.g., verbal processing), information form (characters), and semantic processing (e.g., numbers).

Within the current experiment the actual structure of a single ToH subtask was explored with regard to interruption position. Specifically, interruptions were co-ordinated at points within ToH solution-sequences that usually exhibit very similar move

latencies, and are not associated with large amounts of planning (e.g., Altmann & Trafton, 1999, 2002). These were the third move of the first subtask (if solved error-free)(referred to as *within-subtask* hereafter, and should not be confused with the within subtask condition in Experiment 2) and the last move of the first subtask (renamed *end-subtask*; replacing the within-subtask label within Experiment 2). Figure 2.8 provides examples of each point of interruption. Although these moves exhibit similar latencies in uninterrupted data (2.2 s and 2.1 s respectively, taken from Anderson et al., 1993), strategies used to recover their goal representations may differ, and might be unequivocally affected by interruption. The end-subtask condition provides a disc-cue that is both available and highly salient (i.e., its movement will advance the solution-sequence), whereas the within-subtask condition does not offer such a viably 'obvious' heuristic. Additionally, the end-subtask condition provides only one destination for the to-be-moved disc: Its overall goal destination peg. This is not the case for the within-subtask condition where there are two possible destination pegs; neither of which is the overall goal destination for the to-be-moved disc. So the time required to associatively prime the within-subtask move might be beyond the 500 ms time required for a disc to complete one trajectory, which might also be more adept to the end-subtask move.

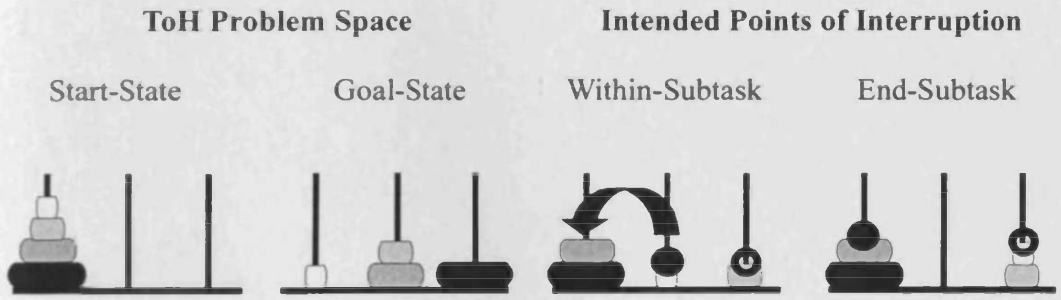

Figure 2.8. A diagrammatic representation of a ToH problem space with examples of the two intended points of interruption in Experiment 3. Interruption would only occur when discs were in the depicted configurations (i.e., after making the last move). Filled circles represent the correct disc cue. Empty circles represent disc cues that might compete with the correct disc cue. The arrow represents the legal

but incorrect move of returning the disc last moved to the peg in which it was moved from, when that disc can be moved to another peg to form the correct move

If participants are relying upon reconfiguration of the problem space as a strategy for re-establishing interrupted state information, resumption latency should be markedly longer for within-subtask conditions compared to end-subtask conditions. This would reflect the time taken to, reconfigure a series of moves (e.g., via re-planning), decide upon the first move to be executed, and configure and deploy the motor action to execute the move. If however, participants are able to retrieve interrupted state information, resumption latencies should be similar irrespective of the point of interruption (e.g., Altmann & Trafton, 2002). Finally, if secondary task similarity were to result in differences between resumption latencies, it might be the case that retrieval as a resumption strategy is disrupted (e.g., Altmann & Trafton, 2002). Having processed goals that are similar to that which has been interrupted might interfere with the representation of that goal when retrieval is attempted. Reconstruction of the goal using the visual array as a source of declarative memory should be void of higher order interference effects (e.g., Anderson & Douglass, 2001).

2.5.2. Method

2.5.2.1. Participants

Thirty-six undergraduates from the School of Psychology at Cardiff University participated in fulfilment of a course requirement. All were naive to the ToH task. Eight participants were removed (leaving 28) for the same reasons specified in section 2.4.2.1.

2.5.2.2. Apparatus and materials

The program used in Experiment 2 was modified for interruptions to occur within or at the end of the first subtask of primary task ToH problems. Secondary tasks were either the 7-move ToH problems used in Experiment 2, or a serial digit-recall task. When initiated the serial digit task displayed 7 text boxes, each containing a single digit from

the range 1 – 9. Each number was presented in a black Times New Roman 16-point font, and each was 2cm apart. None of the digits were repeated in any trial, and they were always constrained not to form a transitional sequence (e.g., '1' '2' '3' or '9' '8' '7'). All digits appeared instantly and remained on the screen for 7 seconds, after which time they disappeared and recalled digits could be recorded in response boxes. A digit had to appear in each response box before the primary task was reinstated.

2.5.2.3. Design

The intended point of interruption was manipulated so that interruptions could occur within or at the end of the first subtask of primary task ToH problems. Additionally, interruption type (similar vs. dissimilar) replaced secondary task set-size as a new independent variable. All dependent measures were the same as those within Experiment 2.

2.5.2.4. Procedure

Participants received instructions that were modified by the experimenter, guiding them on how to perform digit-recall tasks. They were instructed to try and remember as many of the digits as possible and to recall them in the order in which they thought they were presented. For any digits not remembered, they were required to make a guess. They were given three practice problems: One control, one interrupted by a ToH secondary task, and one interrupted by a digit-recall secondary task. The experiment took approximately 45 minutes to complete. Everything else was the same as in Experiment 2.

2.5.3. Results and Discussion

2.5.3.1. Move Latency

Move latency data are depicted in Figure 2.9. Control move latencies are longer within-subtasks compared to end-subtasks. Resumption move latency is longer than control move latency for each condition. However, resumption latencies are longer

within-subtasks, although considerably so when the secondary task is similar, rather than dissimilar, to the primary task. There appears to be no such effect of secondary task similarity on end-subtask resumption latencies.

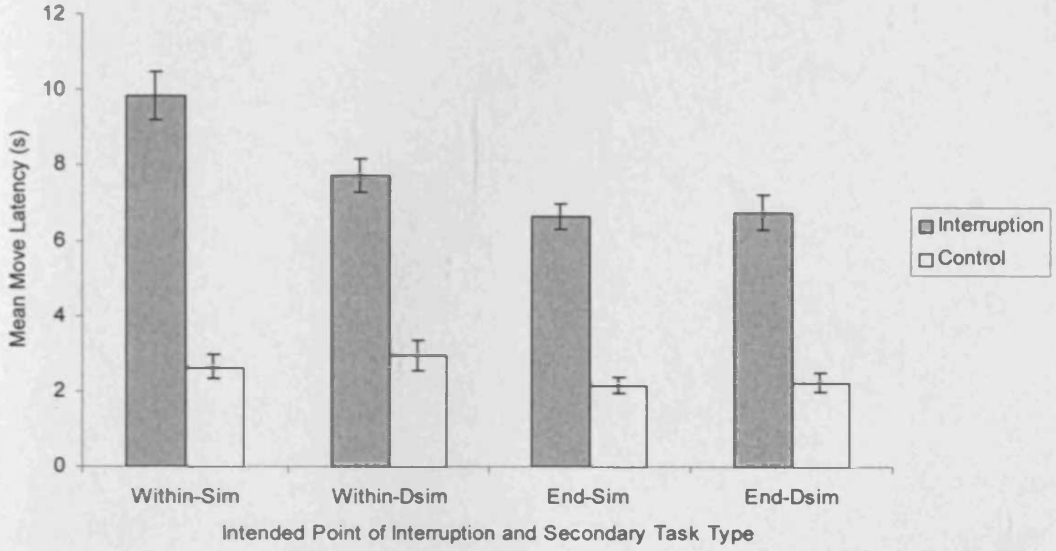

Figure 2.9. Mean move execution latency (±SE) following the intended point of interruption in interruption and control problems of Experiment 3

Move latency data was analysed using a repeated measures 2 (condition) x 2 (intended point of interruption) x 2 (secondary task type) ANOVA. This revealed significant main effects of condition, F (1, 27) = 299.58, MSE = 5.17, $p < .001$, and intended point of interruption, F (1, 27) = 22.33, MSE = 4.51, $p < .001$, but not of secondary task type ($p = .14$). Condition and intended point of interruption interacted significantly, F (1, 27) = 15.30, MSE = 1.99, $p < .001$, as did condition and secondary task type, F (1, 27) = 11.18, MSE = 1.75, $p < .005$, and intended point of interruption and secondary task type, F (1, 27) = 4.26, MSE = 3.38, $p < .05$. There was also a significant three-way interaction between all variables, F (1, 27) = 6.94, MSE = 2.99, $p < .025$. A Bonferroni simple main effects *post-hoc* analysis revealed a significantly longer resumption latency when interrupted within-subtask compared to end-subtask but only when the primary and secondary tasks were similar, F (1, 27) = 28.56, $p < .001$. A

similar effect occurred for dissimilar secondary tasks, although the difference in resumption latency was marginally non-significant ($p = 0.71$). No other variables demonstrated significant simple main effects from this analysis, and thus this was likely to be the cause of the significant three-way interaction.

For each condition, resumption latency is reliably longer than control move latency. However, given the current data, it is suggested that participants are doing something beyond simply reconstructing or reconfiguring suspended goals following un-signalled interruption. The negative impact of secondary task similarity on the length of resumption latencies adheres to the predictions of Altmann and Trafton (2002), that at least some associative priming of the suspended goal has occurred, else retroactive interference should not affect resumption latency.

It is acknowledged however, that resumption latencies within Experiment 3 were substantially longer than those set by the goal-activation model (Altmann & Trafton, 2002; Anderson & Douglass, 2001) and reconstruction parameters shared with another model of goal-directed behaviour (Anderson & Douglass, 2001). This is especially the case following within-subtask interruption. If participants are able to retrieve a suspended goal[6], execution of the move should take somewhere in the region of one to two seconds (e.g., Altmann & Trafton, 2002). In their model, Anderson & Douglass (2001) predict an average time of 1.85 s to reconstruct the representation of a suspended goal, with empirical data in the region of 7.38 s to see the move through to its motoric execution.

Although a reconstruction explanation seems to best fit end-subtask resumption latency in Experiment 3, other factors need to considered carefully when interpreting the processes involved in resuming a goal that has been interrupted. For instance, participants might be able to retrieve more efficiently the representation of a suspended goal in the end-subtask condition compared to the within-subtask condition. The amount of associative priming required to maintain the representation of a single goal is likely to be less than that required to multiple goal representations. For instance, Edwards and

[6]This estimate is just the time in between originally encoding the goal (posting the goal) and the time to execute the action. Thus it does not include the time taken to reconfigure the goal and is an average time based upon all data where encoding, storage and retrieval times were less than 5 s. The average time between reconfiguration and execution for a move corresponding to the within subtask condition of Experiment 3 is 7.38 s in their experimental data.

Gronlund (1998, p. 665) argued that, 'If the primary task lacked associative support among its task components, recovery was more difficult following an interruption that overlapped either completely or partially in the amount of information shared with the primary task'. Perhaps in Experiment 3, the 500 ms move trajectory time prior to interruption initiation might have be sufficient to cause an associative connection between a suspended goal and its context in end-subtask conditions.

However, the confidence in accepting one view over the other will be increased by understanding of the role of ToH priming cues in recovering interrupted state information (e.g., Altmann & Trafton, 2002). Unpacking this debate will form the main motivation for Chapter 3. Additionally, it is noted that the opportunity to prepare for interruption was minimal in Experiment 3, an issue that will be given experimental attention in Chapter 4. Perhaps a longer opportunity to encode a suspended goal prior to interruption (through strengthening and associative priming) might cause faster and more efficient retrieval.

2.5.3.2. Move latency and experience

Figure 2.10 illustrates mean resumption latencies for each intended point of interruption condition for each experimental half. There is an increase in move latencies within control problems at problems 13 – 24 compared to problems 1 – 12. In contrast, resumption latency appears to decrease as a function of experimental half, although this improvement is mostly apparent when a similar interruption fell within a subtask.

For further analysis, a repeated measures 2 (experiment half) x 2 (trial type) x 2 (intended point of interruption) x 2 (secondary task type) ANOVA was conducted. This revealed significant main effects of trial type, $F(1, 27) = 263.11$, $MSE = 12.53$, $p < .001$, and intended point of interruption, $F(1, 27) = 11.29$, $MSE = 9.98$, $p < .005$, although there was a non-significant difference across experimental half ($p = .18$). There were non-significant main effects of experiment half ($p = .18$) and secondary task type ($p = .25$). Experiment half and trial type significantly interacted, $F(1, 27) = 14.02$, $MSE = 11.82$, $p < .001$, as did, trial type and intended point of interruption, $F(1, 27) = 13.73$, $MSE = 4.57$, $p < .001$, trial type and secondary task type, $F(1, 27) = 8.52$, $MSE = 5.83$, $p < .010$, and intended point of interruption and secondary task type, $F(1, 27) = 7.17$, MSE

= 5.21, $p < .025$. There was also a significant three-way interaction between trial type, intended point of interruption and secondary task type, F (1, 27) = 8.86, MSE = 6.29, $p < .010$. A Bonferroni *post-hoc* analysis revealed a significant simple main effect between resumption latency in interruption conditions containing similar secondary tasks when interruption was within-subtask rather than end-subtask, F (1, 27) = 21.52, $p < .001$. The cause of the interaction appears to be secondary task type in interruption conditions: Resumption latency is longer following a secondary task that required similar processing resources to the primary task, but only if interrupted within-subtask rather than end-subtask and reduces more over experimental half.

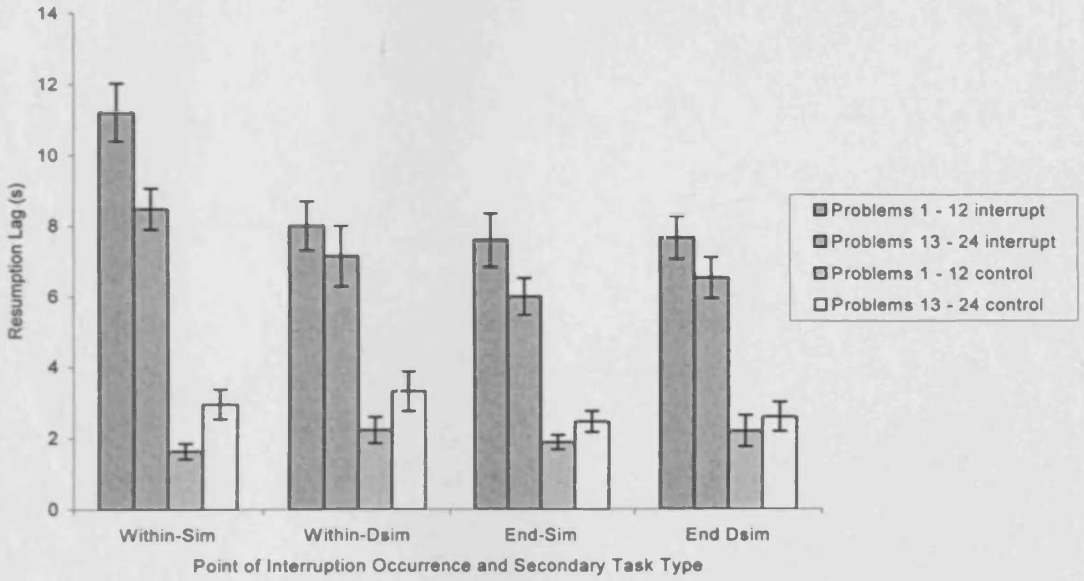

Figure 2.10. Mean move latencies (±SE) for trials 1 – 12 to trials 13 – 24 of the test phase of Experiment 3

Analyses confirmed an effect of experience on move latencies in both interruption and control conditions. The effect was however inverse across conditions in that resumption latency was reduced whereas control latencies were increased. These results lend further support to the proposal that experience in the task domain *and* experience of dealing with interruptions cause a reduction in resumption latency (Trafton et al., 2003). As for the increase in control move latencies, it might be the case that with experience of

interruption, participants learn to anticipate the occurrence of an un-signalled interruption that improves the quality of preparation (e.g., Nagata, 2003). This claim receives support from advocates of rational cognition (Payne et al., 2001; Charman & Howes, 2003). Viewing human behaviour as rational refers to the tendency for "The cognitive system to operate at all times to optimise the adaptation of the behaviour of the organism" (Anderson, 1990, p. 28). Following from this, exposure to interruption might reduce the offloading of memory demands to the environment when participants recognise that faster resumption might ensue from investing resources into trying to remember what they were doing before being interrupted (e.g., Altmann & Trafton, 2002). The emergence of such behaviour gives even more credence to the supposition that participants strengthen to-be-suspended goals, given sufficient opportunity; a claim that will be further explored in Chapter 4.

Importantly, the three-way interaction shows that with both problem solving and interruption experience, resumption latency decreases, although this effect is most pertinent in the within-subtask condition when interrupted by a similar secondary task. This lends more support to the supposition that participants are able to improve the efficiency of resumption, even following an interrupting secondary task that renders retrieval very difficult. This might be due to a refinement of processes used to suspend and resume goals using memorial (e.g., increased strengthening) and/or contextual support (e.g., the use of priming cues). The adoption of more efficient goal maintenance strategies due to increased frequency of interruption has been shown elsewhere (e.g., Monk, 2004), and thus warrants further experimentation.

2.5.3.3. Redundant moves

Two participants were removed from this data set for executing redundant moves that were more than three standard deviations above the mean of one or more intended point of interruption conditions. As illustrated in Table 2.7, more redundant moves were executed following the intended point of interruption in control problems compared to those that were interrupted. Although each condition is affected by high variance, differences are more elevated in end-subtask conditions compared to within-subtask conditions.

Frequency of redundant move mean ranks were calculated for the within-similar condition (interruption Mean Rank = 4.31, control Mean Rank = 4.27), the within-dissimilar condition (interruption Mean Rank = 3.90, control Mean Rank = 5.04), the end-similar condition (interruption Mean Rank = 4.21, control Mean Rank = 4.98), and the end-long condition (interruption Mean Rank = 4.27, control Mean Rank = 5.02). A Friedman test indicated a non-significant difference between scores for redundant moves for the eight conditions $\chi^2(7) = 9.64, p > .05$. So although descriptive data seems to suggest a larger difference between all interruption and control conditions, especially those end-subtask, statistically this is not the case. However, more moves were executed in control conditions overall compared to interruption conditions, and statistical differences may have been masked by the high variance within each condition.

Table 2.7

Mean redundant moves (±SE) within primary task ToH problems following the intended point of interruption in Experiment 3

	Within-Sim		Within-Dsim		End-Sim		End Dsim	
	M	SE	M	SE	M	SE	M	SE
Interruption	0.42	0.15	0.54	0.23	0.35	0.15	0.5	0.2
Control	0.52	0.22	0.85	0.28	0.91	0.26	0.94	0.28

2.5.3.4. Move-error frequency

Table 2.8 presents move-error frequency data for all participants who committed one or more redundant moves out of a possible 416 interruption and control ToH problems. Move-error frequency is higher in all interruption compared to control conditions, in all but the case of the within-similar condition. Overall, participants solved 319 problems error-free and made at least one redundant move in 97 problems, although this difference was marginally non-significant $Q(7, N = 26) = 13.28, p = .067$. Pairwise comparisons using a Bonferroni correct ($p = .0125$) revealed that significantly more errors were made in interrupted problems compared to control problems for, the within-dissimilar condition only, $Q(1, N = 26) = 5.57, p < .025$.

As in Experiment 2 more move-errors were made in interrupted problems compared to control problems, although this was only significantly so for the within-dissimilar condition. However, descriptive data shows that the pattern of move-errors favours a conclusion that their frequency is higher in control problems similar to Experiment 2.

Table 2.8

Move-Error Frequency Following the Intended Point of Interruption in Experiment 3

	Within-Sim	Within-Dsim	End-Sim	End-Dsim
Interruption	11	7	10	9
Control	10	17	17	16

Note. Maximum move-error frequency in each condition is 52 moves

In part, redundant move and error frequency data from Experiment 3 gives more credence to the earlier proposal that interruption might result in more efficient the execution of solution-sequences that are more efficient than that would ensue if not interrupted. Perhaps the resumption phase allows a period to assess the utility that a re-established move will have on the future solution-sequence of the problem (e.g., Speier et al., 2003). However, given that redundant moves and frequency of errors again failed to correlate with resumption latency, there still remains uncertainty about the cost-benefit ratio of such a predicted behaviour. Perhaps during the resumption period, participants simply construct a different move to that which was suspended because of being interrupted, and the new move is occasionally the correct move whereas the suspended move was not.

2.5.3.5. Secondary task performance

Table 2.9 illustrates the mean time taken to complete secondary tasks, and the mean number of moves executed within secondary task ToH problems and the mean number of digits correctly recalled in digit-recall secondary tasks. The time spent completing secondary tasks is very similar, although slightly shorter in the case of digit-recall tasks. Variance was also lower for completion times within digit-recall tasks

compared to ToH problems. The number of moves executed across ToH secondary tasks were very similar, as were the number of digits correctly recalled across digit-recall secondary tasks. However, more moves were executed in ToH secondary tasks (M = 7.22 moves) than digits correctly recalled in digit-recall secondary tasks (M = 5.27 digits).

For secondary task completion time, a repeated measures 2 (intended point of interruption) x 2 (secondary task type), revealed non-significant main effects of both variables, although the main effect analysis of secondary task type almost reached significance ($p = 0.7$). Therefore, a similar amount of time was taken to complete secondary task ToH problems and digit-recall tasks. Due to the differences in data type for performance measures in digit-recall tasks (i.e., parametric) and ToH problems (i.e., non-parametric), these data could not be analysed across secondary task type. A Wilcoxon Signed Ranks test revealed a non-significant difference between mean moves executed in ToH secondary tasks for within-subtask and end-subtask conditions, and a t-test revealed a non-significant difference in the case of digit-recall secondary tasks.

Secondary task completion times were similar, and thus retention interval between primary task suspension and reinstatement was also similar. That variance was lower for secondary task digit-recall tasks compared to ToH problems also supports that performance was more uniform between participants when completing a verbal memory task compared to when problem solving. As previously mentioned, a ramification of secondary task type was a limitation on the number of moves that could be recalled in digit-recall (maximum of seven) tasks whereas secondary task ToH problems may not have always been solved error-free within seven moves. This did not however pose any major concerns for my interpretation of post-interruption performance results, for two reasons. Firstly, secondary task set-size demonstrated no obvious increment in post-interruption performance effect sizes in Experiment 2. Secondly, although participants executed more moves in ToH secondary tasks compared to digits correctly recalled in digit-recall secondary tasks, error rate was reasonably similar in both cases. That is, participants tended to solve ToH problems, on average, 1.22 moves above that required for error-free completion, and they failed to recall correctly, on average, 1.63 digits in digit-recall tasks.

2.5.4. Discussion

Experiment 3 has provided further evidence that interrupting different points of a ToH solution-sequence can have adverse effects on the length of time taken to recover interrupted state information. Additionally, the level of similarity overlap between primary and secondary tasks can have an impact upon the size of these effects, the severity of which seems to be largely influenced by the complexity of the goal structure suspended by an interruption.

Table 2.9

Mean secondary task completion times (±SE), mean number of moves executed in secondary task ToH problems (±SE), and mean digits correctly recalled in secondary task digit recall tasks (±SE) in Experiment 3

	Similar-Within		Similar-Within		Dissimilar-End		Dissimilar-End	
	M	SE	M	SE	M	SE	M	SE
Completion Time	21.89	1.06	21.1	1.28	20.24	0.63	19.92	0.5
Number of Moves or Digits	7.22	0.32	7.21	0.25	5.21	0.19	5.45	0.19

Note. Time is in seconds

The results support the general literature in demonstrating that interruptions are more disruptive when they are co-ordinated within a task activity (i.e., when processing demands are high), compared to the end of an activity (when processing demands are lower) (e.g., Monk et al., 2002, 2004, Miller, 2002). Being interrupted within a subtask of a ToH problem is more indicative of the negative effects of interruption than being interrupted at the end of a subtask. The current results also bolster recent research that has attempted to disencumber the negative impact of interruption, by considering performance degradation as a factor of the exact point of implementation (e.g., Monk et al., 2002, 2004; Hodgetts & Jones, 2003). For instance, Hodgetts (2004) has provided much evidence for the disruptive impact of interruption positioned within the solution-sequence of a goal-directed task (the ToL). She attributes performance degradation to factors such as, the decreased base-level activation of a suspended goal because of

temporal based decay (e.g., Altmann & Trafton, 1999, 2002; Anderson & Douglass, 2001); and the reduced availability of associative activation because having to perform a secondary task (e.g., Altmann & Trafton, 2002). The results of Experiment 3 harmonise with this view, endorsing that both factors might contribute to the efficiency of resuming an interrupted goal.

The current results also reengage the debate from Experiments 1 and 2, as to whether participants are retrieving the representation of a suspended goal, or simply reconfiguring interrupted state information. I believe that the results are more supportive of the former contention. This is based upon the clear disruptive nature of similar secondary tasks on the retrieval performance of goals that are harder to re-establish (within-subtask), compared to those that are easier (end-subtask). If interruption causes participants to reconstruct or reconfigure a goal suspended because of interruption, then why should secondary task similarity have any effect on such processing unless a memorial representation of the goal is in some way involved?

Certainly, ideas used to explain switch costs in the task switching literature suggest that information processed when performing the secondary task might still be active which can potentially impede upon the speed and accuracy of performing the primary task (e.g., Allport et al., 1993, Allport & Wylie, 2000; Mayr & Keele, 2000). Such an inhibition account of previously encoded tasks has been used to try to explain switch costs in task switching studies. However, switch costs are usually within the range of hundreds of milliseconds (see Monsell, 2003). Additionally, these costs are usually higher when switch tasks are dissimilar, with lower costs for those that are similar[7]. It thus seems that an interference account (e.g., Altmann & Schunn, 2001; Nairne, 1990, 2002) provides a more plausible explanation for longer resumption latencies found in Experiment 3. Specifically, the representations of similar secondary task goals compete with and/or mask the representation of a suspended goal (e.g., Altmann & Trafton, 2002; Shulman, 1970), making it more difficult to reactivate when the visual array does not provide an obvious priming cue.

[7] Although, it is noted that similarity effects in task switching studies are usually generated from switching between tasks that demand less complex processing resources (e.g., colour naming)

That participants resume an interrupted task faster with experience, even following a similar secondary task, suggests that they are learning to cope with being interrupted; hence increasing the effort expended in maintaining the representation of a suspended goal (as noted by Altmann & Trafton, 2002, p. 66). Indeed, the current control data (as a function of problem solving experience) suggests that participants invest more time in the encoding of goals that are perhaps associated with moves that might be interrupted. Perhaps, they are strategically forming their own interruption lags by investing more time in individual move execution (e.g., Altmann & Trafton, 2002), given the nature of un-signalled interruption? As the same pattern of data is not shown in interruption conditions (i.e., resumption latencies get shorter), it seems that an improvement in resumption may only be explained by experience of dealing with interruptions (e.g., Hess & Detweiler, 1994; Trafton et al., 2003). Although the additional move data should not be discounted, which shows that less moves are required to complete interrupted problems following their resumption. This might explain, as in Experiment 2, why resumption latencies are longer than would be predicted by current estimates of goal retrieval (e.g., Altmann & Trafton, 2002; Anderson & Douglass, 2001). That is, upon retrieval, participants might engage in a process of evaluating the utility of executing the move to the efficiency of the subsequent solution sequence (e.g., Speier et al., 2003).

Finally the current experiment re-opens the debate regarding how interrupted state information is recovered. Does remembering what one was doing before being interrupted help, and if so, how much is this process reliant upon the utilisation of environmental cues? This question forms the main motivation of Chapter 3.

2.6. General Discussion

2.6.1. Summary of results

It might be useful at this point to provide a general summary of the main findings from the current experimental set. It was shown that interruption appreciably attenuated performance within ToH problems, an effect mainly substantiated by long latencies to make the first move following interruption. In Experiment 1, resumption latencies were

substantially longer for moves that were interrupted earlier within solution-sequences, as were the number of moves required to complete primary task problems (compared to uninterrupted control problems). With primary task complexity reduced, Experiment 2 revealed smaller, but still relatively long resumption latencies compared to control move latencies. Resumption latencies were however longer for moves that usually require longer latencies in uninterrupted ToH solution-sequences (e.g., Anderson et al., 1993). They were also unaffected by secondary task duration (brought about by the manipulation of complexity)(in contrast to Anderson & Douglass, 2001), but might have been influenced by task similarity (in line with Altmann & Trafton, 2002). In contrast to Experiment 1, the number of moves taken to execute problems following interruption were higher than in control problems. Experiment 3 demonstrated longer resumption latencies for moves that would be expected to be harder to reconfigure using the visual array, but only when the secondary task was similar to the primary task. Again fewer moves were required to complete interrupted problems. Experience of dealing with interruption was shown to reduce the negative impact of interruption on resumption latencies within Experiment 2 and Experiment 3.

2.6.2. The status of goals following un-signalled interruption

Experiments within Chapter 2 provide evidence that suspended goals are treated like any other declarative memory items. The representational strength of a suspended goal (assuming that it has been encoded in the first place) seems to be affected by temporal-based decay and/or retroactive interference emanated from intervening goals. Constructs such as the goal-stack (e.g., Anderson, 1990) and the intention-superiority effect (e.g., Marsh et al., 1995) cannot adequately account for the performance deficits caused by interruption in the current experiments. The current results are better supported by the constraints outlined in current models of goal-directed memory (e.g., Altmann & Trafton, 1999, 2002; Anderson & Douglass, 2001). The activation of a goal appears to decay because of interruption, weakening the encoded representation and any associative connection formed with an external priming cue (e.g., Altmann & Trafton, 2002).

The size of resumption latencies within the current experiments can also be explained by retroactive interference caused from processing similar secondary task goals (e.g., Altmann & Trafton, 2002). The interruption-similarity effects observed within the current experiments are revealing about the effects of un-signalled interruption on the fate of a suspended goal in the ToH task. Importantly, suspended goals must be encoded (if only partially) prior to suspension, even with little experience of problem solving in the task domain and of interruption, else resumption latencies would be similar (e.g., Experiment 3). The processing of similar secondary task goals might impinge upon the ability to maintain an active representation of the suspended goal and/or the efficiency of locating an associated priming cue after the primary task has been reinstated. Retroactive interference from intervening goals may decrease the likelihood of reactivating the suspended goal because of similar representational strengths causing activation competition (Altmann & Trafton, 2002). Additionally, such a secondary task may consume a limited supply of associative activation believed to be essential for goal reactivation through contextual priming from a reminder cue (e.g., Lovett & Anderson, 1996). This idea is supported by the results of Experiment 2 where it was found that secondary task content and not duration might explain inflated resumption latencies: Specifically, having performed similar intervening goals, irrespective of their quantity, impinges upon the associative activation required to maintain the representation of a suspended goal. Using the same rationale regarding how a goal might be encoded prior to interruption, experience of problem solving in the task domain (e.g., acquiring domain specific long-term knowledge) and experience of interruption (perhaps of encoding) seems to practically abolish interruption-similarity effects. Specifically, effects are short lasting, and are abated if the interrupted goal-structure requires un-intensive higher-order processing, and are practically abolished with experience in the task domain (Experiment 3). With experience, the remaining resumption latency may consist of costs imposed by interruption other than similarity (processing demands at the point of interruption, insufficient encoding, and so on).

Other factors may slow the rate of temporal-based decay and encourage the use (or refine the efficiency) of preparation strategies (e.g., associative priming using cues). For example, anticipation that retention intervals, or secondary tasks, will be relatively

short might encourage participants to invest more effort in the encoding of goals (e.g., Nagata, 2003). Thus the claim by Anderson and Douglass (2001) that goal storage is stochastically determined by retention interval might only apply to certain contexts, for example, when participants know that the stimulus array is always available to support goal reconstruction (as is not the case when interrupted). Certainly, Experiment 2 demonstrated that the actual length of the retention interval had no effect on the time taken to resume an interrupted goal (in direct contrast to Anderson & Douglass, 2001). Recent work has also revealed that humans opportunistically trade strategies of offloading mnemonic demands to the visual array for more memory-intensive processing, to reduce the cost of processing task relevant information especially when that information is made difficult to access (Fu & Gray, 2004; Gray & Fu, 2004)

2.6.3. Re-establishing a suspended goal

A pertinent purpose of Chapter 2 was to address how participants resume a suspended goal following un-signalled interruption? Do they: Retrieve the representation of the suspended goal? (Altmann & Trafton, 2002); reconstruct the suspended goal using the visual array? (Anderson & Douglass, 2001); or reconfigure the problem space? Perhaps they use a combination of all strategies?

The results of Experiment 1 suggest that when the number of moves between the suspended goal and the overall goal-state is between seven and fifteen, participants choose to reconfigure the problem space, forming a new solution-sequence. In contrast to the predictions of Anderson & Douglass (2001) participants do not seem to reconstruct only the move that is suspended following a retention interval (or in this case a secondary task). However, choice of reconfiguration as a resumption strategy might not result from abandoning suspended goal storage (as would be suggested if the length of the retention interval was predictable, Anderson & Douglass, 2001). Rather, the difficulty of reconstructing a single goal given the complexity of the remaining problem space might encourage the use of reconfiguration, when retrieval might be possible. Note that within Experiment 1, participants must have encoded the suspended goal before being interrupted, given the fact that un-signalled interruption could only follow the attempted execution of that goal-move. In fact, this period of strengthening could be likened to an

interruption lag (Altmann & Trafton, 2002). However, given that resumption latencies for earlier subgoals were longer than their respective initial move latencies, and control move latencies, participants cannot have been retrieving (or efficiently retrieving) the representation of a suspended goal. Additional support for this contention is evident within redundant move data: Redundant moves were higher following interruption for earlier subgoals. When reconfiguring a ToH problem, participants might choose a different – and perhaps less efficient – initial move that consequently leads to a longer solution-sequence. Thus, for Experiment 1 at least, resumption through reconfiguration of the problem space explains the data for earlier ToH problem subgoals.

Experiment 2 showed that when reducing the complexity of primary task ToH problems, and positioning interruptions within subtasks rather than subgoals, resumption latencies are considerably reduced. Resumption latencies however formed a distinctly similar pattern to shorter control move latencies. Perhaps, participants were again attempting to retrieve the representation of suspended goal, but the efficiency of this process is in some way hindered by the interruption context. Given that fewer redundant moves were executed following the intended point of interruption in interruption conditions, it could be argued that participants use resumption as an opportunity to evaluate the utility of performing a retrieved goal (e.g., Speier et al., 2003). However, Experiment 3 showed inflated resumption latency due to similar secondary tasks; an effect moderated by the exact point of interruption. That is, if the visual array did not support efficient primary task resumption through the availability of priming cues, the time taken to make the first move after performing a similar secondary task is increased.

Secondary task duration does not seem to effect resumption performance within the ToH (in contrast to Anderson & Douglass, 2001), and other demanding tasks (e.g., Bailey et al., 2000; Gillie & Broadbent, 1989). Rather the strategies employed to support the recovery of interrupted state information are better determinants of resumption latency.

Some models of strategy selection assume that retrieval is always attempted upon being able to return to suspended item (e.g., Siegler, 1987, 1988, Siegler & Shrager, 1984). If the initial retrieval attempt is met with failure, perhaps because the goal has not been sufficiently strengthened, only then will the deployment of other strategies occur.

Although, Schunn, Reder, Nhouyvanisvong, Richards & Strffolino (1997) partially agree with this view, they also suggest that there must be a feeling-of-knowing for a goal if retrieval is to be attempted at all.

For all experiments, participants might attempt to retrieve the representation of a suspended goal. Retrieval might ensue, or retrieval might fail. Either way, the chance efficiently retrieving the representation of a suspended goal is lower when primary task demands are high, thus encouraging reconfiguration of the interrupted state (e.g., Speier et al., 2003). Simply put: External representations might not efficiently support the reactivation of a goal representation through associative priming, at least when task demands are too high (e.g., Zhang, 1997). Conversely, too much strengthening of goals prior to interruption might increase the interference level (e.g., Altmann & Trafton, 2002). If the activation level of the most active distractor goal (e.g., move two in the strengthened solution-sequence for subgoal one) is higher than that in which is to-be-resumed, the chances of retrieval with the support of associative priming might be reduced (e.g., Altmann & Trafton, 2002). Additionally, resumption of a suspended goal when the reinstated disc configuration represents a complex problem space speaks to the lesioned version of the goal-activation model (Altmann & Trafton, 2002); retrieval failure occurs, but at a cost higher than is predicted by the model in its current form.

Could resumption latencies in excess of that expected if reconstructing a suspended goal reflect an amplified switch cost? As discussed in Chapter 1, switch costs are usually in the region of hundreds of milliseconds rather than the many seconds exhibited here (e.g., Altmann & Trafton, 2004), and are rarely sizeable when switching between two similar although simpler tasks (see Monsell, 2003, for a review). Also, within task switching studies, participants are aware that they will finish one task before starting another, without having to remember anything about the first task to support performance of the second task. A purely task switching account is therefore unlikely.

The goal-activation model (Altmann & Trafton, 2002) best fits the current results. It is chosen over other models given, the lack of support for a negative effect of secondary task duration on post-interruption performance, the importance of primary task cues in supporting resumption (and their availability within the stimulus array), and the negative impact of interference generated by secondary tasks. A problem for the form of

the goal-activation model is that no explicitly cued interruption lag was manipulated, but (1) when a suspended goal could be associatively primed (Experiment 1), retrieval did not always result and (2) retrieval like behaviour seemed to occur when the opportunity to prime was reduced (Experiments 2 & 3). So, the utility of an interruption lag as an essential precedent to interruption to support the retrieval of a suspended goal is questionable, as is shown in related research (e.g., Altmann & Trafton, 2004; Hodgetts & Jones, 2003; Miller, 2002; Trafton et al., 2003). For example, Experiment 1 showed that when participants had the opportunity to strengthen the representation of a suspended goal (through planning at least the first move) and the constraint of having to associate it with an externally based priming cue, they exhibited resumption behaviour that was not characteristic of retrieval. However, when primary task complexity was reduced, and the opportunity to exploit associative priming was reduced immediately before interruption, resumption was much faster (Experiments 2 and 3). This seemed to be a result of learning to recover from interruption by using the environment as a source of declarative support; only when the primary task was easier to solve. Together with evidence for fewer redundant moves, results of Experiment 2 and Experiment 3 demonstrated that retrieval is possible without an interruption lag; the overhead of resumption being the amount of time required to confidently execute the re-established move.

2.6.4. Future directions

The exact point of interruption affects the resumption efficiency of a goal that is suspended because of having been interrupted. I suggest that this is strong evidence that priming cues – mainly their associative strength with the suspended goal – largely influence resumption performance. It is for further experimentation to decipher the exact role of cues in priming a suspended goal for resumption (Chapter 3), and whether resumption efficiency can be increased from allowing an brief period in which a goal may be encoded for retrieval (e.g., Altmann & Trafton, 2002). However, given the exploratory nature of this thesis within these early stages, I want to establish the boundaries of suspended goal resumption in the absence of overtly manipulated resumption aids (the utility of which will be explored in Chapter 4). Finally, although secondary task duration was shown to have no effect on performance of an interrupted

primary task, perhaps the length of secondary tasks in the current experiments was insensitive to such effects. Similarly, Gillie and Broadbent (1989) reported no effect of retention interval on performance within an interrupted task when secondary tasks were between 30-seconds and 2.75 s. Recently, some have noted that longer secondary tasks may not be sensitive to exposing the disruptive effects caused by imposing shorter retention intervals (e.g., Hodgetts, 2004; Monk, Boehm-Davis & Trafton, 2004b).

2.6.5. Summary

Experiments 1 – 3 have provided preliminary evidence that a promptly suspended goal is encoded prior to interruption, but the point of interruption, and to some extent, the secondary task, affects the strength of the representation. Factors that seem to influence the efficient retrieval of a suspended goal are: The difficulty in forming a prompt representation of the suspended goal due to primary task complexity; the similarity (and not the duration) overlap between primary and secondary tasks; and the exact point of interruption within a solution-sequence.

Chapter 3

EMPIRICAL SERIES 2

Priming cues and the recovery of interrupted goals: An exploration of 'cue availability'

3.1. Overview

Altmann and Trafton (2002) argue that for an interrupted goal to be retrieved from memory it must be rehearsed before suspension and associatively primed with a cue to support its reactivation. Four experiments addressed these claims within interrupted ToH problems. Secondary tasks were un-signalled to minimise the opportunity to rehearse, and cue availability was manipulated by changing disc properties in reinstated problems. Maintaining the spatial location of ToH discs between suspension and reinstatement of the primary task caused the fastest resumption (Experiments 4 & 5), but only if a disc prime was distinguishable by an individual colour (Experiments 6 & 7). These effects were augmented with experience of interruption. Retrieval of a suspended goal seemed dependent upon the effective use of disc-colour as a priming cue, although move execution was slower if the problem solution-sequence had changed. ACT-R theory and the *guided-search model* of selective-attention (Wolfe, 1994) offer a conceptual understanding of how associative priming might support the resumption of a suspended goal.

3.2. Introduction

3.2.1. General introduction

Goals often have to be suspended and resumed. The use of priming cues: objects located within the current mental or environmental context (e.g., 'post-it' notes); can serve to remind us of the relevance of a suspended goal. Priming can occur if a cue is selectively associated with a goal (e.g., underlining a sentence in a report) although more permanent priming occurs from learned pairing of the same prime and goal (e.g., the colour green to signal 'go!'). Efficient use of priming cues can increase: the likelihood of resuming a suspended goal (e.g., McDaniel et al., 2004); speed of resumption (e.g., Detweiler et al., 1994); and future use of certain cues (e.g., Goolsby & Suzuki, 2001).

If a goal is interrupted, then linking it to an available priming cue prior to suspension might be the only way to enable its future reactivation (e.g., Altmann & Trafton, 2002). The link between the cue and a goal is mediated by a limited supply of *source* (or *associative*) *activation*: activation generated by the current context (e.g., Lovett & Anderson, 1996). Following secondary task completion, detection of the primed cue can provide source activation to reactivate the suspended goal from its decaying state. Other than cue availability, how and why a priming cue is initially attended, and subsequently chosen and used as a reminder device, is a subject in need of empirical investigation. Chapter 3 offers a starting point for such investigation, using selective-attention theory (e.g., Wolfe, 1994) to explain how cues are selected and processed, and ACT-R theory (e.g., Anderson & Lebiere, 1998) to suggest how a goal is retrieved because of priming. The chapter broadly addresses (1) attributes that might influence cue selection and (2) the role priming has on the speed of resumption and subsequent performance in a task that has been promptly interrupted.

Using the ToH problem solving task, a task susceptible to disruptive effects of interruption (see Chapter 2 of this thesis), the experiments conducted within Chapter 3 speak to the above issues using priming methodology. Interruptions were un-signalled and therefore limited the amount of encoding an interrupted goal could receive. Secondary tasks were cognitively demanding likewise to impede rehearsal of the suspended goal. Disc properties (priming cues in the ToH, Altmann & Trafton, 2002)

were manipulated only after an interrupted problem was suspended. Thus, the effects of strategic or 'selective-priming' could be measured by testing for differences in performance between a condition where *no-change* had occurred to disc properties to conditions where different levels of *change* had occurred. Faster resumption of an interrupted task in the no-change condition compared to change conditions would indicate an effect of associative reactivation using ToH discs as priming cues, and may be augmented further with experience in the task domain.

3.2.2. The use of retrieval aids in prompting a suspended activity

There is an abundant research showing that retrieval aids can cause faster and more efficient recovery of an interrupted activity (e.g., Franke et al., 2003; Harrison et al., 1995). For instance, keeping a primary task in focal awareness when performing a secondary task reduces the time required to resume a goal in an interrupted Tower of London (ToL) problem (Hodgetts, 2004). Memory for intentions is almost faultless with subtle 'on-screen' reminders such a red dot to indicate an unfulfilled intention (McDaniel et al., 2004) and highlighting a suspended goal when it demands response (Chung & Byrne, 2004). Nevertheless, attention-directing cues and control over the use of reminders (e.g., placing of a mouse cursor over relevant information) do not always cause faster resumption of an interrupted task (e.g., Czerwinski et al., 2000). A goal that is difficult to prime (perhaps because of high task demands at the time of interruption) might be more difficult to remember, even with the support of a reminder cue (e.g., Czerwinski et al., 2001b).

Why then do some retrieval aids support more efficient resumption of an interrupted task than others, even when the suspended goal is cognitively demanding? Conceivably, limiting the amount of information to only that required to recover interrupted-state information might cause more efficient retrieval (e.g., van Dantzich et al., 2002). This may cause a 'freeing-up' of resources that can be dedicated to the processing of priming cues (e.g., Lavie & Tsal, 1994), consequently causing faster resumption and less frequent forgetting of suspended goals. Additionally, not all cues have to come from the external visual array. Knowledge of the task domain (i.e., internal representation) and sufficient preparation (e.g., through training) might serve to reactivate

the representation of a suspended goal in the absence of an external cue (e.g., Zhang, 1997).

Conversely, why do some cues fail to be noticed, even when explicitly manipulated to capture attention within the task display? Many tasks by their nature sometimes demand high levels of attention and at others times do not (e.g., Adams et al., 1995; Dismukes et al., 1998). For example, whilst typing an earlier sentence in this chapter, I had to focus attention resources to a computer screen in order to inspect the outcome of my typed keystrokes. Only upon finishing the sentence did I notice a flashing icon prompting me to save earlier changes made to a database file (which from experience was likely to have been flashing for some time). The task of typing the sentence demanded a high level of focused explicit attention (e.g., Neisser, 1976; Posner, 1980; Sanford & Garrod, 1981). The save prompt (although an important feature of the visual array) was set in peripheral focus (e.g., Sanford & Garrod, 1981). Consequently, it was not attended immediately, probably because of the demands associated with the primary task of typing. The sometimes ineffective attention-directing properties of cues are also found within: studies of post-completion error (e.g., Chung & Byrne, 2004); prospective memory (e.g., Guynn et al., 1998), and applied studies of interruption (e.g., Lahlou et al., 2002). Perhaps then, task demands at the point of interruption might interact with the efficiency of an attention-directing cue prime, such that demands will influence the effectiveness of a primary cue.

3.2.3. External cues: a route to internal representations?

Attending to external objects is intrinsic to performance in most goal-directed tasks, with objects serving to guide, constrain and even determine behaviour (e.g., Larkin, 1989; Reisberg, 1987). Rather than being thought of as merely a channel to representations contained within the internal mind, external representations provide knowledge of the identity and structure of the environment, as well as rules, and relations between physical configurations (e.g., Zhang, 1997).

The ToH is an example of a task in which the processing of external cues is required to guide problem solving behaviour, if only to remind oneself of a goal within the solution-sequence (e.g., Anderson & Douglass, 2001; Anzai & Simon, 1979). To

gain knowledge of the solution-sequence, internal representations have to be formed in the first place, which requires attending to external objects to form associative connections. In addition, when external representations are not available, say during performance of an interrupting secondary task, internalisation alone may not be sufficient to remember a goal (e.g., Anderson & Douglass, 2001; Zhang, 1997). In such a case, attending to an external representation of the goal, a cue perhaps that gives the goal identity may cause reactivation of the representation of that goal (e.g., Clifford & Altmann, 2004).

Given the crucial role of external representation in the re-establishment of a previously formed goal, it is important to understand how an object that might serve to reactivate the representation which is processed in the first place. In the following sections, I discuss conceptual and empirical support for the likely role of memory and selective-attention as potential mediators for this process.

3.2.4. The role of rehearsal

Rehearsal is a pervasive strategy for maintaining information in the short-term (e.g., Atkinson & Shiffrin, 1968) and is commonly regarded as a prerequisite to encode the representation of an interrupted goal (e.g., Altmann & Trafton, 2004; Trafton et al., 2003). That rehearsal may mitigate the negative effects caused by interruption is not a new phenomenon (e.g., Gillie & Broadbent, 1989; Storch, 1992), although interpretation of what is being done at rehearsal should sometimes be met with caution (e.g., McFarlane, 2002). For example, it may require encouragement to rehearse to lead to even moderately faster resumption times (e.g., Detweiler et al., 1994). When faced with the opportunity to rehearse, people might instead adapt to task constraints to minimise metal load at the point of interruption, by relying on externalisation to support the recovery of interrupted state information (e.g., Anderson & Douglass, 2001). For instance, participants can adapt to preparation periods by choosing to use memory aids instead (e.g., pointing to interrupted information, Miller, 2002).

3.2.5. Priming cues and ACT-R theory

The role of cues in causing the reactivation of a suspended goal is grounded in ACT-R theory (e.g., Anderson & Lebiere, 1998). Taking a Bayesian approach, history of use and the current context determine the relevance of a particular cue to a goal (e.g., Anderson & Milson, 1989), and whether it will cause the appropriate production rule to fire (e.g., Lovett & Anderson, 1996). A production rule is a mechanism that serves to modify a goal by acting upon condition (e.g., the current context) action (e.g., perform a goal) pairings. The fewer resources available to determine whether a particular goal is relevant to task performance, the less likely production rule firing during the first processing cycle, meaning that the goal will have to undergo a process of conflict resolution. Conflict resolution applies to cueing in goal-directed behaviour, and involves deciding upon which production rule to select based upon: degree of match; production strength; and data refractoriness.

A production is more likely to be selected if the degree of match between a cue and goal representation is high; the process might be abandoned if the match is inadequate. Abandonment might occur, for example, upon encountering an unfamiliar cue. Whether a production is selected might also depend on how often the same condition-action pairing has led to an accurate retrieval in the past: its *history of use* (e.g., Lovett & Anderson, 1996). If one condition (or cue) has successfully caused reactivation of a suspended goal enough times in the past, conflict resolution will respond to that condition causing the production rule to fire more efficiently in the future (e.g., Anderson & Schooler, 1991). Data refractoriness refers to the idea that the same object cannot serve in two sources of information simultaneously. For cueing effects in goal-directed memory, data refractoriness is perhaps best considered by the phenomenon of pattern matching: the general finding that an object is responded to faster if it represents a pattern that was activated previously (e.g., Neely, 1977). Take the example of a single cue presented among a 'busy' array of cues prior to goal suspension (e.g., the disc configuration in a ToH problem). According to the logic of pattern matching, a production is more likely to fire if the cue remains in the same location of the display and maintains the same relative relationships (e.g., distance) to other cues (e.g., Reder & Anderson, 1980).

ACT-R theory makes strong predictions regarding how a goal is retrieved using production rules and conflict resolution, although it might be flawed in assuming a goal is automatically represented before suspension. The idea that goals are stored and retrieved in a serial order from a hierarchical stack, immune to the effects of interference and decay, has recently become a subject of controversy (e.g., Anderson & Douglass, 2001). For instance, there is growing empirical evidence for goal forgetting (e.g., Byrne & Bovair, 1997; Chung, 2004) and goal-selection errors (e.g., Anzai & Simon, 1979; VanLehn, 1991). Recent models of goal-directed behaviour address this issue by suggesting that memory for a goal is like that of all other declarative memories (Altmann & Trafton, 2002; Anderson & Douglass, 2001). Goal memory is not perfect, and is disrupted by factors that also impede retrospective memory performance, such as interference, decay and forgetting. It is to these models that I now turn, to consider why priming of a suspended goal may essential if it is to be reactivated in the future. A sufficiently primed cue may act as resource to reduce conflict within production rule criteria, thus increasing the likelihood of goal retrieval (e.g., Altmann & Trafton, 2002).

3.2.6. Priming cues and the suspension and resumption of interrupted goals

The function of priming cues in the reactivation of a goal suspended because of interruption has surprisingly received little empirical investigation (e.g., Altmann & Trafton, 2002; Trafton et al., 2003) and conceptual deliberation (e.g., Altmann & Trafton, 2002). It is generally accepted that for a suspended goal to be retrieved from memory, it must be identifiable by a perceptually available object, but also has to be prepared for retrieval (e.g., Anderson & Douglass, 2001; Trafton et al., 2003).

In their goal-activation model, Altmann and Trafton (2002) use ACT-R theory (e.g., Anderson & Lebiere, 1998) to conceptualise about the processes involved in goal suspension and resumption. They agree that a goal can only govern behaviour if it is the most active in memory, and that it will decay as a function of the time since it was last processed (e.g., Anderson & Douglass, 2001). As in ACT-R, this is a measure of the goal's base-level activation (BLA). In the case of an interruption, Altmann and Trafton (2002) move away from the ACT-R idea of a hierarchical goal-stack in which goals are encoded and retrieved perfectly in a LIFO order (e.g., Anderson et al., 1993). Instead,

they stress that a cue must be associatively primed with the representation of a suspended goal during a rehearsal period that immediately precedes interruption initiation (the 'interruption lag'); else retrieval will be impossible, and the goal will be forgotten. The priming cue must also be easily detectable upon primary task reinstatement, and should only be associated with the suspended goal and no (or few) others. Upon its detection, the cue (the current context) will utilise source activation to produce a rapid reactivation boost to the decaying BLA of the goal, causing the appropriate production rule to fire and the goal to be retrieved. As such, Altmann and Trafton ensue that the role of priming in an interrupted task is to support the reactivation of a goal that has been rehearsed immediately before its suspension. If rehearsal is prevented, say because of not allowing an interruption lag, then cueing as an associative priming mechanism should be rendered ineffective. Consequently, retrieval will be impossible and reconstruction of the goal using the external visual array will be the only route to resumption.

Similarly, Anderson and Douglass (2001) suggest that retrieval of a suspended goal is possible if it is reactivated from a decaying state, and they too deviate from the goal-stack assumption of a perfect goal memory. Unlike the goal-activation model, Anderson and Douglass (2001) predict little time cost of goal storage; rather the greatest time cost will be exhibited at goal retrieval. They suggest that in anticipation of a long retention interval between goal suspension and resumption, people will choose to abandon goal storage, and pay the cost of reconstructing the goal rather than making a fallible retrieval attempt from memory. Taking a rational analysis approach (e.g., Anderson, 1990), and support from their empirical data, Anderson and Douglass (2001) suggest a high cost of goal storage because of ineffective memory for goals at long retention intervals. Participants invest resources in reconstructing the suspended goal, rationalising that the re-established goal will be more efficient than the output of a potentially flawed retrieval attempt. Taking this view, Anderson and Douglass (2001) suggest little or no effect of priming cues on the speed of resuming an interrupted goal.

3.2.7. Attending to cues in a goal-directed task: what is processed and how?

Attention is arguably the initial process to which information is subject but there is no consensus on how attention operates. William James (1890) classically stated that

'Everyone knows what attention is' (p. 381), but disentangling between attention as a conscious controllable process and one that is obligatory is a locus of continued debate (see Lavie & Tsal, 1994, for a review).

There is however growing agreement on how we selectively attend to visually presented information (see Wolfe & Horowitz, 2004, for a recent review). Visual selective-attention is thought to affect how we initially process information in problem solving tasks (e.g., Zhang, 1997, Zhang & Norman, 1997) and is likely to be highly applicable to the paradigm used for empirical investigation throughout Chapter 3. To unpack visual selective-attention with regards to cue selection in goal-directed memory, early yet still influential ideas of attention in a broader context need to be considered first.

Attention theorists often distinguish between 'early-selection' and 'late-selection' when conceptualising about how information is attended. Early-selection refers to the processing of information based upon physical properties (e.g., colour, loudness, and so on); properties that have to be selectively attended. Late-selection refers to the processing of information after meaning has been derived from objects (e.g., categorising a visual object as a 'chair'). One of the earliest ideas of attention, the well-known 'Filter Theory', maintained that information is attended early on, in parallel, and is held in an unlimited capacity store (Broadbent, 1958). According to Broadbent, physical properties are selectively stored in a temporary buffer, and only some information filters through to higher-level processing. Filter Theory has influenced the development of attention theories since (e.g., Treisman & Gelade, 1980), but it has received criticism due to evidence that information might be attended only late in the selection process (e.g., Duncan, 1980; Deutsch & Deutsch, 1963; Norman, 1968). The idea of late-selection is that all information is processed at the semantic level with only highly relevant information selected for response.

How we selectively attend to visual objects in goal-directed tasks is largely informed by two dominant theories of attention, both heavily influenced by the aforementioned attention theories.

The first, feature-integration theory (Treisman & Gelade, 1980) maintains that information can be attended early on, at a preattentive stage. It is an early-selection theory in the sense that it allows for the generation of basic features during the

preattentive stage, but also a late-selection theory because only certain inputs are allowed further processing after being attended. Although very influential, feature-integration theory has received some criticism because of empirical evidence that object features can be selected preattentively (e.g., Maljkovic & Nakayama, 1994, 2000; Wolfe, Yu, Stewart, Shorter, Friedman, Hill & Cave, 1990).

A more plausible contribution to explaining goal-directed selective-attention is the guided-search model (Wolfe, Cave & Franzel, 1989; guided-search 2.0: Wolfe, 1994). Whilst maintaining a two-stage processing network, guided-search allows more flexibility than most other ideas of selective-attention, in that important objects can be selected for further processing at the preattentive stage. In guided-search, features are selected based upon (1) their salience and uniqueness in distinguishing a target from a distractor and (2) their history of success in resulting in the correct response to a target (Wolfe, 1994). Only features that are selected preattentively are passed onto exclusive feature maps (e.g., one for colour, one for shape, and so on) that allow further processing. The model uses ideas of bottom-up activation and top-down activation to explain how the feature selection process operates. If a target item is unusual compared to distractors, then its activation strength caused by bottom-up perceptual processing should be stronger than that of the distractors, causing its retrieval. However, if the target is similar to the distractors, top-down selectively deployed activation is required for its selection, by tuning a feature channel to respond to one feature more than other feature channels respond to their particular features. In this case, selective-attention can be viewed as a learning process: retrieval success using a particular cue increases the chances of its future use. In both cases, the system learns to respond to a target based upon its history of use: that is, how often its activation level has exceeded that of the mean activation level of all distractors.

The guided-search model makes three solid predictions regarding how cues are selectively attended, and might therefore inform an understanding of how such processes operate in a goal-directed task suspended because of interruption.

Firstly, a number of cue features have high predictive values in terms of their salience (and hence the likelihood they will be selected). Empirical research supports this contention, showing the most salient features to be those such as: colour (e.g., Bauer,

Jolicouer & Cowan, 1996; Maljkovic & Nakayama, 1994, 2000), and size (e.g., Treisman & Gormican, 1988). In contrast less salient features are those such as: novelty (e.g., Wang Cavanagh & Green, 1994) and colour change (e.g., Theeuwes, 1995). For example, it is relatively easy to detect a unique colour-singleton within an array of different yet homogeneously coloured distractors (e.g., Goolsby & Suzuki, 2001). To detect one repeated colour prime (e.g., green) against an array of other different coloured yet heterogeneous distractors (e.g., blue, yellow and red), may however require experience of successful priming using the exact feature cue. Such a prediction lends well to ACT-R theory (e.g., Anderson & Lebiere, 1998), in that history of use can increase the activation of a representation, thus leading to its faster retrieval through reactivation at future encounters.

Secondly, like production rule firing in ACT-R theory, the guided-search model makes predictions regarding the time dedicated to locating a target to reactivate before terminating the search. Only upon finding a target that exceeds an activation threshold within a certain time parameter will the search stop (similar to the idea of an interference level in retrieval of a suspended goal, e.g., Altmann & Trafton, 2002). If during the search, the system realises that past searches have not taken as much time before a successful target 'hit', the search will stop: the target thereafter will be rendered as forgotten. Similarly, if search results in selection of the wrong target, or the selected target is associated with an inappropriate goal, then the correct goal will not be retrieved.

Thirdly, the guided-search model has implications for how cue availability might cause the reactivation of a suspended goal. If a cue is sufficiently attended prior to a retention interval, but is flanked in the feature space by distractors sharing similar features, it is less likely to be attended when a response is required (e.g., Bauer et al., 1996). For example, if a circle is selectively attended because of a particular colour (e.g., green), but is contained amongst an array of similarly or mixed colour circles (e.g., blue, yellow, and so on), detection will be difficult. The chances of retrieving a goal through associative priming are therefore reduced, compared to if the cue was more readily available (e.g., Altmann & Trafton, 2002). If however, the cue is not hidden amongst such distractors and/or is easier to detect; the target-distraction criteria lowers, the cue is

more likely to be detected, and retrieval is more likely to occur (e.g., Treisman & Gelade, 1980).

Finally, the attentional processes involved in selecting a cue as a prime and responding to it upon detection are in many ways like the processes required to remember a suspended goal (e.g., Awh, Jonides & Reuter-Lorenz, 1998; Smyth & Scholey, 1994). In both cases, representations have to be, encoded, stored, and retrieved. The major impact of attentional theory on re-establishing a suspended goal after interruption comes from the nature of external representations (e.g., Zhang, 1997). Specifically, attending to a priming cue should be much faster than trying to re-establish a suspended goal, and may be the only route to its reactivation (Altmann & Trafton, 2002). Choice of cue to prime a suspended goal might involve a process of selecting one that is most likely to capture attention in the future, given the demands imposed by the task domain (e.g., Lavie & Tsal, 1994). When processing a complex goal in a goal-directed task before its prompt suspension, associating that goal with an easily distinguishable external cue might ensure its redetection in the future, thus increasing the likelihood and speed of goal retrieval. Associating the goal with a cue that is, lower in salience, requires complex processing, or one that is more difficult to detect following a demanding secondary task, might make cue based retrieval harder or render it impossible.

3.2.8. The current experiments

Participants were only interrupted after making the second move in the solution-sequence of ToH problems, in order to obtain a larger sample in the maximally informative conditions of Experiment 3. Exercising control over the exact point of interruption ensured that un-signalled interruption would precede only a move of equal difficulty, irrespective of the start-state and goal-state of the ToH problem in which it is embedded. With regards to cue availability, the relative spatial location of discs and the difficulty in using each as a cue was matched in suspended and reinstated ToH problems. For example, the disc most likely to be chosen as a cue for the suspended goal is always disc 4 (the smallest disc) and as such was always on top of a stack of discs (thereby making it accessible). Participants were not trained in the use of extracting cue properties or given instructions to attend to cues. This was done to allow resumption strategy to

develop with the constraints imposed by the task domain, for example, experience of being interrupted.

Within the current experiments, disc properties were only ever manipulated after the primary task was suspended: discs could be in a different configuration or presented in different colours when the task was reinstated. This methodology was chosen so not to explicitly encourage participants to attend to certain cues, but to investigate how cues are used (if used at all) when re-establishing a suspended goal (by measuring speed of resumption).

In all but Experiment 5, conditions were included in which disc properties within only the reinstated current-state or only the reinstated goal-state could have changed. Based upon resumption latency data from experiments presented in Chapter 2, these conditions serve to explore the cost of resumption, above that of time taken to re-establish the suspended goal. If retrieving a goal through associative priming generated by a disc in the current-state of the visual array, there should be no difference in resumption latency for conditions where disc properties in the goal-state are the same or different. Differences would highlight the potential influence of evaluating a re-established move before it is executed: a common behaviour in situations of uncertainty (e.g., Reason, 1990). Implications are that actual retrieval and reconstruction latencies may be masked by this cost of 'goal-directed enquiry', but subtle manipulations may provide some insight into the size of such costs.

The principle motivation was to investigate performance within interrupted ToH problems that had undergone different degrees of change to disc properties in the reinstated problem space, and therefore the visual array did not change at any point within control problems. This was to (1) avoid confounding uninterrupted ToH performance because of having to deal with a change to the visual array and (2) to minimise the expectancy of change to disc properties between suspension and reinstatement of interrupted ToH problems. Nevertheless, following the intended point of interruption, all individual moves remaining within the solution-sequences of interrupted and control problems were identically matched in terms of their difficulty to execute. For example, if solved error-free, the first move following interruption always

required movement of disc 4 (the smallest disc) to a temporary destination peg so that the second LOOP disc could be moved to its goal destination peg.

3.3. Experiment 4

3.3.1. Introduction

Twenty-four ToH problems requiring nine moves for error-free completion were to be solved. Eight were interrupted once by an un-signalled digit-recall task; all were matched with an identical control problem that was not interrupted. Following secondary task completion, the suspended ToH problem was reinstated but the location of discs within the problem space could have changed. Specifically, there were changes to the congruence between disc configurations in the suspended and reinstated current-state and/or the suspended and reinstated goal-state. In two conditions, reinstated disc configurations were either completely congruent (same current-state same goal-state, referred to as SC-SG hereafter) or completely incongruent (different current-state different goal-state, referred to as DC-DG hereafter) with those that were suspended. In a further two conditions, the current-state disc configuration remained congruent and the goal-state disc configuration was incongruent (referred to as SC-DG hereafter) or the goal-state disc configuration remained congruent and the current-state was incongruent (referred to as DC-SG hereafter). Figure B1 of Appendix B provides examples of each condition.

Interruption was un-signalled and thus participants should have little or no opportunity to rehearse the suspended goal prior to interruption. According to Altmann & Trafton (2002), maintaining only the availability of a single priming cue (the suspended move disc) will be sufficient to cause associative reactivation of a suspended goal in the ToH task, but only if rehearsed prior to interruption. Thus, retrieval via associative priming in Experiment 4 should be seriously impaired (perhaps impossible), meaning the goal will have to be reconstructed using the visual array (e.g., Anderson & Douglass, 2001). Time to reconstruct should not be effected by changes to disc location

within the visual array, because the reinstated solution-sequence retains the same level of difficulty between moves. Therefore, all conditions should yield similar resumption latencies.

Alternatively, if associative priming occurs at periods other than immediately before interruption initiation (e.g., during an initial planning period), or is relatively automatic, resumption latency should be shortest in the SC-SG condition. The guided-search model supports this prediction, but only if participants are able to extract a salient cueing feature from the visual array sometime before goal suspension to guide attention toward upon task reinstatement (Wolfe, 1994). Disc-location may serve as such a salient feature (e.g., Tsal & Lavie, 1993), although effects may be moderated by other salient disc properties such as features (e.g., Goolsby & Suzuki, 2001; Maljkovic & Nakayama, 1994, 2001). Conceptual guidance from the guided-search model together with empirical evidence of strategy shifts to extract salient information from an interrupted scene (e.g., Miller, 2002) also impart that resumption performance could improve because of experience of associative priming. That is, resumption latency will reduce if participants learn that attending to a cue property (or properties) results in faster reactivation of a suspended goal.

Finally, whether retrieved or reconstructed, the time to resume an interrupted ToH goal might be effected by the time taken for goal-directed enquiry (e.g., Speier et al., 2003). For example, in conditions where the reinstated current-state disc configuration is the same between problem suspension and reinstatement, resumption latency might be higher in the SC-DG condition. Results to this effect would provide convincing support for a cost of goal-directed enquiry on the time to resume a suspended goal after it has been re-established.

3.3.2. Method

3.3.2.1 Participants

Thirty-two undergraduate students from the School of Psychology at Cardiff University participated in fulfilment of a course requirement. All were naive to the ToH task. Eight participants were removed from analysis for failing to satisfy the strict

criteria of executing the first two moves in interruption and control ToH problems error-free. Thus, data from 24 participants qualified for analysis across all reported dependent measures.

3.3.2.2 Apparatus and materials

The program used in Experiment 3 was modified so that interruptions occurred within the first subtask of primary task ToH problems: *only* after execution of the second move within solution-sequences. In an interruption trial, moving disc 4 to the non-goal peg for the LOOP disc (disc 1) and then moving disc 2 to the goal-peg for the LOOP disc would cause interruption (see Figure B1 in Appendix B). The secondary task was the digit-recall task used in Experiment 3. The program was additionally adapted to respond to secondary task completion by delivering reinstated ToH disc configurations that were either congruent or incongruent with those presented when the task was suspended. Specifically, for reinstated problems: the locations of discs in current-state disc configurations, goal-state disc configurations, or a combination of both, were either the same or different to when the primary task was suspended.

3.3.2.3. Design

The same as in Experiment 3 with some modifications. Firstly, the intended point of interruption was within the first subtask of primary task ToH problems: following execution of the second move. Secondly, ToH problems were only ever interrupted by an un-signalled digit-recall task. Thirdly, reinstated ToH problem disc configurations were manipulated across four levels; each varying in the degree of congruence shared with disc locations in the interrupted problem space. In Condition 1, there was no change to the suspended current-state or to the suspended goal-state (SC-SG). In Condition 2, both the current-state and goal-state had changed (DC-DG). In Condition 3, only the goal-state had changed (SC-DG), and in Condition 4 only the current-state had changed (DC-SG).

All dependent measures for Experiment 3 were included for analysis. Following the intended point of interruption, ToH problems could be solved within seven moves if executed error-free. Redundant moves are therefore any instance when this value is above seven. For secondary tasks, mean completion times and mean digits correctly

recalled did not significantly differ between any interruption condition. Descriptive data for secondary tasks is presented in Table B1 of Appendix B.

3.3.2.4. Procedure

The same as in Experiment 3, except that participants were exposed to one practice trial for each condition; one uninterrupted control; one interrupted SC-SG; one interrupted DC-DG; one interrupted SC-DG; and one interrupted DC-SG.

3.3.3. Results and Discussion

3.3.3.1. Move latency

Figure 3.1 illustrates latency data for the first move following the intended point of interruption in interrupted and control ToH problems. Resumption latencies

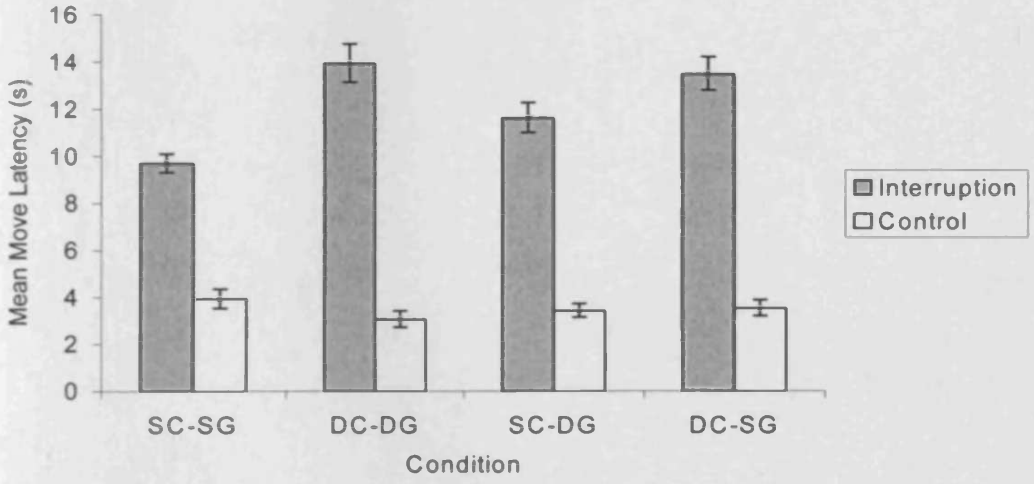

Figure 3.1. Mean move execution latencies (±SE) for interrupted and control trials in Experiment 4

are considerably longer than move latencies in control problems. Additionally, move latency is very similar across all control conditions. For interrupted problems, participants are faster to resume SC conditions compared to DC conditions. However, resumption is slowed in the SC-DG condition compared the SC-SG condition. For DC

conditions, resumption latency is very similar, although slightly shorter in the DC-SG condition.

For further analysis, a repeated measures, 2 (trial type) x 2 (current-state disc configuration) x 2 (goal-state disc configuration) ANOVA was conducted on move latency data. This revealed significant main effects of trial type, $F(1, 23) = 314.87$, $MSE = 11.52$, $p < .001$, and current-state disc configuration, $F(1, 23) = 27.9$, $MSE = 3.95$, $p < .001$, but not of goal-state disc configuration. There was a significant two-way interaction between trial type and current-state disc configuration, $F(1, 23) = 54.15$, $MSE = 2.14$, $p < .001$, with trial type and goal-state disc configuration almost reaching significance ($p = .07$). A marginally significant three-way interaction occurred for all three variables, $F(1, 23) = 4.69$, $MSE = 6.20$, $p < .05$. A Bonferroni *post-hoc* analysis revealed a significant simple main effect between move latency for interrupted SC-SG and interrupted SC-DG conditions, $F(1, 23) = 10.57$, $p < .005$, with no other significant differences between any of the other conditions. Thus, participants resume an interrupted ToH problem faster when the current-state disc configuration is congruent with that suspended, and are even faster when disc-configurations in the entire problem space are completely congruent with those suspended.

Being interrupted during the execution phase of a ToH subtask appreciably increases the time taken to execute the first move following secondary task completion (in support of Experiment 3). Changing the location of discs (and consequently disc-configuration) within the reinstated problem space increases resumption latency, although to different extents. Changing the current-state disc configuration has the most marked effect on increasing resumption latency, even when the goal-state disc configuration remains unchanged (DC-SG). Within DC conditions, the ability to retrieve the representation of a suspended goal or the choice of using a retrieval cue for such a process might thus be attenuated. Consequently, participants seem to reconfigure the problem space using the available visual array (e.g., Anderson & Douglass, 2001).

Resumption latency in the SC-SG condition is to long to consist exclusively of retrieval (e.g., Altmann & Trafton, 2002). However, participants are not just reconfiguring the problem space in the SC-SG condition; else, resumption latency should be the same as in the DC-DC condition. Also, in the DC-SG condition, participants seem

to exhibit reconfiguration like behaviour, although this is not the case in the SC-DG condition. That resumption latency in the SC-DG condition is longer than in the SC-SG but shorter than in the DC-SG condition suggests that resumption is facilitated from having previously processed current-state visual array. That is, something is remembered from the suspended problem, perhaps the location of a priming cue, that supports faster resumption. The difference between resumption latency in the SC conditions suggests that the re-established move undergoes goal-directed enquiry before it is executed, more so in the SC-DG condition. This process could occur due to awareness that the goal-state disc configuration has changed. If this were the case, there is reason to suspect that there should be a cost of goal-directed enquiry in all interruption conditions, including the SC-SG. This idea will be explored further in Experiment 5.

3.3.3.2. Move latency and experience

Figure 3.2 presents mean move latency data for interrupted and control problems across experimental half (i.e., trials 1 – 12 versus trials 13 – 24). There is a general increase in move latency in between control trials after having solved 13 ToH problems, although not for the SC-DG condition. Variance within each control condition is however quite high. For interrupted trials, there is a clear reduction in resumption latency after 13 ToH problems: markedly so for SC conditions. Although a similar reduction is evident for DC conditions, variance is much higher compared to SC conditions.

Control move latencies were removed from further analysis for failing to exhibit significant differences across experimental half. Due to occasional violations of assumptions of linearity, and evident heteroscedasicity, all move latency data were transformed using a logarithmic function that served to normalise the linear relationship between the variability in each condition. Transformed resumption latency data only was analysed using a repeated measures 2 (experimental half) x 2 (current-state disc configuration) x 2 (goal-state disc configuration) ANOVA. This revealed a significant main effect of experimental half, $F(1, 23) = 22.51$, $MSE = .15$, $p < .001$, and goal-state disc configuration, $F(1, 23) = 4.66$, $MSE = .14$, $p < .05$, but not of current-state disc configuration. There was also a significant two-way interactions between, current-state disc configuration and goal-state disc configuration, $F(1, 23) = 36.01$, $MSE = .07$, $p < .001$. A Bonferroni *post-hoc* analysis revealed a significant simple main effect between

resumption latency in the SC-SG and SC-DG conditions, $F(1, 23) = 21.44$, $p < .001$, but not between the DC-DG and DC-SG conditions. Resumption latency was generally shorter in the SC conditions compared to the DC conditions, but there was a larger increase in resumption latency in the SC-SG condition compared the SC-DG condition than there was in DC conditions.

Experience of interruption seemed to result in a larger reduction in resumption latency for SC conditions compared to DC conditions, but this effect was not supported by a significant three-way interaction. Refinement of the use of ToH discs in priming the re-establishment of an interrupted goal might lead to faster resumption. Task constraints, for example the stochastic predictability of changes to the current context, might render associative priming using ToH disc cues marginally ineffective (at least to level of producing a significant difference) at causing improvement in resumption performance. Consequently, selectively attending to priming cues may occur less frequently, and/or trying to locate a disc prime after primary task reinstatement might be met with a higher frequency of terminated searches (Wolfe, 1994). That resumption latency is shorter in the SC-SG condition compared to all other conditions, especially in the second experimental half, does however suggest some benefit from having been exposed to the same problem space prior to being interrupted.

3.3.3.3. Redundant moves

Table 3.1 illustrates mean number of redundant moves executed following the intended point of interruption in interrupted and control ToH problems. Redundant moves are more numerous in control conditions, although are affected by high amounts of variance. Fewer redundant moves are made interruption conditions, more so in SC conditions compared to DC conditions (with an apparent floor effect in the SC-SG interruption condition). Overall, redundant moves are lowest in the SC-SG condition and highest in the DC-DG condition. Variance is high within each condition, although is similarly distributed across trial type. This reflects the range of redundant moves in each case: between 0 and 5.5 for interruption conditions, and between 0 and 7 for control conditions.

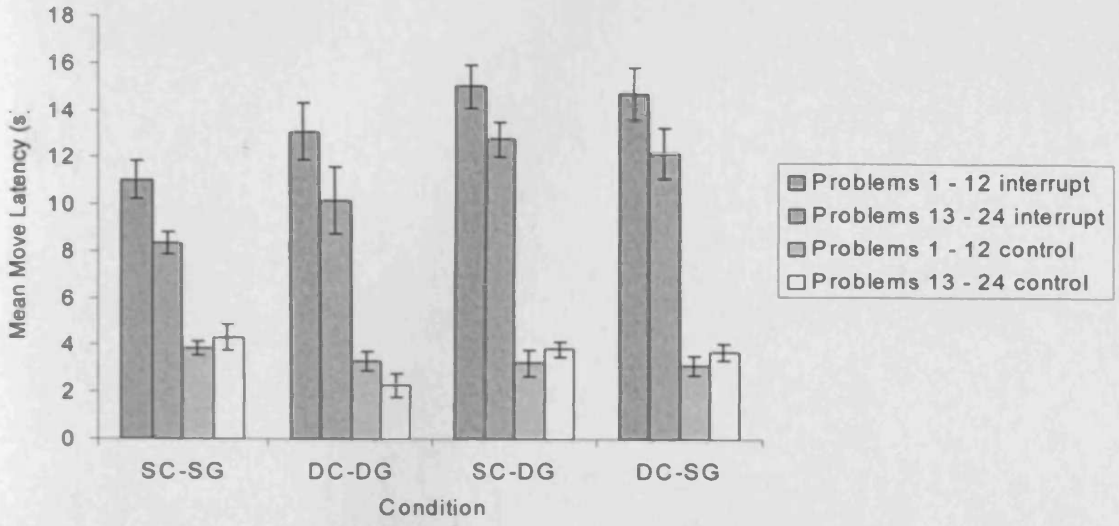

Figure 3.2. Mean move latencies (±SE) for trials 1 – 12 and trials 13 – 24 of the test phase of Experiment 4

Table 3.1

Mean redundant moves (±SE) in Experiment 4

	SC-SG		DC-DG		SC-DG		DC-SG	
	M	SE	M	SE	M	SE	M	SE
Interruption	0.2	0.19	1	0.24	0.45	0.81	0.81	0.29
Control	0.47	0.9	1.27	0.27	1.75	0.41	0.9	0.34

Frequency of redundant move mean ranks were calculated for the SC-SG condition (interruption Mean Rank = 3.23, control Mean Rank = 5.35), the SC-DG condition (interruption Mean Rank = 4.21, control Mean Rank = 5), the DC-SG condition (interruption Mean Rank = 4.1, control Mean Rank = 4.06), and the DC-DG condition (interruption Mean Rank = 4.79, control Mean Rank = 5.25). A Friedman test indicated that a significant difference occurred between scores for redundant moves for each of the eight conditions $\chi^2(7) = 19.15, p < .010$. A series of Wilcoxon Signed Rank tests with a Bonferroni correction ($p = .0.006$) were carried out to provide *post-hoc* comparisons of

mean ranks for each condition. These tests revealed significant differences between frequency of mean ranks for redundant moves in interruption and control conditions for, the SC-SG condition, $Z(1) = -2.11, p < .05$, and the SC-DG condition, $Z(1) = -2.45, p < .10$. No other two-way comparisons reached significance.

Redundant move data shows that a greater number of redundant moves are executed in uninterrupted control ToH problems, but only significantly so for SC conditions. However, this result should be approached with some caution, given the exceptionally high variance evident within those conditions. Nevertheless, fewer redundant moves are executed in SC conditions, so few that this effect was unlikely to have been caused by chance.

3.3.3.4. Move-error frequency

Table 3.2 presents move-error frequency data for all participants that committed one or more redundant moves out of a possible 368 interruption and control ToH problems. Fewer move-errors are made following interruption, with an apparent floor effect in the SC-SG condition. Participants solved 267 problems error-free compared to 101 problems in which at least one redundant move was made, $Q(7, N = 24) = 42.08, p < .001$. Pairwise comparisons using a Bonferroni correction ($p = .0125$) revealed that more move-errors were made in control problems compared to interruption problems in the SC-SG condition, $Q(1, N = 24) = 12.8, p < .001$, and the SC-DG condition, $Q(1, N = 24) = 5.4, p < .025$. Error frequency did not significantly differ across trial type for either of the DC conditions.

Table 3.2

Move-error frequency in experiment 4

	SC-SG	SC-DG	DC-SG	DC-DG
Interruption	2	6	11	17
Control	18	20	7	20

Note. Maximum move-error frequency for each condition is 46 moves

Fewer move errors are made when a ToH problem is reinstated with the discs in the same location as when that problem was suspended. There are two possible reasons for why such effects occurred. Firstly, participants might make an initial error upon resumption. Is this more likely in DC conditions because the removal of priming information causes more frequent re-establishment of an incorrect goal? Closer inspection of the point of move-error showed that this was not the case: move-error was just as likely within the solution-sequence of both SC and DC problems than just at resumption. A second and more plausible explanation is that executing the correct first move leads to fewer move-errors. Executing the correct first move is more likely in SC conditions because participants are able to remember more efficiently that move due to the location of a priming cue. This may lead to better memory for goals further into the solution-sequence of SC ToH problems: a knock-on effect of faster retrieval of other previously encoded goals. The idea here is that after the suspended goal has been reactivated; the goal with the next highest activation value may also be associatively reactivated (as is suggested for uninterrupted performance in the ToH task by Altmann & Trafton, 2002)

3.3.4. Discussion

The results of Experiment 4, along with those in Chapter 2, establish that interruption to problem solving performance in the ToH task produces significant deficits in the time to resume the reinstated problem. The greater the level of incongruence between disc location in suspended and reinstated problems, the greater the level of performance attenuation across all dependent measures. Resumption latency, resumption latency across experimental half, redundant moves, and move-error frequency were always lowest in the SC-SG condition and always highest in the DC-DG condition.

The form of the resumption latency data suggests that resumption strategy could be determined by the amount of associative activation provided by priming cues in the current context, even in the absence of an explicitly cued interruption lag. For SC conditions, participants might be able to use disc location (and perhaps feature properties) as a priming cue to support the reactivation from memory of a previously suspended goal.

For DC conditions, participants seem to reconfigure the problem space, and do not appear to attempt to solve problems based upon cue availability relative to when the problem was suspended. Such behaviour is reflected by resumption latencies far in excess of the time required to retrieve *or* reconstruct the representation of a single goal (e.g., Altmann & Trafton, 1999, 2002; Anderson & Douglass, 2001).

The efficiency of associative priming using the current context may however be affected by the current experimental design. Firstly, there was little (or no) opportunity to form the associative connection immediately before goal suspension, perhaps engendering any associative connection to be weak (Altmann & Trafton, 2002). Successful retrieval may have occurred infrequently, or perhaps partial retrieval was more likely given the difficulty in choosing a production for the appropriate goal (e.g., Siegler, 1987, 1988). Retrieval efficiency may have improved perhaps when participants were more adept at priming strategically a suspended goal with a retrieval cue. Secondly, task constraints might have rendered the utility of associative priming low because of the frequent changes to disc location in reinstated problems. Perhaps participants rationalised that more often than not, using the current context to form a connection between a cue and a suspended goal was generally ineffective at reducing resumption latency (e.g., Anderson, 1990, Anderson & Douglass, 2001). For now, both explanations are given equal credence, although better overall performance from the SC-SG condition leads one to infer that participants did benefit substantially from having been exposed to the same disc configuration prior to primary task suspension. This benefit may have been because of the ability to detect a cue at a familiar location (e.g., Wolfe, 1994), rather than being able to efficiently retrieve enough information to resume the suspended goal in a timely manner.

Although resumption latency data in the SC-SG condition is consistent with modelled estimations of the time required to reconstruct a goal using the visual array (e.g., Altmann & Trafton, 2002), this does not necessarily mean that retrieval does not occur. Participants are significantly faster to resume an interrupted problem in the SC-SG condition compared to the DC-DG condition but are not significantly slower in the SC-DG condition. Following re-establishment of a goal, participants may evaluate the expected gain of executing a re-established move (e.g., Einstein & McDaniel, 1996;

Speier, 2003). In the SC conditions, such a cost of goal-directed enquiry is lowest, perhaps partly because discs are in the same location in the current-state. In the DC-DG and DC-SG conditions, goal-directed enquiry is highest, and is perhaps influenced by the relocation of discs since primary task suspension. The main implication is that resumption latency for the SC-SG condition might not consist solely of the time taken to re-establish the move, but also the time taken up by goal-directed enquiry. Goal-directed enquiry in the SC-SG condition may not be necessary (i.e., nothing has changed between problem suspension and reinstatement), but is encouraged because of frequent changes in other reinstated conditions.

Experiment 5 set about addressing some of the issues that might have encouraged participants not to rely entirely on associative priming as a mechanism for supporting the retrieval of a suspended goal in Experiment 4. Firstly, the uncertainty engendered by reinstated problems with incongruent current-state and/or goal-state disc configurations might have motivated a strategic decision to dedicate fewer resources to associative priming before interruption. That is, participants may have quickly learned that priming cues within most reinstated problems were somehow 'less available' than in a fewer number of other problems, and rationalised that memorising interrupted-state information was futile. In Experiment 5, solution-sequences in reinstated problems were either completely congruent or completely incongruent with those suspended (like in the SC-SG and DC-DG conditions of Experiment 4). With experience of interruption, participants might learn that effective associative priming could improve resumption performance in half of the interrupted problems, causing a shift in willingness to dedicate resources to its deployment. Secondly, resumption latency data in Experiment 4, as in Experiments 2 and 3, seemed to highlight a cost of goal-directed enquiry that might be masking the actual time required to re-establish the representation of a suspended goal. In some conditions of Experiment 5, the distance to goal, in terms of the number of moves required for error-free completion in reinstated problems, was reduced. This may have the desired effect of causing participants to spend less time evaluating the cost of performing a re-established move, which might result in speedier task resumption.

3.4. Experiment 5

3.4.1. Introduction

The central aim of Experiment 5 was to investigate further the role of priming cues in recovering from interruption, specifically how cue location and the difficulty in locating a cue might affect how a suspended goal is re-established. Motivated by results from Experiment 4, Experiment 5 also set out to explore if there is a cost of evaluating the execution of a move after it has been re-established: that is a cost of goal-directed enquiry.

Two main questions are echoed throughout Experiment 5, and both build upon the predictions and findings of Experiment 4. Firstly, does the location of a priming cue affect how an interrupted ToH problem is resumed and subsequently completed? Obtaining a similar pattern of resumption latency data to the SC-SG and DC-DG conditions in Experiment 4 would reinforce the inference that maintaining the exact disc location in a reinstated problem causes faster resumption of the suspended goal. Importantly, shorter resumption latencies in SC conditions compared to respective DC conditions would provide further evidence that participants are not simply reconfiguring the problem space in the former case. Secondly, following goal re-establishment, is there a cost of goal-directed enquiry? Shorter resumption latencies for problems that require fewer moves for error-free completion and longer resumption latencies for those that require more moves would support that at least part of the resumption latency consists of time taken up by goal-directed enquiry. Differences may serve to highlight potential masking effects of goal-directed enquiry on the time to retrieve an associatively primed goal in a suspended ToH problem.

To speak to the issues raised above, Experiment 5 manipulated both opportunity to deploy resumption strategies to support the retrieval of a suspended goal, and the cost of goal-directed enquiry.

Manipulating how a suspended goal could be re-established was achieved in a similar way as in Experiment 4: by changing the level of congruence between disc location in suspended and reinstated current-state configurations. However, following

completion of a secondary task, disc-locations within the current-state configuration were either completely congruent (same-current: SC) or completely incongruent (different–current: DC) with those in the suspended visual array. If possible at all, priming should only occur in SC conditions where disc properties have not changed in the current-state.

Manipulating the cost of evaluating move execution, was achieved by varying the number of moves (and consequently problem difficulty, e.g., Davies, 2003), required to solve ToH problems error-free following their reinstatement. There is evidence to show that the distance to goal can inflate operator selections (e.g., increasing execution choice)(e.g., Altwood & Polson, 1976; Newell & Simon, 1972), but little work has been done on how history of success might effect such selection. Moves to completion in reinstated problems were either: nine (two moves more than before suspension); seven (no additional moves); or five (two moves less than before suspension). Note that the seven-move goal-state condition (referred to as *G* hereafter) imposes no additional load in terms of the number of moves to goal compared to when the problem was suspended. In contrast, the nine-move goal-state condition (referred to as *G+2* hereafter) imposes a higher cost, which might be reflected by longer time taken for goal-directed enquiry. Applying the same rationale, the five-move goal-state condition (referred to as *G–2* hereafter) imposes a lower cost of goal-directed enquiry. In the G-2 conditions, the LOOP disc does not have to be considered in reinstated problems: this might be reflected in a shorter time for enquiry and consequently shorter resumption latency.

Similar performance outcomes to those found in Experiment 4 are expected for SC-SG and DC-DG conditions in Experiment 5. Shorter resumption latency in all SC conditions compared to respective DC conditions would suggest that resumption strategy in the former case does not just involve reconfiguration of the problem space. Resumption latencies that are longer when there are more moves to goal and shorter when there are fewer moves would suggest that resumption latency is affected by a cost of goal-directed enquiry. Finally, problem difficulty might interact with disc location, such that resumption strategy is affected. In SC conditions, more moves to completion might increase perceived problem difficulty, causing participants to abandon retrieval through associative reactivation of the goal. Using the same rationale, fewer moves to completion might encourage goal reactivation via associative priming, but could cause

abandonment of retrieval because the problem is perceived as easier to resume using a different strategy.

3.4.2. Method

3.4.2.1. Participants

Fifty-one undergraduate students from the School of Psychology at Cardiff University participated in fulfilment of a course requirement. All were naive to the ToH task. Fifteen were removed for failing to meet the same criteria stated in section 3.3.1.1, leaving 36 participants.

3.4.2.2. Apparatus and materials

The same program used in Experiment 4 was modified. The number of experimental trails was increased from 24 to 36. Four new ToH problems were introduced, each repeated once, to create 4 additional interruption problems and 4 identically matched control problems. Overall, 12 trials contained an interruption, and after secondary task completion, 4 required seven moves for error-free completion, 4 required nine moves for error-completion, and 4 required five moves for error-free completion. There were also 12 control trials, each requiring nine moves for error-free completion.

3.4.2.3. Design

The same in Experiment 4 aside from some modifications. Firstly, following secondary task completion reinstated ToH disc configurations were manipulated across six levels, each varying in the level of congruence shared between the spatial location of discs prior to and following interruption. Reinstated current-state disc configuration had two levels: (1) same-current (SC), and (2) different-current (DC). Reinstated goal-state disc configuration had three levels: (1) goal (G: same number of moves to error-free completion), (2) goal+2 (G+2: two additional moves required for error-free completion), and (3) goal-2 (G-2: two fewer moves required for error-free completion). Thus, one of the following six reinstated disc configuration conditions was possible following

completion of secondary tasks: (1) SC-SG (2) SC-SG+2 (3) SC-SG-2 (4) DC-DG (5) DC-DG+2 (6) DC-DG-2. Figure B2 in Appendix B provides examples of each condition. All control and filler ToH problems required nine moves for error-free completion based upon the same reasoning outlined in Section 3.3.1.

Only move latency was fully explored analytically across each condition of Experiment 5. Other dependent measures considered in previous experiments in this thesis, were affected by potential confounds caused by manipulating distance to goal in interruption conditions and not in control conditions. That is, redundant moves in interrupted ToH problems were instances when the number of moves executed were more than, seven (G), nine (G+2) or five (G-2). Within control trials however, redundant moves were only instances when the number of moves required to complete ToH problems exceeded the minimum of seven. Given this confound, redundant move data, move-error frequency and adjusted completion times were only compared for SC-SG and DC-DG conditions where interruption and control trials were comparable. Nevertheless, SC-SG and DC-DG exhibited no significant differences across trial type for any measure other than move latency, and were thus received no further consideration.

For secondary tasks, mean completion times and mean digits correctly recalled did not significantly differ between any interruption condition. Descriptive data for secondary tasks is presented in Table B2 of Appendix B.

3.4.2.4. Procedure

The same as in Experiment 4, except that the whole experiment took approximately 1 hour to complete.

3.4.3. Results and discussion

3.4.3.1. Move Latency

Figure 3.3 displays move latency data for the first move following the intended point of interruption in interruption and control ToH problems. Resumption

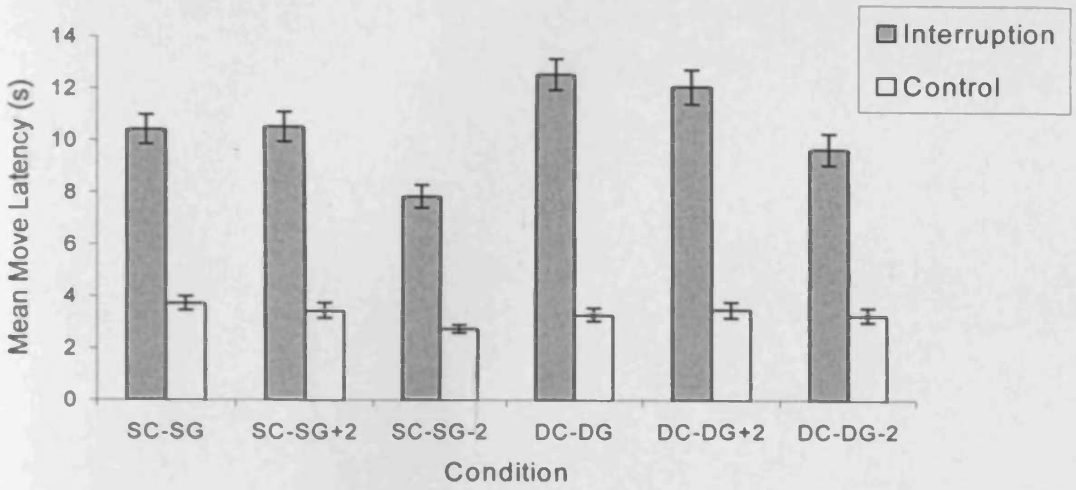

Figure 3.3. Mean move execution latencies (±SE) for interrupted and control trials in Experiment 5

latency is elevated in comparison to control move latency. Within interrupted SC and DC conditions, resumption latency is similar for G and G+2 goal-state conditions, but is considerably reduced in the G-2 condition. For each level of the goal-state, SC conditions yield shorter resumption latencies than in DC conditions.

Move latency data was further analysed using a repeated measures 2 (trial type) x 2 (current-state disc configuration) x 3 (goal-state disc configuration) ANOVA. This revealed significant main effects of, trial type, $F(1, 35) = 357.17$, $MSE = 15.71$, $p < .001$, current-state disc configuration, $F(1, 35) = 16.21$, $MSE = 6.23$, $p < .001$, and goal-state disc configuration, $F(1, 35) = 18.58$, $MSE = 6.20$, $p < .001$. There were also significant interactions between, trial type and current-state disc configuration, $F(2, 70) = 17.69$, $MSE = 5.11$, $p < .001$, and trial type and goal-state disc configuration, $F(2, 70) = 14.77$, $MSE = 3.8$, $p < .001$. All other interactions were non-significant. The significant interaction between trial type and current-state disc configuration was further explored using a Bonferroni *post-hoc* test. This revealed a significant simple main effect between current-state disc configuration for interrupted trials, $F(2, 34) = 20.30$, $p < .001$, but not for control trials ($p = .79$).

Thus, SC conditions resulted in consistently shorter resumption latencies than in DC conditions. The interaction between trial type and goal-state configuration was

similarly explored using a Bonferroni *post-hoc* test. There were significant simple main effects in interrupted trials between resumption latencies in, the SG and SG-2 conditions, $F(2, 34) = 19.08, p < .001$, and the DG and DG-2 conditions, $F(2, 34) = 19.08, p < .001$. There were no such significant differences between move latency for any of the other conditions. In interrupted trials, changing the reinstated goal-state disc configuration so that it requires two moves less than when it was suspended causes shorter resumption latencies compared to when it requires the same moves as when suspended or two moves extra.

Move latency data for Experiment 5 has provided further evidence for shorter resumption latencies when disc location is the same between the suspended and reinstated problem space. A reinstated goal-state that requires two additional moves for error-free completion (i.e., G+2) has no impact on the length of resumption compared to problems that require two fewer moves to completion. However, when the reinstated goal-state requires two less moves for error-free completion, resumption is faster for both the SC and DC conditions.

Perhaps in the G+2 conditions, the change to the goal-state disc configuration is not salient enough to affect resumption time. Take as an example the SC-SG and SC-SG+2 conditions. Since the potential priming cue (disc 4) remains in the same spatial location in the current-state disc configuration of both conditions, this might have led to similar retrieval latencies. Additionally, goal-directed enquiry might be similar for both conditions, because of how many moves participants consider in the solution-sequence before executing a move (e.g., Phillips et al., 2001). For both conditions, the current position of the LOOP disc and its overall goal destination position is the same between suspension and reinstatement. Perhaps then, look-ahead, like planning, stops at or before the LOOP disc (e.g., Davies, 2003) meaning that the expected performance gain of executing the move in both conditions would be the same. Given the strong evidence that move execution time increases linearly with the number of subgoals (e.g., Anderson et al., 1993, Anderson & Douglass, 2001, Altmann & Trafton, 2002; Ruiz, 1987), this seems to be the most plausible explanation for the non-significant difference. The same explanation is also suggested for resumption latencies across DC-DG and DC-DG-2

conditions, although these are elevated compared to SC conditions because of the cost having to reconfigure the suspended move, rather than retrieve it.

Some might argue that a goal-stack account could explain resumption latency in interrupted conditions of Experiment 5 (e.g., Anderson et al., 1993). For example, in G-2 conditions, the suspended LOOP disc is no longer a parent goal in reinstated problems, and so goal-recursion should be faster; thus increasing the speed of resumption through reconfiguration. However, this interpretation fails to explain why there is a discrepancy between resumption latency for interrupted SC-SG-2 and DC-DG-2 conditions. Participants resume reinstated problems faster when they contain discs that have not changed across any property dimension (e.g., location, colour, and so on) since suspension. They seem to be faster at detecting the cue that might lead to the resumption of a suspended goal, an associative link likely to have been formed by selectively attending to the cue prior to interruption (e.g., Wolfe, 1994).

3.4.3.2. Move latency and experience

Figure 3.4 illustrates mean resumption latency for reinstated interrupted trials across experimental half. Figure 3.5 illustrates comparable move latencies for control problems. Resumption latency decreases after 18 trails, whereas control move latency increases. After 18 trials, resumption latency is much shorter across each level of the goal-state in SC conditions, a reduction unmatched in all DC conditions.

Control move latencies were removed from further analysis for failing to exhibit significant differences across experimental half. A logarithmic transformation was applied prior to analysis for the same reasons outlined in Section 3.3.3.2. Transformed resumption latency data only was analysed using a repeated measures 2 (experimental half) x 2 (current-state disc-configuration) x 3 (goal-state disc configuration) ANOVA. There were significant main effects of, experimental half $(1, 35) = 15.68$, MSE = 0.19, $p < .001$, current-state disc configuration, $F(1, 35) = 37.94$, MSE = 0.13, $p < .001$, and goal-state disc configuration, $F(1, 35) = 16.91$, MSE = 0.16, $p < .001$. There was also a marginally significant interaction between experimental half and current-state disc configuration, $F(1, 35) = 4.2$, MSE = 0.16, $p < .05$. A Bonferroni *post-hoc* test revealed

a significant simple main effect between experimental half in SC conditions, F (1, 35) = 14.27, $p < .001$, which was not the case for DC conditions.

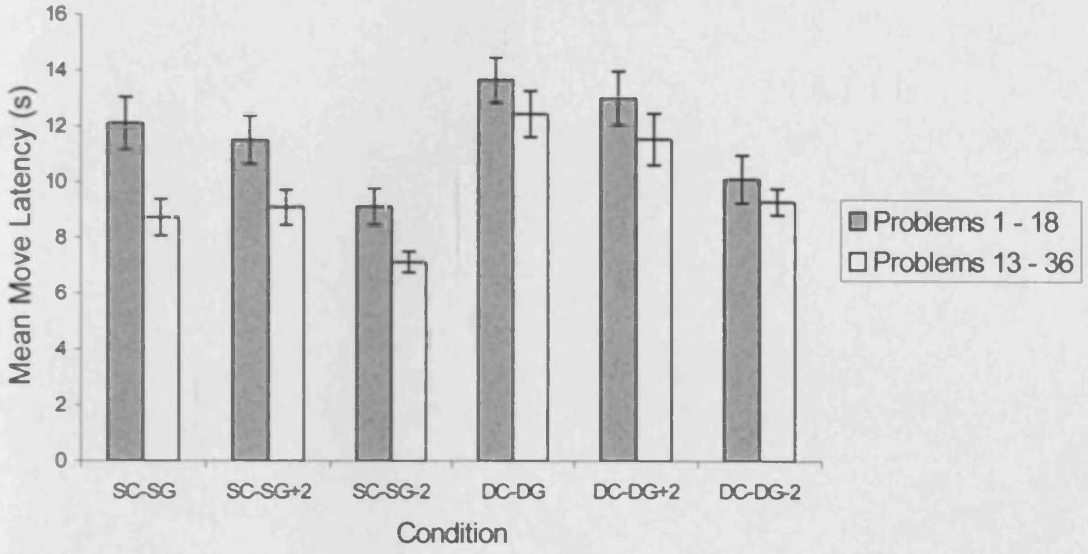

Figure 3.4. Mean un-transformed move latencies (±SE) in trials 1 – 18 and trials 19 – 36 for interrupted ToH problems in Experiment 5

Resumption latency reduced in the second experimental half of Experiment 5 but only when the current-state disc configuration was congruent between the suspended and reinstated visual array. Experience of being interrupted or experience of problem solving in the task domain cannot just account for such a reduction. This is because, a similar reduction in resumption latency should occur for DC conditions, which is not the case. In Experiment 5, compared to Experiment 4, participants seemed to become more adept at efficiently using priming cues to support the resumption of an interrupted ToH problem. For DC conditions, participants may have rationalised that for a new problem space, as indicated by a new current-state and new-goal state, reconfiguration of the problem space would result in the highest performance gain. In Experiment 5, unlike Experiment 4, all interruption DC conditions contained reinstated goal-state disc configurations that were different to those upon primary task suspension.

In sum, SC conditions in Experiment 5 seemed to show that resumption of an interrupted goal was possible using associative priming, if discs maintained the same properties between suspension and reinstatement. This process becomes more efficient with experience of interruption, perhaps because of increased feedback indicating the efficiency of such a strategy on reducing the length of resumption latency (unlike Experiment 4).

3.4.4. Discussion

Experiment 5 has provided further support that resuming an interrupted ToH problem is faster if the properties of ToH disc cues have not changed between problem suspension and resumption. Resumption latency was reduced when the reinstated goal-state required two less moves for error-free completion, but was not increased when the goal-state required two more moves for error-free completion.

Resumption latency data in SC-SG and DC-DG conditions of Experiment 5 is similar to that of the same conditions in Experiment 4. Thus, similar resumption strategies might be deployed for these conditions within each experiment. For the SC-SG conditions, attempted retrieval or reconstruction of the suspended goal, and reconfiguration of the problem space in the case of the DC-DG conditions.

Results from Experiment 5 have shown that re-establishment of a suspended goal using priming cues can account for faster resumption of an interrupted ToH problem. Experiment 5 has also built upon results of Experiment 4 in confirming that participants do not execute a move immediately after its reestablishment. For example, if retrieval of a suspended goal was determined by the time to detect a primed cue, then resumption latency should have been similar across all SC conditions. That resumption latency was shorter when there were fewer moves to goal suggests that participants evaluate the utility of executing a re-established move, with time affected by the distance to goal. Such results impart a non-strategic shift to reconstruction; rather, participants prefer to try to retrieve the goal suspended because of interruption.

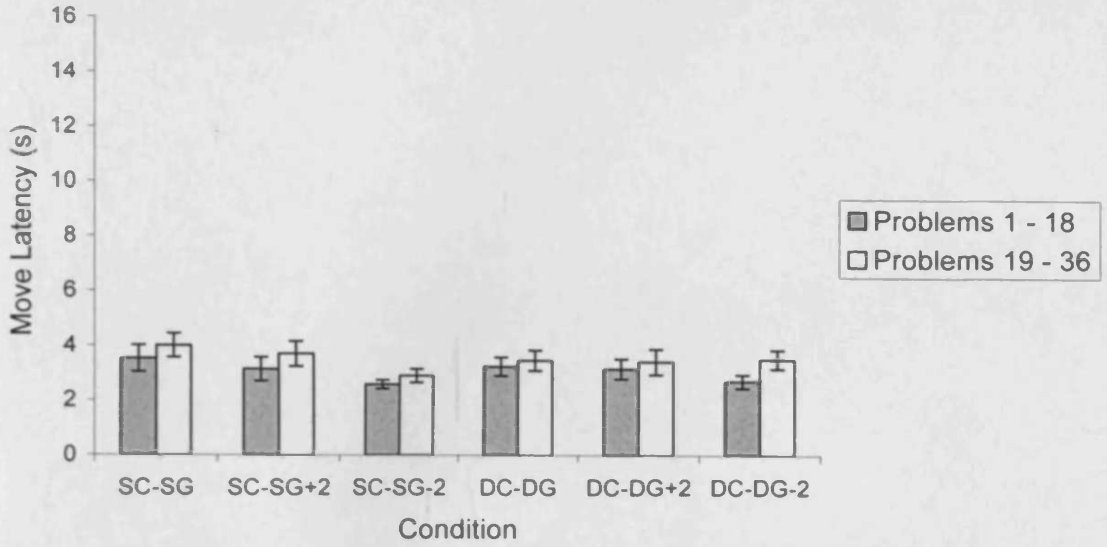

Figure 3.5. Mean control move latencies (±SE) in trials 1 – 18 and trials 19 – 36 for control ToH problems in Experiment 5

Although cue location might support the resumption of an interrupted ToH problem, perhaps other cue-features contribute to this effect. Location alone might not have the greatest expected gain in priming the representation of a suspended goal a dynamic task such as the ToH where discs have to be moved around the visual array in satisfaction of task requirements. Attending to cueing stimuli may be a more selective process based upon permanent properties within the visual space (e.g., Wolfe, 1989, 1994). Specifically, attention might be guided toward a feature-cue(s) that is salient and unique enough to support its identification after a retention interval. A likely candidate for such a feature-cue might be colour (e.g., Sutcliffe, 1995).

Experiment 6 was conducted to assess the role of disc colour in resuming a promptly interrupted ToH task. Simply, disc-colour might serve as a salient priming cue if an interrupted goal is to be re-established through associative priming. Maintaining the location of discs between primary task suspension and reinstatement might also encourage the use of disc-colour as a priming cue. Colour may be regarded as even more salient if participants do not expect a change to disc locations (and consequently the solution-sequence) during interruption (e.g., Wolfe, 1994). Advance knowledge of the

location of an object might improve detection speed, but with additional knowledge of object colour, performance might improve further (e.g., Humphries, 1981).

3.5. Experiment 6

3.5.1. Introduction

Twenty-four ToH problems were to be solved, 8 of which were interrupted by un-signalled digit-recall tasks. Upon primary task reinstatement, discs could have changed colour since they were suspended, although disc-size rules did not change. In the SC-SG condition the colour of discs was the same in the current-state and the goal-state (i.e., red (disc 1), yellow (disc 2), blue (disc 3), and green (disc 4)). In the DC-DG condition, discs in the current-state and goal-state were different combinations of the standard four colours (e.g., green (disc 1), blue (disc 2) and so on). In condition 3, disc-colour in the current-state was the same but was different in the goal-state (referred to as SC-DG hereafter). In condition 4, disc-colour in the current-state was mixed but was the same in the goal-state (referred to as DC-SG hereafter). In conditions where disc colour had changed, no disc was the same colour between problem suspension and reinstatement (see Figure B3 in Appendix B for examples).

Disc-colour did not change during performance of control problems or filler problems (i.e., they were the same as in the SC-SG condition). This was firstly to minimise expectation of disc-colour-change in interrupted problems so that adaptation to the task constraints could be explored (i.e., effects of experience). Maintaining the colour of a target object across trials can increase more frequent and faster responses to that object (e.g., Goolsby & Suzuki, 2001), a finding supported by the guided-search model of selective-attention (e.g., Wolfe, 1994). Given that the same move is always suspended in interruption trials, its corresponding disc-colour (i.e., green) might be regarded as a cue that predicts the next move in the solution-sequence. This could encourage selective attendance to the colour when an interrupted problem is reinstated. Noticing that a different disc contains the primed colour in a reinstated problem might

render the prime ineffective, thereafter terminating the retrieval attempt, and causing reconstruction of the problem space. Alternatively, the goal may be retrieved without having to go through any conflict resolution (e.g., selected by colour), and as such, its associated move was unachievable upon attempted execution. Due to the same rationale, changing disc colour on every trial might have the undesired effect of encouraging participants to abandon associative priming using disc-colour as a cue in interrupted problems (e.g., Maljkovic & Nakayama, 1994).

The main predictions for Experiment 6 are as follows. If disc-colour is a salient enough cue to allow an associative connection to be formed between a disc and a suspended goal, maintaining the same colour-to-disc pattern in the reinstated problem should result in the lowest resumption latency. Changing disc-colours between suspension and reinstatement in the current-state, should abolish the opportunity to associatively reactivate the goal, only if the suspended move is primed by a specific colour cue (e.g., Bauer at al., 1996; Goolsby & Suzuki, 2001). In addition, if goal-directed enquiry is a robust phenomenon following the re-establishment of a suspended goal, resumption latency should be longest in the DC-SG condition, although longer in the SC-DG condition compared to the SC-SG condition. Finally, if disc-location is used to assess the utility of executing a re-established move during goal-directed enquiry, resumption latencies should be shorter for all conditions in Experiment 6 compared to SC conditions of Experiment 4. Resumption latencies in Experiment 6 similar to those in DC conditions of Experiment 4, would suggest that disc-location and *not* colour is a better determinant of whether a goal is re-established in the first place.

3.5.2. Method

3.5.2.1. Participants

Sixteen undergraduate students from the School of Psychology at Cardiff University participated in fulfilment of a course requirement. All participants met the same criteria as detailed in section 3.3.1.1, and none had to be removed.

3.5.2.2. Apparatus and materials

The same program used in Experiment 4 was modified to accommodate the two changes to the design of Experiment 6. The location of discs within the reinstated visual array did not change between primary task suspension and primary task reinstatement. Depending upon the condition, the colour of individual ToH discs contained within the suspended visual array could change between primary task suspension and reinstatement. Prior to primary task suspension, colour always followed the same disc-size rule: red for disc 1; yellow for disc 2; blue for disc 3; and green for disc 4. In disc-colour different reinstated ToH problems, discs were different combinations of these colours (e.g., blue [1], green [2], yellow [3] and red [4]). Figure B3 in Appendix B provides examples of all conditions.

3.2.2.3. Design

The design was the same as in Experiment 4 except for changes to disc-location and disc-colour in reinstated ToH problems. Disc-location did not change between primary task suspension and reinstatement. Disc-colour could change between primary task suspension and reinstatement, with different coloured discs in either the current-state, the goal-state, or both states. In the SC-SG condition, disc-colour did not change. In the SC-DG condition, disc-colour was different in only the goal-state. In the DC-SG condition, disc colour was different in only the current-state. In the DC-DG condition, disc colour was different in the current-state with matched colour in the goal-state.

For secondary tasks, mean completion times and mean digits correctly recalled did not significantly differ between any interruption conditions. Descriptive data for secondary tasks is presented in Table B1 of Appendix B.

Adjusted completion time was removed from further analysis for failing to exhibit significant differences between trail type and across any within trial type condition.

3.5.2.4. Procedure

The same as Experiment 4, except that whole experiment took approximately 45 minutes to complete.

3.5.3. Results and discussion

3.5.3.1. Move latency

Figure 3.6 illustrates mean move latency data for the first move following the intended point of interruption in interrupted and control ToH problems. Resumption latencies are longer than move latencies in control problems. Additionally, move latency is very similar across all control conditions. For interrupted problems, resumption latency is much shorter in the SC-SG condition compared to the DC-DG condition, with similar resumption latencies between the DC-DG and SC-DG conditions. However, resumption latency is noticeably highest in the DC-SG condition, some 3.78 s longer than in the SC-SG condition.

For further analysis, a repeated measures, 2 (trial type) x 2 (current-state disc-colour) x 2 (goal-state disc-colour) ANOVA was conducted on move latency data. There were significant main effects of, trial type, $F(1, 15) = 319.27$, $MSE = 4.8$, $p < .001$, and, current-state disc-colour, $F(1, 15) = 6.68$, $MSE = 5.34$, $p < .025$, but not of goal-state disc-configuration. There was however, a significant three-way interaction between all three variables, $F(1, 15) = 5.6$, $MSE = 7.91$, $p < .05$. Bonferroni *post-hoc* analyses revealed a significant simple main effect between the SC-SG and DC-SG conditions, $F(1, 15) = 10.11$, $p < .005$, with no other significant differences between any of the other conditions.

Resumption latency data is consistent with the original hypothesis that disc colour is a salient priming cue for supporting the reestablishment of a suspended goal. Maintaining the colour of a potentially primed disc between problem suspension and resumption results in the fastest resumption of the goal associated with that disc. In contrast, changing disc colour after problem suspension causes marked performance attenuation in the speed of resuming the same goal. In the former case, it is suggested that upon noticing the same disc-colour pairing, a production rule is fired to retrieve the representation of the suspended goal (e.g., Altmann & Trafton, 2002). In the latter case, noticing that a primed disc-colour pairing does not correspond to the same disc, causes abandonment of goal retrieval, meaning that the goal has to be re-established from the visual array (e.g., Anderson & Douglass, 2001). It is also noted that resumption latency,

especially in the SC-SG condition, is shorter than in the same condition of Experiments 4 and 5, an observation that will be given further deliberation in the following section.

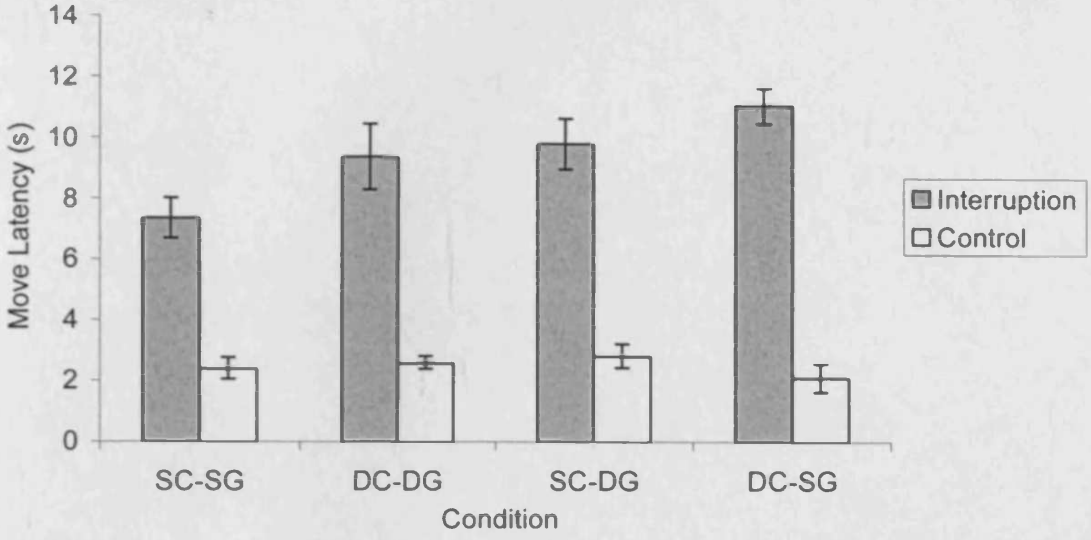

Figure 3.6. Mean move execution latencies (±SE) for interrupted and control trials in Experiment 6

Unlike Experiments 4 and 5, resumption latency was very similar between the SC-DG and DC-DG conditions. A possible explanation of this effect extends the original hypothesis, but builds upon the conclusions regarding resumption strategy in Experiments 4 and 5. Across all three Experiments, participants are able to retrieve the suspended goal in the SC-DC condition, using colour from the current-state as a fundamental priming cue for associative reactivation. The move outcome is evaluated through goal-directed enquiry, which incurs a higher cost because of the incongruent goal-state. However, in the DC-DG condition of Experiment 6, participants are faster to reconstruct a goal than in the DC-DG condition of all other experiments, because the location of discs between suspension and reinstatement is the same. Thus, although the goal has to be re-established using new colour-to-size rules, goal-directed enquiry takes less time. This is because some memory still exists for the suspended solution-sequence, perhaps a plan formed prior to interruption. In Experiments 4 and 5, goal-directed enquiry takes longer because of this inconsistency, a factor that is removed from

Experiment 6. That disc-colour might be responsible for reactivation of a suspended goal, and disc-location might influence goal-directed enquiry will be further deliberated in the general discussion.

3.5.3.2. Move latency and experience

Figure 3.7 illustrates mean move latency for interrupted and control problems across experimental half (i.e., trials 1 – 12 versus trials 13 – 24). Although control move latency seems to increase in the second experimental half, this increase is small, with trials 13 – 24 affected by high amounts of variance. For interrupted trials there is a general decrease in resumption latency in the second experimental half, markedly so in the SC-SG condition. Although a similar reduction in resumption latency is apparent in the SC-DG condition, trials 13 – 24 are affected by higher amounts of variance than within the respective SC-SG condition. Resumption latency also decreases in DC conditions, with a similar reduction in each case. As for the SC conditions, the DC-SG condition is affected by a high amount of variance across both experimental halves, which is not the case for the DC-DG condition, where variance is relatively small.

Control move latency failed to significantly differ across experimental half and was subsequently removed from further analysis. A logarithmic transformation was applied prior to analysis for the same reasons outlined in Section 3.3.3.2. Transformed resumption latency data only was analysed using a repeated measures 2 (experimental half) x 2 (current-state disc-colour) x 2 (goal-state disc-colour) ANOVA. This revealed significant main effects of experimental half, $F(1, 15) = 11.52$, $MSE = .13$, $p < .005$, current-state disc-colour, $F(1, 15) = 4.61$, $MSE = .11$, $p < .05$, and goal-state disc *colour,* $F(1, 15) = 5.07$, $MSE = .28$, $p < .05$. There was also a significant two-way interaction between experimental half and current-state disc-colour, $F(1, 15) = 16.76$, $MSE = .11$, $p < .001$. A Bonferroni post-hoc test revealed a significant simple main effect across experimental half only between resumption latency in SC conditions, $F(1, 15) = 54.93$, $p < .001$, with no such effect across DC conditions. Thus, with experience of problem solving in the task domain and experience of being interrupted, resumption latency significantly reduced only when discs within the current-state did not change colour between suspension and reinstatement.

With experience of problem solving in the task domain and of being interrupted, resumption latency reduces markedly in the SC conditions (to 5.65 s in the SC-SG condition). In the SC-SG condition, this might reflect more successful goal retrieval attempts, because of better use of disc colour as a priming cue. Since

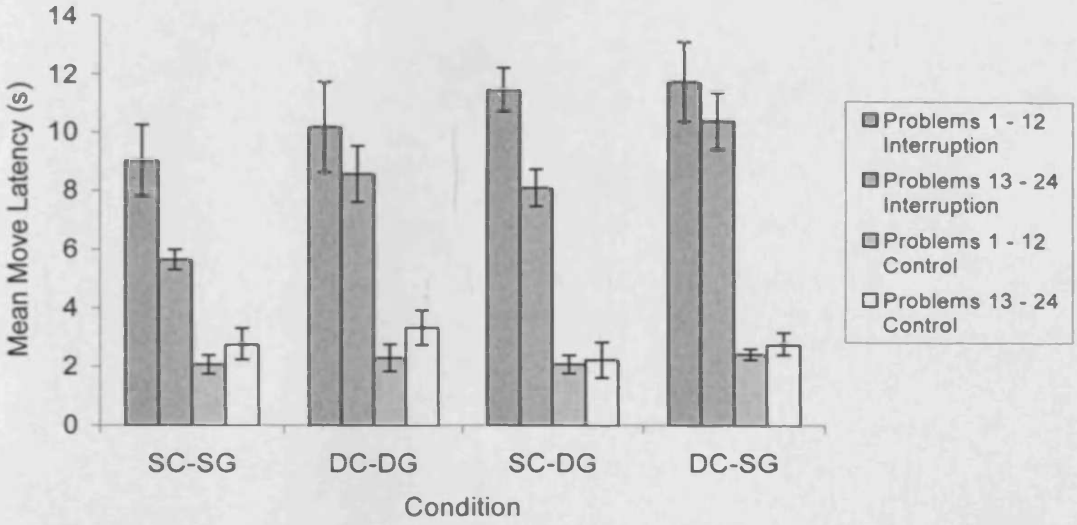

Figure 3.7. Mean move latencies (±SE) for trials 1 – 12 and trials 13 – 24 of Experiment 6

disc-colour in the goal-state is also consistent with that in the current-state, the cost of goal-directed enquiry is reduced markedly compared to all other conditions. This might only occur because of experience of being interrupted: specifically, the cost of enquiry is higher in the first experimental half because of the novelty of being interrupted. In the SC-DG condition, the goal is retrieved, but move execution is met with a higher level of caution, because disc colour in the goal-state is incongruent with disc colour in the current-state. However, with experience of interruption and problem solving in the task domain, this level of caution is reduced.

Resumption latency is similar for the SC-DG and DC-DG conditions in the second experimental half, and as such, may remove some credence from the above claim. However, these resumption latencies may be explainable by the differential costs of goal-directed enquiry. In the SC-DG condition, the suspended goal might still be reactivated from the associative activation provided by the colour priming cue, but following

retrieval, there is a cost of goal-directed enquiry because of the mixed disc-colours in the goal-state. In the DC-DG condition, goal retrieval through associative reactivation of a colour cue is met with failure, but the cost of evaluation is lower compared to the SC-DG condition, because of the congruence between disc-colour in both the current-and goal-state. Simply, participants are faster to execute a reconstructed move when disc colour throughout the problem space is consistent, and are slower to execute a retrieved move when disc colour between the current-state and goal-state is inconsistent. Different weightings of the times to complete both processes in each condition may result in similar resumption latencies. Additional support for this claim comes from resumption latency in the DC-SG condition, which shows little improvement across experimental half. Here, as in the DC-DG condition, goal retrieval through associative reactivation of a colour cue is met with failure, but the cost of goal-directed enquiry is higher because of the incongruence of disc colour in the goal-state. Even with experience of interruption and problem solving in the task domain, a similar level of caution is exhibited.

3.5.3.3. Redundant moves

Table 3.5 illustrates mean redundant moves following the intended point of interruption in interrupted and control ToH problems. Redundant moves are higher in control conditions for both the SC-SG condition and the DC-DG condition, and higher in interruption conditions for both the SC-DG condition and the DC-SG condition. Variance is high within each condition across trial type, again reflecting the range of redundant moves in each case: between 0 and 8 for interruption conditions, and between 0 and 6 for control conditions.

Table 3.5

Mean redundant moves (±SE) in Experiment 6

	SC-SG		DC-DG		SC-DG		DC-SG	
	M	SE	M	SE	M	SE	M	SE
Interruption	0.47	0.17	0.38	0.2	0.6	0.26	1.47	0.58
Control	0.81	0.38	0.81	0.51	0.38	0.2	1.16	0.38

Frequency of redundant move mean ranks were calculated for the SC-SG condition (interruption Mean Rank = 4.34, control Mean Rank = 4.56), the SC-DG condition (interruption Mean Rank = 4.19, control Mean Rank = 3.75), the DC-SG condition (interruption Mean Rank = 5.19, control Mean Rank = 5.34), and the DC-DG condition (interruption Mean Rank = 3.75, control Mean Rank = 4.88). Number of redundant moves did not significantly differ between interruption and control conditions when analysed by a Friedman test ($p = .21$). Participants execute a similar number of redundant moves following interruption compared to when they are not interrupted.

Numbers of redundant moves were quite similar between interruption SC-SG and interruption DC-DG conditions, but somewhat higher in the DC-SG condition compared to the SC-DG condition. Thus, the uniformity of disc-colour might affect performance following the resumption of an interrupted move; specifically if disc-colour is the same in the current-state and goal-state, fewer the redundant moves are executed, irrespective of disc-colour in the suspended visual array.

3.5.3.4. Move-error frequency

Table 3.6 presents move-error frequency data for all participants that committed one or more redundant moves out of a possible 256 interruption and control ToH problems. Move-error frequency appears very similar across trial type (interruption = 29, control = 32), except in the case of SC-DG, where a higher frequency of move-errors occurs in the interruption condition. Participants solved 297 problems error-free compared to 59 problems in which at least one redundant move was made, $Q(7, N = 16) = 14.56$, $p < .05$. However, pairwise comparisons using a Bonferroni correction ($p = .0125$) revealed that move-error frequency did not significantly differ across trial type for any of the conditions. As for redundant moves, move-error frequency is similar between interruption and control conditions.

3.5.4. Discussion

Experiment 6 has provided convincing support that disc-colour is an important cue that may be responsible for the subsequent retrieval of goal that has been promptly suspended by an interruption. Maintaining the colour of discs between the suspended and

reinstated current-state of an interrupted ToH problem leads to more timely resumption of the task. Changing disc-colours in the reinstated current-state of a ToH problem causes appreciable impairments to the time taken to resume, probably because of the

Table 3.6

Move-error frequency in Experiment 6

	SC-SG	DC-DG	SC-DG	DC-SG
Interruption	7	7	5	10
Control	7	3	10	12

Note. Maximum move-error frequency in each condition is 46 moves

abolishment of associative reactivation opportunity (e.g., Altmann & Trafton, 2002; Wolfe, 1994). Conditions in which disc-colour changed only in the current-state or only in the goal-state also caused longer resumption latencies. Such an effect can be explained again by the cost of goal-directed enquiry: it is higher because of the discrepancy between disc-colour in the current and goal states, but is unaffected by the method used to re-establish the suspended goal. Simply, the cost accrues after the suspended goal is retrieved or reconstructed.

In considering the differences between resumption latency data from Experiment 6 to that from Experiment 4, it might be the case that disc-colour compliments the effects of location-based priming. For example, advance knowledge of object location is sufficient to prime a response to that location, but advance knowledge of colour serves to further increase the speed of response (e.g., Humphries, 1981; Tsal & Lavie, 1993). Moreover, a colour feature cue might form a more effective prime in a goal-directed task than object location, because of its higher salience value compared to location (e.g., Sutcliffe, 1995; Wolfe, 1994). Maybe, disc colour was responsible for the effects attributed to location based priming in Experiments 4 and 5. That is, when discs within the reinstated current-state are in the same location as when suspended, but also presented in the same colour, the chances of reactivation of a suspended goal through associative priming are increased. Relocation of discs in the reinstated current-state might have rendered colour cue detection ineffective, because the change to the solution-sequence

caused participants to abandon strategies that might otherwise lead to retrieval (e.g., Anderson & Douglass, 2001). This idea will be revisited in Experiment 7.

Disc-colour, if used most often as a priming cue, is perhaps more effective in supporting the retrieval of a suspended goal, when location is not manipulated within the experimental design. In the SC-SG condition of Experiment 6, resumption latency was shorter compared to resumption latencies for the same conditions in Experiment 4 and Experiment 5. Even more compelling was that with experience, resumption latency in the SC-SG condition was 5.65 s in Experiment 6, compared to resumption latencies of between 8 and 9 s for the same condition in Experiments 4 and 5. These results harmonise with the idea of Wolfe and colleagues that selective-attention occurs to serve the demands of the task domain (Wolfe et al., 1989, Wolf et al., 1990; Wolfe, 1994). That is, selective-attention can be refined to a feature-cue that delivers the highest performance rewards because of its salience and uniqueness compared to other cues. Participants respond to the salience offered by disc colour in a promptly suspended ToH problem, and its use it more frequently than less efficient cues (e.g., location) increasing the chances of successfully retrieving a suspended goal through associative reactivation.

The role of spatial location and disc-colour is tested further in Experiment 7 by removing disc-colour from the visual array, whilst manipulating the location of discs in reinstated ToH problems (as in Experiments 4 and 5). By removing colour features from every trail (even those that are not interrupted), the spatial location of ToH discs may serve as the only available cue to cause resumption of a suspended goal through associative reactivation. If disc location determines how a goal is re-established, and does not just contribute to how long it is assessed before being executed, the pattern of resumption latencies should be similar to those in Experiment 4. The opposite effect is expected if disc-colour is the only route to associative priming, at least in the case of un-signalled interruption.

3.6. Experiment 7

3.6.1. Introduction

Twenty-four ToH problems were to be solved, and interruption occurred as in Experiment 4. However, each ToH problem contained discs that were all presented in the same colour (black), thus removing disc-colour cues that might have been responsible for the priming effects observed in Experiments 4 – 6. Reinstated problems were manipulated across the same 4 manipulations of disc location in Experiment 4, as to assess its role as a priming cue in the absence of disc-colour.

Predictions for Experiment 7 are as follows. If disc-location has a similar or more salient cueing status to disc-colour, resumption latency in all conditions of Experiment 7 should be similar or shorter to those in respective conditions of Experiment 4. On the other hand, increased resumption latencies in Experiments 7 compared to Experiment 6, would suggest that disc-colour is a better priming cue, and may have been responsible for the way that suspended goals were re-established in Experiments 4 – 6. If goal-directed enquiry cost is influenced by the congruence between location cues in the suspended and reinstated problem space, and not disc colour, resumption latencies should be longer for each condition in Experiment 7 compared to those in Experiment 4. Results consistent with this claim would provide compelling support for the role of colour in supporting the reactivation of a suspended goal, and location in partly determining how long a re-established move is evaluated before it is executed.

3.6.2. Method

3.6.2.1. Participants

Forty undergraduate students from the School of Psychology at Cardiff University participated in fulfilment of a course requirement. All were naive to the ToH task. Twelve were removed for failing to meet the point of interruption criteria specified in section 3.3.1.1, leaving 28 participants.

3.6.2.2. Apparatus and materials

The same program used in Experiment 4 was modified to accommodate the change to disc colour. Specifically, all discs were presented in black within each ToH problem

3.6.2.3. Design

The same design as in Experiment 4 except for the modification to disc-colour. For secondary tasks, mean completion times and mean digits correctly recalled did not significantly differ between any interruption condition. Descriptive data for secondary tasks is presented in Table B1 of Appendix B.

3.6.2.4. Procedure

The same as in Experiment 4.

3.6.3. Results and Discussion

3.6.3.1. Move latency

Figure 3.8 illustrates latency data for the first move following the intended point of interruption in interrupted and control ToH problems. Control move latency is very similar across all conditions, and resumption latencies are somewhat longer. Resumption latency is marginally shorter in the SC-SG condition compared to DC-DG condition, and is longer although very similar for SC-DG and DC-SG conditions.

For further analysis, a repeated measures, 2 (trial type) x 2 (current-state disc-configuration) x 2 (goal-state disc-configuration) ANOVA was conducted on move latency data. There was only a significant main effect trial type, $F(1, 27) = 419.86$, MSE $= 9.45$, $p < .001$, although current-state disc-configuration only marginally failed to reach significance ($p = .064$). There was a significant interaction between current-state disc-configuration and goal-state disc-configuration, $F(1, 27) = 9.36$, MSE $= 5.63$, $p < .010$, and a significant three-way interaction between all three variables, $F(1, 27) = 6.94$, MSE $= 7.97$, $p < .025$. A Bonferroni *post-hoc* analysis revealed a significant simple main

effect for resumption latency between interrupted SC-SG and DC-SG conditions, $F(1, 27) = 11.96, p < .005$: the likely cause of the interaction, as there were no other significant differences. Thus, resumption of an interrupted ToH problem when disc-configuration has not changed throughout the entire problem space is faster that when disc-configuration in only the current-state has changed. There is no significant difference between resumption latencies in conditions where disc-configuration in the reinstated problem is completely congruent or completely incongruent with that in the suspended state.

Resumption latencies were longer than respective move latencies in control conditions. Although descriptive data suggested differences between resumption latency across conditions, these were mostly non-significant, and modest in relation to comparable conditions in Experiment 4. In Experiment 7, there was no significant difference between resumption latencies for either of the SC conditions compared to the DC-DG condition suggesting that disc-location was not primed (at least sufficiently so) prior to interruption. The only differences between resumption latency occurred for the SC-SG and the DC-SG conditions. Perhaps participants find it more difficult to configure a move for execution when there is a discrepancy between disc-configuration in the current-state and goal-state. In support of this claim, the difference between move latency in the SC-SG and SC-DG condition almost reached significance ($p = .18$).

If the re-establishment of an interrupted goal is facilitated by the availability of priming cues (e.g., Altmann & Trafton, 2002), disc-location does not seem to fully support this process. Maintaining disc-location in the absence of disc-colour does not significantly support faster resumption of an interrupted task.

3.6.3.2. Move latency and experience

Figure 3.9 illustrates mean move latencies for interrupted and control problems across experimental half (i.e., trials 1 – 12 versus trials 13 – 24). There is a general increase in move latency between control trials after having solved 13 ToH problems. For interrupted trials, there is a reduction in resumption latency after 13 ToH problems, but only marginally so for the SC-SG condition. The reduction in resumption latency is most elevated in DC conditions, markedly so for the DC-SG condition, where resumption

latency after 13 ToH problems is similar to the respective SC-SG condition. Interestingly, resumption latencies are very similar for each interruption condition in the second experimental half, averaging 10.48 s.

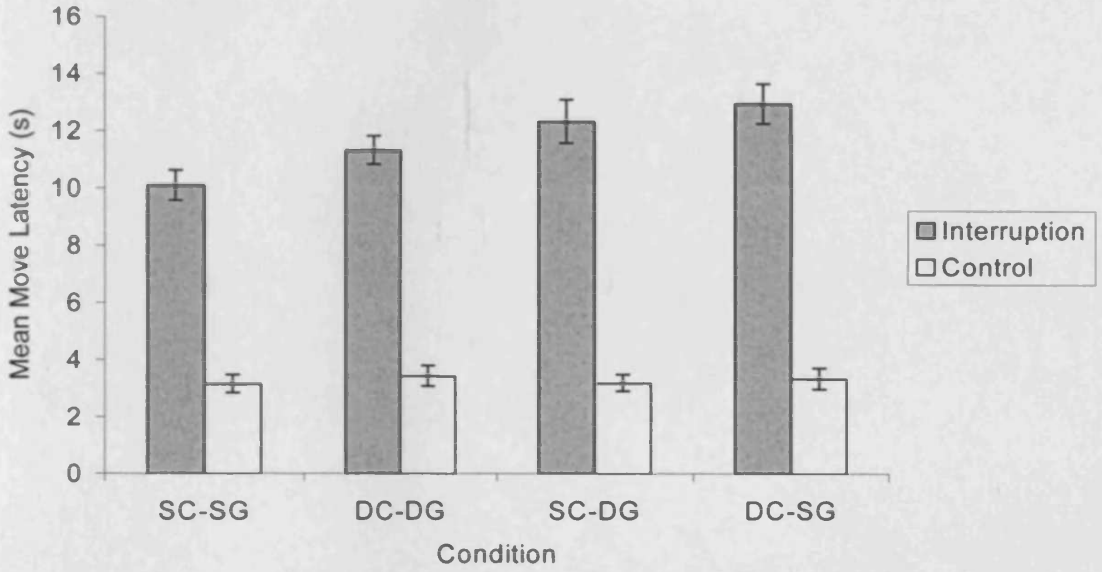

Figure 3.8. Mean move execution latencies (±SE) for interrupted and control trials in Experiment 7

Control move latency failed to significantly differ across experimental half, and was subsequently removed from further analysis. A logarithmic transformation was applied prior to analysis for the same reasons outlined in Section 3.3.3.2. Transformed resumption latency data only was analysed using a repeated measures 2 (experimental half) x 2 (current-state disc-configuration) x 2 (goal-state disc-configuration) ANOVA. This revealed a significant main effect only of experimental half, $F(1, 27) = 25.03$, MSE $= .12$, $p < .001$. There was also a significant two-way interaction between, current-state disc-configuration and goal-state disc-configuration, $F(1, 27) = 5.58$, MSE $= 1.5$, $p < .05$. A Bonferroni post-hoc was carried out to locate the cause of the two-way interaction. A significant simple main effect was found between resumption

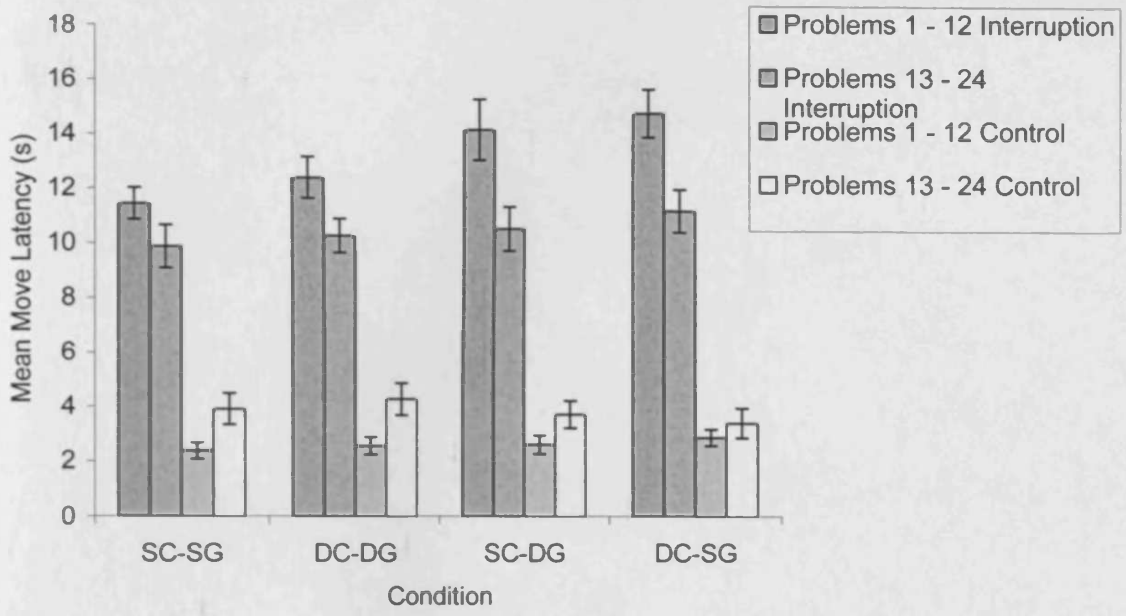

Figure 3.9. Mean move latencies (±SE) for trials 1 – 12 and trials 13 – 24 in Experiment 7

latency in the SC-SG and DC-SG condition, F (1, 27) = 7.38, $p < .025$. This does not reflect an improvement in resumption latency in one condition over the other, rather an extension of the findings in Section 3.6.3.1 that it takes longer overall to resume the DC-SG condition compared to the SC-SG condition.

Resumption latency did significantly decrease across experimental half. Although from descriptive data one might infer a larger reduction in resumption latency for the SC-DG and DC-SG conditions compared to the SC-SG and DC-DG conditions, there was a non-significant three-way interaction. The reduction in resumption latency, to the extent that all conditions are very similar, does however lead one to infer that participants settle upon a strategy of reconfiguring the problem space with experience. Specifically, there is no evidence of associative priming on the time taken to resume a suspended goal. If associative priming were to play a role in the resumption of a suspended goal, not only should resumption latency overall be shortest in the SC-SG condition, improvement in resumption latency should be greatest in this condition also. This was not the case, and

therefore it can be concluded that associative priming using disc location cues did not play a significant role in how a goal was re-established in Experiment 7.

3.6.3.3. Redundant moves

Table 3.7 illustrates mean redundant moves executed following the intended point of interruption in interruption and control conditions. A greater number of redundant moves were executed in all control problems compared to those that were interrupted. Redundant moves are however lowest in the SC-SG condition, but consistently longer in all conditions that contained some change to disc-configuration after they were suspended. Variance is high in each condition, reflecting the range of redundant moves executed across each trial type. These were between 0 and 7 redundant moves in interruption conditions, and between 0 and 10.5 redundant moves in control conditions.

Table 3.7

Mean redundant moves (±SE) in Experiment 7

	SC-SG		DC-DG		SC-DG		DC-SG	
	M	SE	M	SE	M	SE	M	SE
Interruption	0.7	0.28	1.27	0.43	1.32	0.43	1.46	0.51
Control	1.48	0.36	1.32	0.58	1.71	0.43	2.3	0.63

Frequency of redundant move mean ranks were calculated for the SC-SG condition (interruption Mean Rank = 3.73, control Mean Rank = 4.91), the SC-DG condition (interruption Mean Rank = 4.46, control Mean Rank = 4.95), the DC-SG condition (interruption Mean Rank = 4.36, control Mean Rank = 5.30), and the DC-DG condition (interruption Mean Rank = 4.29, control Mean Rank = 4). Number of redundant moves did not significantly differ between interruption and control conditions when analysed by a Friedman test ($p = .13$). Although, participants do seem to execute a greater number of redundant moves in control trials compared to those that are interrupted, there was no significant difference between trial type. It is however noted that redundant moves were markedly fewer in the interruption SC-SG condition compared to all other conditions. In agreement with a contention presented in

Experiment 6, after a goal is re-established, the quality of the solution-sequence is better if the discs are in the same location as when they were suspended throughout the visual array. Perhaps then, the activation of goals contained in a solution-sequence processed prior to interruption is higher in the SC-SG condition, and only re-initiation of this solution-sequence is disrupted if the first move is not primed by a cue.

3.6.3.4. Move-error frequency

Table 3.8 presents move-error frequency data for all participants that committed one or more redundant moves in ToH problems. There are no apparent differences between move-error frequency across trial type, apart from fewer moves executed in the interruption SC condition compared to its respective control condition.

Table 3.8

Move-error frequency in Experiment 7

	SC-SG	DC-DG	SC-DG	DC-SG
Interruption	13	19	19	17
Control	21	22	22	13

Note. Maximum move-error frequency in each condition is 56 moves

Of a possible 448 interrupted and control problems, participants solved 302 problems error-free compared to 146 problems in which at least one redundant move was made. A Cochran's Q test confirmed that move-error versus non-move error did not significantly differ ($p = .29$). As was the suggested by number of redundant moves committed in Experiment 7, move-error frequency did not significantly differ.

3.6.4. Discussion

It was hypothesised that when faced with un-signalled interruption, to support resumption through associative reactivation, participants might selectively extract a cue-property from the visual array that would result in the highest possible performance gain

(e.g., Wolfe, 1994). The results of Experiment 7 show that this might not always be the case, at least when the visual array is starved of a salient feature-cue such as colour. Disc-location, when not supported by a colour that might help distinguish disc-identity, does not sufficiently prime the representation of a suspended goal.

Participants may have quickly adapted to the task constraints by abandoning goal storage, and choosing to reconstruct the suspended goal from the visual array (e.g., Altmann & Trafton, 2002). Such behaviour is predicted by both ACT-R theory of goal-directed behaviour (e.g., Anderson & Lebiere, 1998) and the guided-search model of selective attention (Wolfe, 1994). Both suggest that a history of successful use of a particular object to support retrieval of a target will increase the likelihood that the same object is used in the future. For the current experiment, attempts at identification of an object (perhaps primed before task suspension) were usually met with a failure, thus rendering retrieval impossible (e.g., Altmann & Trafton, 2002). Under such conditions, guided-search predicts that selective-attention will terminate, resulting in no selection of an appropriate cue, and ACT-R predicts that a production rule will not be fired because conflict resolution is not satisfied resulting in a goal being forgotten. Simply, the current context does not provide enough source activation to boost the BLA of the suspended goal to a level for retrieval, and the BLA is not high enough for the goal to be reactivated in any other way.

Experiment 7 has provided compelling evidence that locating a disc-cue following interruption is a process facilitated by the disc-feature colour (e.g., Goolsby & Suzuki, 2001). At least for a ToH problem that is suspended by an un-signalled interruption, cue-location has little influence on priming an object for response, a contrasting result to studies of non-goal-directed selective attention (see Lavie & Tsal, 1994, for a review). Rather disc location may contribute to the time spent evaluating the utility of performing a move after it has been re-established. Evidence in support of this claim comes from shorter resumption latencies in Experiment 6 where disc location was not manipulated, compared to longer resumption latencies in Experiments 4, 5 and 7, where disc location was manipulated.

It is noted that within the Experiment 7, participants could have attempted to use disc features other than colour prime the representation of a goal suspended because of

interruption. For example, object shape, when heterogeneous, has been shown to improve the speed of responding in attention cueing paradigms, but usually only when object location is attended to first (e.g., Nissen, 1985; Tsal & Lavie, 1988, 1993). Such effects are however generally found when the shape of a cue prior to its offset is unique (e.g., a triangle) from other distractor objects (e.g., all ellipses). Disc-shape then, although a salient feature-cue in some tasks, should not have been salient enough to capture attention prior to interruption in the current experiment, because all discs are the same shape (e.g., Schneider & Shiffrin, 1977). Alternatively, given the rule of not being able to place a larger disc on top of a smaller disc in the standard ToH paradigm, disc-size might be a more salient feature-cue to support the associative priming of an interrupted goal. For example, attending to the smallest disc (disc 4) prior to interruption, might prime the representation that the smallest disc is that which is to be moved when encountering the reinstated primary task. Results from Experiment 7 suggest that this is not the case, at least in the case of un-signalled interruption where disc colour cues are not available within the visual array.

Finally, the interruption paradigm employed within the current experiments might not have been powerful enough to extract the effects of location as a priming cue to support associative activation. This conforms to the predictions of a recent model that suggests participants choose to offload memorial demands to the external visual array, if retrieval is perceived as too difficult a resumption strategy (e.g., Anderson & Douglass, 2001). This might be a rational behaviour performed by human beings when cognition and the demands imposed by the task domain do not easily support smooth transitions between tasks (e.g., Anderson, 1990). However, when the task does support resumption, perhaps through the use of a salient priming cue such as colour, retrieval is possible even at long retention intervals (e.g., Altmann & Trafton, 2002).

3.8. General Discussion

3.8.1. Summary of the controversy motivating Chapter 3

Chapter 3 addressed an important conceptual controversy regarding how a suspended goal is resumed when its suspension was caused by an interruption.

Specifically, the role of priming cues in causing reactivation of the goal's representation from memory. It is generally accepted that a suspended goal must be associatively primed before interruption initiation, with a reminder cue that can serve to reactive its decaying representation (e.g., Altmann & Trafton, 2002; Miller, 2002). However, according to one standpoint, priming has to occur in conjunction with rehearsal of the to-be-suspended goal for successful retrieval to ensue (Altmann & Trafton, 2002, 2004; Monk et al., 2002; Trafton et al., 2003). From another standpoint, strategic use of reminder cues can substitute the need to rehearse at a lower cost to overall performance (Chung & Byrne, 2004; Detweiler et al., 1994; Miller, 2002). The experiments presented in Chapter 3 addressed the role of priming cues in the speed of resuming a suspended goal and the quality of the subsequent task performance, in the absence of an opportunity to prepare for interruption

3.8.2. Overview of the main manipulations and key findings

Since empirical data from the ToH task is used to create and test the constraints in recent models of goal suspension and resumption (Altmann & Trafton, 2002; Anderson & Douglass, 2001), it was used as a primary task in Experiments 4 – 7. Cue properties (i.e., ToH discs) that might mitigate the retrieval efficiency of a suspended goal were manipulated following interruption, with little or no opportunity to rehearse the goal immediately before primary task suspension. Upon returning to the primary task after secondary task completion, the location of ToH disc cues (Experiments 4, 5 and 7) or the colour of ToH disc cues (Experiments 6 & 7) could have changed since problem suspension.

Changing only the location of discs, particularly within the current-state, attenuated the time to resume the suspended goal compared to when disc location was held constant (Experiments 4 and 5). This effect was augmented by experience of the task domain, but only when more than one reinstated condition contained discs in the same location of the current-state and the goal-state (Experiment 5). Priming using disc location might have however been mediated from the opportunity to distinguish the prime from an array of distractors because of colour features (Experiments 6 and 7). When disc-colour could be used to identify a cue in the absence of changes to disc location,

resumption latency was considerably shorter than when the colour-cue had changed, especially with experience within the task domain (Experiment 6). Resumption latency seemed to improve at 5.65 s in Experiment 6, compared to resumption latencies of between 8 and 10 seconds in all other experiments. Finally, retrieval of a suspended goal was abolished when the opportunity to selectively adapt to condition-action parings using disc-colour was removed from the task domain, even when a cue to disc location was available (Experiment 7). Throughout Experiments 4 – 6, participants were faster to resume an interrupted goal when cue availability (i.e., location and colour) was congruent between suspension and reinstatement in the current-state disc configuration. With incongurence to the reinstated goal-state disc configuration (Experiments 4 and 6), participants were slower to resume. With a reduction in the number of moves required to reach the reinstated goal-state, participants were faster to resume, but only if there was no change to a cue to guide resumption in the reinstated current-state (Experiment 5).

Overall, the results suggest that the representation of an interrupted goal can be reactivated by the source activation provided by the current context, even in the absence of a period to rehearse the suspended goal. However, participants appear to adapt to the task constraints by extracting the most salient cue as a prime; most likely a feature of the primed object (colour) rather than a property of the solution-sequence (location).

3.8.3. Priming cues and reactivating a suspended goal

Generally, resumption latency indicates that participants are able to remember some aspect of a goal suspended because of interruption, and that this effect is mediated by the use of ToH discs as priming cues (e.g., Altmann & Trafton, 2002). A number of factors may determine the efficiency of resumption using associative priming cues, each of which lends well to some predictions of a selective-attention model and ACT-R based memory models described in detail in Sections 3.2.5 – 3.2.7.

Given that SC-SG conditions (and SC conditions in general) are resumed faster than DC conditions in Experiments 4 – 6, some aspect of the suspended current-state disc configuration must serve to increase the speed of resumption. It is likely upon detection of a salient feature cue (e.g., colour) conveying the identity of a condition-action pairing,

the decaying representation of the suspended goal is reactivated, probably resulting in retrieval (e.g., Altmann & Trafton, 2002; Wolfe, 1994).

The above argument has implications for ideas regarding how an interrupted goal should be prepared for resumption prior to its suspension. Specifically, the proposition of Altmann and colleagues, that associative priming has to occur during a period of focused preparation (e.g., rehearsal) immediately prior to primary task suspension, is not supported by the current results. Instead, it appears that a suspended goal can be prepared for interruption preattentively through selective extraction of a salient feature cue from the visual array (e.g., Wolfe, 1994). That such a strategy is deployed with little experience of being interrupted suggests that it is an automatic response to a task that is promptly suspended because of interruption (e.g., Miller, 2002). Nevertheless, resumption efficiency through associative priming of a salient feature property can develop with experience of the task domain and experience of being interrupted: a result fitting to the history of use component of ACT-R theory (e.g., Lovett & Anderson, 1996).

The implications of the current findings do not condemn the idea that cue availability is an essential prerequisite if a goal is to be efficiently retrieved through associative priming (Altmann & Trafton, 2002, p. 53). Neither do they impart that a priming cue is processed at any time other than before interruption initiation and immediately following secondary task completion (Altmann & Trafton, 2002, p. 66). Rather, they extend such claims to suggest that certain cue features rather than cue availability *per se* can cause the reactivation of a goal suspended because of interruption. Within the ToH task, a colour-cue that distinguishes a to-be-suspended disc from all distractors, seems to be at the core of eliciting associative priming effects. Removing colour from the visual array, and allowing selection to be made upon disc-location (or size), is not sufficient to cause priming (Experiment 7). Priming effects are much more powerful when colour is available to aid this selection process (Experiment 6). For a cue to successfully prime the representation of a suspended goal in the ToH task, it has to be 'available' within the suspended and reinstated visual array, but also has to be easily distinguishable from other potential distractor cues. Specifically, it has to be associated only with the suspended goal and no other goals (e.g., Altmann & Trafton, 2002).

The effective use of colour as a priming cue in the ToH task might be refined by its history of use (e.g., Lovett & Anderson, 1996). For example, the move that was interrupted within all experiments in Chapter 3 was one that can be executed using the perceptual-heuristic, 'one-follows-two'. Specifically, following movement of disc two to a peg, a simple rule of moving disc one to the same peg can be applied, arguably with no need to remember the goal (e.g., Anderson & Douglass, 2001; Simon, 1975). Such a rule has to be learned, else participants tend to approach the move using an inefficient problem solving strategy such as 'hill-climbing' (Anzai & Simon, 1979). For all Experiments in Chapter 3, participants exhibited behaviour to suggest that they could not apply this rule in conditions where disc location and/or disc colour had changed between problem suspension and reinstatement.

Results also impart that when a reinstated problem does not contain a cue that was an original feature of the suspended problem space, either selective search or goal retrieval is abandoned. Selective search for a previously attended feature might be a relatively automatic response in a reinstated task, whereas retrieval of a suspended goal is likely to be determined by the success of the outcome of the search. That is, if the search locates a target that meets the conditional demands of retrieval (an associated priming cue), it should be terminated, although substituted by attempted goal retrieval. If on the other hand, the search terminates with no target selected, or returns the wrong target, retrieval of the correct goal will be seriously impaired. Given that selective search seems to be responsible for shorter resumption latencies (i.e., possible retrieval) in SC conditions, I propose that it is always deployed when a suspended problem is reinstated. Goal retrieval is abandoned when conditions of retrieval are not met: a negative return of conflict resolution because of failed, degree of match, production strength, and data refractoriness.

Reattending to the primed colour (e.g., green) when the task is reinstated, only causes reactivation of the suspended goal if that colour remains a property of the disc corresponding to the move that was interrupted. According to the guided-search model of selective-attention, the colour of an object can be processed preattentively, and may serve to guide attention toward a target when a response is required (Wolfe, 1994). Since within Experiments 4 – 6, the colour of the target disc (green) did not stand out from

heterogeneously coloured distractors, it was selectively attended more frequently with experience of causing the retrieval of a suspended goal. However, if disc colour was selected preattentively before primary task suspension, and reattended automatically upon reinstatement, there should have been no differences between resumption latencies within Experiments 4 and 5. Specifically, in some conditions, only the location of the disc has changed and not the condition-action pairing (i.e., the smallest or green disc always had to be moved to the peg containing the next smallest or blue disc). That participants were much slower to resume in conditions where disc location had changed in the current-state imparts an influence of knowledge of the suspended solution-sequence on the willingness to retrieve a goal through associative priming.

3.8.4. The Resumption latency: too long to exclusively represent retrieval of a suspended goal?

Given the goal-directed nature of the ToH task, it is possible that within the current experiments a direct measure of the time to retrieve a representation of a suspended goal was masked by other factors. Move latency data across Experiments 4 – 7 provided strong support that participants evaluate the expected gain of executing a move that has been re-established in an interrupted ToH problem; perhaps a ramification of goal-directed enquiry. The likelihood that the re-established goal is deployed in the form of a physical move, is dependent upon the time taken to evaluate whether or not the move will have a positive effect on subsequent task performance (e.g., Einstein & McDaniel, 1996). That retrieval like behaviour is still exhibited in SC conditions compared to reconstruction/reconfiguration in DC conditions supports goal retrieval through associative priming. Nevertheless, the time to execute a move was still much higher that retrieval estimates in current models of goal-directed behaviour (e.g., one to three seconds in the goal-activation model, Altmann & Trafton, 2002).

Manipulating the congruence between disc properties in only the suspended and reinstated current-state or only in the suspended and reinstated goal-state appreciably increased resumption latency, and negatively affected the quality of the solution-sequence. Such results were only expected in the former case, where the availability of a priming cue had changed. That they occurred in both conditions suggests that even when

a goal is retrieved through associative priming with a cue in the current-state, elements of the suspended goal-state disc configuration are also remembered, and these impact upon speed of resumption. Strong evidence for this effect comes from resumption latency data in Experiment 4. If participants were only processing the means-ends cue associated with the suspended goal prior to its suspension there is no reason why they should recognise a change to the goal-state disc configuration (e.g., Altmann & Trafton, 2002). It too must have been processed prior to interruption, with a representation strong enough to cause such a delay.

Uncertainty in performing an action following interruption is by no means a novel idea (e.g., Adams et al., 1995; Speier et al., 2003), especially when interruption un-signalled (see McFarlane, 2002 for a review). Within problem solving behaviour, humans often check whether the current-state is on the path to the correct solution (e.g., Atwood & Polson, 1976). Additionally, participants often use 'break' opportunities in the ToH task to serve goal-directed enquiry (usually at subgoal and subtask boundaries (e.g., Anderson & Douglass, 2001; Davies, 2003). Results from the current experiments suggest that if a suspended goal is retrieved through associative priming with a previously processed priming cue, the cost of goal-directed enquiry may be low. Lower at least, than if the goal has to be reconstructed from the visual array. However, This process might be influenced by the time taken to recurse through the goal-stack (Experiment 5)(e.g., Anderson 1983, 1993), although having retrieved rather than reconfigured the suspended goal renders the cost of goal-directed enquiry lower.

Perhaps within an interrupted ToH task, move execution might not immediately follow goal retrieval. Take for example Experiment 5, in which not only the configuration of discs changed between primary task suspension and resumption, but also the number of moves to goal. Increasing error-free completion requirements from seven moves to nine moves had no adverse effect on resumption latency when reinstated disc configurations were the same or different to when they were suspended. By contrast, decreasing the number of moves to goal, from seven to five, largely reduced resumption latency. Let us assume that in the SC-DG condition participants were able to retrieve the representation of a suspended goal. Let us further assume that resumption latency in the second experimental half is more reflective of the efficient use of priming cues, and

hence retrieval. For Experiment 5, mean resumption latency was 1.61 s shorter in the SC-DG-2 condition compared to the 8.75s resumption latency in the SC-SG condition. Presumably, this reduction in move latency occurred because the LOOP disc no longer had to be considered in goal-directed enquiry. If we then apply this logic to the resumption latency for in the second experimental half of the SC-SG condition of Experiment 6, resumption without the need to consider the LOOP disc could be in the region of four seconds. This estimate is more fitting with retrieval parameters set by current models of goal suspension and retrieval (e.g., Altmann & Trafton, 2002; Anderson & Douglass, 2001).

After considering the potential time for goal-directed enquiry, actual re-establishment of the suspended goal in the SC-SC condition may be more in the region of three to four seconds. Goal retrieval time in the ToH task is estimated to take between one to three seconds in the goal-activation model. This is when, then task has not been interrupted, sufficient associative priming has occurred, and is taken from data generated by participants who solved problems error-free (and are very experienced in the task domain). Thus, resumption latencies in SC-SG conditions of Experiments 4 – 6 might consist of retrieval of a suspended goal, but the time for goal-directed enquiry masks such effects. This is likely to be an inherent property of goal directed tasks such as the ToH (e.g., Anzai & Simon, 1979; Simon, 1975), and thus there is no reason to be concerned about its occurrence in problems that have been interrupted.

3.8.5. Future directions

Although results suggest that an explicitly cued interruption lag is not necessary for associative priming to occur, a preparation period might increase (1) the likelihood of attending to a priming cue and (2) the strength of the connection formed (e.g., Altmann & Trafton, 2002). There was no way to be certain within the current experiments that participants were always attending to a priming cue prior to interruption. Participants may have been preoccupied with solving the ToH problem rather than deliberately trying to attend to information that may support the retrieval of a suspended goal (Anderson and Douglass, 2001). In addition, there was no way to be certain that even when participants were attending to a priming cue, the attended object was always the disc corresponding to

the to-be-suspended move. Attending to the wrong cue prior to primary task suspension would have affected, for example, whether the correct goal could be retrieved upon primary task reinstatement. Additionally, during goal-directed enquiry, participants might have realised the futility of performing the originally suspended move: that its execution would not seem to logically advance the solution-sequence. Additionally, there is no way to be sure that provision of an interruption lag immediately prior to interruption would not strengthen the associative connections formed between a priming cue and a suspended goal.

Any of the above factors could have affected resumption latencies in Experiments presented in Chapter 3, but were of no concern to the original questions addressed. Specifically, I was interested if an interrupted goal could be retrieved without an explictly cued preparation period and whether manipulation of disc properties served to mediate the opportunity to retrieve a suspended goal in interrupted ToH problems. Chapter 4 sets out to explore the influence of various types of preparation on performance within interrupted ToH problems. The role of the interruption lag is addressed, and results might serve to compliment or extend those found in Chapter 3.

3.8.6. Summary

Experiments 4 – 7 compliment findings from Chapter 2 in suggesting that a suspended goal is encoded prior to suspension, even when interruption is un-signalled and falls within a complex goal sequence. The encoded representation might not be complete, resulting in partial retrieval of the suspended goal. Through exposure to changes in disc properties prior to and following interruption, particularly colour, participants learn to extract and encode a cue prior to interruption that is most likely to reactivate the decaying representation of a suspended goal. Furthermore, with experience in the task domain and experience of dealing with interruption, participants become more effective at opportunistically encoding a suspended goal using associative activation; behaviour reflected by even shorter retrieval latencies. Linking the suspended goal to an externally represented cue that is both salient (i.e., represents that goal and no others) and predictive of past retrieval success increases the likelihood of faster resumption and future use (e.g., Anderson & Douglass, 1998; Wolfe, 1994).

Chapter 4

EMPIRICAL SERIES 3

Preparing an interrupted goal for resumption: The role of the 'interruption lag'

4.1. Overview

The opportunity to prepare a suspended goal for retrieval during the *interruption lag* – the transitory period between an interruption alert and secondary task initiation (Altmann & Trafton, 2002) – was manipulated across three experiments. With a two-second interruption lag (alerted by a 'blinking' visual dot), the time to resume a ToH problem was reduced only with experience in the task domain (Experiment 8a). A two-second interruption lag (alerted by an auditory tone designed to be less disruptive) markedly reduced resumption latency, and this tendency was yet more marked after experience in the task domain (Experiment 8b). Increasing the length of the interruption lag to four-seconds had no additional impact on resumption latency (Experiment 9). Results accord with the idea that an interruption lag supports faster resumption of a suspended goal (Altmann & Trafton, 2002), although the degree of gain in resumption time relative to the interruption lag is discussed.

4.2. Introduction

4.2.1. General introduction

With the rapid development of communication technology, humans are more than ever faced with a dilemma of how best to manage interruptions. Office-based interruptions (e.g., email, *Instant Messenger*) are often attended to immediately (Jackson et al., 2002, 2003), even though guidelines advise delaying their uptake until task demands are low (McFarlane, 2002). In more safety-critical work environments (e.g., the flightdeck), handling of interruptions cannot be delayed because of the consequences of untimely response (e.g., Dismukes et al., 1998; Loukopolous et al., 2003). Although in some tasks, such as human-human conversing, people can and do often choose to finish what they were doing (e.g., speaking a sentence) before attending to another information source (e.g., Clark, 1996).

In coordinating an interruption into an ongoing activity, the *negotiated* mode (where the person has control over when to take up interruption) usually results in the fewest performance deficits, such as, the efficiency in resuming the suspended task, with *immediate* interruption (un-signalled and delivered at the earliest opportunity) ranked worst. Overall, allowing human operators some control over when to initiate an interruption, for instance through an opportunity to consolidate interrupted-state information, usually leads to better overall performance efficiency in the primary task (McFarlane, 2002).

However, there is often little control over how interruption is co-ordinated into an ongoing task (McFarlane, 1998, 2002). Imagine you are reading an important document and a colleague knocks on your office door. The interruption lag – the time between the onset of the interrupt alert (the first knock) and secondary task initiation (conversing with your colleague) – is a period in which the suspended task might be strategically prepared for resumption (Altmann & Trafton, 2002). Strategic preparation could involve, memorisation of the interrupted goal (e.g., through rehearsal), associating the goal with a contextual reminder cue, or recording information on a physical medium (e.g., a 'post-it' note). Preparation could serve to trigger memory for the suspended goal in the future, reducing the *resumption latency* – the time taken to return to an interrupted task (e.g.,

Trafton et al., 2003). Neither entirely immediate nor negotiated, a brief interruption lag – a period suited to preparing a suspended goal for resumption – might allow humans to deal with interruptions in a timely manner whilst minimising costs to performance efficiency.

Chapter 4 addresses the issue of how the ToH is performed with an interruption lag. To my knowledge, no single study has investigated empirically the impact of the interruption lag – an apparent window of opportunity to encode postponed goals – on retrieval of suspended ToH goals. Would this period be used to better encode the suspended goal, perhaps even allowing time to associatively prime its representation with a reminder cue: both processes increasing the likelihood of faster and more efficient resumption? Conversely, might there be a trade-off in the relation between the additional time spent preparing a goal for resumption and benefits in terms of a reduction in resumption latency?

4.2.2. Memory for goals and interruption management

The classic conceptualisation of memory for goals is the goal stack, which assumes that goals are encoded, stored, and retrieved hierarchically in a last-in first-out (LIFO) order (e.g. Anderson et al., 1993). The time taken to retrieve a goal is determined by where it resides on the stack, with superfluous activation rendering the goal stack immune to decay-based forgetting. This idea of a perfect goal memory system is met with apprehension, given evidence of long resumption latencies for goals suspended because of interruption; latencies that cannot exclusively reflect goal retrieval (e.g., Altmann & Trafton, 2004; Hodgetts & Jones, 2004). There is also abundant research supporting, goal forgetting (e.g., Chung, 2004; Byrne & Bovair, 1997) and goal-selection errors (e.g., Anzai & Simon, 1979; VanLehn, 1991).

The received view is that goals are suspended and resumed like other declarative memory items, and as such, are vulnerable to the effects of decay, and interference from other goals (e.g., Altmann & Trafton, 2002; Anderson & Douglass, 2001). Each time a goal is processed there is an increase in its *base-level activation* – activation determined by how often it has been processed in the past and the time elapsed since it was last

processed. To govern behaviour, a goal has to be sufficiently activated through recurrent strengthening – a process that renders the goal more resilient to decay-based forgetting.

Through empirical investigation of goal-directed memory in the ToH task and a cognitive model, Anderson and Douglass (2001) provide evidence that goals are processed in an opportunistic non-memory intensive manner. They demonstrated that move latency increases as a function of time since a goal was suspended, and provide a compelling case that humans frequently abandon goal storage processes altogether and instead rely on external cues at the point of resumption. Conversely, Altmann and Trafton (2002) introduce a functional role of activation through encoding to cause a suspended goal to exceed the activation levels of competing goals. The form of their goal-activation model also departs from Anderson and Douglass (2001) in suggesting that the only way to reactivate the decaying representation of a suspended goal is through *associative priming*. Priming a goal with a salient reminder cue prior to its suspension will cause rapid reactivation of its decaying representation when the cue is reattended, thereby helping to overcome retroactive interference generated by intervening goals. Associative priming has to occur in synchrony with goal suspension; else, future retrieval will be unlikely. Consequently, there must be a period before goal suspension and secondary task initiation where encoding a suspended goal for future retrieval can occur – the interruption lag (Altmann & Trafton, 2002). In the ToH task, the obvious candidate-priming cue is the disc corresponding to the suspended goal (Altmann & Trafton, 2002). Predictably therefore, there is a tightly coupled relationship between cue availability and goal encoding, both prior to and following interruption.

4.2.3 Empirical investigations of rehearsal and the interruption lag

Rehearsal normally improves memory performance (e.g., Atkinson & Shiffrin, 1968). However, the influence of rehearsal on remembering interrupted information is an issue of continued controversy (see McFarlane, 2002, for a comprehensive review).

Allowing opportunity to rehearse primary task information prior to secondary task initiation can reduce resumption latency (Czerwinski et al., 1991b; Detweiler et al., 1994), especially if the secondary task demands little cognitive effort (Gillie & Broadbent, 1989). Participants were twice as fast to resume a VCR task suspended by

unfilled retention intervals compared to when having to track aircraft, reflecting possibly the greater opportunity to rehearse during the unfilled retention interval (e.g., Monk et al., 2004). Conversely, some studies show no benefit from opportunity to rehearse. For example, rehearsal had no effect on the speed of resuming ones place in a computer-based-shopping task when secondary tasks were similar to those in the primary task or demanded complex processing (Gillie & Broadbent, 1989). Participants may also be reluctant to rehearse unless instructed to do so (e.g., Czerwinski et al., 1991a, 1991b; Detweiler et al., 1994). Failure to use rehearsal to prepare a suspended goal for retrieval has since been interpreted to reflect strategic adaptation to constraints imposed by the context, and not a deficit in remembering. Humans routinely offload mnemonic demands prefer to pay the cost of having to reconstruct suspended goals (e.g., Anderson & Douglass, 2001), and opportunistically use reminder cues instead of memory (Miller, 2002).

The utility of a period in which to prepare resuming a suspended task is also a subject of controversy, one fuelled largely by the variation in method used to promote preparatory behaviour (McFarlane, 2002). With a visual alert, a three-second delay before secondary task initiation, and instructions to remember ones place in mathematical equations, performance was improved only slightly if the suspended goal involved high but not low memory load (Detweiler et al., 1994). At delays of 30-seconds before interruption, memory for a position suspended in a spacecraft-monitoring task did not show significant improvement (Czerwinski et al., 1991a), unless participants were instructed explicitly to use the delay to prepare (Czerwinski et al., 1991b). Because of the perceived cost of encoding and storage of the suspended activity there may be a tendency to rely more on display-related cues that did not impose a burden on memory (e.g., Anderson & Douglass, 2001; Wilson, 2002). For instance, when instructed to rehearse information that could have promoted smoother return to a radar operator task participants often used less memory intensive strategies, such as pointing toward the vicinity of the suspended activity (Miller, 2002).

Recent research has stressed the positive utility of an interruption lag to reducing the time spent resuming an interrupted task (Altmann & Trafton, 2004; Clifford & Trafton, 2004; Trafton et al., 2003). Nevertheless, the apparent benefits to resumption

time of an interruption lag are usually more than offset when the sum of these two intervals is taken into account. Eight-second interruption lags signalled by a flashing visual alert were shown to reduce primary task resumption latency of a complex resource allocation task when interrupted by 30-second tactical assessment secondary tasks (Trafton et al., 2003). At two and four seconds, resumption latency was unaffected if cues were more salient during an interruption lag, with marginal benefits at six seconds, and a significant reduction at eight-seconds (Altmann & Trafton, 2004). Closer examination of data from both studies revealed that the time taken to prepare a goal for retrieval through strengthening and associative priming was longer than the time saved resuming the task.

In general, what research is available has tended to show that the goal-activation model may have underestimated the time required to prepare a suspended goal for retrieval during an interruption lag (Altmann & Trafton, 2002). Empirical data have shown that one to two seconds may be insufficient to prime a suspended goal for resumption (at least in complex tasks demanding high processing resources having few obvious priming cues). However, a longer interruption lag can come at a cost of increasing the time spent handling an interruption; specifically, there is an unequivocal reduction in resumption latency. Nevertheless, the original priming parameters set in the goal-activation framework were modelled on empirical data in the ToH task and perhaps one to two seconds is sufficient to encode such goals for retrieval without imposing an additional cost.

4.2.4 Signalling the interruption lag: making the most of the alert

It is reasonable to assume that a fully articulated interruption lag will be better received and used fuller if it is alerted effectively. A number of studies highlight the importance of effective alerting, particularly warning, to attract attention toward informative information (e.g., Edworthy, Loxley & Dennis, 1991; Patterson, 1990; Posner, Snyder & Davidson, 1980).

Patterson (1990) describes four major stages to processing a warning stimulus: *notification*; *perception*; *understanding*; and *compliance*. Logically, in order to notice a warning signal it has to be distinguishable from other stimuli (e.g., abrupt changes in

luminance, pitch, timbre, and so on) to attract attention when attending to other information. Following notification, a warning has to be perceived, that is, given meaning and relevance in memory, else it might be discarded. Once perceived, the warning must be understood in terms of what information it is trying to convey (e.g., the colour red to convey danger or an indication to 'stop'). Finally, compliance refers to the requirement that the waning signal causes a desired behaviour (e.g., stopping at the traffic light when it is on red).

On signalling an imminent interruption, an effective alert should attract attention whilst minimising disruption to the ongoing task (e.g., Müller & Rabbit, 1989; Nugent & Obermayer, 2000a, b). In their Attention Allocation Subsystem (AAS) designed for multi-modal Navy command and control workstations, Nugent and colleagues provide a number of guidelines for effective warning systems. An effective warning will: be relevant to task goals; capture attention appropriately; support reorientation to a task; allow some control over its handling; and have minimal disruptive properties (Obermayer & Nugent, 2000). For instance, if a warning requires full attention, but its associated action can be delayed, the most effective mode is a *blinking border* – a visually flashing perimeter around the to-be-performed action (Nugent & Obermayer, 2000a, b). Alternatively, an auditory icon (e.g., a tone or speech) may be just as effective as attracting attention whilst allowing the task to continue (Nugent & Obermayer, 2000a, b).

4.2.5 The current experiments

Previous chapters have highlighted factors that might intensify the performance deficits caused by being interrupted, and in the current chapter I wanted to explore if such deleterious effects could be alleviated. In chapter 2, performance deficits were made worse when interruption fell during an uncompleted subtask rather than at the end of subtask, and were worsened further by similar processing requirements between the primary and secondary task. These results showed that primary task processing demands at the point of interruption hinder the formation and maintenance of a suspended goal. This effect was augmented further by interference from the secondary task presumably impinging upon inadequate knowledge of the primary task solution-sequence. In Chapter

3, participants were able to strategically adapt to un-signalled interruption by utilising salient priming cues to remind of the relevance of a suspended goal.

Experiments 8a and 8b were designed to test directly the impact of a two-second interruption lag (alerted by a visual or an auditory cue) on the time taken to resume a ToH problem interrupted at different points within its solution-sequence. Experiment 9 was designed in response to recent findings that a longer interruption lag might be required to encode a suspended goal to a sufficient degree to enable future retrieval.

4.3. Experiment 8a

4.3.1. Introduction

In this experiment, twenty-four ToH problems requiring nine moves for error-free completion were to be solved, and eight were interrupted by ToH secondary task problems. Four interruptions fell during an uncompleted ToH subtask (*within-subtask*) and another four at the end of a ToH subtask (*end-subtask*) – two of each indicated by a two-second interruption lag. In line with a recent study, the onset of an interruption lag was alerted by a visually presented 'blinking' dot to increase the salience of the onset of an interruption lag (Trafton et al., 2003). To promote natural development of strategies used to handle interruptions, participants were not instructed to prepare during the interruption lag; they were told only that primary task controls would be disabled during this time.

There are contrasting views on how a suspended goal might be re-established if its suspension was preceded by an interruption lag. On one view, encoding a goal by strengthening its representation in memory and associating it with a priming cue, should markedly reduce resumption latency (Altmann & Trafton, 2002). Effective association using cue-based identification might abolish the difficulty in distinguishing a suspended goal because of task demands at the point of interruption and should protect against retroactive interference from previously encoded goals. An alternative view is that an interruption lag will not cause faster resumption if participants abandon goal encoding prior to interruption (e.g., Anderson & Douglass, 2001). In this case, resumption latency

might only consist of the time to reconstruct the problem space, and should be largely unaffected by retroactive interference. Reconstructing a within-subtask goal might involve reconsidering intermediate states in the interrupted solution-sequence (i.e., previously processed goals and future goals), and end-subtask conditions might only involve reconsideration of the suspended goal.

4.3.2. Method

4.3.2.1. Participants

Forty-four undergraduates from the School of Psychology at Cardiff University participated in fulfilment of a course requirement. All were naive to the ToH task prior to testing. Eight participants were removed (leaving a total of 36) for failing to produce ToH solution-sequences in interruption and control problems that could cause interruption initiation within 4 moves of the intended point of interruption (see Footnote 2 of Section 2.4.2.1).

4.3.2.2. Apparatus and materials

The same program used in Experiment 3 was modified. Four of the interrupted ToH problems contained a two-second interruption lag (2 within-subtask and 2 end-subtask), whereas interruption was un-signalled in another 4 problems. In interruption lag conditions, a visually presented orange dot (4cm x 4cm in diameter) – the *interrupt alert* – would appear 2 cm underneath the centre peg of the main ToH problem for 600 ms. The dot appeared to blink by inserting 45 ms inter stimulus intervals (ISI) every 170 ms until offset. For the duration of the interruption lag, the primary task remained on the screen, but all primary task controls were inactive. At the end of the interruption lag, the primary task visual display was replaced by a secondary task, which upon completion caused the reinstatement of the primary task and the reactivation of all task operations. Secondary tasks were the ToH problems used in Experiments 1 – 3.

4.3.2.3. Design

The design was the same as in Experiment 3 with some modification. In Experiment 8a, there were twenty-four primary task ToH problems, and eight were interrupted by secondary task ToH problems that could be solved error-free within seven moves. A new repeated measures variable, *interruption lag*, replaced *secondary task type* from Experiment 3. In interruption lag conditions, configuring ToH discs to satisfy the initiation of an interruption would initiate the interruption alert and cause all operations to become temporarily inactive for two-seconds. Two of the four within-subtask interruption conditions and two of the four end-subtask interruption conditions contained an interruption lag. The remaining four interruption conditions (2 within-subtask and 2 end-subtask) did not contain interruption lags; interruption was un-signalled. For ease of explanation hereafter, the four interruption conditions will be referred to as: *within-lag* (within-subtask, interruption lag), *within-no-lag* (within-subtask, no interruption lag), *end-lag* (end-subtask, interruption lag), and *end-no-lag* (end-subtask, no interruption lag). In secondary tasks, mean completion times and moves to completion did not significantly differ between any of the interruption conditions. Descriptive statistics are presented in Table C1 of Appendix C. Dependent variables were the same as in Experiment 3.

4.3.2.4. Procedure

The same procedure as in Experiment 3 was used, with some modification. Participants were instructed that in some ToH problems, a 'blinking' orange circle might appear in the visual array to indicate impending interruption and that thereafter, all primary task operations would be disabled until the onset of a secondary task. One practice trial contained a visually alerted interruption.

4.3.3. Results and Discussion

4.3.3.1. Move latency

Figure 4.1 displays mean latency data for the first move following the intended point of interruption, which is clearly inflated following an interruption. Resumption latency is markedly longer when interruption fell within a subtask compared to

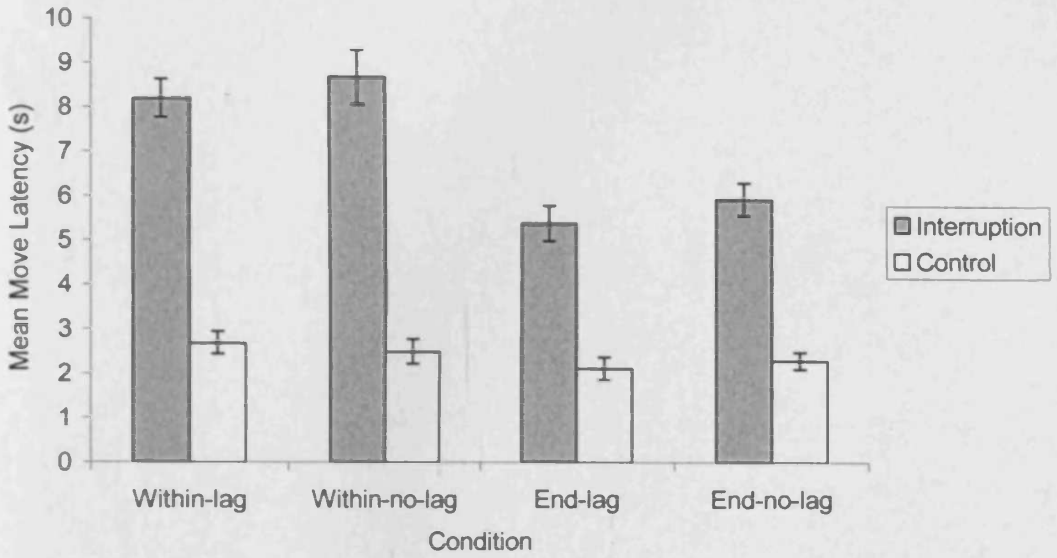

Figure 4.1. Mean move execution latencies (±SE) for interrupted and control trials in Experiment 8a

at the end of a subtask (irrespective of interruption lag), with a similar although less pronounced difference in control conditions. With an interruption lag resumption latency is only 0.5 s shorter in the within-subtask condition and 0.6 s shorter in the end-subtask condition.

A repeated measures 2 (trial type) x 2 (intended point of interruption) x 2 (interruption lag) ANOVA was applied to resumption latency data. In the case of Experiments 8a – 9 Bonferroni *post-hoc* tests were used to investigate significant interactions. The ANOVA revealed significant main effects of trial type, $F(1, 35) = 199.89$, $MSE = 7.8$, $p < .001$, intended point of interruption, $F(1, 35) = 35.82$, $MSE = 4.89$, $p < .001$, with no effect of interruption lag. Interruption resulted in significantly longer resumption latencies than move latencies in comparable control trials; an effect that increased in the within-subtask condition but did not reduce because of an interruption lag at either point of interruption. Trial type and intended point of interruption interacted significantly, $F(1, 35) = 19.07$, $MSE = 5.62$, $p < .001$, due to a significant difference between resumption latency in interruption conditions only, F (1,

35) = 32.21, $p < .001$. A within-subtask move took longer to execute than an end-subtask move; a difference markedly amplified because of interruption.

Interruption caused resumption latencies that were significantly longer than control move latencies; augmented when interruption fell within a subtask, but not alleviated with an interruption lag. A brief opportunity in which a suspended goal could be encoded for future retrieval actually added to the time spent dealing with interruption.

4.3.3.2. Move latency and experience

Figure 4.2 presents mean move latency data for interrupted and control problems across for the first 12 ToH problems and the last 12 ToH problems. Control move latency shows no reliable improvement; in fact, there is a slight increase in the end-no-lag condition. Resumption latency generally decreases across experimental half, with reduction highest in the end-lag condition and lowest in the within-no-lag condition. Overall, the reduction in resumption latency seems to be more amplified in interruption lag conditions compared to conditions with un-signalled interruption.

Control move latency was removed from further analysis for failing to exhibit any significant differences between trial types. Due to occasional violations of assumptions of linearity, and evident heteroscedasicity (especially in within-subtask conditions), all move latency data were transformed using a logarithmic function that served to normalise the linear relationship between the variability in each condition. Transformed data only were analysed using a repeated measures 2 (experimental half) x 2 (point of interruption) x 2 (interruption lag) ANOVA. Unless otherwise stated, all interactions were investigated using Bonferroni post-hoc tests. There were significant main effects of, experimental half, F (1, 35) = 11.94, MSE = .13, $p < .005$, intended point of interruption, F (1, 35) = 26.38, MSE = .33, $p < .001$, and interruption lag, F (1, 35) = 4.93, MSE = .26, $p < .05$. Generally, resumption latency reduced across experimental half, especially with an interruption lag, although overall was still longer in within-subtask conditions. Experimental half and interruption lag interacted significantly, F (1, 35) = 8.89, MSE = 0.1, $p < .010$, due to a significant reduction in resumption latency across experimental half in conditions with an interruption lag, F (1, 35) = 19.47, $p < .001$. Participants benefited only from an opportunity to encode an interrupted goal with experience of an interruption lag with effects exhibited by shorter resumption latency.

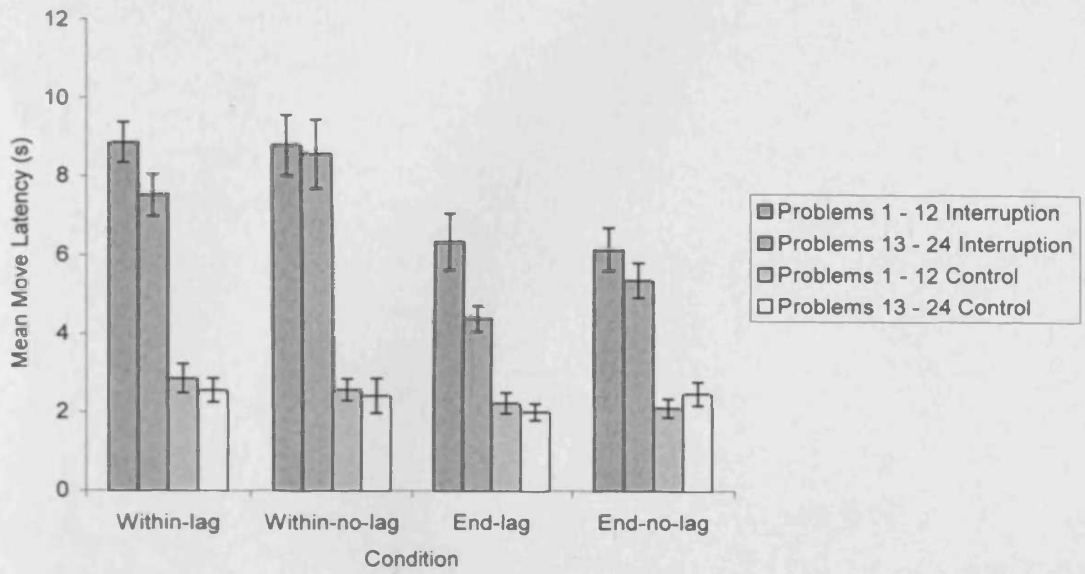

Figure 4.2. Mean move latencies (±SE) for trials 1 – 12 and trials 13 – 24 of Experiment 8a

Speed of resumption was increased with encoding experience during an interruption lag, perhaps because of more efficient preparation strategies, subsequently causing faster resumption. ACT-R's expected gain parameter anticipates such an effect: experience of preparation strategy leading to shorter resumption latencies increased the likelihood of the same strategy being used in the future (e.g., Anderson & Lebiere, 1998). With limited experience in the task domain, the blinking alert might have drawn attention away from the primary task, contaminating the opportunity to encode a suspended goal. However, the reduction in resumption latency was higher in the end-lag condition compared to the within-lag condition. Such an effect may be explained by associative priming using ToH discs as cues to reactivate the representation of a suspended goal together with processing demands at the point of interruption. In contrast to the within-subtask condition, the end-subtask condition offers a cue that is highly salient (the second LOOP disc), and servicing the goal is not dependent on the processing of other intermediate goals. With experience in the task domain, a two-second preparation period may be sufficient to encode a goal for retrieval in the end-subtask condition but not in the within-subtask condition.

4.3.3.3. Redundant moves

Table 4.1 illustrates mean redundant moves executed following the intended point of interruption in interrupted and control ToH problems. Redundant moves are greater in number in interrupted within-subtask conditions than in control conditions, and in contrast are greater in control end-subtask conditions compared to respective interrupted conditions. Variance within each condition is very high, reflecting the range of redundant moves across each trial type: between 0 and 11 moves for interruption conditions, and between 0 and 11.5 moves for control conditions.

Table 4.1

Mean redundant moves (±SE) within primary task ToH problems following the intended point of interruption in Experiment 8a

	Within-lag		Within-no-lag		End-lag		End no-lag	
	M	SE	M	SE	M	SE	M	SE
Interruption	1.06	0.4	0.88	0.26	0.86	0.31	0.67	0.33
Control	0.68	0.2	0.65	0.18	1.21	0.32	1.32	0.50

Frequency of redundant move mean ranks were calculated for the within-lag condition (interruption Mean Rank = 4.43, control Mean Rank = 4.47), the within-no-lag condition (interruption Mean Rank = 4.49, control Mean Rank = 4.49), the end-lag condition (interruption Mean Rank = 4.44, control Mean Rank = 5.26), and the end-no-lag condition (interruption Mean Rank = 4.04, control Mean Rank = 4.38). A Friedman test indicated a non-significant across trial type ($p = .35$). These data suggest a similar number of redundant moves in interrupted and control ToH problems, irrespective of the intended point of interruption and the inclusion of an interruption lag.

4.3.3.4. Move-error frequency

Table 4.2 presents move-error frequency data for all participants that committed one or more redundant moves out of a possible 576 interruption and control ToH problems. Frequencies of move-errors are similar in interruption and control trials. Participants solved more ToH problems error-free (461 problems) compared to problems

in which at least one move-error was committed (115 problems), although this difference was non-significant (Q = 7.83, $p = .35$). Move-error frequency data provides evidence that participants made a similar number of single move errors that may have led to redundant moves following resumption of interrupted ToH problems compared to those that were not interrupted.

Table 4.2

Mean move-error frequency (±SE) within primary task ToH problems following the intended point of interruption in Experiment 8a

	Within-lag	Within-no-lag	End-lag	End-no-lag
Interruption	15	13	14	10
Control	18	13	20	12

Note. Maximum error frequency in each condition is 72 moves

4.3.4 Discussion

The results of Experiment 8a provide further evidence that interrupting a ToH problem leads to slowing in resumption time. Task demands at the point of interruption determine the speed of resumption following an interruption: resumption latency was shorter when interruption occurred at the end of a subtask than when it occurred during a subtask (e.g., Monk et al., 2002). Interestingly, resumption latency was reduced with a brief opportunity to encode a suspended goal, but only with experience in the task domain. The quality of the solution-sequence did not show any such improvement.

That an interruption lag had no impact on resumption latency, at least without experience in the task domain, might seem to contrast against the predictions of the goal-activation model (Altmann & Trafton, 2002). One interpretation of this effect is that a suspended goal is neither sufficiently strengthened nor associated with a priming cue during an interruption lag. However, the goal-activation model suggests that association between a goal and a priming cue is formed by co-occurrence – the goal and the cue have to enter focal awareness at the same time (e.g., Anderson & Reder, 1999). As such,

inexperience of problem solving or interrupt handling should not affect the ability to associate a goal with a cue, given that the cue should have been attended to when encoding the goal prior to primary task suspension.

It is important to take into account other factors before discounting the goal-activation model as a conceptual viewpoint for explaining resumption latency in Experiment 8a. The relationship between encoding a suspended goal and faster retrieval might only be realised after experience of an interruption lag, subsequently causing more frequent and efficient use of goal encoding strategies (e.g., Anderson & Lebiere, 1998). With low expected gain, encoding may be infrequent or even abandoned, resulting in slower and more reconstructive-like behaviours at resumption (e.g., Anderson & Douglass, 2001). Another possibility is that a suspended goal is encoded and subsequently retrieved, but move execution is delayed because of goal-directed enquiry brought on by the uncertainty associated with the event of being interrupted (e.g., Speier et al., 2003). Participants may pause to ask themselves 'where they are' in relation to the longer-term solution-sequence before committing to move execution. This idea may be applied to the current data, where even with experience of an interruption lag, resumption latency was reduced largely when fewer goals required consideration at the point of interruption (i.e., end-subtask conditions). Maintaining an active representation of a suspended goal, especially when primary task demands are high may also be more difficult with retroactive interference from secondary task goals (e.g., Speier et al., 2003). Finally, one to two seconds may not be adequate to sufficiently prime a suspended goal for retrieval (e.g., Altmann & Trafton, 2004; Trafton et al., 2003); a possibility that will be considered in Experiment 9. In sum, the goal-activation model might support resumption latency data from Experiment 8a, although its priming parameters may need fleshing out to support data generated by novice ToH problem solvers.

A shortcoming of Experiment 8a might have been the mode of alerting the onset of the interruption lag. Perhaps the visual blinking dot by its salience impeded on the ability to prepare the suspended goal for retrieval. It was expected that the alert would signal participants toward the onset of an interruption lag, but instead it may have had the undesired effect of drawing attention away from the primary task (e.g., Posner, 1980). That the alert blinked for 600 ms could have disrupted cognitive continuity within the

primary task and impinged upon the time to encode a goal for retrieval. Thus, one to two seconds might be enough time to encode a goal for retrieval, but only if the preparation period itself is uncontaminated by disruptive stimuli. This possibility is considered in Experiment 8b.

4.4. Experiment 8b

4.4.1. Introduction

The design of Experiment 8a was replicated, except that the mode of alerting the onset of an interruption lag was an auditory tone. The rationale for this modification was that the visual alert used in Experiment 8a might have diverted attention away from the primary task for up to 600 ms because of its animated status. This might have had the undesired effect of contaminating the two-second interruption lag; impinging upon the time to encode a suspended goal for resumption. An auditory annunciation signal does not require that participants divert their visual attention from the primary task to a peripheral part of the visual array, and may be useful when the visual system is overloaded (e.g., Nugent & Obermayer, 2000a, b).

If encoding a suspended goal during an encoding opportunity was disrupted by a visual cue in Experiment 8a, shorter resumption latencies should result with a less attention-diverting auditory alert in Experiment 8b. This might be manifested without experience in the task domain, and would provide compelling support for the priming constraint in the goal-activation model (Altmann & Trafton, 2002). Alternatively, if the duration of the alert (i.e., 600 ms), and not the mode of presentation, impinges upon the available time for encoding, resumption latency should be similar to that in Experiment 8a. Finally, if participants choose to not take up the opportunity for encoding, resumption latency might be unaffected by an interruption lag irrespective of the alert mode.

4.4.2. Method

4.4.2.1. Participants

Forty undergraduates from the School of Psychology at Cardiff University participated in fulfilment of a course requirement. All were naive to the ToH task prior to testing. Four participants were removed (leaving 36) for failing to meet the point of interruption criteria specified in Section 2.4.2.1.

4.4.2.2. Apparatus and materials

The program used in Experiment 8a was modified so that an auditory tone alerted the onset of an interruption lag. The tone was a pure tone digitised at 16-bit resolution, at a sampling rate of 48kHz using *SoundForge 5.0* software (*Sonic Foundry, Inc., Madison, WI; 2001*) and edited digitally to have an exact duration of 600 ms. The tone was presented through *Sony MDR CD250* headphones at a comfortable level of roughly 65 dB(A).

4.4.2.3. Design

The same design as in Experiment 8a except that the interruption lag was signalled by a 600 ms pure tone played through headphones. In secondary tasks, mean completion times and moves to completion did not significantly differ between any of the interruption conditions. Descriptive statistics are presented in Table C2 of Appendix C.

4.4.2.4. Procedure

Everything was the same as in Experiment 8a except that participants were instructed that an auditory tone would indicate impending interruption. The tone was played through headphones worn throughout all practice and experimental trials.

4.4.3. Results and discussion

4.4.3.1. Move latency

Figure 4.3 displays mean latency data for the first move following the intended point of interruption, which is greatly inflated in interruption conditions. Resumption latency following un-signalled interruption is only marginally shorter in end-subtask conditions compared to within-subtask conditions. However, with an interruption lag, resumption latency is shorter at both points of interruption, although is appreciably shorter for the end-subtask condition compared to the within-subtask condition. There is a 1.4 s reduction in resumption latency in the within-lag condition compared to the within-no-lag condition, with a 2.74 s reduction in respective end-subtask conditions.

A repeated measures 2 x 2 x 2 ANOVA revealed significant main effects of, trial type, $F(1, 35) = 316.45$, $MSE = 5.06$, $p < .001$, intended point of interruption, $F(1, 35) = 16.24$, $MSE = 4.66$, $p < .001$, and interruption lag, $F(1, 35) = 25.33$, $MSE = 3.13$, $p < .001$. Resumption latency was significantly longer than control move latency; longest following interruption to a goal contained within a subtask, and was reduced with an interruption lag. There was a significant two-way interaction between trial type and interruption lag, $F(1, 35) = 17.73$, $MSE = 4.14$, $p < .001$, due to a reduction in resumption latency only for conditions with an interruption lag, $F(1, 35) = 23.59$, $p < .001$. Intended point of interruption and interruption lag also interacted significantly, $F(1, 35) = 4.28$, $MSE = 2.95$, $p < .05$. With an interruption lag, resumption latency was significantly shorter in the end-subtask condition compared to the within-subtask condition, $F(1, 35) = 20.94$, $p < .001$; a difference only marginally non-significant without an interruption lag ($p = .078$). Overall, the time to resume an interrupted goal was reduced by the inclusion of an interruption lag, but was markedly reduced in the end-subtask condition.

A two-second interruption lag seems sufficient for a suspended goal to be strengthened and associated with a salient priming cue leading to faster resumption of the primary task (in support of Altmann & Trafton, 2002). With an auditory alerted

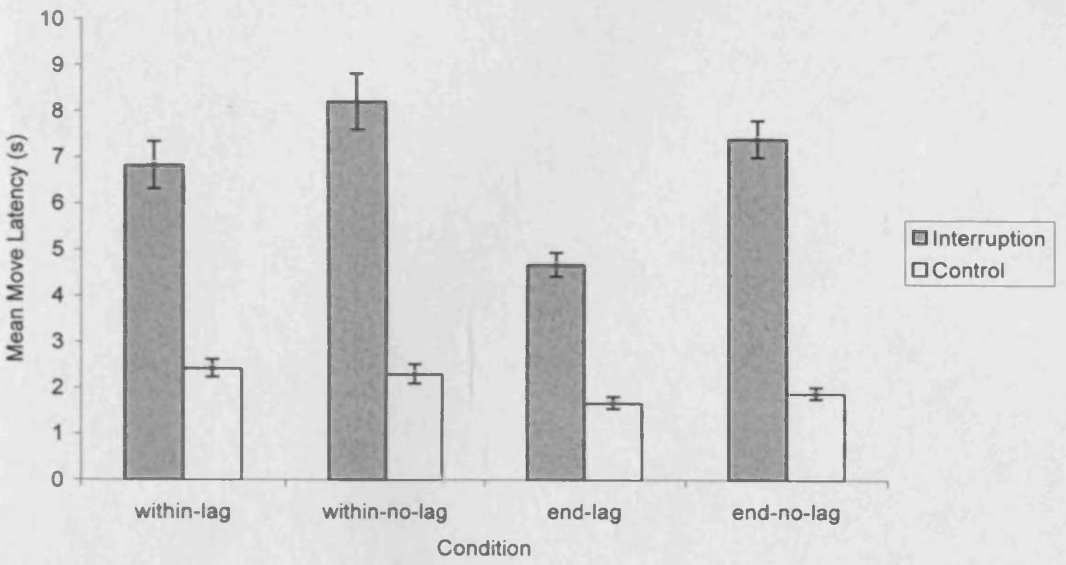

Figure 4.3. Mean move execution latencies (±SE) for interrupted and control trials in Experiment 8b

two-second interruption lag, the time taken to resume an interrupted ToH problem was reduced considerably compared to when interruption was un-signalled. This was especially the case when interrupted fell at the end of a subtask, although the time to resume an uncompleted subtask was also reduced significantly.

Encoding a suspended goal for retrieval during the interruption lag seems to be improved because of the less disruptive mode of alert in Experiment 8b (an auditory tone) compared to Experiment 8a (a visual blinking dot). However, encoding and/or retrieval is still adversely affected by task demands at the point of interruption and/or the similar processing demands between primary and secondary tasks. Resumption latency is longer when interrupted within a subtask, perhaps reflecting a weaker association between the suspended goal and a cue compared to the end-subtask conditions, that might increase vulnerability to retroactive interference from a similar secondary task.

Overall, the cost of the interruption lag is offset by the reduction in resumption latency when interrupted at the end of a subtask but not when interrupted within a subtask. This fits with the idea that a priming cue is more salient in the end-subtask

condition, and is thus more likely to stand-out upon primary task reinstatement causing faster resumption than in the within-subtask condition (e.g., Altmann & Trafton, 2002).

4.4.3.2. Move latency and experience

Figure 4.4 illustrates mean move latency data for interrupted and control problems across experimental half. Control move latency does not diminish, although is consistently shorter in end-subtask conditions. With un-signalled interruption, resumption latency is reduced marginally, although performance varies considerably for each experimental half. In contrast, resumption latency is reduced considerably with an interruption lag, with less performance variation than in un-signalled interruption conditions. For the within-lag condition, resumption latency is shorter by the second experimental half (MD = 2.63 s), and is similarly shorter than the resumption latency in the within-no-lag condition in the second experimental half (MD = 2.17 s). Although resumption latency is reduced to a lesser extent in the end-lag conditions (MD = 1.57 s), it is considerably shorter in relation to resumption latency in the end-no-lag condition by the second experimental half (MD = 2.84 s).

Control move latency was removed from further analysis for failing to exhibit any significant differences between trial types. A repeated measures 2 x 2 x 2 ANOVA revealed significant main effects of, experimental half, $F(1, 35) = 20.03$, $MSE = .17$, $p < .001$, intended point of interruption, $F(1, 35) = 5.04$, $MSE = 3.1$, $p < .05$, and interruption lag, $F(1, 35) = 36.37$, $MSE = 2.17$, $p < .001$. Resumption latency was reduced considerably in the second experimental half, more so for the within-subtask condition with an interruption lag. Experimental half and interruption lag interacted significantly, $F(1, 35) = 7.88$, $MSE = .11$, $p < .10$. An interruption lag caused a larger reduction in resumption latency across experimental half, $F(1, 35) = 18.66$, $p < .001$, although resumption latency significantly improved with no interruption lag, $F(1, 35) = 6.29$, $p < .025$. Intended point of interruption and interruption lag also interacted significantly, $F(1, 35) = 7.55$, $MSE = .22$, $p < .10$, because with an interruption lag, resumption latency was significantly shorter in the within-subtask condition compared to the end-subtask condition, $F(1, 35) = 14.65$, $p < .001$. This distinction between

resumption latency and intended point of interruption was not apparent when interruption was un-signalled.

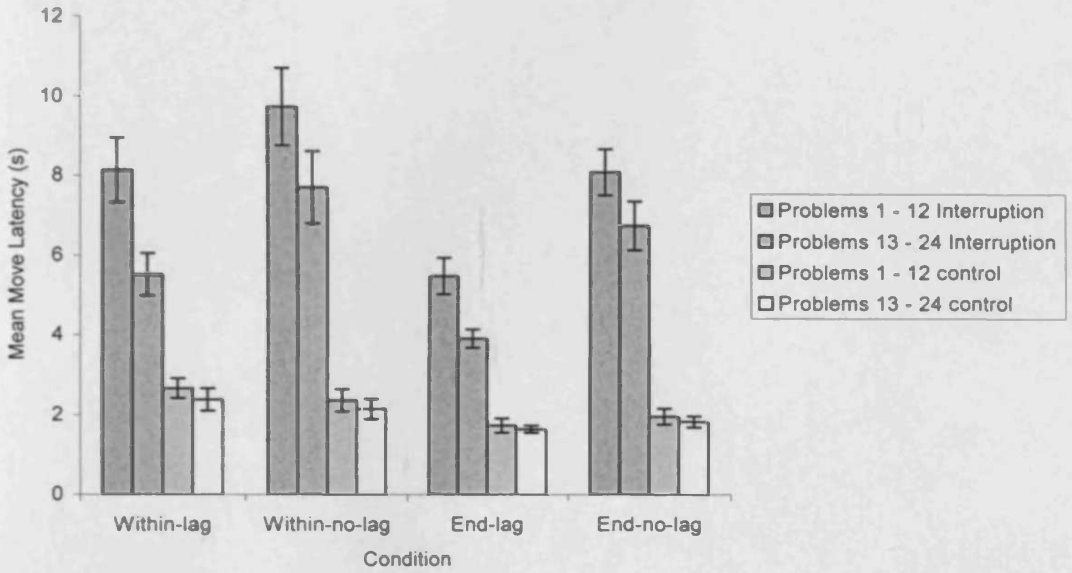

Figure 4.4. Mean move latencies (±SE) for trials 1 – 12 and trials 13 – 24 of the test phase of Experiment 8b

Resumption latency across experimental half was shorter when there was a brief opportunity to prepare a goal for retrieval even before experience in the task domain. Perhaps participants were able to associate a suspended goal with a retrieval cue with little experience of problem solving in the task domain, although better encoding might have been augmented with more experience (e.g., knowledge of task rules). Alternatively, more efficient encoding during an interruption lag may develop with experience of faster resumption of a suspended goal; an expected gain of preparation leading to faster resumption (e.g., Anderson & Lebiere, 1998). This is the more plausible explanation, since resumption latency improves also in conditions where interruption was un-signalled (e.g., Trafton et al., 2003). The idea here is that encoding strategies may be used more often and more efficiently in conditions with an interruption lag, but perhaps the utility of encoding influences performance in conditions in which interruption is un-signalled. That is, participants may be able to encode a suspended goal without an interruption lag, for example, during the 500 ms delay between clicking on the buttons to

execute a previous move and the trajectory of that move. Such an account explains why, with experience, participants become faster at resuming interrupted ToH problems without an interruption lag, but does not dismiss the interruption lag as a better opportunity to prepare.

Overall, an appropriately alerted delay in which a suspended goal can be prepared for interruption results in faster resumption of that goal compared to conditions in which interruption was un-signalled. This interpretation is reinforced by lower performance variability in conditions with an interruption lag. Participants are (1) more likely to produce shorter resumption latencies because of preparation during an interruption lag and (2) less likely to prepare as frequently or as efficiently in conditions with no interruption lag.

4.4.3.3. Redundant moves

Table 4.3 illustrates mean redundant moves executed following the intended point of interruption in interrupted and control ToH problems. Redundant moves are fewer in number for within-subtask interruption conditions compared to respective controls, although are larger in number in end-subtask interruption conditions compared to respective controls. Variance within each condition is very high, reflecting the range of redundant moves across each trial type: between 0 and 8 moves for interruption conditions, and between 0 and 5.5 moves for control conditions.

Table 4.3

Mean redundant moves (±SE) within primary task ToH problems following the intended point of interruption in Experiment 8b

	Within-lag		Within-no-lag		End-lag		End no-lag	
	M	SE	M	SE	M	SE	M	SE
Interruption	0.56	0.26	0.49	0.23	0.93	0.42	1	0.36
Control	0.69	0.18	0.72	0.22	0.75	0.25	0.49	0.18

Frequency of redundant move mean ranks were calculated for the within-lag condition (interruption Mean Rank = 4.49, control Mean Rank = 4.86), the within-no-lag condition

(interruption Mean Rank = 4.25, control Mean Rank = 4.72), the end-lag condition (interruption Mean Rank = 4.32, control Mean Rank = 4.56), and the end-no-lag condition (interruption Mean Rank = 4.50, control Mean Rank = 4.31). A Friedman test indicated a non-significant across trial type ($p = .85$). These data suggest that the number of redundant moves is similar in interrupted and control ToH problems, irrespective of point of interruption or the inclusion of an interruption lag.

4.4.3.4. Move-error frequency

Table 4.4 presents move-error frequency data for all participants that committed one or more redundant moves out of a possible 576 interruption and control ToH problems. Move-error frequency is similar between interruption and control trials. Participants solved more ToH problems error-free (471 problems) compared to problems in which at least one move-error was committed (105 problems), although this difference was non-significant ($Q = 7.93$, $p = .34$).

Taken with redundant move data, move-error frequency data provides convincing evidence that participants made a similar number of mistakes following resumption of all interrupted ToH problems compared to those that were not interrupted.

Table 4.4

Mean move-error frequency (±SE) within primary task ToH problems following the intended point of interruption in Experiment 8b

	Within-lag	Within-no-lag	End-lag	End-no-lag
Interruption	13	11	12	10
Control	16	20	13	10

Note. Maximum move-error frequency in each condition is 72 moves

4.4.4. Discussion

The priming constraint of the goal activation model receives persuasive support from the pattern of resumption latency when interruption was signalled by a two-second interruption lag alerted by an auditory tone. Generally, resumption latency was reduced

appreciably in conditions with an interruption lag compared to conditions in which interruption was un-signalled. With an auditory alert, maximum facilitation of an interruption lag is associated with shorter resumption times, especially when interrupted at the end of a subtask. With un-signalled interruption, there is some reduction in resumption latency with experience in the task domain; possibly reflecting a carryover effect from conditions with an interruption lag (in support of Trafton et al., 2003). Perhaps the number of retrieval successes associated with better encoding of a suspended goal (during lag trials) causes participants to strategically adapt to un-signalled interruption by preparing at other opportunities (e.g., during the move prior to that which is interrupted). Being interrupted, even when complimented with an opportunity to encode the suspended goal, had no effect on improving the quality of the solution-sequence, suggesting no effect of interruption after resumption of a suspended goal.

When interrupted at the end of a subtask, the cost of the interruption lag was offset by shorter resumption latency, an effect augmented by experience in the task domain. Although the interruption lag imposed a cost when interrupted within a ToH subtask, this was compensated for with experience of encoding a suspended goal for retrieval. With limited experience in the task domain, participants are able to form and maintain an association between a salient priming cue, but only with experience in the task domain is such behaviour exhibited when a suspended goal requires complex encoding.

Although resumption latency was reduced with an interruption lag, and reduced further with experience, it was still beyond that predicted by the goal-activation model to retrieve a *sufficiently* primed goal from memory (i.e., one to two seconds). Perhaps, one to two seconds allows better encoding of a suspended goal, but is not sufficient to cause the rapid reactivation boost predicted by the goal-activation model. Consequently, a partial retrieval may occur, causing a timely process of conflict resolution to locate the appropriate production rule to fully reactivate the goal (e.g., Siegler, 1987, 1988). Alternatively, participants may retrieve the representation of a suspended goal, but delay its execution until they are certain of its utility to the longer-term solution-sequence. Since resumption latency in the end-lag condition was taken to only two-seconds above

that in a control trial, the cost may merely reflect the time taken to reconfigure primary task processes (e.g., Monsell, 2003)

One question remains outstanding. Would a longer interruption lag cause a greater reduction in resumption latency (as suggested by Trafton et al., 2003), and if so, would this come at a cost to the overall time spent handling the interruption (e.g., Katz, 1995)? In experiment 9, a parametric manipulation of the length of the interruption lag was employed to address this issue.

4.5. Experiment 9

4.5.1. Introduction

Experiment 9 set out to replicate the findings from Experiment 8b, whilst considering the impact of length of encoding opportunity on resumption performance. It has been shown elsewhere that a minimum of six-seconds might be required for an interruption lag to significantly reduce resumption latency (Altmann & Trafton, 2004; Clifford & Altmann, 2004). However, faster resumption with a longer interruption lag may be a consequence of the task, that is, if priming cues are harder to locate and will take longer to associate with a suspended goal (Trafton et al., 2003). Conversely, the goal-activation model asserts that one to two-seconds should be more than enough time to encode a suspended goal for retrieval in the ToH task – in which a priming cue is realised with each action performed (Altmann & Trafton, 2002).

Within Experiment 9, a parametric manipulation of the length of the interruption lag was employed (2 s vs. 4 s). Even in a controlled task such as the ToH, the utility of encoding through associative priming may be better exposed over longer periods. The goal-activation model may have underestimated the time required to sufficiently encode a suspended goal for retrieval, in which case a longer interruption lag may lead to shorter resumption latency. The benefit for a longer interruption lag might be expressed in within-subtask conditions where cue availability is reduced and processing demands are

high. Nevertheless, a brief two-second encoding opportunity may be sufficient to encode a goal prior to and following interruption. This prediction is based on there being limited cues (four ToH discs), and the assumption that with each move executed, the cue for the next move is automatically attended (Altmann & Trafton, 2002).

4.5.2. Method

4.5.2.1. Participants

One hundred participants from the School of Psychology at Cardiff University took part in fulfilment of a course requirement. All were naive to the ToH task prior to testing. Half were randomly selected to serve in group 1 (two-second interruption lag) and the other half in group 2 (four-second interruption lag). Overall, twelve were removed (see section 2.4.2.1 of Chapter 2 for exclusion criteria) leaving 88 participants considered for further analysis (48 in group 1 and 40 in group 2).

4.5.2.2. Apparatus and materials

Everything was the same as in Experiment 8a except that the interruption lag interval was either two-seconds (for group 1) or four-seconds (for group 2).

4.5.2.3. Design

The same design as in Experiment 8b except for the one modification. A mixed repeated-measures design was employed: one group of participants (*Group 1*) only ever experienced 2 s interruption lags, whereas a second group (*Group 2*) only ever experienced 4 s interruption lags. This was to prevent a 'carryover effect' of preparation interval. With a repeated-measures design, small differences in preparation interval may not be realised, causing participants to overestimate the length of a short preparation period and underestimate the length of a long preparation period (Altmann, 2004).

In secondary tasks, mean completion times and number of moves to completion did not significantly differ across groups or between any of the interruption conditions. Descriptive statistics are presented in Table C2 of Appendix C.

4.5.2.4. Procedure

Everything was the same as in Experiment 8b, except that the third practice trial contained an interruption lag 2 s in duration for group 1 and 4 s in duration for group 2.

4.5.3. Results and discussion

4.5.3.1. Move latency

Figure 4.5 displays mean latencies for the first move executed after the intended point of interruption for group 1 and Figure 4.6 displays comparable move latencies for group 2. For both groups, control move latency is higher in within-subtask conditions compared to end-subtask conditions. Resumption latency is also considerably longer than control move latency. For interruption conditions, resumption latencies follow a similar pattern across group; the only difference being a higher resumption latency in the end-no-lag condition in group 1. Across groups, resumption latency is highest in the within-no-lag condition and is shortest in the end-lag condition. For group 1, resumption latency is reduced with an interruption lag in within-subtask conditions (MD = 1.3 s), although the reduction is higher for end-subtask conditions (MD = 1.92 s). Within group 2, resumption latency is also reduced with an interruption lag in within-subtask conditions (MD = 1.55 s), with a much smaller reduction (compared to group 1) in end-subtask conditions (MD = 0.65 s). In neither group is the gain in resumption time relative to the interruption lag, (although the cost is only marginal for the end-lag condition in group 1).

Resumption latency data were analysed using a mixed-design 2 x 2 x 2 ANOVA, with group (1 vs. 2) serving as the between-subjects variable. There were significant main effects of, trial type, $F(1, 86) = 526.64$, $MSE = 4.69$, $p < .001$, point of interruption, $F(1, 86) = 118.64$, $MSE = 2.19$, $p < .001$, and interruption lag, $F(1, 86) = 37.53$, $MSE = 2.56$, $p < .001$. This confirms that resumption latency was consistently longer than move latency, although participants took longer to execute a move contained within a subtask irrespective of being interrupted. Resumption latency was however reduced with an interruption lag. Trail type and point of interruption significantly

interacted, F (1, 86) = 14.11, MSE = 1.96, $p < .001$, with large differences between resumption latency and move latency for within-subtask conditions, F (1, 86) = 434.42, $p < .001$, and end-subtask conditions, F (1, 86) = 345.21, $p < .001$. Being interrupted greatly increased the time taken to resume a suspended within-subtask goal as well as a suspended end-subtask ToH goal. Trial type and interruption lag also interacted significantly, F (1, 86) = 47.47, MSE = 1.95, $p < .001$, because the time to execute a within-subtask and end-subtask move was greatly reduced with an interruption lag F (1, 86) = 49.3, $p < .001$, with no such difference in control conditions. That there was a difference in resumption latency suggests an effect of interruption lag on the time to resume an interrupted ToH problem, and confirms no such effect (as should

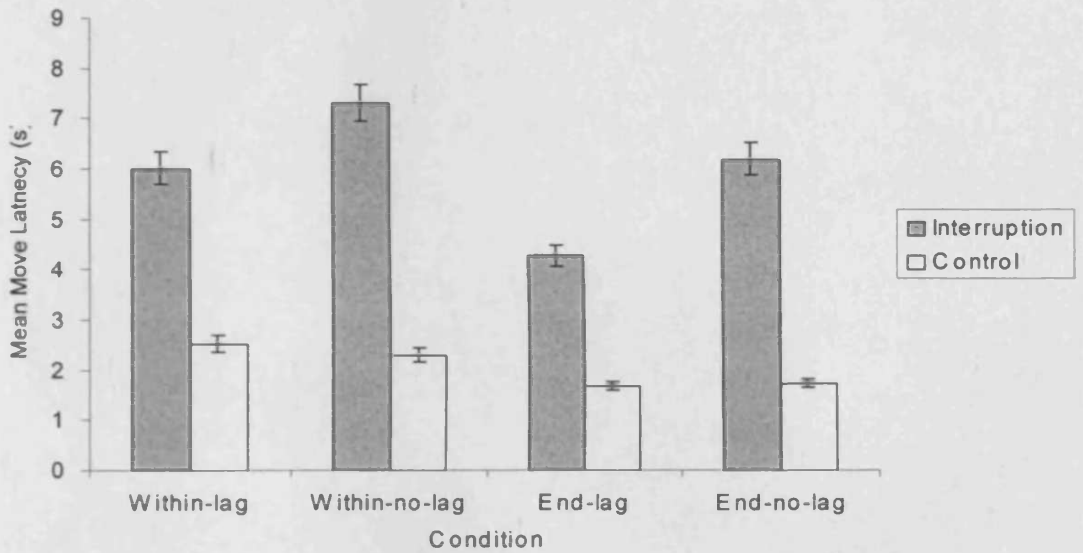

Figure 4.5. Mean primary task resumption latencies (±SE) for group 1 of Experiment 9

be the case) in control conditions. The data for resumption latencies largely replicate those within Experiment 8b, but are similar with two-second *and* four-second interruption lags. A beneficial effect of encoding a suspended goal during a two-second interruption lag is expressed through shorter resumption latencies compared to when interruption was un-signalled, with no additional impact of a longer encoding opportunity. With a longer interruption lag, resumption latency was reduced only in the end-no-lag condition, suggesting a change in encoding strategy in un-signalled interruption conditions. With more experience of encoding leading to faster retrieval, the expected utility encoding

even in un-signalled interruption increases, but is only realised by a shorter resumption latency in end-subtask conditions where there is an obvious priming cue (e.g., Baron, 1986). Nevertheless, for neither group is the gain in resumption latency greater than the time cost imposed by interruption lags. However, the two-second interruption lag is less costly to the overall time spent handling an interruption.

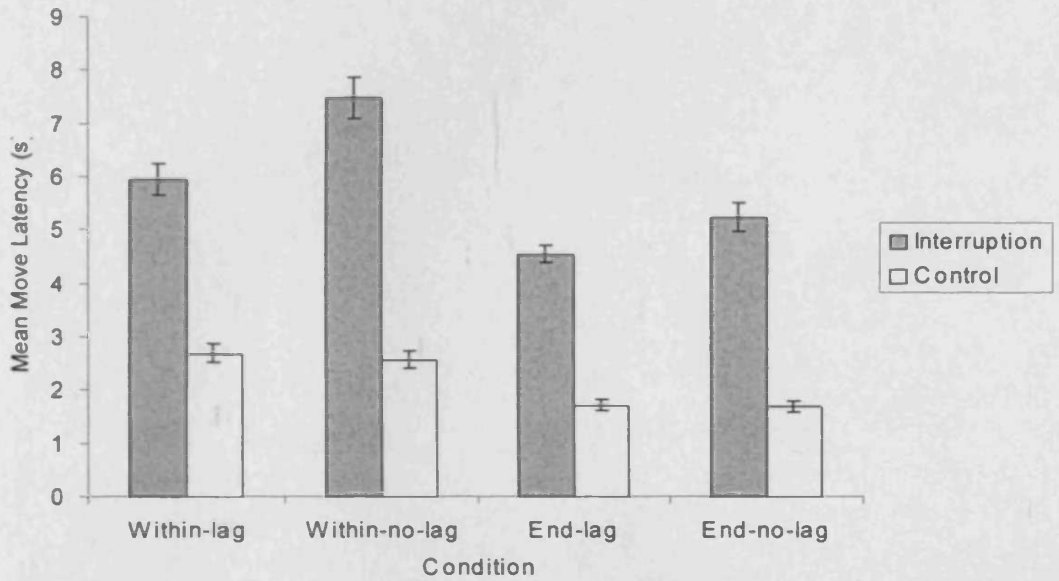

Figure 4.6. Mean primary task resumption latencies (±SE) in group 2 (4 s interruption lag) of Experiment 9

4.5.3.2. Move latency and experience

Figure 4.5 displays mean move latency for interruption and control conditions across experimental half for group 1 and Figure 4.6 illustrates move latencies in comparable conditions for group 2. Control move latency is very similar between groups: higher in within-subtask conditions compared to in end-subtask conditions, with no reliable improvement at either point of interruption. Resumption latency only differs between groups in the case of the end-no-lag condition, with resumption latency higher in group 1 than in group 2, even in the second half of the experiment. Across experimental half there is a reliable reduction in resumption latency in the within-lag condition for group 1 (MD = 1.7 s) and for group 2 (MD = 1.99 s). Resumption latency reduces in end-lag conditions, although to a lesser extent (group 1 MD = 0.67 s, group 2 MD = 0.48

s). For un-signalled interruption conditions, resumption latency shows a marginal reduction across experimental half for group 1 but increases in comparable conditions for group 2. Across groups, there is little difference in resumption latency in the first experimental half between lag and no-lag conditions, at each point of interruption (with the exception of a 1.8 s reduction in the end-lag conditions for group 1). In the second experimental half, resumption latency reduces considerably between lag and no lag conditions for group 1 (within-subtask MD = 1.94 s, end-subtask MD = 2.05 s). This is also the case in the within-subtask condition for group 2 (MD = 2.76 s), although a smaller improvement is found in the end-subtask condition (MD = 0.95 s).

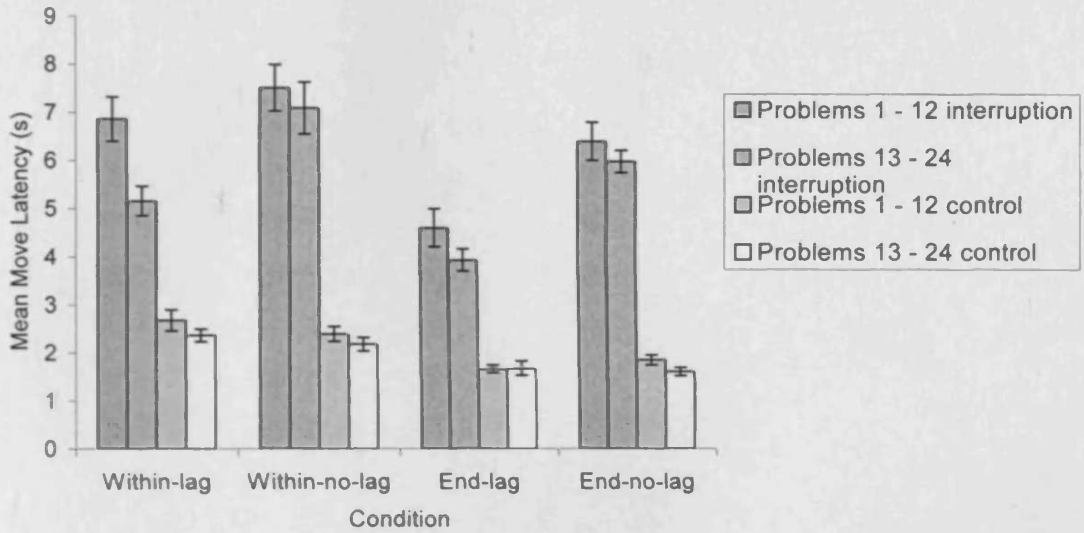

Figure 4.5. Mean primary task resumption latencies across experimental half (±SE) in group 1 (2 s interruption lag)

Control move latency was removed from further analysis for failing to exhibit any significant differences between trial types. A mixed-design 2 x 2 x 2 ANOVA was conducted on resumption latency data with group (1 vs. 2) serving as the between-subjects variable. This revealed significant main effects of, experimental half, $F(1, 86) = 24.10$, $MSE = 6.07$, $p < .001$, point of interruption, $F(1, 86) = 108.63$, $MSE = 5.6$, $p < .001$, and interruption lag, $F(1, 86) = 44.64$, $MSE = 4.89$, $p < .001$. Resumption latency was generally shorter in the second experimental half especially for conditions with an interruption lag, although this was consistently longer when interruption fell within a

subtask compared to when interrupted at the end of a subtask. Although there was a non-significant main effect of group, group and experimental half interacted significantly, although only marginally, F (1, 86) = 4.65, MSE = 6.07, $p < .05$, due to a significant reduction in resumption latency between experimental half for group 1, F (1, 86) = 27.46, $p < .001$; a difference which was marginally non-significant for group 2 ($p = .07$). There was also a significant three-way interaction between, group, experimental half and interruption lag, F (1, 86) = 7.69, MSE = 4.01, $p < .010$. This appeared to be caused by longer resumption latency in the second experimental half following un-signalled interruption for group 2 compared to group 1, even though this difference was marginally non-significant (p = .058). Resumption latency across experimental half did not significantly differ between groups, although the difference in resumption latency between groups for the within-no-lag condition in the second experimental half was nearing significance. There were no such differences between groups in any other interruption condition. Experimental half, point of interruption and interruption lag also interacted significantly, F (1, 86) = 16.5, MSE = 4.3, $p < .001$, due to a non-significant difference between resumption latency across experimental half when interrupted within a subtask by an un-signalled secondary task ($p = .94$), whereas all other conditions differed significantly.

Apart from a significant difference between resumption latencies generated by group 1 and group 2 for the end-no-lag condition, resumption performance following an interruption was highly similar across groups, even with experience within the task domain. Interestingly, with interruption lags at two-seconds and four-seconds, the time to resume suspended goal within a subtask was reduced considerably in the second experimental half, with a smaller reduction in end-subtask conditions. As suggested throughout the current chapter, priming a suspended goal through association with a reminder cue in end-subtask conditions may be easier than in within-subtask conditions, because of the salience of a priming cue. Experience within the task domain may render associative priming of an end-subtask cue more efficient during an interruption lag, but not to the extent that preparation strategy is refined (thus having little effect on reducing resumption latency). Conversely, priming a within-subtask goal with a reminder cue may prove more cognitively demanding with insufficient knowledge of the primary task

solution-sequence; although is refined with experience. A limited amount of associative activation may be distributed among other goals within the subtask as well as that suspended (and secondary task goals), and as such, single-cue based retrieval may be adversely affected. However, with experience in the task domain and a high expected gain of associative priming leading to shorter resumption latencies in end-subtask conditions, participants might be more likely to adopt similar goal preparation strategies in within-subtask conditions. Encoding strategies may be refined, by processing only the goal that is being suspended and not dedicating vital processing resources to other co-dependent goals. This is perhaps a consequence of not having to consider other goals (and their associated cues) prior to and following interruption because of stronger representations of goal chunks in memory (e.g., Gunzelmann & Anderson, 2003). Maintaining an active representation of a suspended goal may be a process more resilient to retroactive interference caused by similar secondary task goals that may impede upon keeping track of a larger goal-structure.

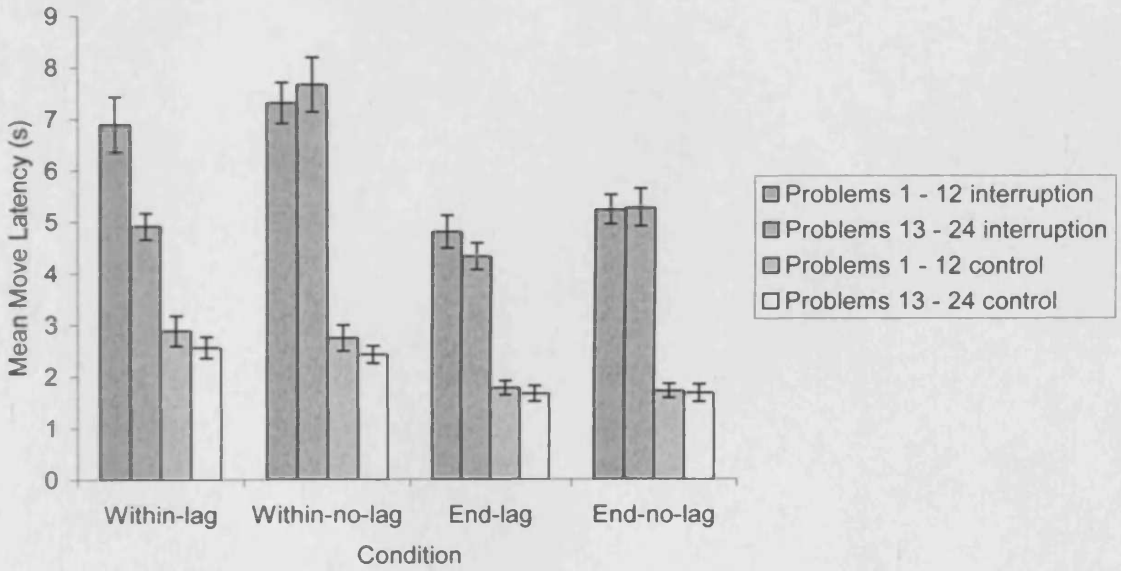

Figure 4.6. Mean primary task resumption latencies across experimental half (±SE) in group 2 (4 s interruption lag)

Better or more frequent encoding of the suspended goal occurs with experience, but only in conditions with an interruption lag. Consequently, having little or no explicit opportunity to prepare a suspended goal prior to interruption (no-lag conditions) results in dedication of fewer processing resources to maintaining an active representation of a suspended goal. Frequency of retrieval successes in no-lag conditions are low, and thus participants seem to offload memorial demands to the visual array (e.g., Anderson & Douglass, 2001). This might be a consequence of the interruption lag (i.e., its success at leading to faster resumption of a suspended goal), and not a general tendency for participants to abandon memory-intensive processing.

4.5.3.3. Redundant Moves

Table 4.5 displays the mean number of redundant moves following interruption for group 1 and group 2 in each condition of Experiment 9. Redundant moves are fewer in number following interruption compared to those executed in control trials (with the exception of the end-no-lag condition for group 2), and are greater in number for within-subtask conditions compared to end-subtask conditions. Redundant moves do not seem to vary considerably across groups following within-subtask interruption, irrespective of the inclusion of an interruption lag. End-subtask conditions demonstrate very few redundant moves, irrespective of group. Variance is however high within each condition reflecting the range of redundant moves. For group 1, number of redundant moves range between, 0 and 21.5 moves in within-subtask conditions, and 0 and 7.5 moves in end-subtask conditions. For group 2, number of redundant moves range between, 0 and 14.5 moves in within subtask conditions, and 0 and 4.5 moves in end subtask conditions.

A Kruskal-Wallis test was applied to redundant move data across groups. The within-no-lag condition significantly differ across group, $H(1) = 13.74$, $p = .001$, with no other significant differences between groups.

These results can be taken to show that number of redundant moves following the intended point of interruption is largely unaffected by the presence of an interruption or the length of an encoding opportunity. For group 2, redundant moves are reduced with un-signalled interruption to an uncompleted subtask compared to the same condition in group 1. This effect is coherent with group 2 resumption latency data for the same

condition, which was higher compared to all other conditions, and actually increased with experience in the task domain. In this condition, participants appear to be reconstructing a suspended goal from the visual array at a cost to resumption latency but to the benefit of the accuracy of the subsequent solution-sequence. Ruminating a suspended task during timely reconstruction of the problem space might improve the chances of selecting a move that would lead to error-free completion compared to when the suspended (and perhaps wrong) move is retrieved and executed.

Table 4.5

Mean redundant moves (±SE) within primary task ToH problems for group 1 and group 2 following the intended point of interruption in Experiment 9

	Within-lag		Within-no-lag		End-lag		End-no-lag	
	M	SE	M	SE	M	SE	M	SE
Interrupt	0.89	0.27	0.82	0.3	0.16	0.06	0.06	0.03
Control	1.26	0.53	1.02	0.34	0.26	0.15	0.33	0.15
Interrupt	0.86	0.4	0.04	0.02	0.03	0.03	0.25	0.13
Control	0.95	0.25	0.71	0.2	0.05	0.04	0.04	0.02

4.5.3.4. Move-error frequency

Table 4.6 displays move-error frequency data for all participants that committed one or more redundant moves in 1408 ToH problems (768 for group 1 and 640 for group 2). Overall, participants solved 1285 problems error free compared to only 123 problems in which one or more redundant moves was made; 73 for group 1 and 50 for group 2. With a parametric manipulation of interruption lag, there seems to be little difference in number move-errors executed between groups (with the exception of large difference between groups in the within-no-lag condition).

A Kruskal-Wallis test was applied to move-error frequency data which revealed a significant difference between group only in within-no-lag interruption conditions, $H(1) = 11.38, p < .001$. Move-error frequency was largely unaffected by the inclusion of an interruption lag prior to interruption initiation, although was reduced substantially in the within-no-lag condition when participants had been exposed to longer interruption lags.

Table 4.6

Mean move-error frequency (±SE) within primary task ToH problems for group 1 and group 2 following the intended point of interruption in Experiment 9

Group		Within-lag	Within-no-lag	End-lag	End-no-lag
1	Interrupt	16	15	4	6
	Control	11	11	3	7
2	Interrupt	11	1	1	6
	Control	10	8	3	2

4.5.4. Discussion

Overall, resumption latency data conform to the predictions of the goal-activation model that one to two seconds is enough time to prime a suspended goal with an available priming cue, increasing the likelihood that it is retrieved in the future (Altmann & Trafton, 2002). It is evident that during a two-second preparation period, a goal's base-level activation is increased through repeated strengthening, during which time the goal is associated with a salient priming cue that will, upon re-attendance, support reactivation of its decaying representation. If more time were required for refinement of either of these processes, resumption latency should have been reduced more so with a longer interruption lag, which clearly was not the case.

As in Experiment 8b, an interruption lag alerted by an auditory tone markedly attenuated the time taken to resume a suspended ToH problem. This was the case when interruption occurred within and at the end of a subtask. Contrary to recent claims (e.g., Altmann & Trafton, 2004), increasing the length of the interruption lag – allowing more time to prepare a suspended goal for re-establishment – had no effect on reducing resumption latency. At four-seconds, an interruption lag seemed only to benefit resumption latency through a residual effect of more efficient resumption of a suspended goal in a condition where interruption was un-signalled. Even with a four-second interruption lag, resumption latency was not reduced appreciably in within-subtask

conditions, at least until sufficient knowledge of the solution-sequence was in place to support cue-based retrieval. Redundant move and move-error frequency data suggest that being interrupted, even if the secondary task is delayed, has little effect on performance of the remaining solution-sequence following resumption of the suspended goal. In sum, it seems to be knowledge of the solution-sequence and encoding that improves resumption latency, and not the length of the encoding opportunity.

4.6. General Discussion

The aim of Chapter 4 was to address whether the interruption lag – the time between the onset of an interruption alert and secondary task initiation – would be opportunistically used to encode a suspended goal for future retrieval. Two models offer differential viewpoints of whether suspended goals are encoded for future retrieval, and manipulations of the opportunity available for encoding allowed further appraisal of their predictions. The goal-activation model asserts that the base-level activation of a goal will decay since it was postponed, but it may be retrieved in the future through an activation boost from an associated cue encoded before its suspension (Altmann & Trafton, 2002). One to two seconds seems a realistic estimate for sufficient encoding time. In contrast, Anderson and Douglass (2001) propose no influence of associative activation in goal memory – in their view participants do not prepare a suspended goal for retrieval, even when there is an opportunity to do so. Instead, they will pay the consequences of forgetting.

The experiments reported in the current chapter provide some support for the priming constraint of the goal-activation model (Altmann & Trafton, 2002). A two-second encoding opportunity prior to secondary task initiation led to a marked reduction in resumption latency for a goal suspended at the end of a subtask with a lesser reduction for a goal suspended within a subtask (Experiments 8b & 9). This effect was always augmented with experience in the task domain (Experiments 8a – 9), suggesting strategic adaptation to an interruption lag from experience of encoding leading to faster resumption of a suspended goal. I suggest that such experience increases the ability to

associate a suspended goal with a salient priming cue. Encoding opportunity was impinged upon if the alert to signal the onset of an interruption lag diverted attention from the primary task (Experiment 8a), but not if the alert allowed uncontaminated preparation (Experiments 8b & 9). Contrary to recent tests of the goal-activation model (e.g., Altmann & Trafton, 2004), a longer interruption lag had no additional impact on reducing resumption latency (Experiment 9), irrespective of the point of interruption. Since resumption was faster in conditions with a more salient priming cue and lower cognitive demands, I propose priming cues are strategically encoded to support the faster retrieval of a goal suspended because of interruption.

The current experiments provide benevolent support for the priming constraint in the goal-activation model (Altmann & Trafton, 2002). Effects of encoding opportunity were amplified when resuming a goal that was suspended at the end of a subtask (shown throughout Experiments 8a – 9). A cue obvious in the sense that execution of its associated goal will advance the solution-sequence of a ToH problem is more salient in end-subtask conditions (Altmann & Trafton, 1999). Shorter resumption latencies result through association of such a cue with a suspended goal (Altmann & Trafton, 2002). Associative activation using priming cues may also explain why resumption latency was reduced when interrupted within a subtask, but only with experience in the task domain. I suggest that participants became more adept at associatively priming a suspended goal with a cue to support its retrieval in within-subtask conditions, with strategic adaptation motivated by the expected gain of encoding leading to faster resumption in end-subtask conditions.

Resumption latency in conditions with a brief opportunity to prepare a goal for faster resumption in Experiments 8a – 9 was too short to consist exclusively of goal retrieval, but too long to consist of re-planning the problem space. Current models of goal-directed memory suggest that retrieval and subsequent execution of a suspended goal within the ToH task should take somewhere in the region of one to three seconds (Altmann & Trafton, 2002; Anderson & Douglass, 2001). With sufficient knowledge of the solution-sequence, reconstruction is also estimated to take between one and three seconds, although actual move execution may be over seven seconds if the goal has to be re-planned (Anderson & Douglass, 2001). In the current experiments, end-subtask

conditions with an interruption lag (excluding Experiment 8a because of the disruptive nature of the alert) produced resumption latencies ranging from 4.69 s (Experiment 8a) to 4.22 s (Group 1, Experiment 9). In within-subtask conditions with an interruption lag, resumption latency ranged from 6.83 s (Experiment 8b) to 5.32 s (Group 2, Experiment 9). However, with experience in the task domain, resumption latency was as little as 3.91 s in the end-lag condition of Experiment 8b and 3.94 s in the same condition for group 1 in Experiment 9. Therefore, resumption latencies in Experiments 8a – 9, although seemingly in support of goal retrieval, might have been affected by other factors.

A number of factors may affect how quickly an interrupted task is resumed, even after a suspended goal has been retrieved from memory. For instance, with limited knowledge of the solution-sequence, there may some uncertainty associated with the execution of that move because of having been interrupted (e.g., Adams et al., 1995; Speier et al., 2003). This may be especially the case in goal-directed tasks, such as the ToH, with behavioural data showing that participants often use break opportunities (e.g., subgoal boundaries), to reconsider at least part of the solution-sequence (e.g., Anderson, 1993). Some studies have shown how effective reminders abolish the forgetting of intentions (e.g., Chung & Byrne, 2004; McDaniel et al., 2004), although effects might involve cueing the memory of when to perform the intention and not what the intention was. In the current experiments, reinitiation of the primary task is enough to remind oneself of an unfulfilled goal, but is clearly not effective at indicating what that goal was. Preparation is also shown to reduce the time cost imposed by a task switch (e.g., Rogers & Monsell, 1995), although such opportunity may not be used entirely (Altmann, 2004), and switch costs may never be abolished completely (e.g., Sohn & Carlson, 2000). Nevertheless, results from the current experiments are encouraging in the sense that the cost of goal-directed enquiry is reduced not only with experience in the task domain, but also with experience of preparing a goal for resumption. For instance, resumption latency in one condition was reduced to 3.92 s (Experiment 8b). This might not even reflect asymptotic performance, which could be potentially encouraged with further experience within the task domain and different manipulations of preparation (e.g., instruction).

It is important to consider the cost imposed to the time spent handling an interruption because of the inclusion of an interruption lag. Although resumption latency was reliably reduced because of a by a brief opportunity to prepare for interruption in Experiments 8a – 9, the time cost imposed by the interruption lag was not always offset by shorter resumption latencies. For instance, in Experiment 8b, a two-second interruption lag resulted in a 2.74 s reduction in resumption latency when interrupted at the end of a subtask, but only a 1.4 s reduction when interrupted within a subtask. Additionally, even when the cost of the interruption lag was offset by resumption latency, this effect was not entirely reliable: an artefact likely to have been influenced by no specific instruction given to prepare a suspended goal for resumption. With a two-second interruption lag prior to the suspension of an end-subtask goal, there was a saving of 0.84 s spent handling an interruption in Experiment 8b; reduced to 0.05 s in the same condition in Experiment 9. The problem of balancing the length of the interruption lag as to avoid imposing an additional time cost spent dealing with interruption is an issue reflected in the interruption literature. For instance, in a recent study, the time to resume a suspended goal when preceded by an eight-second interruption lag was only three seconds less than that in a condition with no interruption lag (Trafton et al., 2003). Thus, the utility of even a brief opportunity to prepare a suspended goal for resumption should be approached with caution given the cost-benefit ratio of interruption lag length on resumption latency. Within a controlled task such as the ToH, a two-second interruption lag may not impose a high cost-benefit ratio to the overall time spent handling an interruption. This cost may be offset by resumption latencies reduced by more than two-seconds. However, it is acknowledged that at longer interruption lags (e.g., four-seconds or more) the likelihood of resumption latency offsetting this cost is reduced; a factor that is proving to be problematic in complex applied settings (e.g., Altmann & Trafton, 2004).

Whilst acknowledging that the ToH is a highly controlled structured laboratory task, the implications of results generated in Experiments 8a – 9 might inform interrupt management in similarly structured real-world tasks (e.g., the flight management system in an aircraft cockpit). An interruption lag might be better used only when speed of resumption is vital to the task goal. Preferably, the interruption lag would be complimented by priming cues in which the suspended goal can be associated with prior

to interruption and reactivated from following secondary task completion (e.g., Trafton et al., 2003). Additionally, if the average time to resume an interrupted task is less than that imposed by the interruption lag, then un-signalled interruption may be more appropriate (e.g., Katz, 1995). Finally, for an interruption lag to be most effective, it should be alerted by a method that is neither attention-directing nor resource demanding, allowing uncontaminated encoding of a suspended goal.

Since an interruption lag was shown to reduce resumption latency in the current experiments, there is much scope for future work to be done to fully unpack the benefits of preparing an interrupted goal for retrieval. Across Experiments 8a – 9, no instruction was given to prepare a suspended goal for resumption (e.g., via rehearsal), to allow natural strategic adaptation to an opportunity to prepare for interruption. However, instruction might encourage participants toward trying to remember a goal suspended because of interruption (e.g., Czerwinski et al., 1991b; Detweiler et al., 1994); when remembering might otherwise be abandoned (e.g., Czerwinski et al., 1991a; Hess & Detweiler, 1994). The goal-activation model states that "...performance in dynamic task environments would be facilitated if operators were taught to react to an alert by searching for a cue and associating it with the goal being suspended" (Altmann & Trafton, 2002, p. 66). Clearly, Experiments 8a – 9 demonstrated that with an opportunity to prepare for interruption, participants do not always choose to offload memorial demands to the task environment (in contrast to Anderson & Douglass, 2001). Given that preparation may be strategically controlled in the ToH task, then instruction, say to remember the disc that was to be moved prior to interruption, might result in even shorter resumption latencies. Preparation may also be augmented even further by experience within the task domain, leading to shorter times to resume an interrupted task. In Experiments 8a – 9, the time to resume an interrupted ToH goal was shown to reliably reduce with experience in the task domain and with experience of interruption. Experience in a task domain, specifically knowledge of the correct solution-sequence, was shown to abolish an interruption-similarity effect in a computerised shopping task even with no explicit opportunity to prepare for interruption (Edwards & Gronlund, 1998). Within the ToH task, knowledge of the solution-sequence is believed to be reflected by the chunking of goal sequences in declarative memory (e.g., Anderson &

Douglass, 1998). It might therefore prove worthwhile to explore the effects of an in interruption lag within tasks in which participants possess expert knowledge; an area in need of empirical investigation if the current results are to withstand tests of ecological validity.

Chapter 5

GENERAL DISCUSSION

5.1. Aims and objectives of the thesis

Task interruption is a subject of applied interest and has long been in need of fine-grained theoretical and experimental support. In line with recent work (Hodgetts, 2004), the current thesis set out to reconcile this paucity in psychological understanding through a cognitive experimental analysis of task interruption in a controlled laboratory setting. The current thesis assessed systematically the effects of task interruption on goal-directed behaviour in the Tower of Hanoi (ToH) task; a task for which there is well-established theoretical support (ACT-R) and hybrid models directly related to the processes involved in goal suspension and resumption (Altmann & Trafton, 2002; Anderson & Douglass, 2001). In this chapter, I will summarise the main findings of the empirical work, focusing on theoretical implications for goal-directed memory and related conceptual domains. I will end by demonstrating the suitability of the methodology for studying the psychological processes involved in task interruption as well as highlighting areas that warrant future research

5.2. Establishing common effects of interruption: A theoretical perspective

Establishing the common effects of task interruption is an important venture for theoretical and applied psychologists, although research in the field has been sparse and inconsistent, partly fuelled by the lack of a developed theoretical framework in which to guide experimental design and interpretation of findings. Using a controlled experimental task with an established theoretical framework, the current experiments have gone some way in cleaning up this disarray, offering scope for future research. Common effects of task interruption were established in Chapter 1, and these were tested experimentally throughout Chapters 2 – 4. These were interruption position and processing complexity; secondary task attributes, and the similarity between primary and secondary tasks. Factors that may on one hand amplify the negative effects of

interruption in their absence, but if controlled appropriately may distinguish such effects were also investigated: Rehearsal and the interruption lag; retrieval cues; practice and expertise.

5.2.1. Interruption position and processing complexity

It is generally accepted that where an interruption falls within a primary task can have serious consequences on performance after the secondary task has been completed (e.g., McFarlane, 2002; Miyata & Norman, 1986). Factors such as, assuming complexity because of task structure (e.g., 'simple' vs. 'complex', Speier et al., 1999) have made it difficult to unpack the reasons for performance losses. There are also very different views on what constitutes processing complexity at different task points (e.g., Wood, 1986). On one view, processing resources may be constrained by a limited working-memory resource (e.g., Anderson, Lebiere & Reder, 1996) such that new goals may exceed capacity limitations (e.g., Just et al., 1996), rendering the encoding of a suspended goal too difficult or beyond cognitive capabilities. Others have recognised that processing complexity may vary because of the relationships between items (or *relational complexity*: Halford, Wilson & Phillips, 1998) and strategy selection in the service of a task (e.g., Reder & Schunn, 1999). For instance, experience at performing a task may cause certain task components to become automated, so much so, that more processing resources are made available to prepare interrupted goals for retrieval.

The current work followed on from recent fine-grained approaches to studying the cognitive processes impinged upon by interruption, by controlling the exact point of interruption (at subgoal, subtask and individual goal levels) and the availability and number of potential reminder cues (see also, Hodgetts, 2004). Un-signalled interruption caused longer resumption latencies and a greater number of redundant moves when it fell before the execution of a complex goal-sequence (Experiment 1). Longer resumption latencies (although shorter than those in Experiment 1) were found when interruption fell within a complex goal-sequence (Experiment 3 & Experiment A), compared to the end of a goal-sequence (Experiment 3 & Experiment A) or in between two such goal-sequences (Experiment 2). Even with a brief opportunity to encode a goal before its suspension (during 2-seconds and 4-second interruption lags), the time to resume a primary task was

appreciably longer following interruption within a goal-sequence (Experiments 8a – 9). In all experiments, resumption latencies became shorter with experience in the task domain, although to a slightly greater extent when provided with an encoding opportunity prior to switching to a secondary task (Experiments 8b and 9). The results support the following assertion: Interruption affects the opportunity to formulate and maintain the representation of a suspended goal, although such processes can operate more effectively with greater encoding opportunity and/or when resources otherwise dedicated to solving problems are made available to process suspended goals.

The findings are incompatible with ideas of perfect goal memory and heightened activation of suspended goals, specifically, ACT-R's goal stack concept (e.g., Anderson & Lebiere, 1998) and the intention-superiority effect (e.g., Goschke & Kuhl, 1993). More compatible, are models of goal-directed behaviour that view goal memory in the same way as other declarative memories, and as such assume that a suspended goal will decay as a power function of time elapsed since it was processed last (Altmann & Trafton, 2002; Anderson & Douglass, 2001). A problematic issue for the model of Anderson and Douglass (2001) is the differential resumption latencies when secondary tasks took more than 18-seconds to complete; which in their view should encourage reconstruction as a resumption strategy. Only when the distance between the suspended current-state and goal-state was large (Experiment 1), and when the task array had changed between suspension and resumption (Experiments 4 – 7), did participants seem to exhibit only reconstructive-like behaviour when resuming a primary task.

A more plausible account of the current findings is offered by Altmann and Trafton (2002) who assume that a goal can be retrieved following interruption as long as it is effectively encoded prior to suspension (through rehearsal and associative priming). Their model predicts longer resumption latencies following suspension of a complex goal-sequence compared to a less complex goal sequence, because of a limited amount of source activation causing difficulty in locating, processing and relocating a salient priming cue. Their account falls short however in explaining why such processing seemed to occur without a brief preparatory period prior to interruption (although was amplified with an interruption lag), and why it was augmented with experience in the task domain. They stress that one to two seconds is essential to sufficiently encode a

suspended goal for retrieval, and without this time, retrieval should be impossible. The results of the current experiments suggest revision of the time parameters of the goal-activation model to include strategic adaptation to interruption. Consideration should be given to possible interactions between task demands at the point of interruption (when some goals may require fewer processing resources than others may) and knowledge of the task domain and experience of being interrupted.

5.2.2. Secondary task attributes

Secondary tasks vary in their degree and kind across studies. Often studies tended to use simple values (e.g., '30 s of mental arithmetic', Gillie & Broadbent, 1989, p. 244) and assume task difficulty (e.g., complex text editing tasks, Speier et al., 2003), causing variability across studies. The current experiments adopted a parametric and systematic approach using secondary tasks that could be assessed in terms of, similarity to the primary task (ToH problems vs. digit-recall tasks), number of required actions (7 move vs. 4 move ToH problems) and duration (controlled by number of component parts). When secondary tasks were similar to the primary task, the time to resume a suspended goal increased; more so when interruption fell during a complex move sequence (Experiments 3, 8a – 9, & Experiment A), although there was no additional impact of increasing the number of component parts and duration (Experiment 2).

That secondary task attributes had any impact on the time to resume an interrupted goal is problematic for a theory assuming a 'perfect goal memory' system (e.g., ACT-R's goal stack) and the view that suspended goals reside in a level of heightened activation (e.g., Marsh et al., 1998). The findings also provide further support that goals are encoded (to some extent) before long retention intervals (in contrast to Anderson & Douglass, 2001). Factors such as encoding opportunity (Experiments 8a – 9) and the difficulty in encoding a suspended goal for retrieval (e.g., Experiment 3) seem to be more predictive of the disruptive effects caused by interfering secondary tasks, with secondary task similarity impinging upon the efficiency of the outcomes of these processes. The goal-activation model seems to offer the most plausible account for many of the findings linked to secondary task manipulations in the current experiments, although some effects are still left unexplained. For instance,

Altmann and Trafton (2002) predict interruption-similarity effects when similar secondary task goals interfere with the encoded representation of a suspended goal, impacting upon retrieval efficiency. They model cannot however accommodate the abolition of such effects when primary task processing demands are low at the point of interruption and when there was little (or no) opportunity to encode a goal before its suspension (Experiment 3 & Experiment A). The current findings suggest that secondary task demands impinge upon the processes used to maintain the representation of a suspended goal after it has been formed, but only if such maintenance requires complex processing. For instance, secondary tasks that are similar to primary tasks are less disruptive to resumption efficiency when participants can inhibit the status of other goals (perhaps through long-term memory chunking) to protect the strength of an associative connection with the suspended goal.

Sensitivity to the potential effects of interruption interval might have not been dissociable from using secondary tasks that were not particularly brief (i.e., those longer than five to six seconds in duration). In a recent study, Gray et al. (2004) found longer resumption latencies following unfilled 5 second interruption intervals compared to when intervals were 250 ms and 1-second (although there was no difference between the shorter interval conditions). Secondary task duration has also been linked with longer resumption latencies with interruption intervals of less than 18-seconds but not more than 6 seconds (Hodgetts & Jones, 2003). Perhaps then, there are effects of interruption interval manifested only at shorter intervals, and the longer secondary tasks used in the current experiments masked such effects. Nevertheless, much of the interruption literature has been in the tradition of using realistic secondary tasks (e.g., tracking aircraft, taking a telephone call and so on) in which duration ranges from many seconds to minutes. The experiments presented in the current thesis follow-on from such a tradition, providing perhaps greater scope for generalisation to applied domains.

5.2.3. Similarity between primary and secondary tasks

There are very different conceptual views on how task similarity might effect retrieval of a suspended goal. On one view, similar items are expected to interfere retroactively with one another, weakening the representational status of a goal that has

been suspended, consequently impinging upon retrieval efficiency (e.g., Altmann & Trafton, 2002). By contrast, some believe that switching between dissimilar tasks will degrade retrieval efficiency because of the necessity to change or update processing resources (e.g., Adams et al., 1995). Different again, is the view that task similarity will have no effect on retrieval efficiency, given that the suspended goal might not be encoded in the first place (Anderson & Douglass, 2001). In the literature, incoherent effects may be explained by methodological problems such as placing insufficient control on the similarity overlap between tasks (see, Latorella, 1996, for a discussion). Interruption-similarity effects are reported in a number of studies (e.g., Gillie & Broadbent, 1989; Hess & Detweiler, 1994), are diminished when the steps to achieve a primary task are highly learned (Edwards & Gronlund, 1998) and are completely abolished when secondary tasks are undemanding (Latorella, 1996).

With fine-grained control over the level of similarity overlap the secondary tasks used in the current experiments were either completely similar to primary tasks (ToH problems sharing information form, content and semantics) or completely unlike (digit-recall tasks). Secondary tasks that were similar to primary tasks were more disruptive when initiated within a complex goal sequence but not when interruption fell at the end of the same goal sequence (Experiment 3 & Experiment A). With a brief encoding opportunity before switching to a similar secondary task, resumption latency was shorter in the complex goal sequence (Experiment 8b – 9). However, performance deficits were reduced considerably with experience in the task domain when interrupted within a subtask (especially with an encoding opportunity prior to secondary task initiation). Taken together, results imply that similar secondary task goals interfere with the representation of a suspended goal, but if a salient priming cue is sufficiently encoded prior to interruption – a process more likely with experience in the task domain – such effects are abolished.

5.3. Effects of preparation

Three methods of preparing an interrupted goal for resumption were considered throughout the current thesis: Rehearsal and the interruption lag; reminder cues; and practice and expertise. In his doctoral dissertation, Jambon (1996) proposed three periods

within the timeframe of an interruption where a suspended activity may be encoded for retrieval: The prologue (i.e., the period immediately preceding interruption); the interruption body (i.e., the secondary task); and the epilogue (i.e., during resumption of the primary task). In their goal-activation model, Altmann and Trafton (2002) agree that the temporal structure of an interruption can be exploited to prepare a suspended goal for resumption, but make specific predictions regarding the cognitive mechanisms underlying such processes. Moreover, the time to resume an interrupted task may be reduced with, experience in the task domain (e.g., Edwards & Gronlund, 1998); augmented with experience of dealing with interruption (e.g., Trafton et al., 2003). Similarly, Anderson and Douglass (2001) accept that humans can strategically prepare suspended goals for retrieval, but by contrast stress that participants will not encode a suspended goal for retrieval with preparation opportunity; instead paying the cost of reconstruction from the task array.

5.3.1. Rehearsal and the interruption lag

Rehearsing a goal to increase its activation might protect it against retroactive interference experienced from newer competing goals (e.g., Gillie & Broadbent, 1989; Detweiler et al., 1994), but may not be the only cognitive mechanism that could be useful to prepare a suspended goal for future retrieval. The received view is that a goal will decay as a temporal function since it was last processed; and unless reactivated during the retention interval (perhaps through rehearsal), it will be forgotten because of retroactive interference generated from intervening goals (Altmann & Trafton, 2002; Anderson & Douglass, 2001). Some regard the interruption lag as a critical period to encode a suspended goal future retrieval (e.g., Altmann & Trafton, 2002; Hodgetts, 2004; Trafton et al., 2003). Whilst endorsing rehearsal as a prerequisite for increasing the representational strength of a suspended goal, Altmann and colleagues add that the goal must be prospectively linked with a reminder cue through associative priming if it is to be retrieved following an interruption (e.g., Altmann & Trafton, 2002; Trafton et al., 2003). According to advocates of the goal-activation model, if a goal can be rehearsed following its suspension, introducing a primed cue upon reinstatement of the primary task will draw

upon associative activation to reactivate the goal so that it is retrieved (see Hodgetts, 2004 for a review).

As the literature shows, researchers in the field are still engaged in debates as to the effectiveness of rehearsal and an interruption lag on reducing the disruptive effects caused by interruption (e.g., Altmann & Trafton, 2004; Detweiler et al., 1994; Miller, 2002). Experiments presented in Chapter 4 of the current thesis investigated the role of the interruption lag on performance in an interrupted ToH problem and in doing so added further fuel to the debate. On a positive note, allowing a brief opportunity to prepare a goal for resumption (with two-second and four-second interruption lags) saw a significant reduction in the time taken to resume an interrupted goal (Experiments 8b & 9). Presumably, participants used this opportunity to rehearse a suspended goal(s), although it is just as plausible that they were encoding cues located in the task visual array. That participants self-generated strategies to prepare a goal for retrieval (even over brief delays) is encouraging given previous recommendations that such behaviour has to be instructed (e.g., Czerwinski et al., 1991; Detweiler et al., 1994). On a slightly negative note, there was a cost of an explicit opportunity to prepare for interruption (Experiment 8a & Group 2 of Experiment 9), unless the opportunity to prepare was brief (i.e., two-seconds) and not impinged upon by a disruptive alert to signal the onset of this period (Experiment 8b). A visual alert to mark the onset of an encoding opportunity spoilt the quality of preparation probably because of the visual processing demands imposed by the primary task were too high (Experiment 8a). The degree of gain in resumption time relative to the time spent preparing will be considered in more detail in Section 5.4.1.3.2 of the current chapter).

5.3.2. Retrieval cues

Reminder cues may alleviate the effects of goal forgetting, although their role as effective retrieval aids in an interruption context is not so clear cut. Limiting presented information to only that required to resume an interrupted activity after completion of a secondary task reduces performance decrements (Cypher, 1986; Field, 1987). By contrast, marking the suspended activity supports faster resumption, but usually only when task demands at the point of interruption are low (e.g., Cutrell et al., 2001). The

current experiments set out to test a number of theoretical claims about the role of retrieval cues made by Altmann and Trafton (2002) in their goal-activation model. In their view, a good retrieval cue: Must be salient and perceptually available; will prime the goal and few (or no) others; and should be encoded immediately prior to interruption and be available when the primary task is reinstated. The current experiments provided compelling evidence that participants do use retrieval cues to support faster resumption of a suspended goal, but that level of salience is determined by how striking features are (e.g., disc colour) rather than the association between the goal and the cue formed when encoding the move.

The current experiments affirm the prediction of Altmann and Trafton (2002, p. 53) that a logical candidate to serve as retrieval cue in the ToH task is the disc associated with the suspended goal. As such, simple manipulations of disc properties could be assessed in terms of the effects on resumption efficiency. The time to resume an interrupted task was consistently shorter if interruption fell at a point within the solution-sequence where the suspended goal was easily distinguished by a salient disc cue (Experiments 3, 8a – 9, & Experiment A). By salient, I mean that in such conditions, the goal was to move the second LOOP disc to its goal destination peg, and thus the cue was striking in the sense that it could be considered independently of other cues. Conversely, when the goal was to un-block the second LOOP disc by formulating and executing a series of indirect moves, cue salience was reduced, and resumption was slowed (Experiments 4 – 9, & Experiment A). In accordance with the goal-activation model, the current experiments showed that with experience in the task domain, "At goal-encoding time the system has to process the disk anyway, in the course of applying the goal recursion algorithm to decide where the disk should go" (Altmann & Trafton, 2002). Without such experience, associative priming using disc cues is limited to contexts where a priming cue is salient and when the current does form part of an abstract move sequence (in contrast to Altmann & Trafton, 2002).

Convincing support that ToH discs support goal memory was demonstrated throughout Experiments 4 – 7. With little opportunity to exclusively encode priming cues prior to interruption, participants seemed to nevertheless encode the most salient feature property – colour – (e.g., Goolsby & Suzuki, 2003; Wolfe, 1994), a process

augmented with experince of encoding that feature cue caused by changing disc properties within some reinstated problems. Goals were retrieved and resumed faster when colour had not changed between problem suspension and resumption (Experiments 4 – 6), but only when the disc configuration in the current-state had not changed (Experiments 7). I believe that such adaptation was the result of a trade-off between applying the least memory-intensive preparation strategy (i.e., encoding a salient feature-cue) against a more memory-intensive preparation strategy (e.g., remembering disc location in a dynamic task context).

5.3.3. Practice and expertise

Experience of problem solving in the task domain and experience of interruption was shown to attenuate the time to resume an interrupted ToH problem in all experimental conditions where disc properties had not changed between problem suspension and reinstatement. I broadly propose that through experience of problem solving in the task domain (e.g., organising goals into chunks that demand fewer processing resources), and experience of associative priming using salient disc cues (especially with a brief opportunity to encode these cues: Experiments 8a – 9), participants become yet more adept at associative priming. This may reflect strategic adaptation brought upon by the expected gain of encoding (e.g., Anderson & Douglass, 1998), perhaps experience of effective associative priming leading to faster resumption. For display-based problem solving, Howes and Payne (2001) suggest that a successful operator (i.e., one that is 'familiar' in that its application has advanced the solution-sequence on past attempts) is more likely to be used on future encounters of similar problem states. Such an effect may also occur from higher-order processing of an interrupted goal: Declarative knowledge generated from frequently attaching mental tags to goals considered in the short-term may serve to facilitate a more active internal long-term representation of the goal (e.g., Fisk & Rogers, 1995).

Without sufficient experience in the task domain or of being interrupted, resumption was generally faster for a suspended goal that could be encoded independently of other goals (e.g., end-subtask conditions). Specifically, I am referring to ToH moves supported by a single salient priming cue available for associative priming.

Such 'uninhibited' priming was further augmented with experience, although to a lesser extent than in conditions where without experience, encoding through associative priming may have been drawn upon by resources taken up by having to consider other goals (i.e., within-subtask conditions). Experience of faster goal resumption because of associative priming in end-subtask conditions, and experience of problem solving in the task domain and may have encouraged a strategic decision to adopt similar preparation strategies in within-subtask conditions (Experiments 3, 8a – 9, & Experiment A).

With experience in the task domain, reductions in resumption latency were found in Experiments 4 – 6 when disc properties (i.e., colour and location) were congruent between primary task suspension and reinstatement. In these experiments, disc properties frequently changed between problem suspension and reinstatement (sometimes in 75% of the interruption conditions), with participants demonstrating reconstruction-like behaviours in such conditions. Perhaps, with experience of problem solving in the task domain, there is a general freeing-up of a limited supply of associative activation that may subsequently be dedicated to associatively encoding a single goal-cue. Because of not having to consider other goals (and their associated cues) prior to and following interruption (perhaps due to learned knowledge of the solution-sequence, e.g., Edwards & Gronlund, 1998), a greater supply of associative activation is available to associatively prime a suspended goal. The results from Experiment A harmonise with this idea, where having considered the solution-sequence prior to goal suspension causes shorter resumption latencies even with little experience of problem solving in the task domain.

5.4. Theoretical implications

Task interruption has received some empirical attention, mostly in applied domains, although much of the existing research lacks theoretical support. A range of theories and computational models have been proffered as to how goals are suspended and resumed, and seem to be ideally suited to the study of task interruption. Three models that are based on the well-established ACT-R theory of cognition are reviewed in the context of the current work. They differ fundamentally in their predictions of the processes involved in goal suspension and resumption, and as such offer an interesting base for comparative investigation. In the following sections I will review in turn,

Anderson and Lebiere (1998), Anderson and Douglass (2001), and Altmann & Trafton (2002), with the aim of (1) summarising their predictions for performance in an interrupted task, and (2) assessing their compatibility with the results generated by the current experiments.

5.4.1. ACT-R and the goal stack (e.g., Anderson & Lebiere, 1998)

ACT-R's goal stack guarantees information about goals. In ACT-R, the goal stack has privileged status: Items posted on the stack are not construed by memory capacity, and do not have decaying activation so therefore cannot be forgotten (e.g., Anderson & Lebiere, 1998). Goals are arguably stored in a LIFO order and as such are pushed down the stack when they are created or when they cannot be acted upon (e.g., when forming an intermediate state in a subgoal), and are popped from the stack when satisfied.

Findings from the current experiments exposed problems for this widely held oversimplification of goal-directed memory. The first issue raised is the idea that an encoded goal has a privileged status in memory such that its activation maintains a heightened level of activation compared to other non-goal memory items, and is immune to decay-based forgetting (e.g., Goschke & Kuhl, 1993; Marsh et al., 1998). Second, is that the goal stack assumes that goals are protected against the potential disruption caused by retroactive interference (e.g., Anderson & Lebiere, 1998), with completed goals removed from working-memory by active inhibitory processes (Mayr & Keele, 2000). A final issue for a stack-based concept is the assumption that goals are hierarchically encoded in a non-idiosyncratic manner to satisfy sophisticated decomposition strategies (such as subgoaling, e.g., Anderson et al., 1993). The goal stack in my view offers an excellent conceptual framework in which to model goal-directed memory, but only if the processor - human or machine – is capable of such faultless performance. The experiments reported in the current thesis suggest that for humans at least, this is not the case. Suspended goals are encoded opportunistically with processing effort governed by knowledge of the task domain and strategy deployment, and decay in activation as well as being susceptible to interference generated from other goals.

In the first experiment, the time to resume a suspended goal increased with the number of moves left to problem completion, irrespective of encoding opportunity prior to interruption. When interruption fell before the first move in a fifteen-move ToH problem, resumption latency was almost six-seconds longer than the time to encode the problem prior to interruption, and over four-seconds longer than the time to initiate an identical control problem (in contrast to the predictions of Anderson & Lebiere, 1998). Shorter resumption latencies when problems required fewer moves for error-free completion is taken as evidence that participants abandon goal storage when rationalising that the problem space may be easier to reconstruct using the visual array (e.g., Anderson & Douglass, 2001) or suffer retrieval failure (e.g., Siegler, 1987). By contrast, there was evidence for goal retrieval (although fallible without experience in the task domain): When task demands were low (subgoal 3 in Experiment 1, Experiments 2, 3 & Experiment A); if the suspended goal was primed with a salient reminder cue (Experiments 4 – 6); and, if participants were given a longer encoding opportunity (Experiments 8a – 9). The goal stack assumption is yet further weakened for failing to account for the extra disruption caused after similar secondary tasks, where newer goals seemed to interfere retroactively with the process of reactivating a suspended goal contained in a complex structure (Experiments 1 – 3, Experiments 8a – 9 and Experiment A).

In all, results did not fit with the 'perfect goal memory' assumption of the goal stack, and instead exposed the fallibility of goal memory. Arguably, the goal stack has been conveniently used for too long to gloss over the variability in human performance (see Meyer & Keiras, 1997), with empirical data such as that provided in the current thesis supporting the new wave of research exposing its limitations (e.g., Hodgetts & Jones, 2003). I agree that encoding, storing, maintaining and retrieving the representation of a suspended goal is an effortful process and is seriously impaired because of a task interruption (Altmann & Trafton, 2002). Similarly, I am not the first to suggest that working memory capacity largely contributes to goal selection errors in the ToH task (Just et al., 1996) although I am the first to show that this might affect encoding prior to and following task interruption in the ToH task. The two models reviewed next are able to better deal with the findings from the current work, and although they cannot explain

all of the results, they might be easily revised to accommodate the effects of task interruption.

5.4.2. A model of goal suspension and retrieval (Anderson & Douglass, 2001)

Anderson and Douglass (2001) question the traditional ACT-R view that goal memory is perfect, and suggest instead that goals are processed in the same way as other declarative memory items. They harmonise with the view that the base-level activation of a suspended goal will decrease as a power function of the delay since it was last processed, and suggest that humans deploy non-memory intensive strategies to maintain representational strength. Simply put, at long retention intervals (e.g., those over 18-seconds), humans are likely abandon goal storage altogether relying instead on reconstructing the goal from the perceptual task array. Findings from some conditions in the current experiments do dismiss reconstruction as a resumption strategy, but suggest its use not because of the length of the retention interval (Experiment 2) but rather the difficulty associated with formulating and maintaining the representation of a suspended goal (e.g., Experiment 3). Resumption latencies were longer, for instance, when an interruption fell at a point within the solution-sequence where maintaining awareness other goals was necessary to reactivate a meaningful representation of an abstract suspended move, longer again when the secondary task required the processing of similar goals and goal structures.

Within their model, Anderson and Douglass (2001) may rightly infer goal decay as a function of time, but at no point do they accept that, given goals can be forgotten, there must be processes to overcome or lessen the impact of such an inferred cognitive processing limitation. For instance, in their experimental design, the primary task visual display is always available for inspection, and thus might encourage participants to use less memory-intensive processing, instead opting to more often offload mnemonic demands to the task array. Second and following from the first point, participants were instructed to 'post' immediately unachievable goals to an invisible goal stack that could be inspected at any point in the future. In most interruption contexts, the primary task array is temporarily removed, a constraint shown by the current experiments to encourage more intensive encoding of a task goal before its suspension. Such strategic processing is

refined with experience in the task domain, especially for goals that are harder to encode prior to interruption, although might be a natural response with a brief opportunity to encode interrupted-state information prior to interruption (Experiments 8b & 9).

Experiments 4 – 7 provided convincing evidence of goal reconstruction when the reinstated visual array did not boast a previously encoded priming cue or when properties of that cue (mostly colour) had changed. That participants were faster to resume interrupted primary tasks when such cues had not been manipulated between task suspension and resumption suggests that goal encoding is not abandoned prior interruption (as suggested should be the case by Anderson & Douglass, 2001), rather that goal retrieval is abandoned or terminated when that goal is too difficult to remember.

It seems that participants attempt to form a representation of a suspended goal, the fate of which is determined by factors such as retroactive interference from secondary task goals, and inexperience of associative priming, rather that the delay between goal formation and expected retrieval.

Overall, participants did not appear to abandon goal storage prior to task suspension (in contrast to Anderson & Douglass, 2001), but instead deployed opportunistic encoding; the fate of which was determined by factors other than secondary task duration (e.g., having a brief encoding opportunity). A diagnostic issue for Anderson and Douglass (2001) to consider is the impact of removing the visual array from perceptual awareness on goal suspension and resumption behaviours. I propose, as in the current experiments, that participants would adapt to loosing an external memory source by dedicating greater effort to processing goals before they are suspended. It would be interesting and informative to model the trade-off between memory-intensive and environmental-intensive processing in the ToH task, which are encoding strategies shown elsewhere to be sensitive to even small changes in task constraints (e.g., Fu & Gray, 2004). For instance, the more costly it becomes to access information (called an *information access cost*), the more likely participants are to trade knowledge 'in-the-environment' for fallible knowledge 'in-the-head' (Gray & Fu, 2004).

5.4.3. The goal-activation model (Altmann & Trafton, 2002)

The goal-activation model seems the most plausible framework for explaining the findings of the experiments presented in the current thesis. It moves away from the mechanisms of the goal stack in suggesting that a suspended goal will decay as a power function of the time since it was last processed (in support of Anderson & Douglass, 2001), and puts forward associative priming as a protective mechanism for retrospective interference caused from intervening goals. In their model, Altmann and Trafton adhere to the view that associative activation is in limited supply (e.g., Lovett & Anderson, 1996); consequently a suspended goal must be linked with a salient cue that primes it and few (or no) other goals.

The third empirical series investigated a core feature of the goal-activation model: That associative priming of a suspended goal with an external cue is a prerequisite for retrieval, and that such processing had to occur during the interruption lag – a brief period between the interruption alert and secondary task initiation. In line with the goal-activation model, an interruption lag caused a reduction in the time to taken to resume the primary task (Experiments 8b & 9) and two-seconds was enough time to encode a suspended goal for retrieval, with no effect of slightly increasing such encoding opportunity (Experiment 9). Nevertheless, an interruption lag did not guarantee retrieval, and when retrieval-like behaviour was exhibited (usually with experience in the task domain), the time to execute a suspended move was still far beyond move latencies in control trials.

5.4.3.1. Limitations of the goal-activation model

On one extreme the goal stack assumes infallible goal memory, and on another extreme Anderson and Douglass (2001) believe that humans abandon trying to remember goals altogether. The goal-activation model provides a middle ground of sorts, in suggesting that memory for goals is fallible but not impossible, but in doing so may have underestimated human cognitive capabilities and the ability to strategically adapt to interruption. The benefit of an interruption lag is questioned for similar reasons.

5.4.3.1.1. Underestimates human cognitive capability

Assuming that a suspended goal is sufficiently strengthened, Altmann and Trafton (2002) stress that retrieval will be impossible unless the goal was associatively primed with a retrieval cue prior to interruption: A process believed to take in the region of one to two seconds (preferably during an interruption lag). Without such an encoding opportunity, the representation will decay and eventually become redundant. This premise may be limited for three reasons. First, Altmann and Trafton are vague about the effect of goal processing prior to the interruption lag. Processing a to-be-suspended goal at task points other than immediately before the onset of an interruption is likely to boost its base-level activation (Anderson & Douglass, 1998), forming a stronger representation (Lovett & Anderson, 1996) which may improve subsequent performance in the ToH task. Findings from some of the current experiments, in which interruption was un-signalled, support encoding at different task points, causing faster resumption of a suspended goal. Secondly, the goal-activation model seems to underestimate the structure of the goal that is being interrupted, and how this might affect resumption performance. For instance, when interruption fell at the end of subtask (involving the processing of a single goal) resumption was faster than when interruption fell within a subgoal (a multiple goal structure). Specifically, the more complex the goal structure being interrupted, the more efficient encoding has to be to ensure retrieval of a suspended goal. Finally, the goal-activation model assumes that associative priming during an interruption lag is essential for retrieval of a suspended goal even when participants are highly trained at the primary task. By contrast, findings from some of the current experiments suggest that experience in the task domain and experience of being interrupted leads to resumption latencies that are more fitting with retrieval parameters, even without an opportunity to prepare (e.g., Experiment 6 & Experiment A).

The priming constraint in the goal-activation model, I feel, underestimates human ability to strategically adapt to constraints imposed by the task domain (e.g., interruption). Preparatory strategies, such as planning, might be sufficient to activate certain goals to a level that they are retrievable without having to rely upon associative priming. However, it is conceded that such retrieval may only be effective if the goal is of high priority (e.g., where timely recovery is encouraged) and if it supports uninhibited

performance (e.g., a single goal structure). In any case, I propose that task expertise (e.g., knowledge of a perceptual problem solving strategy) will have an influence on the importance of using an interruption lag to support suspended goal recovery.

5.4.3.1.2. A cost-benefit utility of the interruption lag and retrieval time?

Whilst the experiments presented in Chapter 4 illustrate that an interruption lag provides an opportunity for more efficient goal encoding prior to secondary task initiation (in support of Altmann & Trafton, 2002), the gain in resumption latency was not always equivalent to the cost imposed by such a preparation delay. In only one experiment did the interruption lag have an overall positive cost-benefit ratio with resumption latency and this with a non-distracting auditory alert and when interruption fell before the execution of a simple move (Experiment 8b). The cost-benefit ratio showed was more positive with experience in the task domain and experience of interruption, but only when employing a shorter interruption lag signalled by an auditory alert (Experiment 8b and the end-lag condition for Group 1 in Experiment 9). Thus, there is clearly a problem of balancing a delay to encode a suspended goal so there are no additional costs accrued form distracting primary task performance (i.e., with an attention directing alert) and by the overall time spent dealing with interruption (e.g., Group 2 in Experiment 9).

These findings are more supportive of the original predictions of the goal-activation model than the subsequent research which generally shows a negative cost-benefit ratio of the length of an interruption lag to the reduction in resumption latency (Altmann & Trafton, 2004; Clifford & Altmann, 2004; Hodgetts & Jones, 2003; Trafton et al., 2003). Under certain task constraints, an interruption lag can cause even longer resumption latencies (Miller, 2002). Nevertheless, it seems that Altmann and Trafton (2002) may have underestimated the time required to sufficiently encode a suspended goal for retrieval in one to two seconds following primary task reinstatement. The current experiments have identified points that need consideration, including, quantification of the differential impact of experince in the task domain, accounting for experience and inexperience of dealing with interruption, and modelling parameters with task constraints in mind (e.g., the cost of goal-directed enquiry in the ToH task). As a

future research direction, instructing participants to prepare during the interruption lag, especially using priming cues, might result in a more positive cost-benefit ratio.

5.4.4. Summary of models of goal-directed memory

The current empirical work has shown that ACT-R's goal stack cannot accommodate the effects of interrupting performance in the ToH task and that the ACT-R model of goal suspension and retrieval may largely underestimate human processing capabilities in such a context. Whilst the goal module of ACT-R theory has undergone an important change in that it allows for decay by not treating goals with a special processing status (Anderson & Douglass, 2001), further work is needed to quantify the effects of interruptions. The goal-activation model offers the most compatible account of the findings from the current experiments, largely because it accepts that (1) since goals decay, there must be strategies for protecting against forgetting and (2) that decaying representations are susceptible to the effects retroactive interference caused by intervening goals. It too requires further work, and hopefully the current research provides some helpful guidance for appropriate revision.

5.5. Implications of the current work on related conceptual domains

I now turn to a discussion of the implications of the current work for conceptual domains related to goal directed memory. In doing so, I hope to show how the interaction of knowledge across domains may benefit all concerned in the ongoing research effort as well as demarcating task interruption as an area of cognitive experimental psychology worthy of future attention.

5.5.1. Intention superiority and the Zeigarnik effect

Advocates of the intention superiority effect assert that unfulfilled goals or intentions reside in a heightened and constant level of activation until such time as they are executed or described (e.g., Goschke & Kuhl, 1993; Marsh et al., 1998; Zeigarnik, 1927). Intention superiority cannot be completely ruled out by the results of the current experiments, although it is certain that other factors such as decay and interference impinge upon the ability to maintain the status of suspended actions. Even with

knowledge in the task domain, experience of being interrupted, and having to perform secondary tasks that require a dissimilar processing resource to the primary task, memory for a suspended goal does not seem to be superior, as evidenced in long resumption latencies.

Participants may automatically deploy preparatory processes before turning to the interrupting secondary task, with appropriate opportunity (Experiments 8a – 9) and even without such opportunity (Experiments 4 – 6). In Zeigarnik's experiments, such preparation, for instance the rehearsal of goals and/or the use of reminder cues, might have contributed to the finding of superior memory for tasks that were interrupted.

Uncompleted actions may be represented in a higher form of activation than completed actions (e.g., Anderson & Lebiere, 1998), although this activation may rapidly dissipate when dealing with complex intervening tasks such as those used in the experiments presented in the current thesis. Reconstructing the content of an uncompleted intention when probed may also be easier than actually executing the goal immediately after being interrupted. For example, in Experiment 4, one participant commented: "I know that I had to get the little disc out of the way of the next largest disc, but I just couldn't quite remember where exactly I wanted to move it, not to mention where the next one had to go". In this example, the participant remembered the goal, although experienced problems in reforming the rest of the task so that the goal could be performed. Clearly, intention superiority comes into play in many everyday activities (e.g., automatically knowing that keys are required to activate a vehicle), but fulfilling such actions can be disrupted by interruptions (e.g., taking a phone call and forgetting where one left the keys), with consequences to performance efficiency.

5.5.2. Prospective memory

It is of interest to compare interruption in goal-directed behaviour to the research on *prospective memory*. The two paradigms are similar in many ways: They require setting, retrieval and execution of a suspended intention at some point in the future. The main difference is that within a prospective memory task a participant has to remember to perform the intention at a specific time or in response to a particular cue. The current experiments provide evidence that the processes involved in both paradigms are similar,

although performance after a task interruption is regulated by how effectively the interrupted goal is prepared for suspension. A suspended goal may or may not be encoded (before the ongoing task is postponed) which might be retrieved when the ongoing task is reinstated (at a certain time), arguably in response to a retrieval cue (e.g., Altmann & Trafton, 2002). In this sense the interrupted task 'becomes a prospective memory task' (Dodhia & Dismukes, 2003).

I assert that a task interruption paradigm, especially one as controlled as that employed in the current thesis, tells us more about the cognitive processes used to maintain the representation of a suspended goal and 'how' we go about performing that goal after it has been re-established. By contrast, prospective memory research remains more informative about human ability to maintain an intention of 'when' to perform a suspended intention during an ongoing activity, usually in response to a certain cue or after a certain time interval.

5.5.3. Task Switching

Having studied task interruption one needs to consider the potential effects of switching between tasks, an area of cognitive psychology in which there is a strong base of empirical and theoretical work. The typical paradigm employed to study task switching involves alternating between two tasks, and measuring the 'switch cost' (usually time and accuracy) of starting each task. Switch costs are usually higher when switching from one task to another task that requires completely different processing resources (see Monsell, 2003 for a comprehensive review). Inevitably, switching from one task to another will impinge upon task start times, even when such switching is due to interruption.

The current experiments have gone some way in distinguishing the study of task interruption from that of task switching, although there seem to be some commonalties. For instance, when returning from a secondary task brought about because of interruption, participants seem to require time to reconfigure their place within the primary task; an effect echoed in a number of task switching studies (e.g., Rogers & Monsell, 1995; Sohn & Carlson, 2000). Also, longer resumption latencies because of having performed a secondary task may consist of inhibitory processes regarded by some

as essential to suppress the mental representations of items processed within such a task (e.g., Allport & Wylie, 2000; Mayr & Keele, 2000).

Task interruption is however distinguishable from task switching because of a number of different processes (or differences in the deployment of these processes) involved in performing tasks. It is important to note that the main distinction between the two paradigms is that participants are required to remember a suspended activity only following an interruption. The paradigms also differ because of, processes prior to the interruption/switch (i.e., different forms of preparation), the cognitive demands of retaining a representation of a suspended goal (required only when interrupted), and the differences in the time to resume a suspended task and start a new task (which are usually longer when interrupted). For example, in a typical task switching cueing paradigm, participants are given time (also called an 'inter-cue interval') to reconfigure processing resources from the current task to the second task, usually resulting in reduced switch costs (e.g., Koch, 2000, 2003; Meiran, 1996; Rogers & Monsell, 1995). By contrast, in a typical interruption cueing paradigm (such as that employed in Experiments 8a – 9 of the current thesis) participants seem to strategically encode a suspended goal for future retrieval rather than preparing to switch to the secondary task, resulting in faster times to resume the primary task. Another difference unearthed by the current experiments is that alternating between similar tasks because of a task interruption can result in greater performance deficits than when alternating between different tasks: this is the inverse effect found typically in the task switching literature (Monsell, 2003), although it is acknowledged that task similarity is usually defined and manipulated at a somewhat simpler level in typical task switching studies.

Overall, task interruption seems to affect performance by impinging upon the opportunity to encode a suspended goal for future retrieval as well as disrupting the processes used to maintain the representation of that goal. This may explain why resumption latencies in many task interruption studies are of the order of seconds rather than the millisecond costs typically demonstrated in task switching studies (e.g., Altmann & Trafton, 2004). It is acknowledged that switch costs may contribute to such resumption latencies, and a future venture for interruption researchers will be to isolate these costs and find methods of eradication.

5.5.4. Long-term working-memory

The current experiments have addressed an important question raised by advocates of a long-term working-memory account (e.g., Ericsson & Kintsch, 1995) and ACT-R theory of expected gain (e.g., Anderson & Lebiere, 1998). That is, do human memory strategies evolve in service of systematic learning of problem solving strategy development (e.g., Gunzelmann & Anderson, 2001; VanLehn, 1991) as a function of speed versus accuracy within the task domain (e.g., Fum & Del Missier, 2001)? I believe that the findings from the current experiments go some way in providing an answer to this question. Memory strategies do evolve but in a functional way. That is, deploying processing resources to retain a suspended goal that would be otherwise difficult and time consuming to reconstruct is a strategic decision, and is largely governed by experience in the task domain. With experience, cognition is better equipped to activate the long-term memory representation of the suspended goal using pointers that are located in the visual array (Ericsson & Kintsch, 1995).

Associative priming is not as effective when the suspended goal forms part of a complex goal sequence, perhaps because (1) that goal is too abstractly represented to memorise and/or (2) because it has no long-term memory representation with which a cue can be linked. If a to-be-suspended goal is supported by an attention-grabbing retrieval cue, one that has been used in the past to successfully ensure retrieval, it may be strategically encoded in the future increasing the likelihood of faster and more efficient retrieval (e.g., Anderson & Douglass, 1998; Wolfe, 1994). Future research might want to assess the role of long-term knowledge on strategy selection in the service of preparing, maintaining and retrieving the representation of a suspended goal. It would be interesting to establish at what level performance asymptotes, and whether this occurs at a suboptimal level.

5.6. Applied implications

The existing research effort is locked into the development of technological tools to best support the synchronisation of interruption (see McFarlane, 2002) with little more than common-sense interpretation of implementation costs instead of informed systematic theory and experiment. Interruption is linked with productivity loss in office-

work (e.g., Jackson et al., 2002, 2003), performance deficits in safety-critical work contexts (e.g., process-control systems, Bainbridge, 1984) and even fatal accidents in the aviation domain (e.g., NTSB, 1988). Among the common effects of interruption in applied domains are, extended times to fully recover an interrupted activity (e.g., up to 15 minutes, DeMarco & Lister, 1987) and performance errors (e.g., Loukopoulos et al., 2003); with similar effects found in the current experiments. The experiments presented in the current thesis employed a well-structured problem solving task (the ToH) for which there has been much empirical attention and theoretical deliberation (see Anderson & Douglass, 2001). Although many applied tasks are far more complex (i.e., involve multiple activities), less structured, and uncontrollably dynamic (e.g., having to monitor system changes), the current cognitive experimental approach may prove informative, if only to find ways to reduce time and performance efficiency costs.

Notification technologies are a part of a knowledge-worker's daily activities in the modern office environment (e.g., *Instant Messenger*, *MSN*). Research has shown notable performance losses from interrupting memory-demanding office-type tasks (e.g., Czerwinski et al., 2000b), with little alleviation of these effects from counter-measures such as reminders (Cutrell et al., 2001; Lahlou et al., 2002). Findings from the current experiments highlight comparable damaging effects of the point of interruption, particularly, interruption is more disruptive when implemented within a complex goal-sequence than when implemented at the beginning or end of such a sequence (in accordance with Czerwinski et al., 2000a, b). Theoretically, this effect is expected: Interruption may impinge upon the opportunity to rehearse an interrupted goal to form an active memory representation (Altmann & Trafton, 2002). Noteworthy from the current experiments is the alleviation of the disruptive effects of interruption with, salient reminder cues, knowledge of the task domain and experience of interruption, and a brief opportunity to encode a suspended goal prior to secondary task initiation. These results fit well with another key prediction of the goal-activation model, particularly that participants may benefit from associating a suspended goal with a priming cue, with advantages exhibited in faster resumption times and fewer performance errors. Previous attempts at using reminder cues to support more efficient transition back to an interrupted task, such using a mouse cursor (e.g., Cutrell et al., 2000a) and leaving track changes on

in *Microsoft Word* (Lahlou et al., 2002), might have proved ineffective in supporting the formation of such associations. For instance, Lahlou et al. (2002) highlighted changes only after secondary task completion, meaning that associative connections could not be formed prior to interruption.

Given the pervasive disruptive effects of interruption reported in the current experiments and reflected throughout the interruption literature, office-workers are encouraged to minimise the occurrence of avoidable interruption. Simple measures such as setting a web browser to decline Internet 'pop-ups' and turning off instant messaging systems when engaged in cognitively demanding tasks are examples of such counter-measures.

In safety-critical work environments such as the flightdeck, the timing of interruption is often unpredictable, and because of time pressures, crews have little opportunity to encode effective cues to interrupted activities. It has recently been noted that, "...time pressure and task demands may prevent the pilot from deeply encoding in memory the intention to resume the deferred task, thereby reducing the likelihood of retrieval. Retrieval of intentions hinges on the individual noticing cues in the environment that are associated with the stored intention that can trigger retrieval of the intention from memory" (Loukopolous et al., 2003, p. 17). These problems are considered in a leading model of goal-suspension and resumption (e.g., Altmann & Trafton, 2002), and are echoed in the findings of the current experiments. Human goal memory is susceptible to the effects of decay and interference. Performance can however be improved, when a salient priming cue is encoded with a suspended goal both prior to and following interruption, with improvement augmented further with a brief encoding opportunity (i.e., a momentary delay before interruption) and experience in the task domain as well as that of being interrupted. During training protocols (e.g., flight simulations), aircrew may benefit from being exposed to interruptions, and I encourage engineers to design interrupt-driven communicative systems that are sensitive to the task demands at the point of interruption; systems capable of detecting salient priming cues and briefly delaying secondary task initiation whenever possible.

5.8. Evaluation of the methodology and experimental design

Interruption researchers have been in the tradition of using an experimental methodology that addresses address a specific problem of applied interest (e.g., military resource allocation tasks, computer notification applications and so on). Results generated from these tasks may be informative to the applied domains in which they are addressing, but generalising their use to mainstream experimental psychology is problematic (e.g., McFarlane, 2002). To add to this disarray, there has been a tendency in laboratory studies to use tasks that are very different in terms of structure causing problems for drawing firm theoretical implications.

5.8.1. The advantages of the Tower of Hanoi task

Finding a suitable laboratory-based methodology in which a comprehensive study of the effects of interruption on goal-directed behaviour could be tested was achieved from looking to the main models and theories of goal-directed memory (e.g., Altmann & Trafton, 2002; Anderson & Lebiere, 1998; Anderson & Douglass, 2001). Within all of these frameworks, performance data from the ToH task was used to constrain model parameters and speak to the cognitive processes involved in the suspension and resumption of goals. As such, the ToH task was chosen as a primary task for all of the experiments reported in the current thesis, primarily because its richness in providing the fine-grained control lacking in many previous studies of task interruption. The well-structured form of task also allowed inferences to be made regarding how goals are processed prior to interruption, particularly during planning phases and during the interruption lag.

The ToH task also proved fruitful in providing a number of performance measures, some of which produced reliable effects across experiments and others that were not so sensitive to the current manipulations.

5.8.2. Move latency

Move latency at the point of interruption was used as the main dependent variable in all of the experiments presented in the current thesis; allowing for comparison of resumption latency in interrupted problems to comparable inter-move latency in control

problems. Control move latency was sensitive to differences between ToH goal structures that required different amounts of processing effort (i.e., moves 1, 3, 4, 5, 9, and 13 in problems solved error-free), supporting the general literature on problem solving in the ToH task (e.g., Anzai & Simon, 1979). The same applied to the time to resume interrupted goals that were suspended at different points within the solution-sequence, although latencies were much longer. Without an interruption lag, resumption latency was six to eight seconds longer to make the third move in an interrupted ToH problem compared to a control problem, a difference of three to five seconds for the fourth move. Move latency in control problems also correlated with the number of moves considered during the initial planning phase of a ToH problem, an effect echoed in interruption conditions when secondary tasks were similar and interruption fell within a complex move sequence (Experiment A).

5.8.3. Move latency and experience in the task domain and experience of being interrupted

The ToH is regarded as a well-structured problem solving task, but in the current experiments performance variation was evident without rich experience in the task domain. Indeed, there is growing support that humans adapt to interruption, evidenced by faster or more efficient resumption of a suspended task, especially with experience of being interrupted in the task domain (e.g., Trafton et al., 2003). In all but Experiment 1, move latency was assessed in each condition after 12 problems; ample opportunity to have gained experience in the task domain. There was little improvement in control move-latency, but significant improvements in resumption latency. It seems that experience of problem solving in the task domain equips participants with declarative representations of suspended goals and goal sequences, and experience of interruption encourages a strategic shift to encoding priming cues located within the task environment to support the reactivation of these representations. Given the fruitfulness of these findings, further research might want to address the role of experience on strategic adaptation to interruption, especially when interruption lags are short or unavailable.

5.8.4. Redundant moves and move-error frequency

Error was assessed by measuring the extent to which participants strayed from an error-free ToH solution-sequence following the intended point of interruption in interruption and control problems. In Chapter 2, error was greater in interrupted problems compared to control problems, suggesting reconstruction as a dominant resumption strategy. Reconsidering the problem state after making the initial moves within a solution-sequence (i.e., up until the point of interruption) might increase the likelhood of reconstructing a move that is more likely to lead to error-free completion than when move execution is allowed to continue without a forced break. In Chapter 3, fewer errors were made in interrupted conditions where disc properties in the current-state had not changed between problem suspension and reinstatement, showing that encoding prior to interruption (even without an interruption lag), leads to a better quality solution-sequence. In this case, participants might have encoded a suspended goal based upon a salient priming cue, with that priming mapping onto declarative rule-based problem solving strategy (e.g., the green disc always follows the blue disc), which if retrieved and executed is more likely to lead to error-free performance.

With the exception of Experiment 1, perhaps the ToH problems used in the current experiments were too simple to expose reliable differences in accuracy decrements between interruption and control conditions. Future research might want to address this issue, although care needs to be taken not to make problems so difficult that participants refrain from using retrieval as a resumption strategy.

5.8.5. Less informative dependent measures

The main cost of interruption seemed to be exhibited when retrieving the first move after the intended point of interruption. Elsewhere, the time to complete a primary task has been shown to increase because of interruption (e.g., Edwards & Gronlund, 1998; Gillie & Broadbent, 1989), but was not evident in the current experiments (with the exception of subgoal 1 in Experiment 1). Completion times in the current experiments were however aggregated across redundant moves, and as so, a precautionary measure was employed: The time to execute the move immediately following that which was interrupted (interruption position plus 1 or $IP+1$)(e.g., move five if interrupted at move

four). There were no differences in IP+1 between interruption and control conditions in any of the current experiments, reinforcing the conclusion that interruption mainly affects the retrieval of a suspended goal.

5.8.6. Dissociating retrieval from other costs imposed by the task

Given the goal-directed form of the ToH task, it is possible that the resumption latencies generated in the current experiments contained processing costs other than that required to retrieve a suspended goal. Current models of goal-directed memory set the time to retrieve a goal at one to two seconds, if the goal was encoded appropriately before it was suspended (Altmann & Trafton, 2002; Anderson & Douglass, 2001). In none of the current experiments did the time to resume an interrupted goal come anywhere near this estimate (the shortest resumption latency was 3.91 s with experience in the task domain in Experiment 8b).

The chief resumption strategy used by participants in the current experiments seemed to be retrieval, with the exception of Experiment 1 where participants demonstrated reconstructive-like behaviour when the distance between the suspended-state and the goal-state was large. I assert that resumption could not have consisted of only reconstruction of a suspended goal, else resumption speed should have been similar irrespective of, task similarity, changing the location of discs in the reinstated display, and allowing a brief opportunity to encode a goal prior to interruption.

A number of factors may have been responsible for longer resumption latencies including the because of the opportunity to assess the utility of executing a retrieved goal through goal-directed enquiry of the goal-state disc configuration and the uncertainty in performing an interrupted action (e.g., Speier et al., 2003). In the prospective memory literature, it has been suggested upon 'noticing' a retrieval cue participants might 'search' the visual display for alternative and perhaps better cues (Einstein & McDaniel, 1996). A commonly held view is that when switching between tasks, participants have to reconfigure processing resources; a process that comes at a cost to starting to switched task (Monsell, 2003). Removing the goal-state from the reinstated visual array may not only assess whether participants are able to retrieve a suspended goal in the ToH task (as recently shown in the ToL task, Hodgetts & Jones, 2005, in revision), but may also factor

out the cost of goal-directed enquiry. Reducing the uncertainty of performing a retrieved goal may only result from experience of handling interruptions in the task domain (e.g., Trafton et al., 2003).

5.9. Future directions

5.9.1. Tracking visual gaze during encoding and resumption

Eye-tracking data would provide a useful insight into the direction of visual gaze prior to interruption and when the primary task is reinstated. Prior to interruption (especially during an interruption lag), attention might be spread across more discs in the current-state disc configuration when engaged in a complex goal sequence and fewer discs when at the end of a goal sequence. Allocation of attention may be more confined the current-state with experience in the task domain and experience of interruption, reflecting the selection and encoding of only a salient priming cue(s)(e.g., Wolfe, 1994). The time cost imposed by inspecting the goal-state disc configuration may also be assessed further through measuring gaze direction when a suspended task is reinstated. When there is a tendency to inspect the goal-state disc configuration (e.g., when interruption falls within a subtask, if disc properties have changed, and so on), gaze might frequently fluctuate between the current-state and goal-state disc configuration. By contrast, when there is a low cost of goal-directed enquiry (e.g., when interruption falls at the end of a subtask, after accumulation of problem solving experience, and so on) gaze might be more fixed on the current-state disc configuration. Eye-tracking data might supplement behavioural data, but caution would be taken to avoid misinterpretation of gaze direction as a direct measure of retrieval (e.g., Anderson, 2004): Avoiding dismissal of the possibility that participants will revisit pre-attended sites (e.g., Aimson, Bothell; Douglass & Anderson, 2005, in press; Horowitz & Wolfe, 1998; 2001).

5.9.2. Disrupt encoding

Findings from Experiments 8a – 9 harmonise with the view of Altmann & Trafton (2002) that a suspended goal may be encoded (through rehearsal) for future retrieval, but provide evidence mainly for an associative priming mechanism underlying such effects.

For instance changing cue properties between primary task suspension and reinstatement markedly slows resumption latency (Experiments 4 – 7). To assess the relationship between the strength of an encoded representation prior to interruption and associative priming generated from cues, a future direction might be to disrupt encoding using irrelevant speech. Irrelevant speech is shown to disrupt rehearsal, even of spatial items (e.g., Jones, Farrand, Stuart & Morris, 1995), and could be presented during an interruption lag. If speed of resumption is dictated by the strength of a rehearsed representation, irrelevant background speech would be expected to disrupt the integrity of the representation, impinging on future retrieval.

5.9.3. Training and expertise

The direction of the current thesis was to establish how novice problem solvers cope with interruption, and how encoding and retrieval strategies develop with some experience in the task domain. Outstanding questions are as follows. Are people able to find the optimal strategy for dealing with interruption or does performance asymptote at a suboptimal level? How are interruption-preparation/recovery strategies adapted to cope with problem solving context/expertise? How does the cost structure of the environment, of knowledge, and of the human cognitive architecture combine in order to determine strategies for coping with interruption? Each of these questions can be addressed by training participants in the task domain (e.g., to use sophisticated problem solving strategies such as goal-recursion) as well as instructing them of strategies that may improve encoding strategy prior to being interrupted (e.g., to extract priming cues based upon salient properties). Only then might retrieval performance come close the parameters estimated by current models of goal-directed memory (Altmann & Trafton, 2001; Anderson & Douglass, 2001).

References

Adams, M. J., Tenney, Y. J., & Pew, R. W. (1995). Situation awareness and the cognitive management of complex systems. *Human Factors, 37,* 85-104.

Adamczyk, P. D. & Bailey, B. P. (2004) If not now, when? The effects of interruption at different moments within task execution. In *Human Factors in Computing Systems: Proceedings of CHI'04.* New York: ACM Press, 271-278.

Allport, D. A., Antonis, B., & Reynolds, P. (1972). On the division of attention: A disproof of the single-channel hypothesis. *Quarterly Journal of Experimental Psychology, 24,* 225–235.

Allport, D. A., Styles, E. A., & Hsieh, S. (1994). Switching intentional set: Exploring the dynamic control of tasks. In C. Umilta & M. Moscovitch (Eds.), *Attention and performance, XV* (pp. 421-452). Hillsdale, NJ: Erlbaum.

Allport, A., & Wylie, G. (2000). Task-switching, stimulus-response binding, and negative priming. In S. Monsell & J. S. Driver (Eds.), *Attention and performance XVIII: Control of cognitive processes* (pp. 35-70). Cambridge, MA: MIT Press.

Altmann, E. M. (2002). Functional decay of memory for tasks. *Psychological Research, 66,* 287-297.

Altmann, E. M. (2004). Advance preparation in task switching: What work is being done? *Psychological Science, 15,* 616-622.

Altmann, E. M., & Gray, W. D. (1999). Serial attention as strategic memory. *Proceedings of the 21st annual meeting of the Cognitive Science Society* (pp. 25-30). Hillsdale, NJ: Erlbaum.

Altmann E. M., & Gray W. D. (2000). Managing attention by preparing to forget. *Proceedings of the IEA 2000/HFES 2000 Congress* (pp. 152-155). Santa Monica: Human Factors and Ergonomics Society.

Altmann E. M., & Gray W. D. (2002). Forgetting to remember: The functional relationship of decay and interference. *Psychological Science, 13,* 27-33.

Altmann, E. M., & Schunn, C. D. (2002). Integrating decay and interference: A new

look at an old interaction. *Proceedings of the 24th annual meeting of the Cognitive Science Society* (pp. 65 –70). Hillsdale, NJ: Erlbaum.

Altmann, E. M., & Trafton, G. J. (1999). Memory for goals: An architectural perspective. *Proceedings of the 21st annual meeting of the Cognitive Science Society* (pp. 19-24). Hillsdale, NJ: Erlbaum.

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 conference of the Cognitive Science Society.* Hillsdale, NJ: Erlbaum.

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

Anderson, J. R. (1976). *Language, memory, and thought.* Hillsdale, NJ: Erlbaum.

Anderson, J. R. (1993). *Rules of the mind.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review, 111,* (4). 1036-1060.

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

Anderson, J. R., Kushmerick, N., & Lebiere, C. (1993). Navigation and Conflict Resolution. In J. R. Anderson (Ed.), *Rules of the Mind* (pp. 93-119). Hillsdale, NJ: Lawrence Erlbaum Associates.

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

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

Anderson, J. R., Reder, L. M., & Lebiere, C. (1996). Working memory: Activation limitations on retrieval. *Cognitive Psychology, 30,* 221-256.

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science, 2,* 396-408.

Arrington, C. M., Altmann, E. M., & Carr, T. H. (2003). Tasks of a feather flock

together: Similarity effects in task switching. *Memory & Cognition, 31*, 781-789.

Arrington, C. M., & Logan, G. D. (2004). The cost of a voluntary task switch. *Psychological Science, 15*, 610-615.

Arroyo E., & Selker T. (2003). Arbitrating multimodal outputs: Using ambient displays as interruptions. In J. Jacko and C. Stephanidis (Eds.), *Human-Computer Interaction: Theory and Practice (Part II) - Proceedings of HCI International 2003, Vol. 2* (pp. 591-595). Mahwah: Lawrence Erlbaum Associates.

Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation: Advances in research and theory, Vol. 2* (pp. 89-195). New York: Academic Press.

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

Avrahami, D., Gergle, D., Hudson, S. E., & Kiesler, S. (under review). Improving the accuracy of cell phone interruptions: A study on the effect of contextual information on the behavior of callers. Paper submitted to *Behaviour and Information Technology*.

Awh, E., & Jonides, J. (2001). Overlapping mechanisms of attention and spatial working memory. *Trends in Cognitive Sciences, 5 (3),* 119-126.

Awh, E., Jonides, J., & Reuter-Lorenz, P. A. (1998). Rehearsal in spatial working memory. *Journal of Experimental Psychology: Human Perception and Performance, 24,* 780-790.

Baddeley, A. D. (1968). A three-minute reasoning test based on grammatical transformation. *Psychonomic Science, 10,* 341-2.

Baddeley, A. D. (1986). *Working memory.* Oxford: Oxford University Press.

Baddeley, A. D., & Hitch, G.J. (1974). Working memory. In G.A. Bower (Ed.), *Recent advances in learning and motivation* (Vol. 8, pp. 47–90). New York: Academic Press.

Bailey, B. P., Konstan, J. A., & Carlis, J. V. (2000). Measuring the effects of interruptions on task performance in the user interface. *IEEE Conference on Systems, Man, and Cybernetics.* (pp. 757-762). Nashville, TN.

Bailey, B. P., Konstan, J. A., & Carlis, J. V. (2000). Measuring the effects of interruptions on task performance in the user interface. In *IEEE Conference on Systems, Man, and Cybernetics 2000 (SMC 2000)*, IEEE (pp.757-762).

Bailey, B. P., Konstan, J. A., & Carlis, J. V. (2001). The effects of interruptions on task performance, annoyance, and anxiety in the user interface. *Proceedings of Human-Computer Interaction – INTERACT 2001* (pp. 593-601). IOS Press: IFIP.

Bainbridge, L. (1984). Analysis of verbal protocols from a process control task. In E. Edwards and F. P. Lees (Eds.), *The human operator in process control* (pp. 146-158). London: Taylor and Francis.

Barnes. V. E., & Monan, W. P. (1990). Cockpit distractions: precursors to emergencies. *Proceedings of the Human Factors Society 34th Annual Meeting*, Santa Monica, CA, 1142–1144.

Baron, R. S. (1986). Distraction-conflict theory: Progress and problems. *Advances in Experimental Social Psychology, 19,* 1-39.

Bauer, B., Jolicouer, P., & Cowan, W. B. (1996). Visual search for targets that are or are not linearly-separable from distractors. *Vision Research, 36,* 1439-1466.

Berg, P. T. van den, Roe, R. A., Zijlstra, F. R. H., & Krediet, I. (1996). Temperamental factors in the execution of interrupted editing tasks. *European Journal of Personality, 10*, 233-248.

Bishop, D. V. M., Aamodt-Leeper, G., Creswell, C., McGurk, R., & Skuse, D. H. (2001). Individual differences in cognitive planning on the Tower of Hanoi task: Neuropsychological maturity or measurement error? *Journal of Child Psychiatry, 42 (4),* 551-556.

Botvinick, M. M., & Bylsma, L. M. (in press). Distraction and actions slips in an everyday task: Evidence for a dynamic representation of task content. *Psychonomic Bulletin and Review.*

Brandimonte, M., Einstein, G. O., & McDaniel, M. A. (1996). *Prospective memory: Theory and applications.* Mahwah, NJ: Erlbaum.

Broadbent, D. E. (1958). *Perception and communication.* Oxford: Pergamon.

Brown, J. (1958). Some tests of the decay theory of immediate memory. *Quarterly Journal of Experimental Psychology, 10,* 12-21.

Burmistrov I., & Leonova A. (1996). Effects of interruptions on computerised clerical

task performance. In J. Gornostaev and S. Overmyer (Eds.), *Human-Computer Interaction: Human Aspects of Business Computing. Proceedings of EWHCI'96.* (pp. 21-29). Moscow: ICSTI.

Burmistrov I., & Leonova A. (2003). Do interrupted users work faster or slower? The micro-analysis of computerized text editing task. In J. Jacko and C. Stephanidis (Eds.), *Human-Computer Interaction: Theory and Practice (Part I) - Proceedings of HCI International 2003, Vol. 1.* (pp. 621-625). Mahwah: Lawrence Erlbaum Associates.

Byrne, M. D., & Bovair, S. (1997). A working memory model of a common procedural error. *Cognitive Science, 21,* 31-61.

Cellier, J., & Eyrolle, H. (1992). Interference between switched tasks. *Ergonomics, 35,* 25-36.

Cellier, J., & Eyrolle, H. (2000). The effects of interruptions in work activity: field and laboratory results. *Applied Ergonomics, 31,* 537-543.

Charman, S. C., & Howes, A. (2003). The adaptive user: an investigation into the cognitive and task constraints on the generation of new methods. *Journal of Experimental Psychology: Applied,* 9, 236-248.

Chase, W. G., & Ericsson, K. A. (1982). Skill and working memory. *The Psychology of Learning and Motivation, 16,* 1-58.

Chung, P. H., & Byrne, M. D. (2004). Visual cues to reduce errors in a routine procedural task. In *Proceedings of the 26th Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Coates, J. 1990. It is legitimate to be unavailable. *Industrial and Commercial Training, 22(5),* 8-11.

Clifford J. D. and Altmann E. M. (2004). Managing multiple tasks: reducing the resumption time of the primary task. *Poster presented at the 26th annual conference of the Cognitive Science Society,* Chicago.

Covet, S. R. (1989). The seven habits of highly effective people. New York: Simon & Schuster.

Cubesmart, Inc. (2002). *Social Interruption and the Loss of Productivity.* Retrieved October 16, 2003, from http://www.interruptions.net/literature/CubeSmart-productivity-wp1.pdf.

Cutrell, E., Czerwinski, M., & Horvitz, E. (2000). Effects of Instant Messaging

Interruptions on Computing Tasks. *Extended Abstracts of CHI '2000, Human Factors in Computing Systems* (pp. 99-100). ACM press.

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). IOS Press, IFIP.

Czerwinski, M. P., Chrisman, S. E., & Rudisill, M. (1991). *Interruptions in multitasking situations: The effects of similarity and warning* (Technical Report JSC-24757): National Aeronautics and Space Administration, Lyndon B. Johnson Space Center, Houston, Texas.

Czerwinski M., Chrisman S., & Schumacher B. (1991). The effects of warnings and display similarities on interruptions in multitasking environments, *SIGCHI Bulletin*, 23 (4), 38-39

Czerwinski, M., Cutrell, E., & Horvitz, E. (2000). Instant messaging and interruption: Influence of task type on performance. In C. Paris, N. Ozkan, S. Howard, and S. Lu (Eds.), *OZCHI 2000 Conference Proceedings,* (pp. 356-361). Sydney, Australia, Dec. 4-8.

Czerwinski, M., Cutrell E., & Horvitz, E. (2000). Instant messaging: Effects of relevance and time. In S. Turner and P. Turner (Eds.), *People and Computers XIV: Proceedings of HCI 2000, Vol. 2*, British Computer Society, 71-76.

Czerwinski M., Horvitz E. & Wilhite S. (2004). A diary study of task switching and interruptions. In *Human Factors in Computing Systems: Proceedings of CHI'04*. (pp. 175-182). New York: ACM Press.

Dabbish, L., & Kraut, R. (2003). Coordinating communication: Awareness displays and interruption. *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI'03): Extended abstracts* (pp. 786-787). New York: ACM Press.

Dahms A. R. (1988, March-April) Time management and the knowledge worker, *Industrial Engineering*, 27-29.

Davies, S. P. (2000). Memory and planning processes in solutions to well structured problems. *The Quarterly journal of Experimental Psychology, 53(A), 3,* 896-927.

Davies, S. P. (2003). Initial and concurrent planning in solutions to well-structured

problems. *The Quarterly Journal of Experimental Psychology, 56A(7),* 1147-1164.

Damos, D. L., & Tabachnick, B. G. (2001). *The Effect of Interruptions on Flight Crew Performance: ASRS Reports.* Los Angeles: Damos Research Associates.

Degani, A., & Wiener, E. L. (1990). *Human Factors of Flightdeck Checklists: The Normal Checklist.* NASA-CR-177549. Washington D.C.: National Aeronautics and Space Administration.

De Jong, R. (2000). An intention-activation account of residual switch costs. In S. Monsell & J. Driver (Eds.), *Attention and performance XVIII: Control of cognitive processes* (pp. 357–376). Cambridge, MA: MIT Press

DeMarco, T., & Lister, T. (1999). *Peopleware: Productive Projects and Teams (2nd Ed.)*, New York: Dorset House.

Detweiler, M. C., Hess, S. M., & Phelps, M. P. (1994). *Interruptions and working memory.* Unpublished technical report. University Park, PA: The Pennsylvania State University.

Deutsch, J. A., & Deutsch, D. (1963). Attention: Some theoretical considerations. *Psychological Review, 70,* 80-90.

Duncan, J. (1980). The demonstration of capacity limitation. *Cognitive Psychology, 12,* 75-96.

Dismukes R. K. (2003) Managing interruptions, distractions and concurrent task demands. *Paper presented at the 13th Annual Air Transport Association (ATA) Advanced Qualification Program (AQP) Working Group Conference*

Dismukes, K., Young, K., & Sumwalt, R. (1998). Cockpit interruptions and distractions: Effective management requires a careful balancing act. *ASRS Directline, 10,* 4-9.

Dodhia, R. M., & Dismukes R. K. (2003, November). *A task interrupted becomes a prospective task: Encoding and retrieval manipulations.* Poster presented at the 44th Annual Meeting of the Psychonomic Society, Vancouver, Canada.

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

Edworthy, J., & Stanton, N. (1995). A user-centered approach to the design and evaluation of auditory warning signals: 1. Methodology. *Ergonomics, 38(11),* 2262-2281.

Edworthy, J., Loxley, S., & Dennis, I. (1991). Improving auditory warning design: relationship between warning sound parameters and perceived urgency. *Human Factors, 33,* 205-231.

Einstein, G. O., Holland, L. J., McDaniel, M. A., & Guynn, M. J. (1992). Age-related deficits in prospective memory: The influence of task complexity. *Psychology and aging, 7,* 471-478.

Einstein, G. O., & McDaniel, M. A. (1990). Normal aging and prospective memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 717-726.

Einstein, G. O., & McDaniel, M. A. (1996). Retrieval processes in prospective memory: Theoretical approaches and some new empirical findings. In M. Brandimonte, G. O. Einstein, & M. A. McDaniel (Eds.), *Prospective memory: Theory and applications* (pp. 115-142). Hillsdale, NJ: Erlbaum.

Einstein, G. O., McDaniel, M. A., Manzi, M., Cochran, B., & Baker, M. (2000). Prospective memory and aging: Forgetting over short delays. *Psychology and Aging, 15,* 671-683.

Einstein, G. O., McDaniel M. A., Williford C. L., Pagan J. L., & Dismukes, R. K. (2003). Forgetting of intentions in demanding situations is rapid. *Journal of Experimental Psychology: Applied,* 9, 147-162.

Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review, 102,* 211-245.

Eyrolle, H., & Cellier J. M. (2000). Some effects of interruptions in work activity: Field and laboratory results. *Applied Ergonomics, 31,* 537-543.

Field, G. E. (1987). Experimentus interruptus. *ACM SIGCHI Bulletin, 19,* 42-46.

Field, G. & Spence, R. (1994). Now, where was I? *New Zealand Journal of Computing, 5,* 35-43.

Fischer, B., & Glanzer, M. (1986). Short-term storage and the processing of cohesion during reading. *Quarterly Journal of Experimental Psychology, 38A,* 431-460.

Fisk, A. D., Rogers, W. A. (1991). Toward an understanding of age-related memory and visual search effects. *Journal of Experimental Psychology: General, 120,*131-149.

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: Erlbaum.

Flynn, E. A., Barker, K.N., Gibson, J. T., Pearson, R. E., Berger, B. A., & Smith, L. A. (1999). Impact of interruptions and distractions on dispensing errors in an ambulatory care pharmacy. *American Journal of Health System Pharmacy, 56,* 1319-1325.

Fu, W, -t., & Gray, W. D. (2004). Resolving the paradox of the active user: Stable suboptimal performance in interactive tasks. *Cognitive Science, 26(8),* 359-382.

Fum, D. & Del Missier, F. (2001). Adaptive selection of problem solving strategies. *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society.* Mahwah, NJ: Lawrence Erlbaum Associates, 313-318.

Garcia-Ogueta, I. (1993). Internal attentional switching: Effects of predictability, complexity and practice. *Acta Psychologica, 83,* 13-32.

Gilhooly, K. J., Phillips, L. H., Wynn, V., Logie, R. H., & Della Sala, S. (1999). Planning processes and age in the five-disc Tower of London task. *Thinking and Reasoning, 5,* 339-361.

Gilhooly, K. J., Wynn, V., Phillips, L. H., Logie, R. H., & Della Sala, S. (2002). Visuo-spatial and verbal working memory in the five-disc Tower of London task: An individual differences approach. *Thinking and Reasoning, 8,* 165-178.

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

Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review, 91,* 1-65.

Goel, V., Pullara, D., & Grafman, J. (2001). A computational model of frontal lobe dysfunction: Working memory and the Tower of Hanoi task, *Cognitive Science, 25 (2),* 287-313.

Goolsby, B. A., & Suzuki, S. (2001). Understanding priming of color-singleton search: Roles of attention at encoding and "retrieval". *Perception & Psychophysics, 63 (6),* 929 – 944.

Gopher, D., Armony, L., & Greenspan, Y. (2000). Switching tasks and attention policies. Journal of Experimental Psychology: General, 129, 308–339.

Goschke, T., & Kuhl, J. (1993). Representation of intentions: Persisting activation in memory. *Journal of Experimental Psychology: Learning, Memory, &*

Cognition, 19, 1211-1226.

Goschke, T., & Kuhl, J. (1996). Remembering what to do: Explicit and implicit memory for intentions. In M. Brandimonte, G. Einstein, & M. McDaniel (Eds.), *Prospective memory: Theory and applications* (pp. 53-91). Hillsdale, NJ: Erlbaum.

Gray, W. D., & Fu, W, -t. (2004). Soft constraints in interactive behaviour: The case of ignoring perfect information in-the-world for imperfect knowledge in-the-head. *Cognitive* Science 28*(3),* 359-382.

Gray, W. D., Schoelles, M. J., & Sims, C. R. (2005). Adapting to the task environment: Explorations in expected value. *Cognitive Systems Research, 6,* 27 – 40.

Greene, R. L. (1992). *Human Memory: Paradigms and Paradoxes.* Hillsdale, NJ: Erlbaum.

Griffon-Fouco, M., & Ghertman, F. (1984). Recueil de données sur les facteurs humains à l'Electricite de France. *Operational Safety of Nuclear Power Plants* (pp. 157-172). Vienna: International Atomic Energy Agency.

Guynn, M. J., McDaniel, M. A., & Einstein, G. O. (1998). Prospective memory: When reminders fail. *Memory & Cognition, 26,* 287-298.

Gunzelmann, G., & Anderson, J. R. (2001). An ACT-R Model of the Evolution of Strategy Use and Problem Difficulty. *Proceedings of the Fourth International Conference on Cognitive Modeling,* (pp. 109-114). Mahwah, NJ: Lawrence Erlbaum Associates.

Gunzelmann, G., & Anderson, J. R. (2003). Problem solving: Increased planning with practice. *Cognitive Systems Research, 4,* 57-76.

Halford, G.S., Wilson, W.H. & Phillips, W. (1998) Processing capacity defined by relational complexity: Implications for comparative, developmental and cognitive psychology. *Behavioral Brain Sciences, 21(6),* 803-831.

Handley, S.J., Capon, A., Copp, C., & Harper, C. (2002). Conditional reasoning and the Tower of Hanoi: The role of verbal and spatial working memory. British Journal of Psychology, 93, 501-518.

Harrison, B. L., Ishii, H., Vicente, K. J., & Buxton, W. A. S. (1995). Transparent layered user interfaces: An evaluation of a display design to enhance focused and divided attention. In *Human Factors in Computing Systems: Proceedings*

of CHI'95 (pp. 317-324). New York: ACM Press.

Hayes, J. R. (1978). *Cognitive Psychology: Thinking and Creating.* Homewood, Ill.: Dorsey Press.

Hayes-Roth, B., & Hayes-Roth, F. (1979). A cognitive model of planning. *Cognitive Science, 3(4),* 275-310.

Hess, S. M., & Detweiler, M. C. (1994). Training to reduce the disruptive effects of interruptions. *Proceedings of the Human Factors and Ergonomics Society 38th Annual Meeting, HFES.* Santa Monica, 1173 – 1177.

Hicks, J. L., Marsh, R. L., & Russell, E. J. (2000). The properties of retention intervals and their affect on retaining prospective memories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26,* 1160-1169.

Hirst, W., Spelke, E.S., Reaves, C.C., Caharack, G., & Neisser, U. (1980). Dividing attention without alternation or automaticity. *Journal of Experimental Psychology: General,* 109, 98-117.

Hockey, G. R. J. (1970). Effects of Loud Noise on Attentional Selectivity. *Quarterly* Journal of Experimental Psychology, 22, 28-36.

Hodgetts, H. M. (2004). Task interruption: The effects of length, complexity, preparation and retrieval cues. *Unpublished doctoral dissertation,* Cardiff University.

Hodgetts, H. M., & Jones, D. M. (2003). Interruptions in the Tower of London task: Can preparation minimise disruption? *Proceedings of the 47th Annual Meeting of the Human Factors and Ergonomics Society (HFES)*, pp. 1000-1004.

Hodgetts, H. M. & Jones, D. M. (2005). Interruptions of the Tower of London task: Support for a goal activation approach. *Submitted to Journal of Experimental Psychology: General.*

Horowitz, T.S. & Wolfe, J.M. (1998). Visual search has no memory. *Nature, 394,* 575-577.

Horowitz, T.S. & Wolfe, J.M. (2001). Search for multiple targets: Remember the distractors, forget the search. *Perception & Psychophysics, 63*, 272-285.

Horvitz, E., Koch, P., Kadie, C. M. & Jacobs, A. (2002). Coordinate: Probabilistic forecasting of presence and availability. In J. Breese, & D. Koller (Eds.), *Proceedings of the 18th Conference in Uncertainty in Artificial Intelligence* (pp. 224-233). Morgan Kaufmann Publishers.

Horvitz, E., Jacobs, A. & Hovel, D. (1999). Attention- sensitive alerting. *15th Conference on Uncertainty and Artificial Intelligence (UAI '99)*, Stockholm, Sweden. Morgan Kaufmann Publishers.

Howes, A & Payne, S.J. (2001). The strategic use of memory for frequency and recency in search control. In J. D. Moore & K. Stenning (Eds.), *Proceedings of the Annual Conference of the Cognitive Science Society*, Edinburgh, 435-443.

Humphreys, G. (1981). Flexibility of attention between stimulus dimensions. *Perception & Psychophysics, 30,* 291-302.

Hunt, A. R., & Klein, R. M. (2002). Eliminating the cost of task set reconfiguration. *Memory and Cognition, 30,* 529-539.

Iqbal, S. T., & Bailey, B. P. (2005). Investigating the effectiveness of mental workload as a predictor of opportune moments for interruption. In *Human Factors in Computing Systems: Proceedings of CHI'05: Extended Abstracts* (pp. 1489-1492). New York: ACM Press.

Iqbal, S. T., Adamczyk, P. D., Zheng, X. S., & Bailey B. P. (2005). Towards an index of opportunity: Understanding changes in mental workload during task execution. In *Human Factors in Computing Systems: Proceedings of CHI'05.* (pp. 311-320) New York: ACM Press.

Jackson T. W., Dawson R. J., & Wilson D. (2001). The cost of email interruption, *Journal of Systems and Information Technology, 5 (1),* 81-92

Jackson T. W., Dawson R. J., & Wilson D. (2002). Evaluating the effect of email interruptions within the workplace. *Conference on Empirical Assessment in Software Engineering* (pp. 3-7).

Jackson, T. W., Dawson, R., & Wilson, D. (2003). Understanding email interaction increases organizational productivity. *Communications of ACM, 46,* 80-84.

Jambon, F. (1996). Formal modelling of task interruptions. *Human Factors in Computing Systems: Proceedings of CHI'96, Conference Companion* (pp. 45-46). New York: ACM Press.

James, W. (1890). *The principles of psychology.* Holt and Co: New York.

Jersild, A. T. (1927). Mental set and switch. *Archives of Psychology, 89,* 1-81.

Johansson, G., and Aronsson, G. (1984) Stress reactions in computerized administrative work. *Journal of Occupational Behaviour, 5,* 159-181.

Jones, D. M. (1999). The 1997 BPS Broadbent Lecture: The cognitive psychology of auditory distraction. *British Journal of Psychology, 90,* 167-187.

Jones, D. M., Farrand, P. A., Stuart, G. P., & Morris, N. (1995). The functional equivalence of verbal and spatial information in short-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 1008-1018.

Jones, J. W., and McLeod, R. "The Structure of Executive Information Systems: An Exploratory Analysis, *Decision Sciences (17),* 220-249

Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review, 99,* 123-148.

Kahneman, D. (1973). *Attention and effort.* Englewood Cliffs, NJ: Prentice Hall.

Katz, R. (1995). Automatic versus user controlled methods of briefly interrupting telephone calls. *Human Factors, 37,* 321-34.

Kirmeyer, S. L. (1988). Coping with competing demands: Interruption and the Type-A pattern. *Journal of Applied Psychology, 73,* 621- 629.

Koch, I. (2001). Automatic and intentional activation of task sets. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 27,* 1474-1486.

Koch, I. (2003). The role of external cues for exogenous advance reconfiguration in task switching. *Psychonomic Bulletin & Review, 10 (2),* 488-492.

Kirsh, D. (2002) Interruptions in the Workplace. Unpublished manuscript. Available: *http://adrenaline.ucsd.edu/edf/Position_paper.htm*

Koriat, A., Ben-Zur, H., & Nussbaum, A. (1990). Encoding information for future action: Memory for to-be-performed versus memory for to-be-recalled tasks. *Memory & Cognition, 18,* 568-578.

Kotovsky, K., Hayes, J. R., & Simon, H. A. (1985). Why are some problems hard? Evidence from Tower of Hanoi, *Cognitive Psychology, 17,* 248-294.

Kreifeldt, J. G., & McCarthy, M. E. (1981). Interruption as a test of the user-computer interface. *Proceedings of the 17th Annual Conference on Manual Control* (pp. 655-667). JPL Publication 81-95, Jet Propulsion Laboratory, California Institute of Technology.

Kvavilashvili, L. (1987). Remembering intention as a distinct form of memory. *British Journal of Psychology, 78,* 507-518.

Lahlou, S., Kirsh, D., Rebotier, T., Reeves, C., & Remy, M. (2002). Interruptions in the workplace. Retrieved January 19, 2003 from:*http://www.adrenaline.ucsd.edu/edf/Experiment2.htm.*

Larkin, J. (1989). Display-based problem solving. In D. Klahr & K. Kotovsky (Ed.), *Complex Information Processing: The Impact of Herbert A. Simon.* Hillsdale, NJ: Erlbaum.

Latorella, K. A. (1996). *Investigating Interruptions: Implications for Flightdeck perfromance.* Unpublished Doctoral dissertation. State University of New York at Buffalo.

Latorella, K. A. (1998). Effects of modality on interrupted flight deck performance: Implications for data link. *Proceedings of the 42nd Annual Meeting of the Human Factors and Ergonomics Society, HFES*, Santa Monica, 249-253.

Latorella, K. A. (1999). *Investigating Interruptions: Implications for Flightdeck Performance*, NASA/TM-1999-209707. Washington: National Aviation and Space Administration (also published in 1996 as Doctoral Dissertation, State University of New York at Buffalo).

Lavie, N., & Tsal, Y. (1994). Perceptual load as a major determinant of the locus of selection in visual attention. *Perception & Psychophysics, 56 (2),* 183-197.

Lewin, K. (1926). Vorsatz, Wille und Bedürfnis [Intention, will and need]. *Psychologische Forschung, 7,* 330-385.

Linde, C., & Goguen, J. (1987). *Checklist interruption and resumption: A linguistic study.* NASA Technical Report NASA-CR-177460. Moffett Field: NASA Ames Research Center.

Logan, G. D., & Bundesen, C. (2003). Clever homunculus: Is there an endogenous act of control in the explicit task cuing procedure? *Journal of Experimental Psychology: Human Perception and Performance*, *29*, 575-599.

Loukopoulos, L. D., Dismukes, R. K., & Barshi, I. (2001). Cockpit interruptions and distractions: A line observation study. In R. Jensen (Ed.), *Proceedings of the 11th International Symposium on Aviation Psychology.* Columbus, OH: Ohio State University. Columbus: Ohio State University Press. Retrieved March 30, 2003, from http://human- factors.arc.nasa.gov/flightcognition/pubs.html

Lovett, M. C., & Anderson, J. R. (1996). History of success and current context in problem solving: Combined influences on operator selection. *Cognitive*

Psychology, 31, 168-217.

Lovett, M.C., Reder, L.M., & Lebiere, C. (1999). Modeling Working Memory in a Unified Architecture: An ACT-R Perspective. In Miyake, A. and Shah, P. (Eds), *Models of Working Memory* (pp. 135-182). Oxford: Oxford University Press.

Maljkovic, V., & Nakayama, K. (1994). Priming of pop-out I: A role of features. *Memory and Cognition, 22,* 657-672.

Maljkovic, V., & Nakayama, K. (2000). Priming of pop-out: III. A short-term implicit memory system beneficial for rapid target selection. *Visual Cognition, 7,* 571-595.

Mäntylä, T. (1996). Activating actions and interrupting intentions: Mechanisms of retrieval sensitization in prospective memory. In M. Brandimonte, G. O. Einstein, & M. A. McDaniel (Eds.), *Prospective Memory: Theory and Applications* (pp. 93 – 113). Mahwah, NJ: Erlbaum.

Mäntylä, T., & Sgaramella, T. (1997). Interrupting intentions: Zeigarnik-like effects in prospective memory. *Psychological Research, 60(3),* 192-199.

Marsh, R. L., Hancock, T. W., & Hicks, J. L. (2002). The demands of an ongoing activity influence the success of event-based prospective memory. *Psychonomic Bulletin & Review, 9,* 604-610.

Marsh, R. L., Hicks, J. L., & Bink M. L. (1998). The activation of completed, incomplete and partially completed intentions. *Journal of Experimental Psychology: Learning, Memory and Cognition, 24,* 350-361.

Mayr, U., & Keele, S.W. (2000). Changing internal constraints on action: The role of backward inhibition. *Journal of Experimental Psychology: General, 129,* 4-26.

Mayr, U., & Kliegl, R. (2003). Differential effects of cue changes and task changes on task-set selection costs. *Journal of Experimental Psychology: Learning, Memory & Cognition, 29,* 362 – 372.

McDaniel, M. A., & Einstein, G. O. (1993). The importance of cue familiarity and cue distinctiveness in prospective memory. *Memory, 1,* 23-41.

McDaniel, M., Einstein, G. O., Stout, A., & Morgan, Z. (2003). Aging and maintaining intentions over delays: do it or lose it. *Psychology and Aging, 18,* 823-835.

McDaniel, M, Einstein, G. O., Graham, T., & Rall, E. (2004). Delaying execution of

intentions: Overcoming the costs of interruptions. *Applied Cognitive Psychology, 18(5),* 533-547.

McFarlane, D. C. (1997). *Interruption of people in human computer interaction: A general unifying definition of human interruption and taxonomy.* (NRL formal report NRL/FR/5510-97-9870). Washington: US Naval Research Laboratory.

McFarlane, D. C. (1999). Coordinating the interruption of people in human-computer interaction. In M. A. Sasse & C. Johnson (Eds.), *Human-computer interaction-INTERACT '99* (pp. 295-303). Edinburgh, UK:IOS.

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.

Meacham, J. A., & Leiman, B. (1982). Remembering to perform future actions. In U. Neisser (Ed.), *Memory observed: Remembering in natural contexts* (pp. 327-336). San Francisco: Freeman.

Media Metrix (1999). Report on digital media. Retrieved December 05, 2003, from http://www.mediametrix.com.

Meiran, N. (1996). Reconfiguration of processing mode prior to task performance. *Journal of Experimental Psychology, 22,* 1423–1442.

Meiran, N., Chorev, Z., & Sapir, A. (2000). Component processes in task switching. *Cognitive Psychology, 41,* 211-253.

Merced, D. (1991). *The effects of task interruption on user performance in a multitasking environment: Implications for computerized reminders.* Unpublished master's thesis, Faculty of the Virginia Polytechnic Institute and State University.

Microsoft PressPass (2002). Retrieved November 12, 2002, from http://www.microsoft.com/presspass/press/2002/nov02/1108MSN8GlobalTimeLine.mspx.

Miller, S. L. (2002). Window of opportunity: Using the interruption lag to manage disruption in complex tasks. *Proceedings of the 46th Annual Meeting of the Human Factors and Ergonomics Society.* Santa Monica, CA: Human Factors

and Ergonomics Society. Retrieved 20 July, 2003, fromhttp://interruptions.net/literature/Miller-HFES02.doc

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

Miyata, Y., & Norman, D. (1986). Psychological issues in support of multiple activities. In D. A. Norman and S. W. Draper (Eds.), *User Centered Systems Design: New Perspectives on Human-Computer Interaction* (pp. 265-284). Hillsdale: Lawrence Erlbaum Associates.

Monan, W. P. (1979). *Distraction: A human factor in air carrier hazard events.* NASA- TM- 78608. Washington, D.C.: National Aeronautics and Space Administration.

Monk, C. A. (2004). The effect of frequent versus infrequent interruptions on primary task resumption. *Proceedings of 48th Annual Meeting of the Human Factors and Ergonomics Society, HFES,* Santa Monica, 295-299.

Monk, C., Boehm-Davis, D. A., & Trafton, J. G. (2002). The attentional costs of interrupting task performance at various stages. *Proceedings of the 2002 Human Factors and Ergonomics Society Annual Meeting.* Santa Monica, CA: HFES.

Monk, C. A., Boehm-Davis, D. A. and Trafton J. G. (2004, October). Very brief interruptions result in resumption cost. *Poster session presented at the 26th annual conference of the Cognitive Science Society, Chicago.*

Monk, C. A., Boehm-Davis, D. A., & Trafton, J. G. (2004). Recovering from interruptions: Implications for driver distraction research. *Human Factors, 46 (4),* 650-663.

Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences, 7,*134-140.

Monsell, S., Sumner, P., & Waters, H. (2003). Task-set reconfiguration with predictable task switches. *Memory & Cognition, 31 (3),* 327-342.

Morgan, S. E., Reichart, T. & Harrison, T. R. (2002). *From numbers to words: Reporting statistical results for the social sciences.* Boston: Allyn & Bacon.

Muter, P. (1980). Very rapid forgetting. *Memory & Cognition, 8,* 174-179.

Nagata, S. F. (2003). Multitasking and interruptions during mobile web tasks. *Proceedings of the 47th Annual Meeting of the Human Factors and Ergonomics Society HFES,* Santa Monica, CA, 1341-1345.

Nairne, J. S. (1990). Similarity and long-term memory for order. *Journal of Memory and Language, 29,* 733-746.

Nairne, J. S. (2002). Remembering over the short-term: The case against the standard model. *Annual Review of Psychology, 53,* 53-81.

Nakajima, Y., & Sato, K. (1989). Distractor difficulty and the long-term recency effect. *American Journal of Psychology, 102,* 511-521.

National Transportation Safety Board. (1973). *Aircraft Accident Report: Eastern Airlines Inc. Lockheed L-1011, N310EA, Miami International Airport, Miami, Florida, December 29, 1972.* Washington, D.C.: National Transportation Safety Board.

National Transportation Safety Board. (1988). *Aircraft Accident Report- Northwest Airlines, McDonnell Douglas DC-9-82, N312RC, Detroit Metropolitan Wayne Co., Airport, Romulus MI, August 16, 1987.* Washington, D.C.: National Transportation Safety Board.

Navon, D., & Gopher, D. (1979). On the economy of the human processing system. *Psychological Review, 86(3),* 214-255.

Neisser U. (1976). *Cognition and reality: principles and implications of cognitive psychology.* San Francisco: Freeman.

Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: roles of inhibitionless spreading activation and limited-capacity attention. Journal of Experimental Psychology: General, 106, 226-254.

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

Norman, D. (1968). Toward a theory of memory and attention. Psychological Review, 75, 522-536.

Obermayer R. W. and Nugent W. A. (2000) Human-computer interaction for alert warning and attention allocation systems of the multi-modal watchstation. In P. Hamburger (Ed.), *Integrated Command Environments: Proceedings of SPIE, Vol. 4126, Bellingham: SPIE – The International Society for Optical Engineering,* 14-22

O'Conaill, B., & Frohlich, D. (1995). Timespace in the workplace: Dealing with interruptions. *Companion proceedings of CHI'95* (pp. 262-263). New York: ACM SIGCHI.

Oulasvirta, A., & Saariluoma, P. (2004). Long-term working memory and interrupting messages in human-computer interaction. *Behaviour and Information Technology, 23(1),* 53-64.

Ovsiankina, M. (1928). Die Wiederaufnahme unterbrochoter Handlungen. *Psychologische Forschung, 11,* 302-379.

Patalano, A. L., & Seifert, C. M. (1997). Opportunistic planning: Being reminded of pending goals. *Cognitive Psychology, 34,* 1-36.

Patterson, R.D. (1990). Auditory warning sounds in the work environment. *Philosophical Transactions of the Royal Society of London, B,* 327, 485-492.

Paquiot, J., Eyrolle, H., & Cellier, J. (1986). Conséquences de l'interférence entre les tâches de saisie sur la durée des traitements et l'occurrence des erreurs. In R. Patesson (Ed.), *L'homme et L'écran* (pp. 199- 209). Editions de l'Université de Bruxelles.

Pashler, H. E. (1998). *The Psychology of Attention.* Cambridge, MA: MIT Press.

Pantalano, A. L., & Seifert, C. M. (1994). Memory for impasses in problem solving. *Memory & Cognition, 22,* 234-242.

Payne, S. J., Richardson, J. & Howes, A. (2000). Strategic use of familiarity in display-based problem solving. *Journal of Experimental Psychology: Learning, Memory & Cognition, 26,* 1685-1701.

Perlow, L. A. (1999). The time famine: Toward a sociology of work time. *Administrative Science Quarterly, 44,* 57-81.

Perry, J. (1998). *Effects of interruptions on the cognitive management of complex tasks. Unpublished* master's thesis, University of Oklahoma Graduate College.

Phillips, W. A., & Christie, D. F. M. (1977). Components of visual memory. *Quarterly Journal of Experimental Psychology, 29,* 117-133.

Phillips, L. H., Wynn, V. E., Gilhooly, K. J., Della Sala, S., & Logie, R. H. (1999). The role of memory in the Tower of London task. *Memory, 7,* 209- 231.

Phillips, L. H., Wynn, V. E., McPherson, S., & Gilhooly, K. J. (2001). Mental planning and the Tower of London task. *Quarterly Journal of Experimental Psychology, 54A,* 579-597.

Pitney Bowes (1998). Pitney Bowes study finds messaging creates greater stress at work. Retrieved January 10, 2002, from http://www.exn.ca/Stories/1998/05/20/58.asp.

Posner, M. I., & Konick, A. W. (1966). On the role of interference in short-term retention. *Journal of Experimental Psychology, 72,* 221-231.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology, 32,* 3-25.

Raaijmakers, J. G. W., & Shiffrin, R. M. (1980). SAM: A theory of probabilistic search of associative memory. In G. H. Bower (Ed.), *The psychology of learning and motivation, vol. 14* (pp. 207-262). New York: Academic Press.

Reason, J. (1990). *Human Error.* New York: Cambridge University Press.

Reber, P. J., & Kotovsky, K. (1992). Learning and problem solving under a memory load. *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Lawrence Erlbaum.

Reder, L. M., & Anderson, J. R. (1980). A partial resolution to the paradox of interference: the role of integrating knowledge. *Cognitive Psychology, 12,* 447-472.

Reder, L. M. & Schunn, C. D. (1999) Bringing Together the Psychometric and Strategy Worlds: Predicting Adaptivity in a Dynamic Task. In Gopher, D. & Koriat, A. (Eds.), *Cognitive regulation of performance: Interaction of theory and application. Attention and Performance XVII,* MIT Press, pp. 315-342.

Reisberg, D. (1987). External representations and the advantages of externalizing one's thoughts. *Proceedings of the Cognitive Science Society.* Seattle, WA. Hillsdale, NJ: Erlbaum.

Renaud, K. (2000). Expediting rapid recovery from interruptions by providing a visualization of application activity. *Proceedings of OZCHI 2000: Interfacing Reality in the New Millennium* (pp. 348-355). Sydney: Academic Press.

Richard, C. M., Wright, R. D., & Ward, L. W. (2003). Goal-driven modulation of stimulus-driven attentional capture in multiple-cue displays. *Perception & Psychophysics, 65 (6),* 939-955.

Roe R. A., Zijlstra F. R. H. and Leonova A. B. (1996) *Interruptions in mental information work: Effects on work activity and psychological well-being,* WORC Report 96.06.003/4, Tilburg: Tilburg University Press

Rogers, R. D., & Monsell, S. (1995). The cost of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General, 124,* 207-231.

Rubenstein, J. S., Meyer, D. E., & Evans, J. E. (2001). Executive control of cognitive

processes in task switching. *Journal of Experimental Psychology: Human Perception and Performance, 27,* 763-797.

Ruiz, D. (1987). Learning and problem solving: What is learned while solving the Tower of Hanoi (Doctoral dissertation, Stanford University, 1987). *Dissertation Abstracts International, 42,* 3438B.

Ryle, G. (1949). *The Concept of the Mind.* London: Hutchinson.

Sanford, A. J., & Garrod, S. C. (1981). *Understanding Written Language.* New York: Wiley Press.

Schneider, W., & Detweiler, M. (1988). The role of practice in dual-task performance: Toward workload modeling in a connectionist/control architecture. *Human Factors, 30,* 539–566.

Schneider, W., & Schiffrin, R. (1977). Controlled and automatic human information processing I: Detection, search and attention. *Psychological Review, 84,* 1–66.

Siegler, R. S. (1987). The perils of averaging data over strategies: An example from children's addition. *Journal of Experimental Psychology: General, 116,* 250-264.

Siegler, R. S. (1988). Strategic choice procedures and the development of multiplication skill. *Journal of Experimental Psychology: General, 117,* 258-*275.*

Shaffer, L. H. (1975). Multiple attention in continuous verbal tasks. In P. M. A. Rabbitt and S. Dornic (Eds.), *Attention and performance V.* Academic Press, London,

Shiffrin, R., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review, 84,* 127-190.

Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society of London B, 298,* 199-209.

Simon, H. A. (1975). The functional equivalence of problem solving skills. *Cognitive Psychology, 7,* 268-288.

Sohn, M. H., & Carlson, R. (2000). Effects of repetition and foreknowledge in task-set reconfiguration. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26,* 1445-1460.

Smith, R. E. (2003). The cost of remembering to remember in event-based

prospective memory: Investigating the capacity demands of delayed intention performance. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 29,* 347-361.

Smith, S. L., & Mosier, J. N. (1986). *Guidelines for designing user interface software* (Report ESD–TR–86–278). Bedford, MA: MITRE.

Smyth, M. M., & Scholey, K. A. (1994). Interference in immediate spatial memory. *Memory and Cognition, 22,* 1-13.

Spector, A., & Biederman, I. (1976). Mental set and mental shift revisited. *American Journal of Psychology, 89,* 669-679.

Speier, C., Valacich, J. S., & Vessey, I. (1999). The influence of task interruption on individual decision making: An information overload perspective. *Decision Sciences, 30,* 337-360.

Speier C., Vessey I. and Valacich J. S. (2003) The effects of interruptions, task complexity, and information presentation on computer-supported decision-making performance, *Decision Sciences*, 34 (4), 771-797.

Stanton, N. A., Booth, R. T., & Stammers, R. B. (1992). Alarms in human supervisory control: A human factors perspective. *International Journal of Computer Integrated Manufacturing, 5(2),* 81–93.

Storch, N. A. (1992). *Does the user interface make interruptions disruptive? A study of interface style and form of interruption (Report UCRL-JC-108993).* Springfield: Lawrence Livermore National Laboratory.

Stringer, M., Eldridge, M., & Lamming, M. (2000). Towards a deeper understanding of task interruption. *Proceedings of CHI 2000 Workshop # 14 on situated interaction in ubiquitous computing,* 26-27.

Sudevan, P., & Taylor, D. A. (1987). The cuing and priming of cognitive operations. *Journal of Experimental Psychology: Human Perception and Performance, 13*, 89–103.

Sutcliffe, A.G. (1995). *Human-Computer Interface Design.* London: Macmillan Press Ltd.

Tétard, F. (2000). Fragmentation of working time and SMARTER IS-solutions. *Proceedings of the 33rd Annual Hawaii International Conference on System Sciences, Vol. 7* (p. 7047). Los Alamitos: IEEE Computer Society.

Theeuwes, J. (1995). Abrupt luminance change pops out; abrupt colour change does

not. *Perception & Psychophysics, 57,* 637-644.

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.

Treisman, A., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12,* 97-136.

Treisman, A., & Gormican, S. (1988). Search asymmetry: A diagnostic for preattentive processing of separable features. *Journal of Experimental Psychology: General, 114,* 285-310.

Tsal, Y., & Lavie, N. (1988). Attending to color and shape: The special role of location in selective visual processing. *Perception & Psychophysics, 44,* 15-21.

Tsal, Y., & Lavie, N. (1993). Location dominance in attending to color and shape. *Journal of Experimental Psychology: Human Perception and Performance, 19,* 131-139.

Turner, J. W., & Huntley Jr., M. S. (1991). *The use and design of flightcrew checklists and manuals.* FAOE2-A0180. Washington, D.C.: U.S. Department of Transportation.

Van Bergen, A. (1968). *Task Interruption.* Amsterdam: North-Holland.

Van Dantzich, M., Robbins, D., Horvitz, E., & Czerwinski, M. (2002). Scope: Providing Awareness of Multiple notifications at a glance. *Proceedings of AVI 2002, ACM Conference on Advanced Visual Interfaces,* (pp. 22-24). ACM Press.

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

VanLehn, K., & Ball, W. (1991). Goal reconstruction: How Teton blends situated action and planned action. In K. VanLehn (Ed.), *Architectures for intelligence* (pp. 147-189). Hillsdale, NJ: Erlbaum.

VanLehn, K., Ball, W., & Kowalski, B. (1989). Non-LIFO execution of cognitive procedures. *Cognitive Science,* 13, 415-465.

Vera, A., Howes, A., McCurdy, M., & Lewis, R. L. (2004). A constraint satisfaction approach to predicting skilled interactive performance. In E. Dykstra-Erickson

and M. Tscheligi (Eds.), *Proceedings of ACM HCI 2004, Conference on Human Factors in Computing Systems* (pp. 121-128). Vienna.

Vernon, W. 1990. Basic management for engineers: Part 1 personal management. *Engineering Management Journal, 8,* 107-111.

Vessey, I. (1991). Cognitive fit: A theory-based analysis of the graphs versus tables literature. *Decision Sciences, 22,* 219-240.

Von Wright, J. M. (1968). Selection in visual immediate memory. *Quarterly Journal of Experimental Psychology, 20,* 62-68.

Wang, Q., Cavanagh, P., & Green, M. (1994). Familiarity and pop-out in visual search. *Perception & Psychophysics, 44,* 81-93.

Ward, G., & Allport, A. (1997). Planning and problem-solving using the five-disc Tower of London task. *Quarterly Journal of Experimental Psychology, 50A,* 49-78.

Ware, C., Bonner, J., Knight, W., & Cater, R. (1992). Moving icons as a human interrupt. *International Journal of Human-Computer Interaction, 4,* 341-348.

Wickens, C. D. (1984). Processing resources in attention. In R. Parasuraman and D. R. Davies (Eds.), *Varieties of attention.* New York: Academic Press.

Wickens, C. D. (1992). *Engineering psychology and human performance.* New York: Harper Collins.

Williams, C. L. (1995). Potential conflicts of time sharing the flight management system control display unit with data link communications. *Proceedings of the Eighth International Symposium on Aviation Psychology,* 341–347.

Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review, 9 (4),* 625 – 636.

Wood, R. E. (1986). Task complexity: Definition of the construct, *Organizational* Behavior and Human Decision Processes, 37, 60-82.

Wolfe, J. M. (1994). Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review, 1 (2),* 202-238.

Wolfe, J. M. (2004). What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience, 5,* 1-7.

Wolfe, J.M., Cave, K. R., & Franzel, S. L. (1989). Guided Search: an alternative to the Feature Integration Model of visual search. *Journal of Experimental Psychology: Human Perception and Performance, 15,* 419-443.

Wolfe, J. M., Yu, K. P., Stewart, M. I., Shorter, A. D., Friedman-Hill, S. R., & Cave, K. R. (1990). Limitations on the parallel guidance of visual search: Color and Orientation conjunctions. *Journal of Experimental Psychology: Human Perception & Performance, 16,* 879-892.

Zeigarnik, B. (1927). Das Behalten erleigter und unerledigter Handlungen. *Psychologische Forschung, 9,* 1-85.

Zhang, J. (1997). The nature of external representations in problem solving. *Cognitive Science, 21 (2),* 179-217.

Zijlstra, F. R. H., Roe, R. A., Leonora, A. B., & Krediet, I. (1999). Temporal factors in mental work: Effects of interrupted activities. *Journal of Occupational and Organisational Psychology, 72,* 163-185.

Appendix A1 – Tower of Hanoi Interface and Example Disc-Configurations

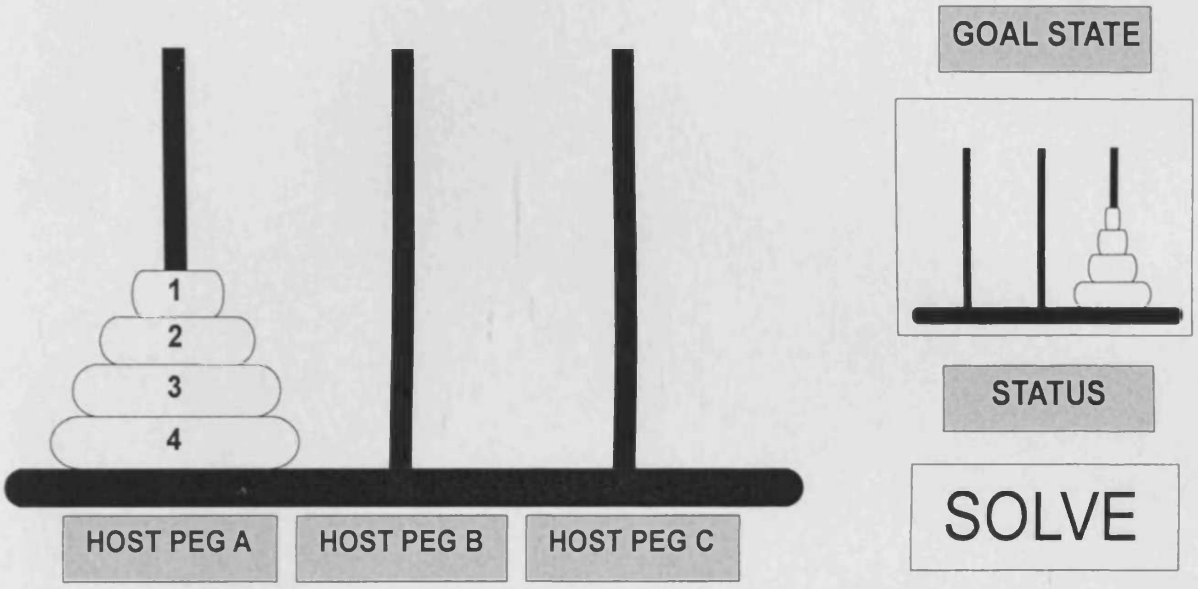

Figure A1. An example of the main ToH interface at the beginning of a fifteen-move problem. Discs are numbered 1 – 4 for descriptive purposes only and were coloured in all experiments other than Experiment 7 so that: Disc 1 was green; disc 2 was blue; disc three was yellow; and disc 4 was red.

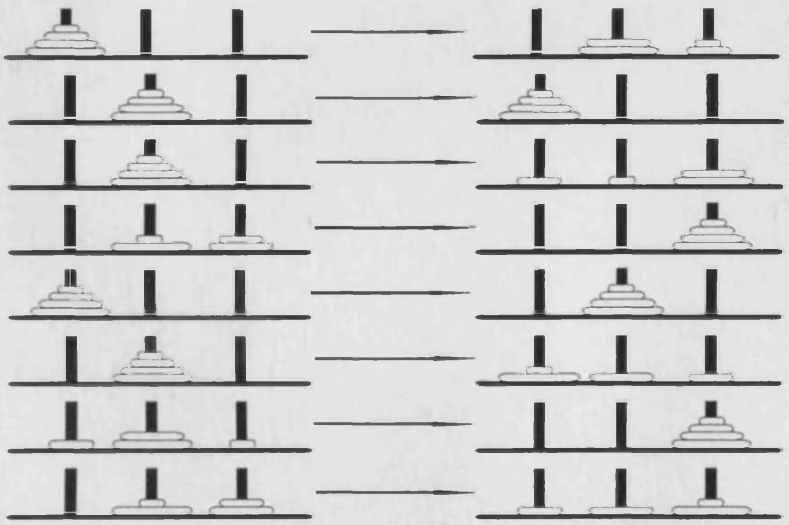

Figure A2. Examples of four-disc ToH problems that are 15 moves apart

Appendix A2

The Effects of Forming an Initial Plan on resuming promptly suspended goals

---

Introduction

A suspended goal is expected to decay in activation since it was last processed, and associating it with a salient retrieval cue may the only way to protect against forgetting (Altmann & Trafton, 2002). Associative priming should be difficult in the case of unsignalled interruption, especially when the goal intruded upon forms part of a complex move sequence for which there are a number of potential priming cues. Might other task phases allow preparation of a to-be-suspended goal to minimise the disruptive impact of unsignalled interruption on task performance?

The current experiment assessed whether planning during the initial phase of a ToH problem impacts upon the speed of resuming a promptly suspended goal. Perhaps a planned goal will reside in a heightened level of activation meaning that its future retrieval will be certain (e.g., Goschke & Kuhl, 1993). Alternatively, planning and consequently activating a goal during an initial phase of the problem might reduce the amount of encoding required prior to suspension. Given that a goal will have been considered at least once before, more processing resources might be available to associate the goal with a salient priming cue prior to interruption, increasing the chances of retrieval.

Evidence on the effectiveness of planning on improving perfromance in uninterrupted well-structured problem solving tasks is however largely mixed. Classically, it was shown that humans formulate partial plans, execute initial moves, and then fall back on a less memory demanding strategy of opportunistically planning further moves as they become apparent (Hayes-Roth & Hayes-Roth, 1979). Recent work has largely focused on planning in the Tower of London (ToL) task, a variant of the ToH, in which participants have to move a number of same-sized coloured discs

(or balls) from an initial pattern to match a goal pattern (e.g., Shallice, 1982). Typically, participants are instructed to initially plan a full series of moves before execution; the quality of this plan arguably leads to better performance during the execution phase (e.g., Shallice, 1982). However, the role of planning on perfromance efficiency in the ToL task has recently undergone careful scrutiny (e.g., Gilhooly et al., 1999; Phillips et al., 2001). Although participants are able to plan up to seven moves ahead (or two subgoals), there is often a discrepancy between planned and executed moves (e.g., Gilhooly et al., 2001) with planning shown to be ineffective at reducing performance time and move-errors (e.g., Phillips et al., 2001). By contrast, recent work on planning in the ToH task has proved more promising in showing improvements to execution perfromance. Encouraging initial planning and discouraging online or concurrent planning (by restricting individual move time to 2.5 s) has been shown to reduce move errors in 3-disc and 4-disc ToH problems but not in more complex 5-disc problems. This is believed to reflect a trade-off between memory-intensive planning strategies against perceived problem complexity, such that participants plan-ahead more moves in the boundary of working-memory limitations (Davies, 2003).

In the current experiment, planning was divided into two temporal phases; one classified as the time spent *mentally planning* ToH goals and the other classified as the time spent *posting* ToH moves. Mental planning was taken as the time spent encoding goals before an executable solution-sequence was derived (see Anderson & Douglass, 2001, for a discussion of the processes involved in planning ToH moves). Posting was the time spent indicating moves that would be executed if ToH discs could be moved. Mental planning was assumed to always precede goal posting as a consequence of observed planning behaviour in the ToH task where goals and subgoals are usually considered before they are executed (e.g., Anzai & Simon, 1979; Davies, 2000, 2003; VanLehn, 1991). Nevertheless, to tighten the distinction between the two phases, participants were instructed to begin posting moves only after they had planned mentally as many moves as they believed would be sufficient to advance the solution-sequence (as should be the case in other experiments presented in the thesis). Specific instructions to encourage such behaviour were as follows:

"It is your task to solve each main ToH problem in the minimum number of moves

possible. During an initial planning phase you will be allocated 15-seconds to (1) assess the difference between the start-state disc configuration and the goal-state disc configuration (2) mentally consider a series of moves that you think will advance the solution-sequence toward the goal state, and (3) indicate as many of these moves as possible before the 15-seconds expires. Please try to start indicating moves (i.e., phase 3) before the 15-seconds expires."

Predictions for Experiment A are as follows. Primarily, it is hypothesised that move latency data will replicate that found in Experiment 3, with similarities showing that participants plan an initial sequence of moves even without encouragement to do so. By contrast, longer inter-move latencies should correlate with insufficient planning and when participants stray from original plans when executing moves. Straying from original plans is likely to encourage concurrent planning, with such behaviour coming at a cost of longer inter-move latencies. Predictions for interrupted move latencies are by and large *ad hoc*. If maintaining a representation of a suspended goal is effected by whether that goal was encoded during the planning phase, resumption latency is expected to be shorter. Such effects would provide additional support that participants are able to retrieve suspended goals from memory, even following unsignalled interruption, and would suggest that some of the preparatory encoding is established at task points other than immediately prior to interruption initiation (in contrast to Altmann & Trafton, 2002).

Method

Participants

Thirty-two undergraduates from the School of Psychology at Cardiff University participated in fulfilment of a course requirement. All were naive to the ToH task. Four participants were removed (leaving 28) for failing to meet the inclusion criteria specified in Section 2.4.2.1.

Apparatus and materials

The same program used in Experiment 3 was modified to incorporate a 15-seconds planning and posting period prior to the solve prompt in all primary task ToH

problems. This period was signalled in the same way as in Experiment 3 with the command 'plan' appearing for 15-seconds in the primary task visual display. Host pegs and destination pegs could be selected during the planning period, although all other primary task functions were disabled, meaning that discs could not move. For the duration of the planning and posting period, the program recorded when goal posting was initiated, and each time a host peg and destination peg was selected following this period (i.e., when a to-be-performed move was *posted*). When the 15-seconds expired, the command 'solve' appeared after which time all primary task functions were re-enabled.

Design

Move latency was measured in the same way as in Experiment 3. A 15-second planning and posting protocol was introduced at the beginning of primary task ToH problems. As in Experiment 3, participants were expected to mentally plan as many moves as they could, but were also required to 'act-out' planned moves through posting move operations. Participants posted goals to an invisible goal-stack (that could not be accessed) by selecting a disc and a destination peg using the mouse cursor. Posted moves where larger discs were moved on top of smaller discs were given a value of zero: Unless undone, any moves planned afterwards were not considered in further analyses. Following the plan and post phase, participants were free to engage in a different or modified solution-sequence to that planned. Planning behaviour was measured by taking the total number of moves posted irrespective of being executed, and the total number of moves subsequently executed in the order in which they were posted (referred to as *posted-executed* hereafter).

Procedure

The same as in Experiment 3 except that participants were given the instructions in the design section and were shown how to post goals.

Results and discussion

Planning behaviour

All 28 participants spent less than 15-seconds planning moves and thus

measures of planning and posting time could be taken for the complete sample. Figure A3 displays mean time spent mentally planning the solution-sequence and the mean time spent posting moves during the 15-second plan-post period. Clearly there was little effect of condition on the time spent mentally planning moves, with participants spending between six and eight seconds before posting moves.

Figure A3. Mean time spent planning (±SE) and posting moves during the 15-second plan and post period preceding the execution phase of ToH problems in Experiment A. Grey bars represent planning latency and white bars represent posting latency.

A repeated measures 2 (planning vs. posting) x 2 (trial type) x 2 (point of interruption) x 2 (secondary task type) ANOVA was conducted on mental planning and posting data. This revealed that more time was spent posting moves than mentally planning moves, although this difference was only marginally significant, $F(1, 27) = 4.65$, $MSE = 11.58$, $p < .05$. There were no other significant main effects and none of the variables significantly interacted. Even with a 15-second planning period, participants spend less than 10-seconds mentally planning moves; well within the fixed planning periods incorporated at the beginning of each primary task ToH problem in Experiments 2 – 9.

Table A1 illustrates mean number of legal moves posted in interrupted and

control ToH problems across each condition. Mean numbers of moves posted are very similar for each condition. Approximately four moves were posted during the initial planning phase of interruption and control problems; more than the number required to reach the point of interruption in within-subtask conditions (three) and almost as many required for end-subtask conditions (four).

Table A1

Mean number (±SE) of posted moves in interrupted and control ToH problems

Interruption				Control			
Wth-Sim	Wth-Dsim	End-Sim	End-Dsim	Wth-Sim	Wth-Dsim	End-Sim	End-Dsim
4.11 (0.3)	4.11 (0.3)	4 (0.2)	3.84 (.22)	3.9 (.29)	3.64 (.21)	4 (.27)	3.93 (0.3)

A Friedman test indicated a non-significant difference between any of the conditions ($p = .93$) suggesting that the number of moves posted were very similar irrespective of condition. With a 15-seconds planning opportunity, participants seem to be able to post more than the required number of moves required to reach the intended point of interruption in within-subtask conditions, although are marginally short of four moves in end-subtask conditions (M = 3.87 moves).

Table A2 displays the mean number of moves posted and subsequently executed in interrupted and control ToH problems across each condition. Moves considered were only those posted-executed in a serial fashion: The first move posted was executed first, the second move posted was executed second, and so on. Numbers of posted-executed moves are less than the total numbers of moves posted, suggesting that participants occasionally strayed from planned solution-sequences. On average, 3.24 moves were posted and subsequently executed; more than the number required to reach the point of interruption in within-subtask conditions (i.e., three moves) but not quite as many required for end-subtask conditions (i.e., four moves).

A Friedman test indicated a non-significant difference across trial type ($p = .86$) suggesting that the number of moves subsequently executed because of posting were very similar irrespective of condition. With a 15-second opportunity at the beginning of ToH problems, participants are able to plan, indicate and subsequently execute almost as many moves required to reach the intended point of interruption in each condition.

Planning data from the current experiment lends well to original hypotheses and supporting literature in showing that participants can plan and subsequently execute the initial steps of a ToH problem (e.g., Anderson & Douglass, 2001; Davies, 2003). Participants usually encoded the within-subtask goal during the plan and post period, and the same move was subsequently executed, even following interruption. Although they found it difficult to post more than four moves, I suspect that this was a confound of the overall time constraint on planning and posting, such that with more time more moves would have been posted. The reason for this suspicion is that the majority of participants were shown to still be posting moves right up until the solve command appeared in the display.

Table A2

Mean number (±SE) of posted-executed moves in interrupted and control ToH problems

Interruption				Control			
Wth-Sim	Wth-Dsim	End-Sim	End-Dsim	Wth-Sim	Wth-Dsim	End-Sim	End-Dsim
3.2 (.37)	3.4 (.23)	3.4 (.42)	3.5 (.28)	3 (.37)	2.9 (.26)	3.3 (.35)	3.2 (0.4)

Move latency

Figure A4 illustrates latency data for the first move following the intended point of interruption in interrupted and control ToH problems. Control move latencies are much shorter than resumption latencies in interrupted problems and are very similar across conditions. Although resumption latencies are very similar across three of the interruption conditions, the within-similar interruption condition exhibits markedly longer resumption latency.

A repeated measures 2 x 2 x 2 ANOVA revealed significant main effects of trial type, $F(1, 27) = 253.46$, $MSE = 3.82$, $p < .001$, intended point of interruption, $F(1, 27) = 10.29$, $MSE = 1.13$, $p < .005$, and of secondary task type, $F(1, 27) = 9.39$, $MSE = 2.54$, $p < .010$. Resumption latencies were longer than move latencies in comparable control problems; longest when interruption fell within a subtask and the secondary task was similar. There was a significant two-way interaction between trial type and intended point of interruption, $F(1, 27) = 8.09$, $MSE = 1.69$, $p < .010$. There were also marginally significant interactions between, trial type and secondary task type, $F(1, 27) = 4.24$. $MSE = 1.44$, $p < .05$, and intended point of interruption

and secondary task type, F (1, 27) = 4.24, MSE = 1.73, $p < .05$. All three variables interacted significantly, F (1, 27) = 4.28, MSE = 2.44, $p < .05$. Bonferroni *post-hoc* analyses revealed that resumption latency was longer if interrupted within a subtask by a secondary ToH problem than if the same secondary task was performed at the end of a subtask, F (1, 27) = 12.21, $p < .001$. There were no other significant simple main effects between any other conditions. The cause of the significant three-way interaction was the significantly longer resumption latency in the within-similar condition compared to the end-similar condition, an effect that was not exhibited in any other comparison.

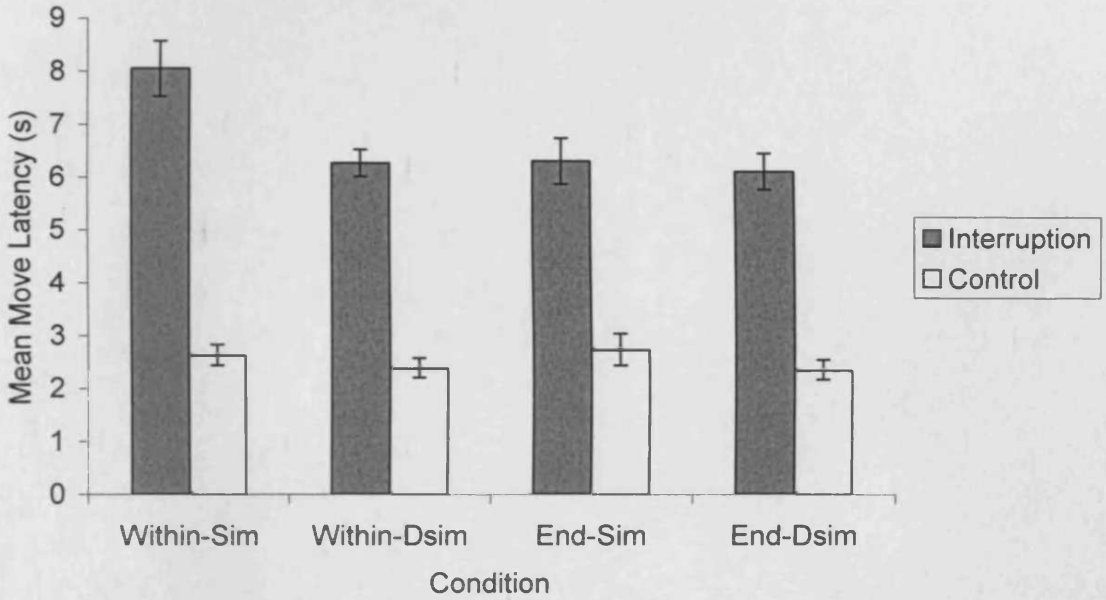

Figure A4. Mean move execution latencies (±SE) for interrupted and control trials in Experiment A

Point of interruption interacted with secondary task type, reinforcing the idea that when processing demands are high at the point of interruption (within-subtask), performing a similar secondary task seriously impairs resumption performance. Enforced planning might have allowed extra encoding that could have abolished the difference in resumption latency for the dissimilar secondary task. Without an interruption lag and with an initial planning phase, goals do seem to be encoded prior to interruption, although representational strength is more susceptible to the activation competition generated from processing similar goals.

On a similar vein, associative activation – the limited supply of activation generated by the current context – might explain the within-subtask interruption-similarity effect. Specifically, a similar secondary task might utilise associative activation that could otherwise be used to support the retrieval of the suspended goal; a resource utilised to a lesser extent when performing a digit-recall task. Simply, in the within-dissimilar condition, there may be more available associative activation to serve the reactivation of a goal that would otherwise be difficult to retrieve (e.g., Altmann & Trafton, 2002).

Finally, resumption latency was almost two seconds shorter for interrupted within-subtask conditions in the current experiment compared the same conditions in Experiment 3. This may have been an effect of the manipulation of enforced posting of planned moves in the current experiment such that representational strength of a goal suspended within-subtask was more resilient to forgetting.

Move latency and planning behaviour

Table A3 displays Pearson's r two-tailed bivariate correlations (with significance values) between number of posted-executed moves and move latency in interruption and control problems. Of particular interest are the entries running from the top left of the table to those diagonally downwards at the bottom right of Table A3. Apart from the within-dissimilar control condition, there were strong significant

Table A3

Pearson r bivariate correlation coefficients and significance values (two-tailed) between number of planned-executed moves and resumption latencies

Move Latency	Number of planned moves							
	Wth-Sim I	Wth-Sim C	Wth-Dsim I	Wth-Dsim C	End-Sim I	End-Sim C	End-Dsim I	End-Dsim C
Wth-Sim I	-.414*	-.239	.258	-.124	-.493*	-.629**	-.335	-.562*
	.029	*.220*	*.185*	*.530*	*.008*	*.000*	*.081*	*.002*
Wth-Sim C	-.221	-.577**	-.168	.178	-.318	.492*	-.208	-.222
	.527	*.000*	*.392*	*.365*	*.099*	*.008*	*.288*	*.257*
Wth-Dsim I	-.262	-.390*	.006	.014	-.340	-.237	-.147	.009
	.178	*.040*	*.976*	*.946*	*.076*	*.225*	*.455*	*.964*

Wth-Dsim C	-.276	-.473*	.249	.091	-.363	-.261	-.259	-.255
	.170	*.011*	*.202*	*.646*	*.057*	*.179*	*.184*	*.190*
End-Sim I	-.405*	-.239	.340	.166	-.361	-.223	-.118	-.408*
	.032	*.221*	*.077*	*.398*	*.059*	*.255*	*.549*	*.031*
End-Sim C	-.222	-.324	-.151	-.038	.306	-656**	.420*	-.425*
	.257	*.093*	*.444*	*.848*	*.113*	*.000*	*.026*	*.025*
End-Dsim I	-.368	-.506*	-.035	.022	-.320	-.2.67	-.222	-.245
	.054	*.006*	*.858*	*.911*	*.097*	*.170*	*.257*	*.209*
End-Dsim C	-.108	-.268	.371	.104	-.306	-394*	-.163	-.601**
	.584	*.168*	*.052*	*.559*	*.113*	*.038*	*.406*	*.000*

Note. I refers to interruption and C to control. Correlations marked with an asterisk (*) were significant at $p < .05$, and those with a double asterisk (**) were significant at $p < .001$.

negative relationships between number of posted-executed moves and shorter move latencies. When interruption involved a ToH secondary task, there was a strong significant negative relationship between number of posted-executed moves and resumption latencies in the within-subtask condition, a relationship only marginally non-significant for the end-subtask condition ($p = .06$).

So why does executing more planned moves lead to faster resumption when interrupted by a similar secondary task but not when interruption entails the completion of a dissimilar secondary task? As was the case for move latency, such relationships may be explained through ACT-R assumptions regarding the availability and distribution of a limited supply of associated activation (e.g., Lovett & Anderson, 1996). Given that a comparable number of moves are posted and subsequently executed when interrupted by either secondary task type, the base-level activation of a suspended goal should be similar in each case. However, more of a limited supply of associative activation may be demanded by a similar secondary task compared to one that is dissimilar. Perhaps then, having explicitly planned the interrupted move (and possibly co-dependent moves) increases the resilience to retrospective interference caused from processing similar goals in the ToH secondary task. The claim here is not one that assumes participants are better at retrieving a suspended goal because of efficient planning when interrupted by similar secondary tasks. Rather strengthening the representations of goals in the planning phase might provide resilience to retroactive interference caused by similar secondary task goals, so much so, that

resumption latency is reduced. In contrast, dissimilar secondary tasks might not demand as much associative activation as similar secondary tasks, causing less reliance on a memory-intensive retrieval strategy and more reliance on a strategy that uses contextual priming cues. In dissimilar secondary task conditions, an activation boost reclaims a suspended goal from its decaying state: A process that is more efficient than trying to remember the goal based upon the strength of its mnemonic representation (e.g., Altmann & Trafton, 2002).

Discussion

First I will summarise the novel findings from Experiment A. Second, I will discuss the results of Experiment A in relation to findings from Experiment 3. Finally, I will consider how the current results impact upon a core prediction of the goal-activation model: Specifically, if a suspended goal is to be retrieved from memory it must be associated with a priming cue immediately prior to interruption, and this process should take between one and two seconds.

Participants are able to plan between four and five moves ahead in ToH problems, although they do not always adhere to these plans when executing moves (in support of the work on the ToL task by Phillips and colleagues). However, planning a sequence of moves at the beginning of a ToH problem and subsequently executing those moves in the same order, leads to shorter move latencies, as evidenced in most control problems, and some interrupted problems.

Planning and subsequently executing a series of moves was related with shorter resumption latencies, but only when interrupted by a secondary task that required similar processing resources to the primary task. This provides evidence that participants are likely to be encoding up to at least the fourth move during the initial planning period, and such encoding protects against the retroactive interference generated by secondary task goals. In the case of the end-subtask condition, a salient priming cue is available before the primary task is interrupted and encoding at an initial planning phase might allow a stronger associative connection; one that is resilient to retroactive interference. The same rationale applies for the within-subtask condition, although the retroactive interference generated by the secondary task seems to impede upon maintaining the associative connection between the suspended goal and its disc cue. If this connection is strong, as evidenced through having planned the

initial steps to the solution-sequence, the impact of disruptive impact of retroactive interference is nevertheless reduced.

The pattern of results from Experiment 3 are partially replicated by Experiment A; the main similarity being the inflated resumption latency in the within-similar condition, and the main difference being the shorter resumption latencies in both of the within-subtask conditions in Experiment A. Mentally planning and indicating a sequence of moves seems to abolish the disruptive impact of processing complexity at the point of interruption (as found in Experiment 3), as long as the secondary task is dissimilar to the primary task. Perhaps, in the within-dissimilar condition participants are able to prepare a suspended goal for retrieval using associative priming, a process rendered less effective in the within-similar condition. Even with experience in the task domain, maintaining the activation of the suspended goal, perhaps the integrity of the associative link between the goal and the priming cue, is impinged upon by retroactive interference generated by secondary task goals when interruption falls within a complex subtask. The introduction of the posting protocol seemed to impact upon how quickly suspended goals were resumed; perhaps a small change in task constraints affected processing strategy such that more memory-intensive processing was deployed (e.g., Fu & Gray, 2004).

The current experiment provides some evidence that seems problematic for the priming constraint in the goal-activation model: Specifically, the assumption that a goal must be encoded immediately before it is suspended if it is to be retrieved in the future. The findings of the current experiment show that strengthening goals during an initial planning phase can reduce the time taken to resume an interrupted goal. Even in the case of unsignalled interruption, a goal might be strategically and opportunistically prepared for interruption at task points other than immediately before interruption. I am not stating that encoding a goal during an initial planning phase leads to equal or better perfromance than would be expected if such encoding were to occur during an interruption lag. Rather I am suggesting that encoding, particularly strengthening (or rehearsing) a to-be-suspended goal can occur at different task points, a finding that warrants future research to assess ways of best synchronising unsignalled interruption into complex tasks to minimise subsequent perfromance deficits.

Appendix B

Disc Configurations to Satisfy Interruption and Secondary task Descriptive Statistics for Empirical Series 3

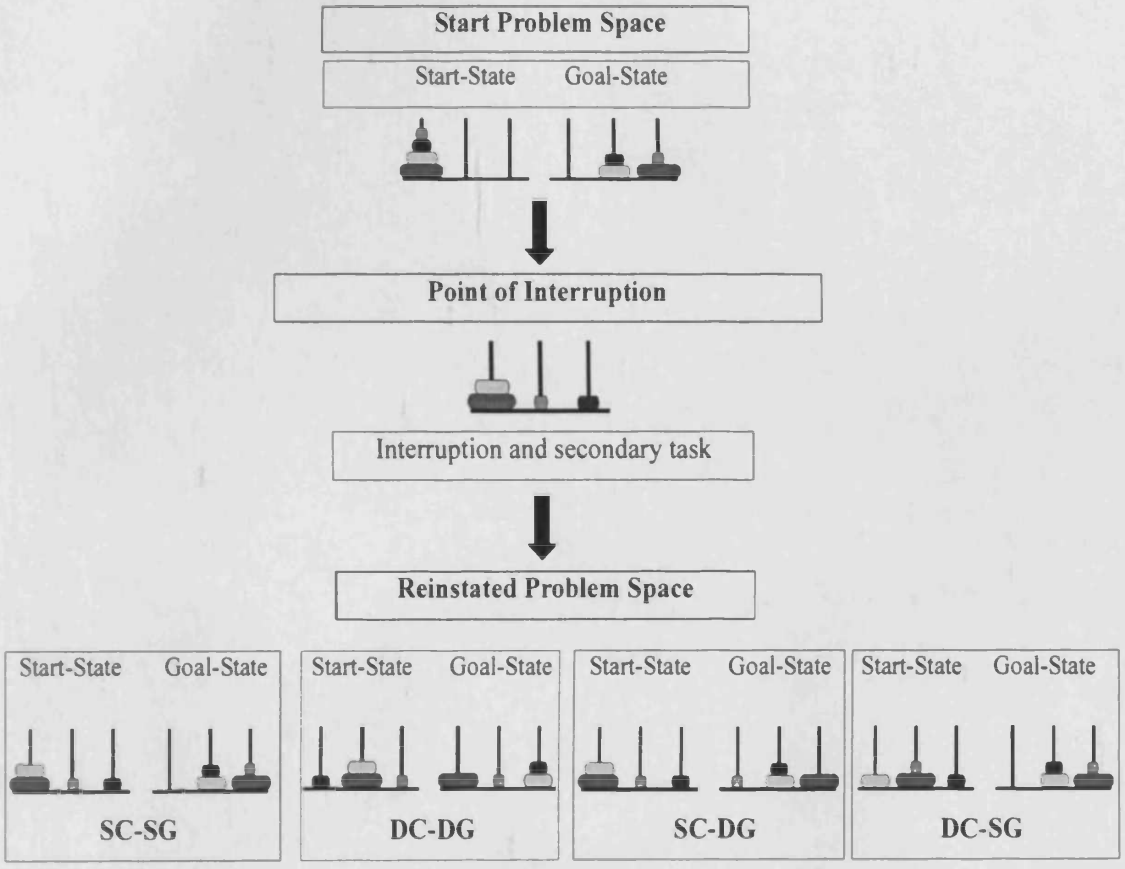

Figure B1. An example of an interrupted ToH problem showing all possible reinstated current-state disc configuration conditions and goal-state disc configuration conditions in Experiment 4

Table B1.

Mean completion time (±SE) and digits correctly recalled (±SE) within secondary tasks of Experiments 4, 5 & 7

Experiment	Secondary Task	SC-SG		SC-DG		DC-SG		DC-DG	
	Performance	M	SE	M	SE	M	SE	M	SE
4	Completion Time	15.67	0.66	16.60	0.77	16.24	0.59	16.02	0.7
	Digits Recalled	5.54	0.3	5.21	0.36	5.08	0.36	5.63	0.23
6	Completion Time	13.09	0.71	13.93	0.77	13.64	0.83	13.64	0.84
	Digits Recalled	4.25	0.35	4.97	0.31	4.22	0.34	4.72	0.34
7	Completion Time	16.86	0.73	16.98	1.03	16.32	0.66	17.07	0.77
	Digits Recalled	4.89	0.24	4.73	0.25	4.75	0.26	4.94	0.28

Note. Time is in seconds

Table B2.

Mean completion time (±SE) and digits correctly recalled (±SE) within secondary tasks of Experiment 5

Secondary Task	SC-SG		SC-SG+2		SC-SG-2		DC-DG		DC-DG+2		DC-DG-2	
Performance	M	SE	M	SE	M	SE	M	SE	M	SE	M	SE
Comp. Time	16.93	1.01	16.56	0.97	16.21	0.8	16.06	0.67	16.67	0.77	17.07	0.85
Digits Recalled	4.15	0.22	4.31	0.22	3.97	0.21	3.92	0.18	4.25	0.2	3.9	0.19

Note. Comp is completion and time is in seconds

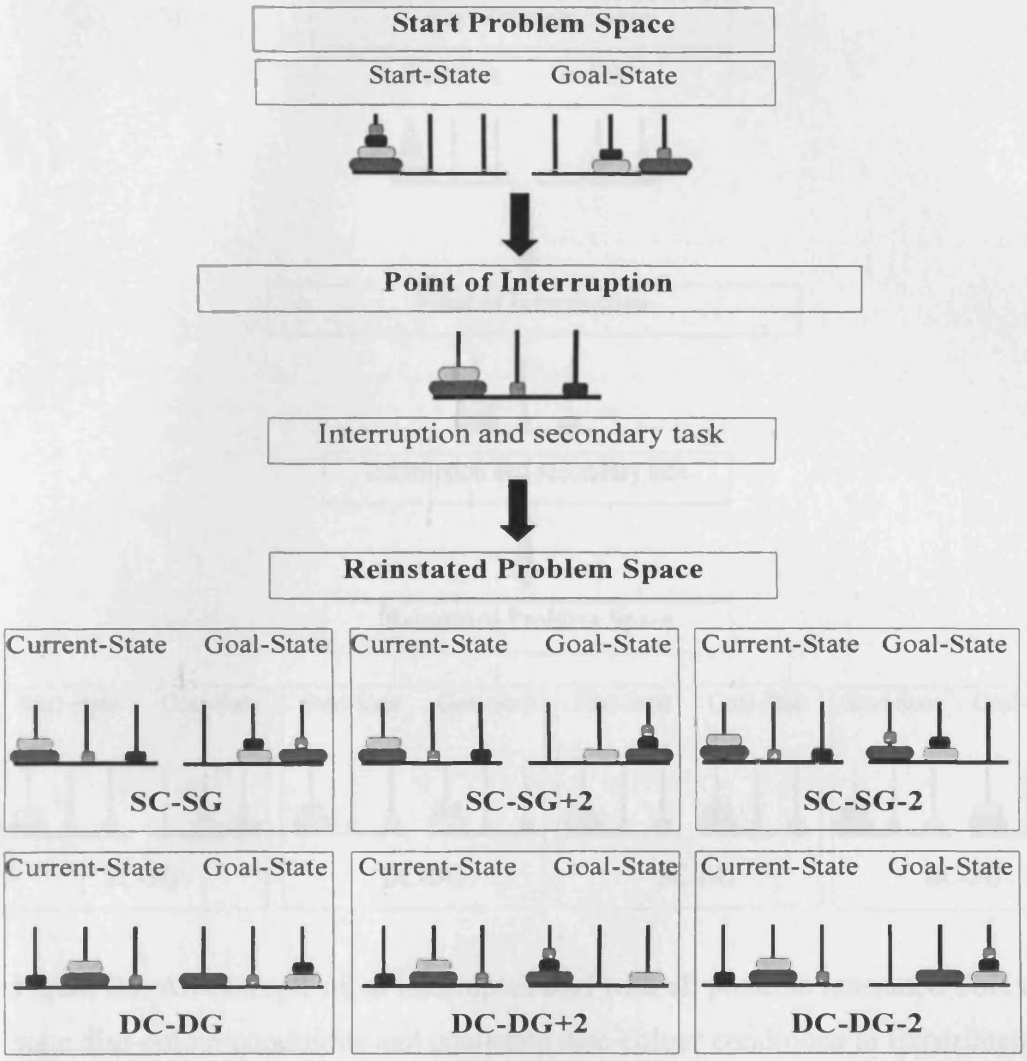

Figure B2. An example of an interrupted trial with the possible reinstated ToH current-state disc configuration conditions and goal-state disc configuration conditions in Experiment 5

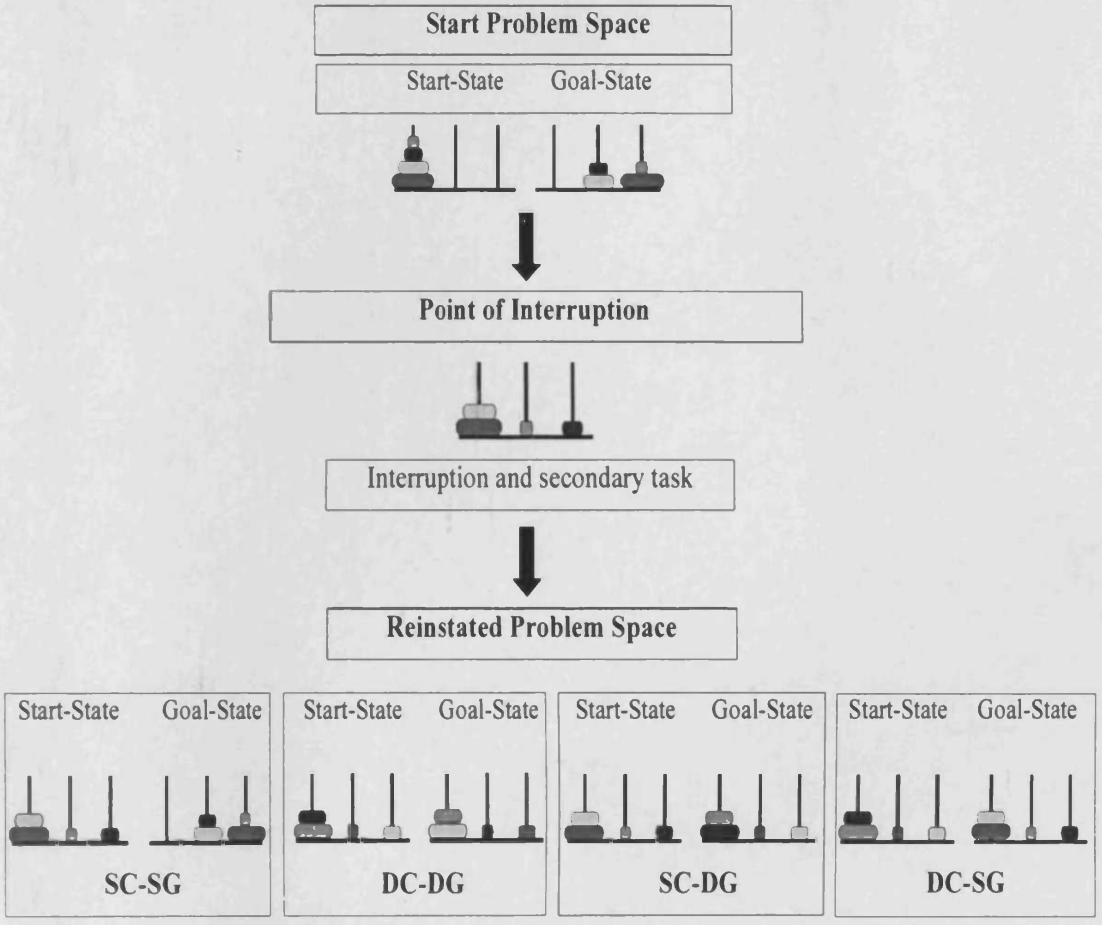

Figure B3. An example of an interrupted trial with all possible reinstated ToH current-state disc-colour conditions and goal-state disc-colour conditions in Experiment 6

Appendix C

Secondary Task Descriptive Statistics for Empirical Series 4

Table C1

Mean Secondary Task Completion Times (±SE) and Mean Number of Moves (±SE) Executed Within Secondary Tasks in Experiment 8a

	Within-Lag		Within-no-Lag		End-Lag		End-no-Lag	
	M	SE	M	SE	M	SE	M	SE
Completion Time	20.65	1.05	21.28	1.39	19.84	0.93	21.28	1.15
Moves	7.86	0.26	8.04	0.31	7.67	0.17	8.03	0.27

Note. Time is in seconds

Table C2

Mean Secondary Task Completion Times (±SE) and Mean Number of Moves (±SE) Executed Within Secondary Tasks in Experiment 8b

	Within-Lag		Within-no-Lag		End-Lag		End-no-Lag	
	M	SE	M	SE	M	SE	M	SE
Completion Time	22.33	1.46	21.38	0.96	21.93	1.29	21.61	1.12
Moves	7.99	0.33	7.88	0.31	8.2	0.32	7.68	0.34

Note. Time is in seconds

Table C3

Mean Secondary Task Completion Times (±SE) and Mean Number of Moves (±SE) Executed Within Secondary Tasks by Group 1 and Group 2 in Experiment 9

Group		Within-Lag		Within-no-Lag		End-Lag		End-no-Lag	
		M	SE	M	SE	M	SE	M	SE
1	Comp. Time	19.05	0.8	18.96	1	18.76	.94	19.46	1.07
	Moves	7.47	.17	7.38	.11	7.49	.18	7.63	.23
2	Comp. Time	20.06	.96	18.73	1.02	19.25	.89	18.18	.89
	Moves	7.6	.26	7.33	.12	7.42	.18	7.67	.19

Note. Comp. is completion and time is in seconds