

Reminders, Alerts and Pop-ups: The Cost of Computer-Initiated Interruptions

Helen M. Hodgetts and Dylan M. Jones

School of Psychology, Cardiff University, Cardiff, CF10 3AT, UK
{hodgettshm, jonesdm}@cardiff.ac.uk

Abstract. Responding to computer-initiated notifications requires a shift in attention that disrupts the flow of work. The degree of cost associated with resuming the original task following interruption may be dependent upon such factors as the transition between tasks (was the worker able to consolidate his/her place in the main task before engaging in the interruption?) as well as the nature of the interrupting task itself (e.g., length or complexity). The current paper reviews a number of studies from our laboratory that investigate the effects of brief interruptions to the execution phase of computer-based 5-disk Tower of London problems. The results are interpreted within the theoretical framework of the goal-activation model [1] and suggestions are made for practical applications that may help to minimize the disruption caused.

Keywords: interruption, Tower of London, goals, activation, memory.

1 Introduction

Emails, instant messages, office assistants, ‘save’ reminders, printing notifications, internet pop-ups and software update alerts, are all distracting for a few moments before we must refocus back on the task in hand. Even such brief unscheduled breaks in task can incur a cost to ongoing performance, not only in terms of the time spent dealing with the actual interruption, but also in terms of the time needed to retrieve previous task goals and resume the interrupted activity. Most existing research has examined the effect of interruptions of several minutes in duration such as telephone calls [2] or walk-in visitors [3], but in the current paper we examine the extent to which very simple computer-based interruptions can be disruptive to an ongoing task. From a theoretically-informed position we examine those factors that may increase or decrease the disruption caused and highlight practical applications of this work.

1.1 The Goal-Activation Model

As a theoretical foundation for our study of interruptions we use the goal-activation model (G-AM) [1], a computational model that derives from the principles of ACT-R [4]. It makes predictions about the processes involved in the suspension and resumption of goals, and as such can be applied to task interruption. The model proposes a number of active goals in memory, the most active of which is the one to govern behavior. Activation is determined by two components: history of use (how

frequently and how recently a goal has been retrieved) and its relevance to the current context (the influence of environmental cues). For a new or interrupting goal to become more active than the others and to exceed the interference threshold, it must be repeatedly sampled or *strengthened* within a short period of time in order to build up base-level activation. Once selected, the rate of sampling decreases and the goal gradually decays to a level below that of other newer goals. Sometimes – in the case of interruption for example – goals are temporarily suspended and must be resumed later. G-AM makes this possible through a process of *priming*. The base-level activation of a suspended goal will have decayed through lack of use, but retroactive interference can be overcome if a goal is deemed relevant by the current context. A goal's associative activation is boosted by environmental cues: Effective cues are formed by co-occurrence and so must be present both at the time of goal suspension and when the goal is to be retrieved. In terms of an interruption, Altmann and Trafton propose that these cues can be encoded during the 'interruption lag', the time between the interruption alert (e.g., the telephone starting to ring) and the actual interruption (e.g., engaging in the telephone conversation). G-AM is one of the few theoretical models to have been explicitly applied to the issue of task interruption and it provides a useful basis for the exploration and interpretation of interruption effects.

1.2 The Tower of London Task

Our experiments used a computer-based version of the 5-disk ToL problem [5] as a primary task because it provides a controlled task environment and allows for the assessment of performance at a sufficiently fine-grained level of detail. Participants are presented with a starting array of different colored, equal-sized disks mounted on three pegs (Figure 1). The aim is to achieve a given goal configuration by moving the disks one at a time, from peg-to-peg. The task involves the formulation, retention and execution of a planned series of actions – processes that are not dissimilar to many everyday computer-based activities (e.g., first formulating and then typing sentences for a report). In the current work it is the execution phase that is interrupted, and we investigate what factors affect retrieval of a planned sequence of action when execution of this plan is unexpectedly broken.

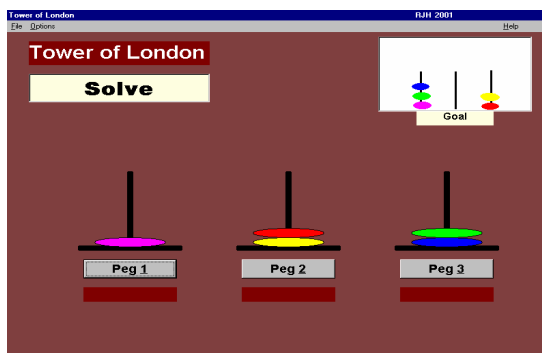


Fig. 1. Screen display during the Tower of London task

Participants were to complete 25 ToL trials, of which either 6 or 8 trials were interrupted (depending on the particular experiment). These interrupted trials were all equivalent in terms of complexity, and required 6 moves to solution. Each interrupted trial was always matched to a control trial which required exactly the same solution path but the colors of the disks were changed. Disks were moved by clicking on the buttons below the pegs, and a pop-up box indicated when the goal state had been achieved.

Interruption always occurred in the middle of a trial, after the participant had made their third move. The main ToL display would then be replaced by the interrupting task which in most cases was a mood checklist. A list of six statements along a mood continuum (e.g., “extremely happy,” to “extremely sad”) were presented in the centre of the screen, and participants were asked to select the one that best applied to them by clicking on the statement with the mouse. This task was irrelevant to the main ToL but provided a plausible means with which to interrupt the primary task in the laboratory context. Interruption was always brief and took around 5 s to complete before the participant then continued with the main task at the exact point at which it had been left. Using this general methodology, we have conducted a number of experiments to investigate the effect of brief on-screen interruptions to ongoing task performance.

2 Interruption Length

G-AM incorporates ACT-R’s base-level learning equation, according to which a goal’s base-level activation is dependent upon how recently and how frequently it has been retrieved. It is this constraint that determines the initial rapid build-up of activation of a new goal with repeated sampling, but also the decay of a suspended goal as a power function of time delay. In accordance with the model it was predicted that a goal suspended for longer would be subject to greater decay and therefore be more time consuming to retrieve after interruption. Existing studies have provided evidence both for [6] and against [7] an effect of interruption duration, comparing interruption intervals in the region of 30 s to two minutes. In our experiments we used a more fine-grained approach than previously, investigating the effect on individual move times of on-screen interruptions that were relatively short in duration (less than 30 s).

Interruptions of either 5 s or 15 s in duration were introduced to the execution phase of ToL problems [8]. The short interruption required the participant to complete one 5 s mood checklist and the longer interruption required the completion of three checklists that each changed automatically every 5 s. Times taken to make the fourth move are shown in Figure 2¹. Goal retrieval following interruption incurred a time cost relative to the control condition in which solution execution was continuous. Furthermore, the cost of goal retrieval was greater following the 15 s rather than the 5 s interruption as goals suspended for longer were subject to greater decay.

¹ The original experiment in Hodgetts & Jones (2006a) included a further factor (whether or not participants were aware of how long the interruption would be). However, this manipulation had no effect so the task resumption time data shown here are collapsed across conditions.




Fig. 2. Time to make the fourth move (in seconds) according to interruption condition

In a further experiment [9] interruption length was manipulated by the number of mood checklists to be completed – either one, three, or five – but the precise timing of these was under the control of the participant rather than being dictated by the computer program. Participants clicked a ‘continue’ button after completion of each checklist, which then either displayed a further checklist or returned the participant to the ToL task. The average lengths of short, medium and long interruptions were 4.48 s, 12.65 s and 20.25 s respectively. Task resumption times following short interruptions were around one second quicker than in the medium condition, a result that parallels the findings of the previous experiment. However, introducing the longer interruption interval did not result in any further increase in resumption time, suggesting that goals suffer a more rapid decline in activation when they are first suspended with a less marked decrease thereafter. G-AM predicts that activation decreases as a power function over time, and the current data seem consistent with this proposal.

2.1 Post-interruption Performance

Interruption incurs a time cost in resuming previously suspended task goals, but does an unexpected break in task exert a more general negative effect on post-interruption performance, beyond that of the resumption lag? In order to gain a general idea of the pattern of move times throughout an entire trial, a sample of three problems were chosen from Experiment 1 of Hodgetts and Jones [8] and subjected to further scrutiny for which all six move times were calculated (Figure 3). The problems selected were the first interruption trial (trial 4), the last interruption trial (trial 25), and an interruption trial from the middle of the experiment (trial 12), as well as their matched controls (problems 13, 9, and 20 respectively).

Only move times for perfect trials were used, that is, those that were completed using the correct six-move solution path (approx 85% of the data). This therefore allowed for a direct comparison between move times in each condition as each move was equivalent. First it is obvious to note that participants take longest to make their first move as this also incorporates planning time. Participants became quicker at planning with practice as plan times were greater for trial 4 than for any of the other later trials shown in Figure 3. Other than this, there were no appreciable differences between conditions (short, long or no interruption) in the first half of the trials

because before interruption these problems were all equivalent. A reliable effect at move 4 is apparent in all three of the problems, with a marked difference between interruption and no interruption trials, and a difference also between short and long interruptions. Although no formal analyses were carried out on these data, the graphs allow us to see at a glance whether interruption may have had a more general effect on post-interruption performance, perhaps slowing moves 5 and 6 as well as move 4. It appeared that this was not the case however, and that participants bear the cost of interruption only at the point of goal retrieval (i.e., the ‘resumption lag’).

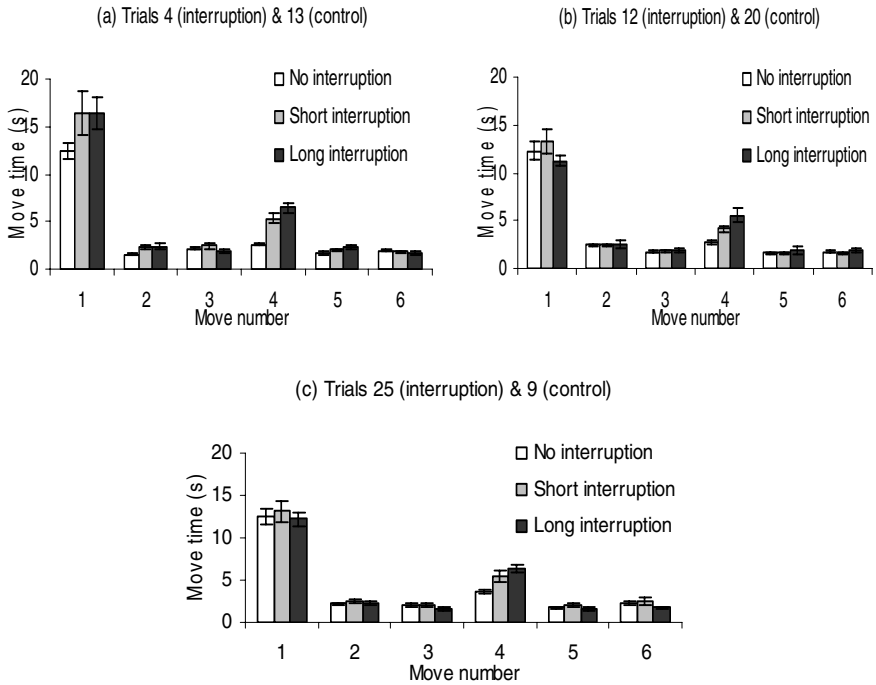


Fig. 3. Move times for perfect trials in Experiment 1, Hodgetts and Jones (2006a). Error bars show standard error.

3 Interruption Complexity

Intuitively we may think that an interruption that is more cognitively demanding would be more disruptive, and our research provides some support for this notion [8]. We find that participants are quicker to resume the primary task following a simple mood checklist task than following a more complex verbal reasoning task of the type “A follows B - - AB” [false] [10], even though they take the same length of time to complete. Furthermore, this effect of interruption complexity is more marked at points of high memory load, when participants are interrupted after their first rather than their third move [11]. Unlike the mood task, the reasoning problem involves several

elements that must each be retained and verified in order to assess the validity of the statement. In ACT-R, associative activation is limited such that the more elements that are active the less associative activation each one receives; it is possible that the number of elements involved in this more complex task leaves less activation available for the maintenance of the suspended goal. A further experiment using mental arithmetic problems found increased task resumption times when participants were interrupted to complete one complex sum (e.g., $37 + 48$) rather than a series of four simple sums (e.g., $1 + 3$, $2 + 4$) [8]. It seems therefore that it is not simply the *number* of subgoals that are activated during the interruption that is critical for disruption, but rather the processing demands imposed by each and the coordination of multiple task goals in the case of the more complex interruption.

4 Expectation and Preparation

Interruptions generally consist of two parts: an alert and the interruption proper. The period between these, termed the *interruption lag*, is thought to be a critical time for 'preparing to resume' a to-be-suspended task, perhaps through retrospective rehearsal of the current state or by prospectively encoding a future goal [12]. We manipulated the opportunity for preparing task goals before the onset of interruption by varying the transition between tasks: Either the interrupting mood checklist task covered the whole screen so that the ToL task could not be seen (as in previous experiments) or the mood checklist appeared in a box in the top left hand corner of the screen (although the ToL was visible, no further moves could be made on the main task during the interruption period; Figure 4) [13]. Participants were quicker to resume the ToL task in this latter condition because the less abrupt transition between tasks allowed a chance to prepare to-be-suspended task goals before attention was fully drawn to the mood checklist. When the checklist covered the whole screen, there was no specific time for the efficient encoding of contextual cues and so reactivation of the suspended goal after interruption was a more difficult and time consuming process. In further experiments, it was found that a 3-s pause before the onset of a full-screen interruption resulted in similar benefits to primary task resumption as with the corner-screen condition, as this provided a brief time lag for the preparation of

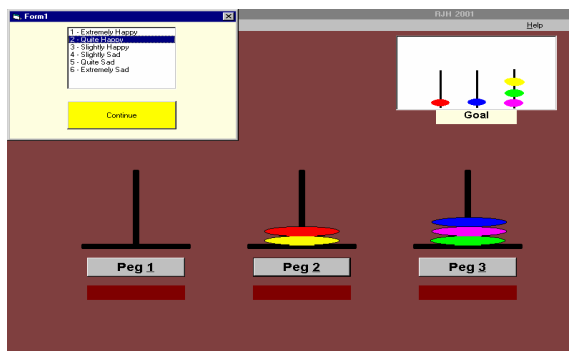


Fig. 4. Mood checklist interruption in the corner-screen condition [13]

task goals before the interruption ensued. This was the case for both mood checklist [13] and verbal reasoning task [14] interruptions. Rearranging the colors of the discs following interruption (but still retaining the original solution path) was found to increase task resumption times because it disrupts the contextual cues that help to prime retrieval of the suspended goal [13].

Some interruptions require immediate attention before any work on the original activity can proceed (e.g., 'save' reminders on spreadsheet programs), while some interruptions provide an alert but allow the user to choose a convenient time to switch tasks (e.g., a flashing instant message icon at the bottom of the computer screen). Given the proposed advantages of the interruption lag as a time for preparation before the onset of an interruption, one might expect that an opportunity to control the timing of the interruption lag would be of particular benefit. In one exploratory experiment we found no difference between 'immediate' and 'negotiated' interruptions in terms of task resumption times [14]. However, it seemed that participants chose to engage in the 'negotiated' interruptions immediately around half of the time anyway, despite the opportunity for the secondary task to be deferred. Nevertheless, the idea of negotiated interruptions deserves further investigation, especially the question of whether it would be possible to train participants to take full advantage of this preparatory period which may in turn help to mitigate the effects of interruption.

5 Practical Applications

5.1 Even Brief Interruptions Are Disruptive

The standard mood checklist interruptions were brief and undemanding in content, comparable to many types of pop-up notifications that increasingly invade our computer screens. Such interruptions may seem trivial, but the current work demonstrates that even these may be impacting upon worker efficiency, particularly given their frequency throughout the working day. Having recognised that even these seemingly inconsequential interruptions may be affecting performance, the first obvious practical recommendation would be to minimize such computerised intrusions. For example, instant-messenger systems should be turned off whilst engaging in tasks that require a lot of planning or concentration, or at least set to 'busy' so that colleagues are aware that unimportant interruptions are not welcome. Similarly, email alerts could be turned off or a priority system used, e.g., onscreen notifications are only given for those emails that the sender tags specifically as being high priority (or at least alerts are not received for incoming messages that the system detects as being 'spam').

5.2 Interruption Length and Complexity

Our work has shown that the effects of interruption are exacerbated by longer or more complex interruptions. Although the current work is limited to particularly short interruption intervals, this may still have implications for the design of computer-initiated alerts with the recommendation that the necessary information is displayed clearly in order to minimize the time spent reading them. For example, email alerts sometimes provide a number options (i.e., clear, read, save, delete), but reducing this

to just *read* or *clear* may decrease the time spent attending to the notification and therefore minimize disruption. If workers anticipate that a particular interrupting event will require a lot of time or effort to deal with, then they should think about rescheduling for a more convenient time (e.g., responding to an email can often be delayed until the worker reaches a more suitable break point in the current task).

5.3 The Importance of the Interruption Lag

Our work has also highlighted the importance of the temporal structure of interruptions, and in particular the benefit of a brief time lag before the onset of the secondary task. This time can be used to rehearse one's place in the task, or to prospectively encode cues that will later help to prime retrieval of the suspended goal. Based on the findings of the current experiments, it is recommended that the task context is such that it enables the encoding of these cues before the participant must attend to the interrupting activity. If the interruption is particularly intrusive or attention-grabbing, it may capture attention immediately without allowing the worker a chance to consolidate his or her place in the task. In light of this, moving icons for example should perhaps be avoided as research has found them to be particularly disruptive [15].

Email packages can differ in the types of notification that they use to alert the worker to incoming mail. Some use relatively discreet alerts that appear in the top corner of the screen, while others can be large and intrusive, appearing centrally on the screen and obscuring much of the original activity. We have shown that the latter type of alert may be particularly disruptive, as there is no 'window of opportunity' for consolidating primary task goals before the interruption ensues [13]. Email alerts and similar pop-up messages, should therefore be as small as possible while still conveying the appropriate information. Another difference between email packages is that some alerts allow the computer-operator to continue typing whilst they appear on the screen, whilst for others, the alert receives the focus and the original activity is put to the background. In this case, the worker has no choice but to deal with the interrupting alert immediately before he or she is able to continue with the ongoing task, but if a chance is available to finish off the current subgoal during the interruption lag (e.g., complete writing a sentence), then the worker may be able to arrive at a more convenient cognitive breakpoint before dealing with the intruding task. Furthermore, it would also be recommended that the alert actually disappears after a few seconds if not responded to, as then the worker can be made aware of the incoming information without explicitly needing to interrupt the ongoing task to respond to it.

Acknowledgment. Paper supported by the UK's Economic and Social Research Council grant no. RES-062-23-0101 and by an Overseas Conference Grant awarded by the British Academy (ref. OCG-46404).

References

1. Altmann, E.M., Trafton, G.J.: Memory for goals: An activation-based model. *Cognitive Science* 26, 39–83 (2002)
2. DeMarco, T., Lister, T.: *Peopleware: Productive projects and teams*, 2nd edn. Dorset House, New York (1999)

3. Storch, N.A.: Does the user interface make interruptions disruptive? A study of interface style and form of interruption (Report UCRL-JC-108993). Lawrence Livermore National Laboratory, Springfield (1992)
4. Anderson, J.R: Rules of the mind. Lawrence Erlbaum Associates, Hillsdale, NJ (1993)
5. Ward, G., Allport, A.: Planning and problem-solving using the five-disc Tower of London task. *Quarterly Journal of Experimental Psychology* 50A, 49–78 (1997)
6. Lahlou, S., Kirsh, D., Rebotier, T., Reeves, C., Remy, M.: Interruptions in the workplace (2000) Available: <http://adrenaline.ucsd.edu/edf/Experiment1.htm>
7. Gillie, T., Broadbent, D.: What makes interruptions disruptive? A study of length, similarity, and complexity. *Psychological Research* 50, 243–250 (1989)
8. Hodgetts, H.M., Jones, D.M.: Interruption of the Tower of London task: Support for a goal activation approach. *Journal of Experimental Psychology: General* 135, 103–115 (2006a)
9. Hodgetts, H.M., Jones, D.M.: Resuming an interrupted task: Activation and decay in goal memory. In: Sun, R., Miyake, N. (eds.) *Proceedings of the 28th Annual Conference of the Cognitive Science Society (CogSci 2006)*, p. 2506. Lawrence Erlbaum Associates, Mahwah, NJ (2006b)
10. Baddeley, A.D.: A three-minute reasoning test based on grammatical transformation. *Psychonomic Science* 10, 341–342 (1968)
11. Hodgetts, H.M., Jones, D.M.: Interrupting problem solving: Effects of interruption position and complexity. In: Katsikitis, M (ed.) *Past Reflections, Future Directions: Proceedings of the 40th Australian Psychological Society Annual Conference*, Melbourne, Australia pp. 128–132 (2005)
12. Trafton, J.G., Altmann, E.M., Brock, D.P., Mintz, F.E.: Preparing to resume an interrupted task: Effects of prospective goal encoding and retrospective rehearsal. *International Journal of Human-Computer Studies* 58, 583–603 (2003)
13. Hodgetts, H.M., Jones, D.M.: Contextual cues aid recovery from interruption: The role of associative activation. *Journal of Experimental Psychology: Learning, Memory & Cognition* 32, 1120–1132 (2006c)
14. Hodgetts, H.M., Jones, D.M.: Interruptions in the Tower of London task: Can preparation minimize disruption? In: *Proceedings of the 47th Annual Meeting of the Human Factors and Ergonomics Society*, pp. 1000–1004. HFES, Santa Monica CA (2003)
15. Ware, C., Bonner, J., Knight, W., Cater, R.: Moving icons as a human interrupt. *International Journal of Human-Computer Interaction* 4, 341–348 (1992)