Temporal Constraints on Human Information Processing
Kimron Shapiro (Ed.), pp. 100-123
Oxford University Press

Perceptual and Central Interference in Dual-Task Performance

Eric Ruthruff
NASA - Ames Research Center
and
Harold E. Pashler
University of California, San Diego

Humans often experience difficulty when asked to perform multiple tasks at the same time. Two of the better-known forms of dual-task interference are the attentional blink **(AB)** effect and the Psychological Refractory Period **(PRP)** effect. These phenomena have traditionally been studied independently, using divergent methodologies and different dependent measures. The purpose of this chapter is to explore the possibility that these dual-task phenomena might reflect the same underlying processing limitation – a central bottleneck. We also discuss how AB and PRP effects are related to other phenomena such as repetition blindness and movements of spatial attention across visual space.

Most observers can reliably find and report a single target object (e.g. a digit) embedded within a rapid serial visual presentation (RSVP) display, even when each object is presented for only a tenth of a second. After detecting one target, however, observers frequently have difficulty detecting targets that appear within the next several hundred ms (e.g. Broadbent & Broadbent, 1987; Chun & Potter, 1995; Raymond, Shapiro, & Arnell, 1992; Shapiro, Raymond, & Arnell, 1994). Because this dual-task interference is not attributable to sensory masking, it has been dubbed the “attentional blink” (**AB**) effect.

This phenomenon at least superficially resembles a type of dual-task interference known as the Psychological Refractory Period (**PRP**) effect. The PRP effect occurs when subjects are asked to make speeded responses to two different stimuli. Whereas subjects usually respond quickly to the first stimulus, they tend to respond slowly to the second stimulus when it appears shortly after the first (e.g., Vince, 1948; Welford, 1952; see Pashler & Johnston, 1998 for a review).

The PRP and AB designs differ in several respects. The typical PRP design requires only fairly simple perceptual discriminations involving stimuli presented well above threshold, often in different sensory modalities. The stimulus-response mappings are usually arbitrary, however, and subjects are pressured to produce their responses very quickly. In contrast, the typical AB design involves perceptual discriminations that are made difficult by visual backward masking, but without any pressure to select and produce a response quickly. Thus, whereas interference in the PRP design appears to arise from response selection and other decision-making processes, interference in the AB design appears to arise in perceptual processing. One might suspect, therefore, that these phenomena have different causes. On the other hand, AB and PRP effects are similar in that performance of the first task is relatively unaffected by task overlap, but performance of the second task is impaired when presented within a few hundred ms of the first. This surface similarity raises the possibility that the two phenomena might be related. It has been recently suggested, in fact, that they might be two different manifestations of a single underlying processing limitation (e.g., a central bottleneck; see Jolicoeur and Dell'Aqua, 1998, 1999).

A unified account of AB and PRP effects would obviously be very attractive by virtue of its parsimony. The purpose of this chapter, therefore, will be to consider whether such an account is viable. We begin by briefly reviewing the AB and PRP literatures. We then present several experiments using hybrid AB/PRP designs to evaluate the unified account of these phenomena. Finally, we will relate our work to the closely related studies conducted by Pierre Jolicoeur and colleagues (of which we were unaware of when we began the present experiments) and to work on other attentional phenomena, such as the effects of spatial attention and repetition blindness.

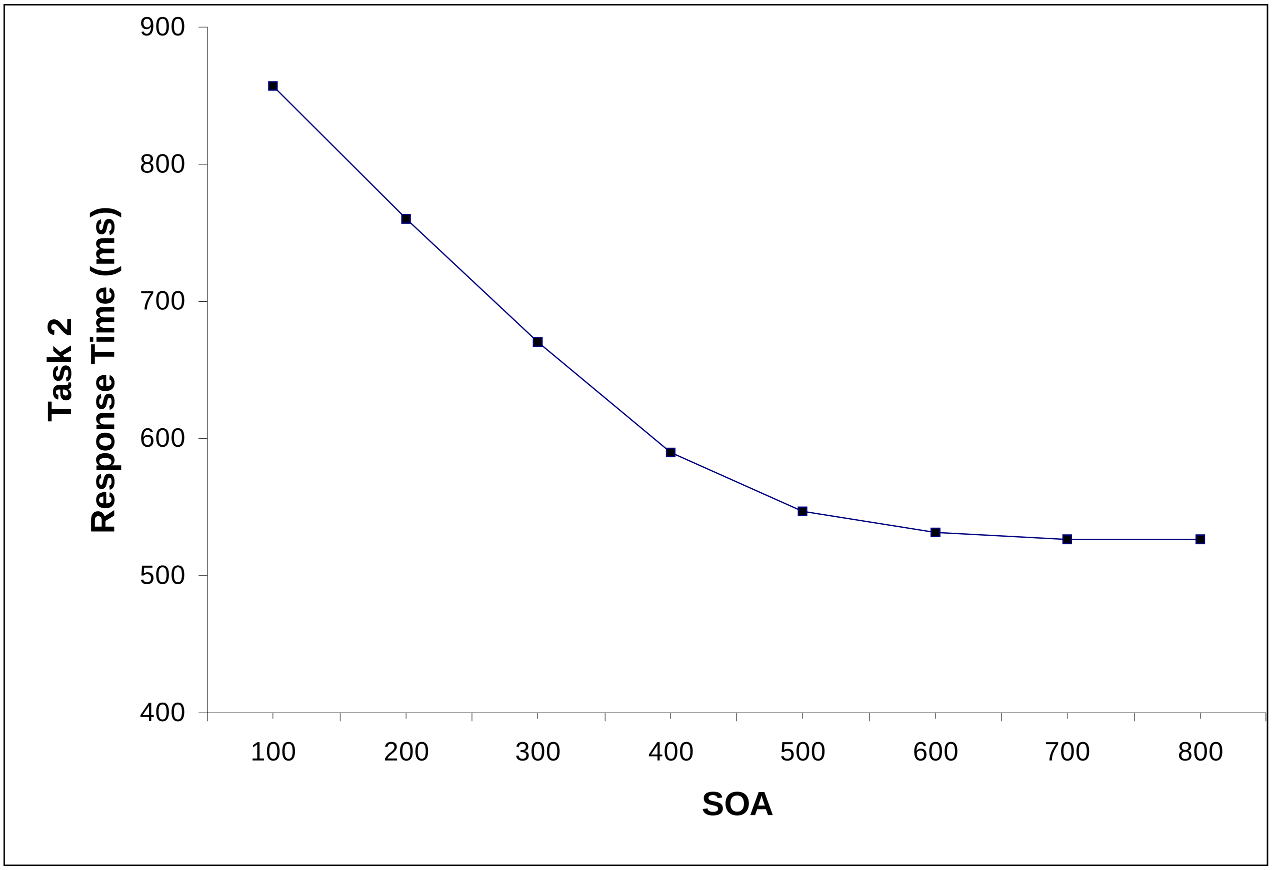

The Psychological Refractory Period (PRP) Effect

In the PRP design, subjects are presented with two tasks (Task 1 and Task 2), each requiring a separate speeded response. The key independent variable is the time between the stimulus onsets, better known as the stimulus onset asynchrony (**SOA**). Although the response time (RT) to Task 1 usually does not depend much on the SOA, the RT to Task 2 can be elevated by 300 ms or more at short SOAs (i.e. when task overlap is high). This phenomena, shown in Figure 1a, was labeled the "Psychological Refractory Period" (or PRP) effect on the mistaken assumption that it is a type of recovery phenomenon, similar to the refractory period of a neuron.

The PRP effect has been found using a very wide range of speeded tasks, including some very simple ones[1]. The effect has been found even when the responses are made with different output modalities (e.g. vocal versus manual) and even when the inputs are presented in different input modalities (e.g. auditory versus visual). In fact, the PRP effect occurs even when subjects respond to two different attributes of the same visual object (Fagot and Pashler, 1992), which presumably serves to minimize input interference (see Duncan, 1984).

The fact that the PRP effect does not depend on any obvious input or output conflicts led Welford (1952) to propose the Central Bottleneck Model. This model asserts that central mental operations (e.g. response selection, planning, decision-making) can proceed on only one task at a time. As shown in Figure 2, while the Task 1 central stage is underway, the Task 2 central stage must wait. This waiting time (a.k.a. postponement, or “bottleneck delay”) is the primary cause of the PRP effect[2].

Central Bottleneck Model

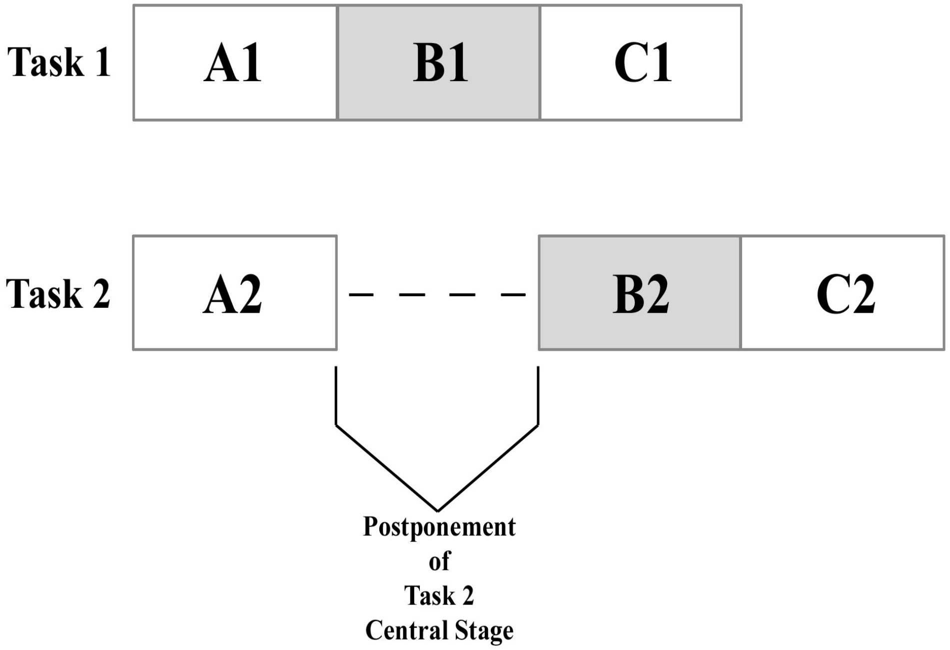

The Central Bottleneck Model makes a number of distinctive predictions that have been confirmed in a wide range of PRP experiments (see Pashler & Johnston, 1998, for a review). For example, the model correctly predicts that the effects of prolonging prebottleneck stages of Task 2 (e.g. by reducing stimulus contrast) should be greatly reduced at short SOAs (e.g. Pashler & Johnston, 1989). Consequently, it seems clear that a central bottleneck, or something very close to it, is chiefly responsible for the PRP effect.

The Attentional Blink (AB) Effect

In the standard AB design subjects view a stream of visual items (usually at a rate of about 10 per second), each presented in the same location on a computer monitor; this presentation technique is known as rapid serial visual presentation (or just “**RSVP**”). The subjects’ task is to report two targets (T1 and T2) within the RSVP stream. For example, subjects might be asked to identify two digits embedded within a stream consisting mostly of letters. Subjects are allowed to report these identities at their leisure; the primary dependent measure in this design is accuracy, not response time. As in the PRP design, task overlap is controlled by varying the SOA between T1 and T2. However, rather than express the task overlap variable in terms of the SOA, many investigators refer instead to the stimulus ‘lag,’ defined as the number of frames between the onsets of T1 and T2.

Typically, subjects can accurately report T1 but tend to miss T2 when it appears within 300-400 ms of T1. This phenomenon, known as the attentional blink, is shown in Figure 1b. The AB effect does not occur when subjects are told to ignore T1; because the mere presence of T1 is the RSVP stream is not sufficient to elicit an AB effect, the effect cannot be due to sensory masking (Raymond et al., 1992). The AB effect also does not occur reliably at lag 1. This “lag-1 sparing” effect provides one potentially important clue to the cause of the AB effect.

Several different theoretical interpretations of the AB effect have been proposed. Raymond et al. (1992) postulated an “attentional gate” that closes following detection of T1, preventing subsequent items in the RSVP stream (except perhaps for the item immediately following T1) from entering visual short-term memory (VSTM). Once T1 has been encoded from VSTM into a more durable form of short-term memory, the gate can then be reopened. This putative attentional gate would help preserve the clarity of T1’s representation in VSTM at the cost of missing T2 when presented shortly after

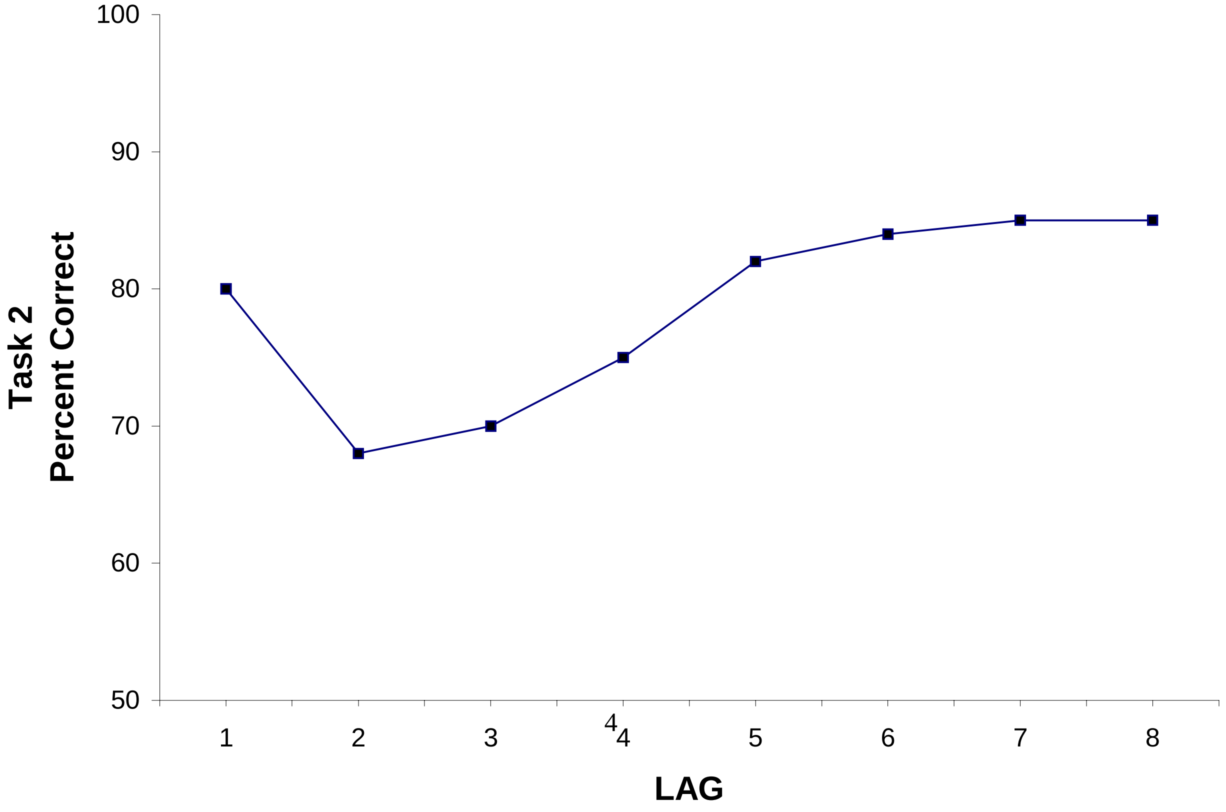

T1.

The attentional gate model can explain lag-1 sparing by adding the additional assumption that the gate does not close immediately following detection of T1. However, this model is inconsistent with several converging lines of evidence that T2, even though often unavailable for report, is nonetheless identified on most trials. Shapiro, Driver, Ward, and Sorensen (1997), for example, showed that the identity of T2 primed the processing of a later target (T3) even when subjects could not reliably report T2 (see also Maki, Frigen, and Paulson, 1997; Shapiro, Caldwell, and Sorensen, 1997). Further, the electrophysiological consequences of stimulus identification (e.g. the ‘N400’ peak of the event-related potential waveform) are largely preserved at short lags, even when T2 is missed (Luck, Vogel, and Shapiro, 1996; Vogel, Luck, and Shapiro, 1998). In addition, Johnston, Ruthruff, and McCann (1999) provided chronometric evidence using locus-of-slack logic (see McCann & Johnston, 1992) that T2 identification is not suppressed during the attentional blink. It seems likely, therefore, that the failure to report T2 arises in postperceptual processing stages.

Several different accounts of the AB effect are consistent with the evidence that T2 can be identified but not reported. Chun and Potter (1995), for example, hypothesized that all RSVP items pass through a categorization stage (i.e. are identified) but only a select few then pass through a limited-capacity stage that forms a reportable perception (e.g. that consolidates categorical representations into short-term memory [STM]). At short lags, the consolidation of T2 will be delayed by the consolidation of T1. Because of this delay, the categorical representation of T2 is likely to become unavailable due to rapid decay and/or overwriting by subsequent RSVP items before it can be consolidated. Hence, T2 will often be missed at short lags. This “two-stage” model can account for lag-1 sparing by adding the auxiliary assumption that T1 and T2 can be concurrently subjected to consolidation if they arrive in close temporal proximity.

Shapiro et al. (1994) proposed a somewhat different account of the AB effect based on interference in visual short-term memory (VSTM). According to this model, items that match the template for T1 and/or T2 are granted entry into VSTM. T1 and T2 are therefore likely to enter VSTM and so are the items immediately following T1 and T2, due to sluggishness of the selection mechanism. Upon entering VSTM the items are assigned a weighting that determines the order in which their representations (which are assumed to decay rapidly) will be passed on to a limited-capacity reporting stage. These weightings are based jointly on (a) how well the item matches the target templates and (b) the amount of remaining “resources.” T1, and perhaps the item immediately following T1 in the RSVP stream, will use up much of the available resources. Thus if T1 has not cleared VSTM by the time T2 arrives (i.e. at short lags), T2 will be assigned a low weighting and will likely be missed.

The VSTM interference model of Raymond et al. (1994) differs in a few details from the two-stage model of Chun and Potter (1995). However, both models agree that consolidation of T1 into a reportable percept interferes (directly or indirectly) with the consolidation of T2 into a reportable percept.

A Potential Unified Central Bottleneck Model of AB and PRP Effects

The primary purpose of this chapter is to ask whether the AB and PRP effects can reasonably be attributed to the same central bottleneck. According to the unified central bottleneck (UCB) model considered in this paper, there is some constraint that prevents the simultaneous occurrence of any two operations that belong to a set of demanding central operations. By hypothesis, at least one of these central operations is required to select and produce a speeded response (i.e. the typical PRP task); hence a processing bottleneck occurs when two such tasks are presented in close temporal synchrony, causing the PRP effect. Likewise, finding and remembering a target embedded in an RSVP stream requires at least one type of operation that also belongs to the set of demanding central operations; thus, a processing bottleneck occurs when two RSVP tasks are presented in close temporal synchrony, causing the AB effect.

Note that the demanding central operation(s) responsible for the AB effect need not be the exact same operation(s) responsible for the PRP effect. In fact, the central bottleneck in the PRP design appears to be due in large part to the process of response selection (see Pashler & Johnston, 1998), whereas the AB design allows subjects to respond at their leisure and thus does not even appear to require on-line response selection[3]. Although it is premature to definitively specify which central operation(s) are responsible for the bottleneck in the AB design, one plausible candidate is the consolidation of categorical representations into STM (see Chun & Potter, 1995; Jolicoeur, 1999a, 1999b; Jolicoeur & Dell'Aqua, 1998).

There is reason for skepticism about the prospects for the UCB model. As noted above, the AB effect appears to reflect a difficulty in determining what stimuli were presented, whereas the PRP effect appears to reflect a difficulty in determining how to respond to stimuli once they have been identified. Further, Pashler (1989) has provided evidence that the cause of the PRP effect is different from the cause of interference between simultaneous visual discriminations. Pashler presented subjects with a tone requiring an immediate speeded response, followed 50, 150 or 650 ms later by a masked visual array of characters. In one experiment, subjects looked for the highest number in an array of eight digits; in another they searched for a green T among green O's and red T's. In both cases, performance on the visual task depended very little on the degree of overlap with the speeded tone task. In contrast, a related set of experiments showed that performance on the visual task was severely impaired by overlap with another visual task (whether speeded or unspeeded). Based on this evidence, Pashler (1989) argued that dual-task interference has (at least) two different causes: (1) there is a constraint requiring that demanding central processes (including but not limited to response selection) be performed one at a time, and (2) visual processes may proceed simultaneously but suffer mutual interference dependent upon the complexity of the visual operations required. On this view, the PRP effect would be due to the first component, whereas the AB effect might be attributed to the second component.

Testing a Unified Account: Experimental Logic

According to the UCB model, speeded-response tasks (used in the PRP design) and RSVP tasks (used in the AB design) both require demanding central operations that can take place on only one task at a time. Thus, this model predicts that one speeded task should interfere with another speeded-response task at short task lags (causing the PRP effect). Likewise, one unspeeded RSVP task should interfere with another unspeeded RSVP task at short task lags (causing the AB effect). Furthermore, the UCB model predicts that an unspeeded RSVP task should interfere with a speeded-response task. In other words, while demanding central operations are underway on an RSVP Task 1, the critical central operations on a speeded-response Task 2 will be postponed. Experiment 1 was designed to test this prediction. Experiment 2 was designed to test the converse prediction, namely that a speeded-response task should interfere with an unspeeded RSVP task.

Table 1 summarizes the relations between the AB design, the PRP design, and the designs of Experiments 1 and 2. Each of the four cells represents an experimental design resulting from pairing together speeded-response and/or unspeeded RSVP tasks as Task 1 and Task 2. When Task 1 and Task 2 are both speeded-response tasks, the experiment is a pure PRP design. When Task 1 and Task 2 are both unspeeded RSVP tasks, the experiment is a pure AB design. When Task 1 is an unspeeded RSVP task and Task 2 is a speeded-response task (Experiment 1), the experiment is a "PRP/AB" design (where the first set of initials refer to the type of Task 1 and the second set of initials refer to the type of Task 2). When Task 1 is a speeded-response task and Task 2 is an unspeeded RSVP task (Experiment 2), the experiment is an "AB/PRP" design. If the UCB model is correct, then we should observe substantial dual-task interference in the hybrid designs.

As will be seen, the hybrid designs of Experiments 1 and 2 do in fact produce substantial interference effects, consistent with the UCB model. The purpose of Experiments 3 and 4 was to provide a more direct test for the involvement of central interference in the AB effect. Finally, Experiment 5 looked for evidence that the AB effect is due to both central interference and perceptual interference.

Experiment 1

(AB/PRP design)

Experiment 1 paired a typical AB Task 1 (i.e. unspeeded report of a target in an RSVP stream) with a typical PRP Task 2 (i.e. speeded report of an unmasked letter). We refer to this as the AB/PRP design. The goal is to see if central operations engendered by a typical AB Task 1 prevent the simultaneous performance of central operations on a typical PRP Task 2, as predicted by the UCB model. If so, then we should observe substantial dual-task interference, in the form of Task 2 slowing at short lags (similar to the PRP effect).

Task 1 (unspeeded) was to report a target digit (1, 2, 3, or 4) presented within an RSVP stream of letter distractors. Task 2 was to rapidly respond to a letter (E or F) by pressing the corresponding key on the keyboard. The E or F was always the last item in the RSVP stream (and therefore unmasked). Further, it was slightly brighter than the other RSVP items and remained present until the subject responded to it. Thus as with most PRP experiments, the Task 2 stimulus was fairly easy to perceive.

Method

Subjects. Twenty-one students at the University of California - San Diego participated in return for partial course credit. All reported normal or corrected-to-normal vision. No subject participated in more than one experiment reported in this chapter.

Stimuli. Stimuli were alphanumeric characters, which subtended approximately 0.5 degrees horizontally by 0.7 degrees vertically from a typical viewing distance of 60 cm. They were gray (IBM VGA color code #8) against a black background, displayed on NEC Multisynch monitors connected to IBM PC-compatible computers. Characters were displayed using rapid serial visual presentation (RSVP). Each RSVP display consisted of fifteen characters in the middle of the screen, presented strictly sequentially for 85 ms each. The non-target characters in the RSVP sequence were letters drawn randomly from the alphabet, excluding the letters E, F, I, and O.

Procedure. The Task 1 stimulus (S1) was a digit (1, 2, 3, or 4) presented in the 4th frame of the RSVP stream. Subjects were asked to identify the digit and report it at the end of the trial (i.e. sometime after they had made their speeded response to Task 2). The digits 1-4 were mapped in numerical order onto the ‘N’, ‘M’, ‘<’, and ‘>’ keys. The Task 2 stimulus (E or F) appeared at a lag of two to eight frames after S1. We avoided lag 1 because this condition does not reliably produce an AB effect (e.g. Raymond et al., 1992; Chun & Potter, 1995). Each lag occurred equally often. Subjects were asked to press the key (‘E’ or ‘F’) corresponding to the Task 2 stimulus as fast as possible, without making too many mistakes.

Subjects completed 6 blocks of 66 trials. The first two blocks were considered practice and therefore not analyzed. Similarly, the first two trials within a block were considered warm-ups. Feedback on Task 1 accuracy, Task 2 accuracy and RT2 was provided at the end of each block.

The sequence of events within a trial was as follows. A fixation point was presented in the middle of the screen for 500 ms followed by a blank screen for 250 ms. Next, the RSVP stream was presented at a rate of 85 ms per character. Once both responses had been made, they were echoed to the screen. The word ‘WRONG’ was displayed adjacent to any incorrect response(s) for 1 sec. The next trial began 1 sec later.

Analyses. We excluded any subject who did not achieve an average of at least 65% correct on Task 1 and 85 % correct on Task 2. Further, we excluded from RT2 analyses all trials on which the Task 1 response was incorrect.

Results and discussion

The results are shown in Figure 3 as a function of lag. Task 1 percent correct, which did not depend on lag, is shown in Table 2. The main purpose of this experiment was to determine if an unspeeded RSVP Task 1 would interfere with a speeded Task 2, as predicted by the UCB model. Task 2 accuracy (mean = 98%) was affected relatively little by lag, however mean Task 2 RT was 124 ms longer at short lags than at long lags, F(6,120) = 54.7, $p < 0.001$. Similar findings have been obtained by Johnston, Ruthruff, and McCann (1999) and Arnell and Duncan (1999) using other variants of the AB/PRP design.

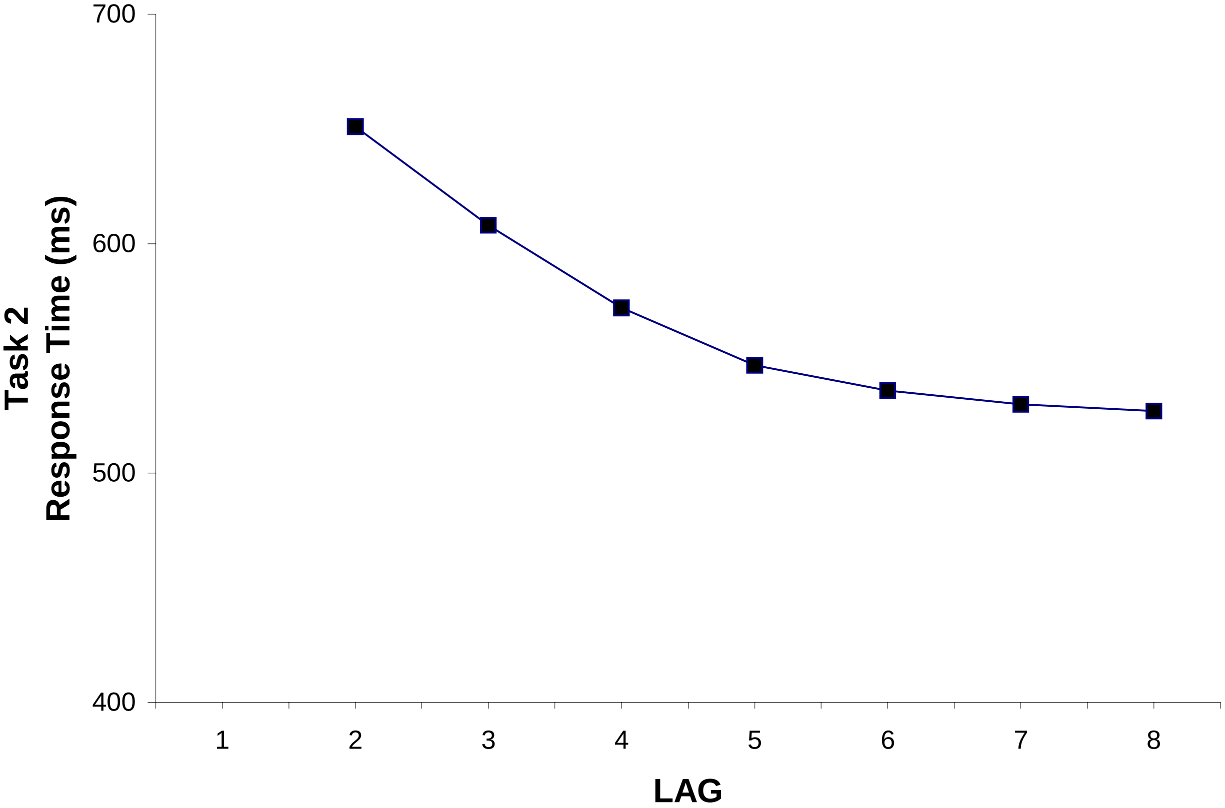

Second-task slowing in this experiment occurred even though the Task 2 stimulus was unmasked and therefore relatively easy to perceive. In fact, we have informally replicated this finding even when the Task 2 stimulus is auditory (see also Arnell and Duncan, 1999; Jolicoeur and Dell'Aqua, 1998, 1999). This pattern of results is consistent with a central locus of interference.

How does the size of this interference effect compare to that typically found in PRP experiments? First, it is important to note that the present experiment used a somewhat restricted range of lags (2 to 8) compared to that used in most PRP experiments (the equivalent of lags 0 and 9). Even after accounting for the difference in lags, the slowing observed here (124 ms) is somewhat less than 170+ ms PRP effects observed across comparable task lags (e.g. McCann & Johnston, 1992; Pashler & Johnston, 1989; Ruthruff, Miller, & Lachmann, 1995). According to the UCB model, this means that the Task 1 central operations required by a typical AB task finish sooner than the central operations required by a typical PRP task. This conclusion is plausible given that the AB task requires only that subjects remember the target identity (presumably a highly-practiced operation), whereas PRP tasks usually require subjects to negotiate a novel and arbitrary stimulus-response mapping.

Experiment 2

(PRP/AB design)

In Experiment 1 an unspeeded RSVP Task 1 (like that used in the AB design) interfered with the processing of a speeded, unmasked Task 2 (like that used in the PRP design) at short lags. This outcome is consistent with the UCB model, which says that central operations on an unspeeded RSVP Task 1 can postpone central operations on a speeded Task 2. In Experiment 2 we reversed the order of these two tasks (the "PRP/AB" design) to test the converse claim: central operations on a speeded Task 1 can postpone the central operations critical to performing an unspeeded RSVP Task 2. If so, then Task 2 accuracy should be reduced at short lags relative to long lags (similar to the AB effect).

Task 1 required a speeded response to the pitch of a pure tone, which was presented on half of the trials and absent on the other half (selected at random). The target-absent trials provide a baseline against which we can performance in the target-present condition. The tone was presented for 85 ms (the same duration as the characters in the RSVP stream) and was not masked. Subjects were instructed to press the 'N' key if they heard the low frequency tone (300 Hz) and press the 'M' key if they heard the high frequency tone (900 Hz). Task 2 was to make an unspeeded response indicating whether the 'E' or the 'F' was present in the RSVP stream.

Method

Except where noted, the method was identical to that of Experiment 1.

Subjects. Fifty-four students at the University of California - San Diego participated in exchange for partial course credit.

Procedure. The Task 1 tone, when present, coincided with the fourth frame of the RSVP stream, which contained a random distractor letter. Subjects were instructed to respond immediately to the tone. After a lag of 1 to 8 frames, the Task 2 stimulus (an 'E' or an 'F') was presented. Subjects were allowed to respond at their leisure to this stimulus.

Analyses. We excluded any subject who did not achieve an average of at least 65% correct overall on Task 1 and 65 % correct on Task 2 at the two longest lags. Further, we excluded from analyses of Task 2 RT and Task 2 accuracy any trials on which the Task 1 response was incorrect.

Results and Discussion

The results are shown in Figure 4. Mean Task 1 RT (620 ms) did not depend on lag, $F(1,31) < 1$; this indicates that subjects did not interrupt Task 1 in order to perform Task 2. The main purpose of this experiment was to see if a speeded Task 1 (like that used in the PRP design) would interfere with an unspeeded RSVP Task 2 (like that used in the AB design), as predicted by the UCB model. As shown in Figure 4, the Task 2 accuracy was in fact much lower overall when the Task 1 stimulus was present than when it was absent, $F(1,31) = 3.66$, $p < 0.01$. Further, Task 2 accuracy declined substantially at short lags when the Task 1 stimulus was present, $F(7,217) = 5.78$, $p < 0.001$.

Using methods similar to those of the present experiment, Jolicoeur (1999) found essentially the same pattern of results and arrived at a similar conclusion. We have also replicated this result in several different pilot experiments. One key finding is that PRP/AB interference occurs even when the Task 1 tone judgment is very easy. For example, we have found this effect when the high and low tone frequencies are even more distinct (200 and 2000 Hz) and the tone is presented for three times as long (i.e. 255 ms). The fact that interference in the PRP/AB design can be obtained even when the stimuli are cross-modal and relatively easy to perceive is certainly consistent with models that propose a central (rather than perceptual) locus for the AB effect.

Experiment 3

(Task-1 Carryover)

It appears that dual-task interference occurs regardless of how speeded-response tasks and RSVP tasks are combined (i.e. in all four cells of Table 1). Performance of two speeded-response tasks produces the robust PRP effect and performance of two unspeeded RSVP tasks produces the well-documented AB effect. Experiment 1 (the AB/PRP design) showed that dual-task interference also occurs when an RSVP Task 1 is paired with a speeded-response Task 2. Further, Experiment 2 (the PRP/AB design) showed that interference occurs when a speeded-response Task 1 is paired with an RSVP Task 2. These results are consistent the UCB model, which says that both of these types of tasks (speeded-response and unspeeded RSVP) require central operations that can proceed on only one task at a time.

Although the results from Experiments 1 and 2 are consistent with the UCB model, they are consistent with a few other models as well. For example, the results are consistent with a model in which the PRP effect is due primarily to a

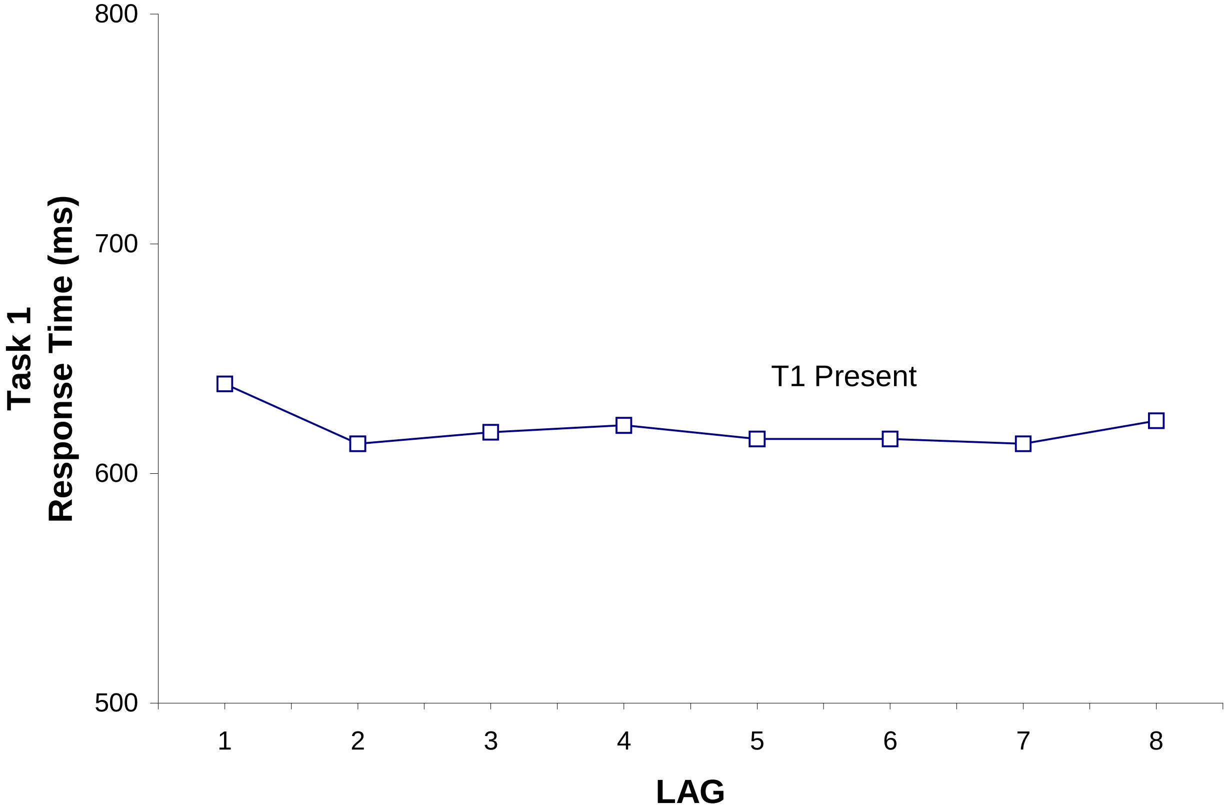

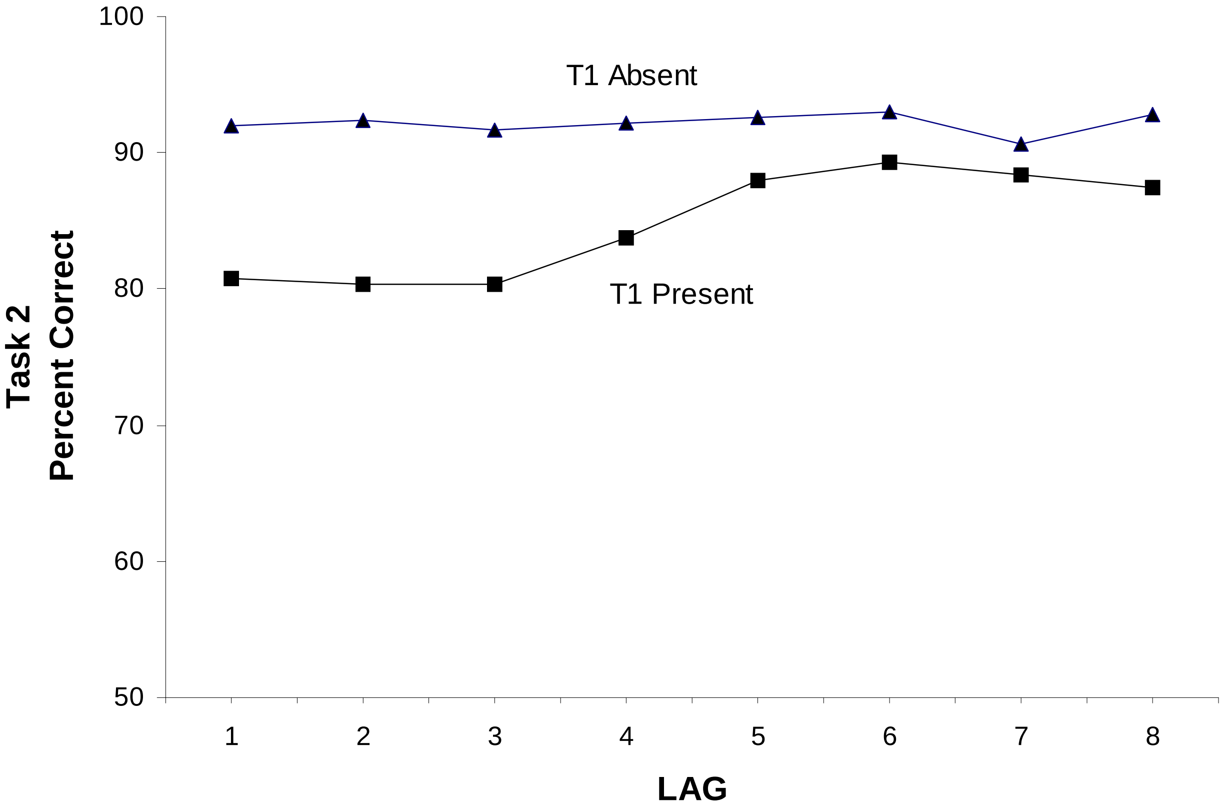

central bottleneck whereas the AB effect is due to an earlier bottleneck that is completely separate from the central bottleneck. The locus of this early bottleneck need not be perceptual processing, per se. It might be STM consolidation, though there are other plausible candidates as well, such as stimulus classification[4]. To explain the interference observed in both Experiment 1 (the AB/PRP design) and Experiment 2 (the PRP/AB design), this dual-bottleneck account need only assume that both speeded tasks and RSVP tasks require the same early stage that causes the AB bottleneck.

If the AB effect is due to an early bottleneck rather and not to a central bottleneck, then the amount of interference in the AB design should not depend on the time required to complete the central stages of Task 1. Instead, the size of the AB effect should depend only on the time required to complete the early stages of Task 1 (i.e. up to and including the stage that causes the early bottleneck). According to the UCB model, on the other hand, dual-task interference occurs primarily because the Task 1 central stage postpones the onset of the Task 2 central stage. At short lags, any manipulation that delays the completion of the Task 1 central stage will cause an additional postponement in the onset of the Task 2 central stage. In the AB design, where Task 2 is an RSVP task, the extra postponement of the Task-2 central stage might cause the categorical representation of T2 to decay even further before it can be consolidated. Consequently, T2 would be especially unlikely to be reported when the Task 1 central stage is prolonged. More precisely, when the Task 1 central stage is prolonged the AB effect should be deeper and more prolonged (i.e. recovery from dual-task interference should occur at a later lag). Henceforth we refer to this as the "Task-1 carryover prediction" because the effects of prolonging the Task 1 central stage carry over onto Task 2.

To test between the UCB model and the dual-bottleneck model, Experiment 3 evaluated the Task-1 carryover prediction. We manipulated the duration of the Task 1 central processing in an AB design by varying whether the response to the Task 1 stimulus was speeded or unspeeded. We reasoned that the requirement to perform an immediate response selection should increase the duration of the central stage relative to a condition where subjects could defer response selection until after T2 had been processed. If the UCB model is correct, the decrement in Task-2 accuracy should be deeper and longer-lasting when Task 1 is speeded. If the early bottleneck account is correct, however, the decrement in Task-2 accuracy should be roughly the same for both the speeded and unspeeded conditions.

Method

Except where noted, the method was identical to that of Experiment 2.

Subjects. Forty-two students at the University of California - San Diego participated in exchange for partial course credit.

Procedure. Task 1 was to identify a digit (1, 2, 3, or 4) embedded in an RSVP stream. Half of the subjects were instructed to respond immediately to Task 1 (the speeded condition), and the other half were instructed to respond only after all 15 characters had been displayed (the unspeeded condition). Task 2 (unspeeded) was to indicate whether the E or the F had been presented.

Results and Discussion

The results are shown in Figure 5. Task 2 performance was worse overall when Task 1 was speeded than when Task 1 was unspeeded, $F(1,40) = 18.3$, $p < 0.001$, although it is debatable whether the recovery from dual-task interference was also prolonged. These results are broadly consistent with the UCB model, but appear to contradict the dual-bottleneck model.

Jolicoeur (1998) has also compared speeded and unspeeded Task-1 judgments and found a very similar pattern of results. Further, Jolicoeur (1999b) manipulated the number of stimulus-response alternatives on Task 1. In some blocks of trials subjects performed a two-alternative forced-choice Task 1 based either on letter size (small or large) or letter identity (H versus S), while in other blocks of trials subjects performed a four-alternative forced-choice Task 1 based on both letter identity and letter size. This manipulation also produced the same basic Task-1 carryover effect observed in the present experiment and in Jolicoeur (1998).

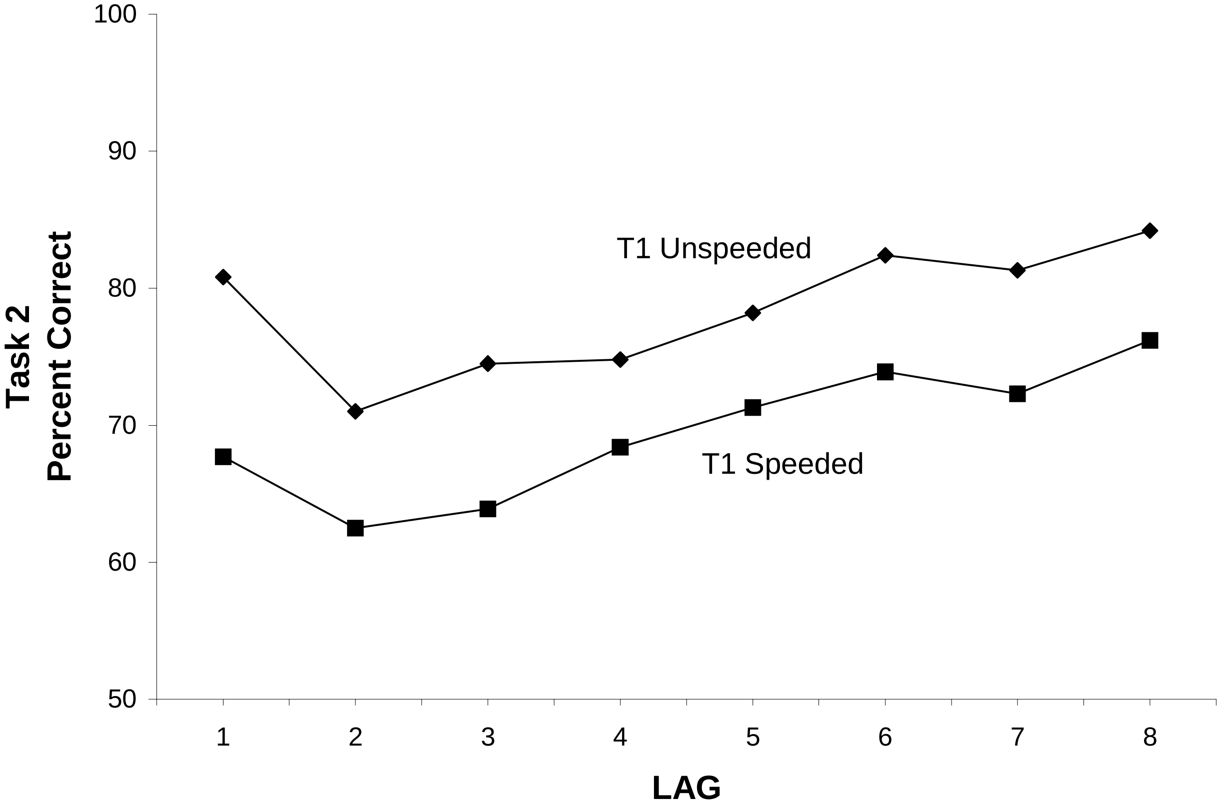
100
90
80
70
60
50
Task 2
Percent Correct
T1 Unspeeded
T1 Speeded
1
2
3
4
5
6
7
8
LAG

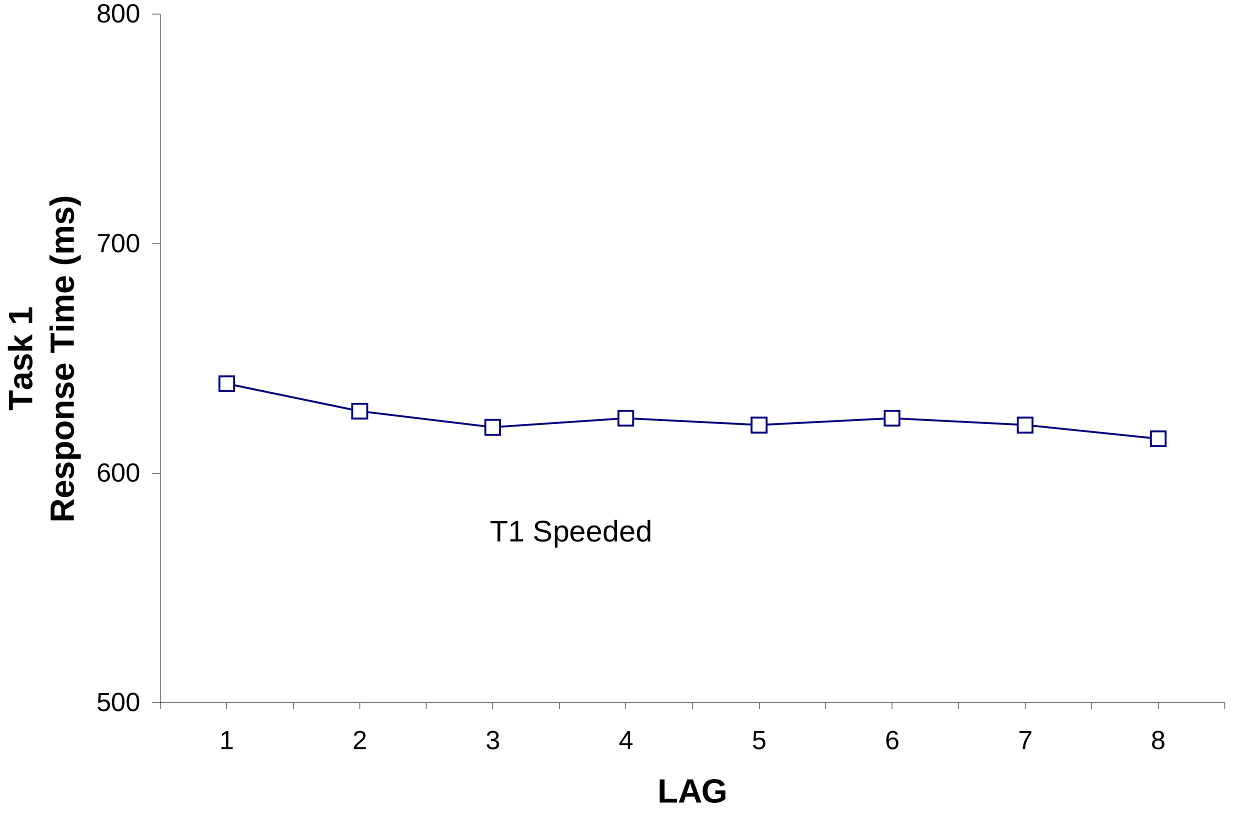
800
700
600
500
Task 1
Response Time (ms)
T1 Speeded
1
2
3
4
5
6
7
8
LAG

Experiment 4

(Task-1 Carryover Using a Within-Block Manipulation)

Varying whether subjects need to select an on-line response to Task 1 (as in Experiment 3) would seem to be a natural way to control the duration of the Task 1 central stage; however, this manipulation has several serious drawbacks. First, although it seems likely that Task 1 central stage is prolonged in the speeded condition relative to the unspeeded condition, this is not necessarily the case. It is conceivable, for example, that it takes about as long to consolidate T1 into memory when the task is unspeeded as it does to select an immediate response to T1 when the task is speeded. Because the unspeeded Task-1 condition does not produce an observable response time, there is no easy way to confirm that central operations were in fact prolonged in the speeded condition relative to the unspeeded condition.

An even more serious problem with Experiment 3 (and with the studies reported by Jolicoeur, 1998, 1999b) is that the speeded and unspeeded Task-1 conditions were not run within the same block of trials. Consequently, subjects might have adopted different response criteria or task strategies in the speeded and unspeeded blocks. Furthermore, subjects might have experienced higher levels of arousal, or simply exerted more effort, in one condition or the other. In addition, more pre-trial preparation might be required for a speeded Task 1 than an unspeeded Task 1, leaving subjects less prepared for Task 2 (see Gottsdanker, 1980). This preparation effect alone could account for the poorer Task 2 performance observed in Experiment 3 when the Task 1 response was speeded.

In Experiment 4 we addressed these problems by varying the stimulus-response compatibility of a speeded Task 1 within blocks of trials. The Task 1 stimulus was a digit drawn randomly from the set {1, 2, 3, 4, 5, 6, 7, 8}. As in Experiment 3, the digits 1-4 were mapped in numerical order (i.e. compatibly) onto four response keys. However, the digits 5-8 were mapped in reverse order (i.e. incompatibly) onto the same response keys. These compatible and incompatible Task-1 conditions were mixed randomly within a block of trials to ensure that these conditions did not differ in arousal, response criteria, the degree of preparation for Task 2, etc. Furthermore, note that both the compatible and incompatible conditions produce an observable RT, allowing us to determine whether we were in fact successful in prolonging the duration of the Task 1 response selection stage and to estimate the amount of prolongation. An additional advantage of this design is that it seems especially clear that the Task 1 stage affected by the present response-compatibility manipulation is central rather than perceptual.

Method

Except where noted, the method was identical to that of Experiment 3.

Subjects. Twenty-two students at the University of California - San Diego participated in exchange for partial course credit.

Procedure. The Task 1 stimulus was chosen at random (with replacement) from the set {1, 2, 3, 4, 5, 6, 7, 8}, with the exception that the same stimulus was never presented twice in a row. Subjects responded by pressing either the ‘N’, ‘M’, ‘<’, or ‘>’ key. The digits 1-4 were mapped in numerical order onto the four responses keys from left to right (i.e. compatibly), while the digits 5-8 were mapped in reverse numerical order from left to right (i.e. incompatibly). Subjects were instructed to respond to the digit as fast as possible without making too many errors. As in the previous experiment, Task 2 (unspeeded) was to indicate whether an E or an F had been presented in the RSVP stream.

Results and Discussion

The results are shown in Figure 6. The compatible Task 1 condition (M = 720 ms) produced much faster responses than the incompatible Task 1 condition (M = 846 ms), F(1,21) = 7.6, $p < .01$. This suggests that we were successful in increasing the duration of Task 1 central operations, making it possible to meaningfully evaluate the Task-1 carryover prediction. Relative to the compatible condition, the incompatible Task 1 condition produced a deeper decrement in Task 2 accuracy, F (1,20) = 13.2, $p < 0.01$. The incompatible condition also appears to have produced a more prolonged decrement in Task 2 accuracy. These observed Task-1 carryover effect contradicts the dual-bottleneck account. However, they provide compelling evidence for the involvement of central operations in the AB effect, as proposed by the UCB model.

Experiment 5

(Within-Modal versus Cross-Modal interference)

The interference observed in the hybrid designs of Experiments 1 and 2 along with the Task-1 carryover effects observed in Experiments 3 and 4 clearly suggest that central interference plays a role in the AB effect. As discussed next, however, there is reason to suspect that a central bottleneck is not the only cause of interference in the AB design. The hybrid AB/PRP design of Experiment 1 produced only 124 ms of dual-task slowing between lags 2 and 8, whereas comparable PRP experiments typically show 170 ms or more of Task-2 slowing between the same

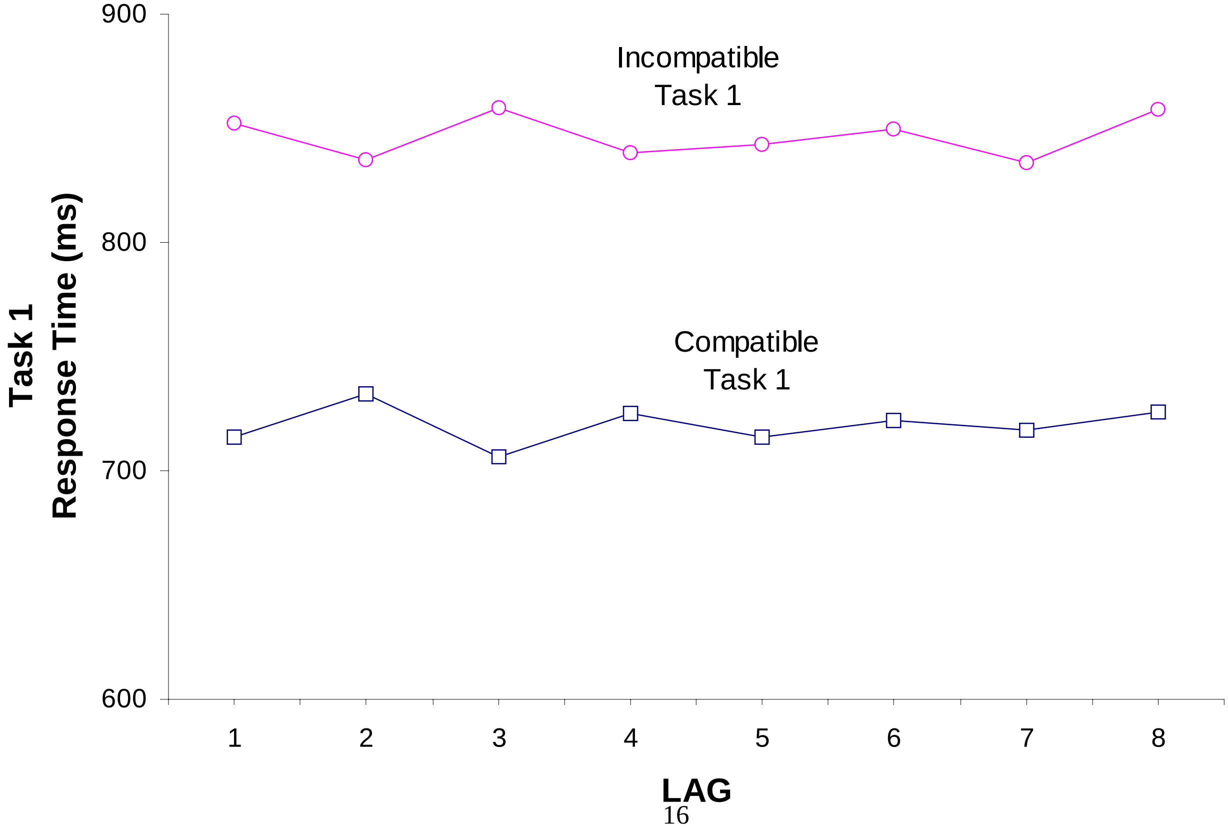

lags (e.g. McCann & Johnston, 1992; Pashler & Johnston, 1989; Ruthruff et al., 1995). In other words, a speeded-response Task 2 appears to suffer less interference when Task 1 is an RSVP task (as in the AB/PRP design) than when Task 1 is another speeded-response task.

According to the UCB model this difference occurred because central operations simply finish earlier when Task 1 is an RSVP task than when Task 1 is a speeded-response task, resulting in less postponement. This explanation is plausible given that an RSVP task merely requires that the subject remember a single target character, whereas most speeded-response tasks require subjects to rapidly carry out a novel and arbitrary mapping of stimulus onto response. However, the UCB must then predict that an RSVP Task 1 should also produce less interference than a speeded-response Task 1 when Task 2 is an RSVP task. In other words, the PRP/AB design should produce less interference than the pure AB design. Contrary to this prediction, the PRP/AB design appeared to produce somewhat <u>less</u> interference than is typically found in pure AB experiments.

One explanation for this pattern of results is that a central bottleneck is not the only cause of interference in these designs. It is possible, for example, that the AB effect is due in part to a central bottleneck and in part to perceptual interference. Experiment 5 was designed to evaluate this possibility. Specifically, the goal was to determine if the AB effect is due to perceptual interference in addition to any central interference.

Our approach was to see if an unspeeded RSVP Task 2 suffers more interference when Task 1 is another RSVP task (the within-modal condition) than when Task 1 requires a speeded response to a pure, unmasked tone (the cross-modal condition). The within-modal condition was essentially a replication of the standard AB effect. The cross-modal condition was explicitly designed to require the same amount of Task 1 central processing as the within-modal condition but with greatly reduced potential for perceptual interference (both because the inputs were cross-modal and because the tone was unmasked). If the AB effect is due in part to perceptual interference, then the AB effect should be substantially reduced in the cross-modal condition relative to the within-modal condition.

To ensure that the within- and cross-modal conditions did not differ with respect to arousal, effort, or task preparation, these conditions were randomly mixed within blocks of trials. Furthermore, to estimate the duration of central processing on Task 1, we required a speeded response to the Task-1 stimulus in both the within- and cross-modal conditions. The task parameters were chosen (based on pilot studies) so that the auditory and visual Task 1 judgments would produce roughly the same RT. Because these judgments required the exact same responses (pressing the ‘L’ or ‘H’ key), we can then infer that they required roughly the same amount of central processing. Therefore, the UCB model predicts that the two conditions should produce roughly equal amounts of Task 2 postponement, which in turn should lead to roughly equal amounts of dual-task interference. However, if the AB effect is due in part to perceptual interference then we should observe greater amounts of interference in the within-modal condition.

Method

Except where noted, the method was identical to that of Experiment 3.

Subjects. Twenty-one students at the University of California - San Diego participated in exchange for partial course credit.

Procedure. S1 on each trial was either a tone (300 Hz or 900 Hz) or a green letter (L or H). When S1 was a letter, it occupied the fourth frame of the RSVP stream. When it was a tone, it was presented simultaneously with the fourth frame of the RSVP stream, which was a gray distractor character selected at random. Subjects responded to the low-pitched tone and to the letter L by pressing the key marked 'L'; subjects responded to the high-pitched tone and to the letter H by pressing the key marked 'H'. Subjects were instructed to respond quickly without making too many errors. Task 2 (unspeeded) was to determine whether an E or an F had been presented in the RSVP stream.

Results and Discussion

The results are shown in Figure 7. Subjects responded somewhat more slowly to Task 1 when S1 was auditory (551 ms) than when it was visual (491 ms), $F(1,20) = 9.5$, $p < 0.01$. Because subjects made the exact same physical responses to both Task 1 judgments, the duration of response processes should have been similar. It therefore seems likely that the central stage of the tone task tended to finish at least as late as the central stage of the letter task (if not later). Consequently, the UCB model predicts at least as much interference when the Task 1 stimulus was a tone (i.e. in the cross-model condition) as when the Task 1 stimulus was a letter (i.e. the within-modal condition). Contrary to this prediction, we actually found less interference in the cross-modal condition than in the within-modal condition. We observed less effect of lag in the cross-modal condition than in the within-modal condition, $F(2,40) = 2.2$, $p < 0.05$. Furthermore, Task 2 performance at short lags (2, 3 and 4) was significantly better in the cross-modal condition than in the within-modal condition, $F(1,20) = 5.7$, $p < 0.05$.

Similar results were obtained in a recent study by Arnell and Duncan (1999). Each task in their experiment (i.e. Task 1 and Task 2) could be either a speeded-response to an auditory stimulus or a unspeeded response to a visual stimulus. As in the present Experiment 5, these conditions were mixed together randomly within a block of trials. Consistent with the present results, Arnell and Duncan found greater interference in both of the pure conditions (i.e. speeded tone task followed by another speeded tone task or unspeeded digit task followed by another unspeeded digit task) than in the hybrid conditions (i.e. speeded tone task followed by unspeeded digit task or unspeeded digit task followed by speeded tone task).

In summary, the present results and those of Arnell and Duncan (1999) suggest that the AB effect is due not only to central interference (see Experiments 1-4) but also to perceptual interference resulting from the requirement to simultaneously perform two demanding visual discriminations. One specific possibility is that within-modal conditions produce relatively weak categorical representations of T2 relative to the cross-modal condition. Because the representation of T2 is especially weak to begin with, it is unlikely to survive the period of Task 2 postponement.

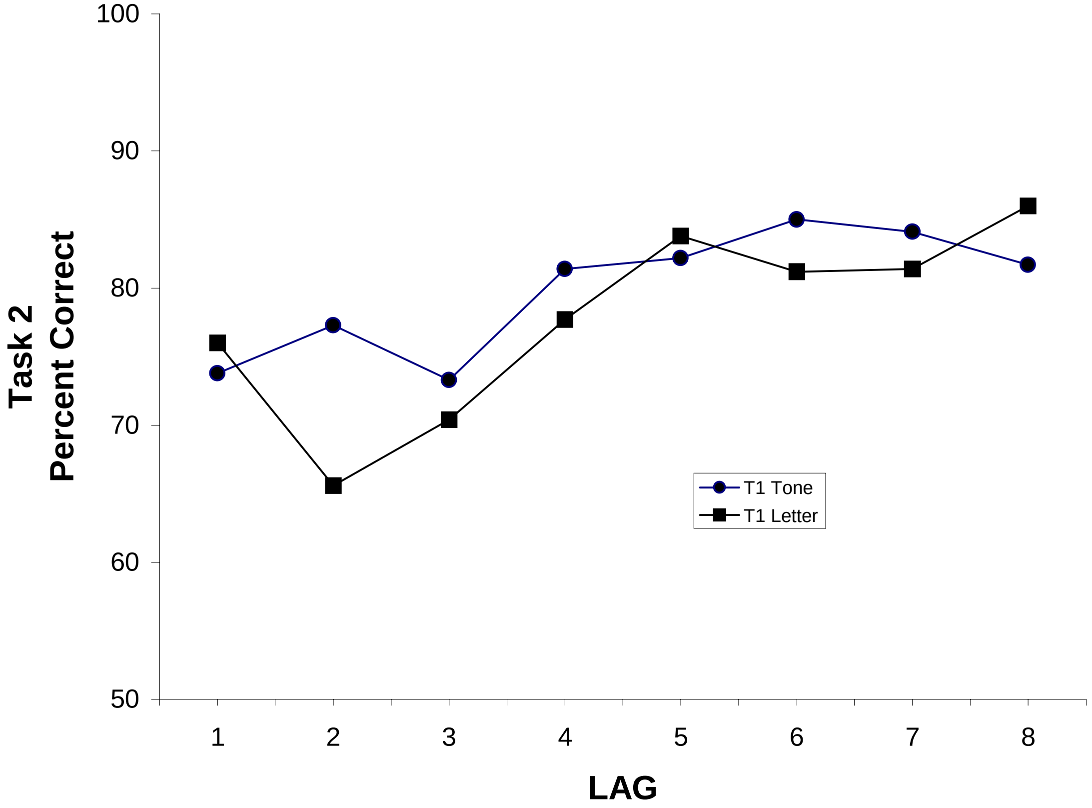
Task 2
Percent Correct
100
90
80
70
60
50
1
2
3
4
5
6
7
8
LAG
T1 Tone
T1 Letter

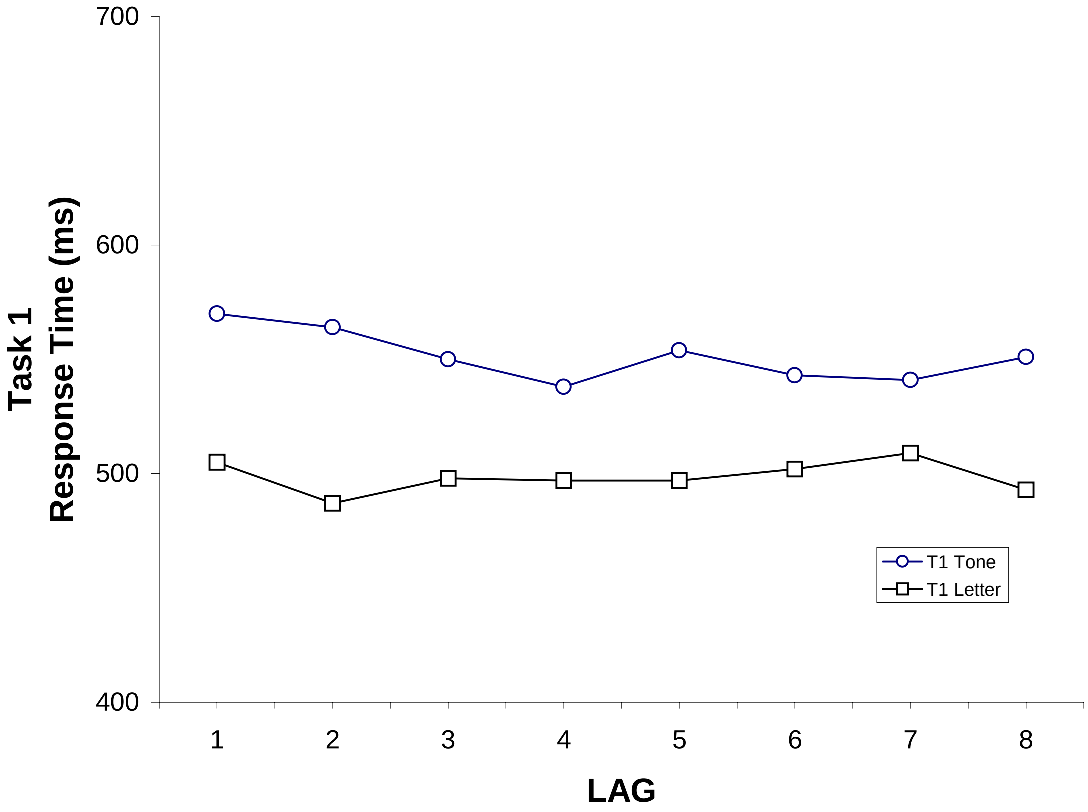
Task 1
Response Time (ms)
700
600
500
400
1
2
3
4
5
6
7
8
LAG
T1 Tone
T1 Letter

General Discussion

The purpose of this chapter is to explore the relationship between two well-studied forms of dual-task interference: the Psychological Refractory Period (PRP) effect and the Attentional Blink (AB) effect. The PRP effect is a tendency to respond slowly to one speeded-response task when it is presented shortly after another speeded-response task. The AB effect is a tendency to miss one target embedded in an RSVP stream when it is presented shortly after another target.

One might suspect that these two phenomena are caused by distinct attentional limitations. The PRP effect arises with two speeded tasks even when input conflicts are minimized. In fact, it occurs even with unmasked supra-threshold stimuli presented in different modalities (e.g. one auditory and one visual stimulus). The AB effect, in contrast, would seem to reflect a sort of input interference limited to rapid presentation rates. It is observed even in the complete absence of speed pressure on either response. The experiments reported in this chapter, however, suggest that AB and PRP effects may nonetheless be due, at least in part, to the same attentional limitation, commonly termed the "central bottleneck."

We conducted several tests of the unified central bottleneck (UCB) model, which says that (a) there is some constraint that prevents the simultaneous occurrence of any two operations that belong to a set of demanding central operations and that (b) speeded-response task and RSVP tasks require at least one of these demanding central operations (though not necessarily the same ones). Thus, both AB and PRP effects occur because the demanding central operations on Task 2 are postponed until the demanding central operations on Task 1 have been completed. In the PRP design, this postponement manifests itself as a slowing of the Task-2 response; in the AB design, the postponement manifests itself as a tendency to miss T2.

If the UCB model is correct, then a central bottleneck should occur not only in the pure AB and PRP design, but also in hybrid designs that pair a speeded-response task with an RSVP task (see Table 1). Experiment 1 verified that an unspeeded RSVP Task 1 does in fact slow processing on a speeded-response Task 2 (see also Johnston et al., 1999). Similarly, Experiment 2 verified that a speeded response Task 1 impaired performance of an unspeeded RSVP Task 2. A very similar finding was also reported by Jolicoeur (1999b).

Experiments 3 and 4 tested an alternative account of Experiments 1 and 2, which says that the PRP effect is due to a central bottleneck whereas the AB effect is due entirely to a separate bottleneck that occurs before the central bottleneck. If this dual-bottleneck account is correct, then the accuracy of an RSVP Task 2 should depend on the duration of early Task-1 stages, but should not depend on the duration of Task-1 central stages. In contrast, the UCB model says that Task-2 central operations cannot begin until Task-1 central operations have finished; hence the accuracy of T2 at short lags should depend very much on the duration of the Task 1 central stage. Specifically, the decrement in Task 2 accuracy should be deeper and longer-lasting when Task-1 central operations are prolonged. We refer to this as the "Task-1 carryover" prediction of the UCB model.

To evaluate the Task-1 carryover prediction, we manipulated whether or not Task 1 was speeded (Experiment 3). We reasoned that the requirement to make an immediate speeded response to Task 1 should add to the duration of the central operations relative to a condition where subjects could defer response selection. As predicted by the UCB model, Task 2 performance was indeed much worse when Task 1 was speeded than when Task 1 was unspeeded. Similar results have also been demonstrated by Jolicoeur (1998; 1999a).

The results of Experiment 3 clearly support the UCB model, but are open to several objections. Because the speeded and unspeeded Task 1 conditions of Experiment 3 were run in different blocks of trials, they may have differed in terms of arousal, effort, response criteria, or task preparation. The same criticism applies to Jolicoeur (1998, 1999a) as well. Differences in task preparation are especially serious, since greater preparation for a speeded Task 1 could result in less preparation (and thus worse performance) on Task 2. To eliminate these possible confounds, Experiment 4 replicated the Task-1 carryover effect using a within-block manipulation of Task-1 difficulty. The digits 1-4 were mapped in numerical order (i.e. compatibly) onto the four response keys, but the digits 5-8 were mapped in reverse numerical order (i.e. incompatibly) onto the same four response keys. Relative to the compatible condition, the incompatible Task 1 condition caused both a deeper and more prolonged decrement in Task-2 performance (i.e. a substantial Task-1 carryover effect). These results clearly contradict the dual-bottleneck account, which predicted that the duration of Task-1 central stages should not affect Task 2 accuracy. These Task-1 carryover effects, however, provide direct support the UCB model.

Locus of the Central Bottleneck

The UCB implies the existence of a set of demanding central processes that cannot operate concurrently. Although it clearly suggests that very early input processes and very late output processes do not belong to this set, the model does not specify exactly which operations in between do belong to this set. This issue is therefore left as an empirical matter. In the PRP design there is clear evidence that character identification does not belong to the set of demanding central operations (e.g. Pashler & Johnston, 1989), but that later processes such as response selection (e.g. Pashler, 1984; McCann & Johnston, 1992), memory retrieval (Carrier & Pashler, 1995), and mental rotation (Ruthruff et al. 1995) do belong to this set.

It is less clear what mental operations in the AB paradigm belong to the set of demanding central operations. As noted earlier, however, there are several lines of evidence indicating that character identification is not suppressed during the attentional blink (Johnston, Ruthruff, and McCann, 1999; Luck, Vogel, and Shapiro, 1996; Maki, Frigen, and Paulson, 1997; Shapiro, Caldwell, and Sorensen, 1997; Shapiro, Driver, Ward, and Sorensen, 1997; Vogel, Luck, and Shapiro, 1998). Therefore, the earliest operation that does belong to the set of demanding central operations must come after the process of character identification. Consistent with this inference, Chun and Potter (1995) have proposed that the AB effect is due to interference between the operations that make the categorical representations of the targets available for subsequent report (e.g. by consolidating those representation into short-term memory [STM]).

Jolicoeur and Dell'Acqua (1998; 1999) have provided particularly clear-cut support for the conjecture that consolidation into STM belongs to the set of demanding central processes. In one experiment (not unlike the AB/PRP design of Experiment 1), they presented a masked display of 1-3 characters, which subjects were to remember for a later report, followed after a variable SOA by a pure tone requiring an immediate speeded response. Jolicoeur and Dell'Acqua found that processing of the visual stimulus slowed the response to the tone at short SOAs. More importantly, the magnitude of the tone-task slowing was greater when three letters needed to be consolidated than when only one letter needed to be consolidated. This result is consistent with the claim that central operations required to make a speeded response to the tone are postponed until short-term consolidation of the visual stimulus has been completed. As noted by Jolicoeur and Dell'Aqua, the hypothesis that STM consolidation cannot proceed on two tasks at the same time provides a ready explanation for interference in the AB design as well (see also Jolicoeur, 1998, 1999a, 1999b).

Multiple Sources of Interference

The fact that we observed substantial interference in a hybrid AB/PRP design (Experiment 1) and a hybrid PRP/AB design (Experiment 2) supports the UCB model. However, the interference observed in these hybrid designs appears to be somewhat smaller than interference effects typically found in the standard PRP and AB designs. This finding raises the suspicion that a central bottleneck cannot fully account for interference in both designs. Consistent with this conjecture, Experiment 5 showed that performance of an RSVP Task 2 was impaired more when the Task 1 stimulus was presented in the same input modality. This effect cannot easily be explained by the UCB model, because the two tasks were explicitly designed so that they would require roughly equal amounts of central interference. Results similar to those of the present Experiment 5 have also been obtained by Arnell and Duncan (1999).

The fact that interference effects were greater in the within-modal condition than in the cross-modal condition of Experiment 5 suggests that RSVP tasks interfere at a central stage and at some earlier perceptual stage. For example, perceptual interference at short lags might result in categorical representations that are relatively weak. These representations might then decay even further because short-term consolidation of T2 must wait until short-term consolidation of T1 is complete. Thus, perceptual interference reduces the initial activation of the categorical representation of T2 and central interference then causes the activation to decay even further.

Further work is needed to better specify the nature of the non-central sources of interference in the AB paradigm and how they interact with the sources of central interference. For the present, however, it seems safe to conclude that although a central bottleneck plays an important role in both the AB and PRP paradigms, it is probably not the only source of interference.

Relation of the Central Bottleneck to other Attentional Limitations

The present data and those of Jolicoeur and colleagues suggest that AB and PRP effects are both due in large part to a central bottleneck. In this section we briefly discuss whether a central bottleneck might play a role in other well-known phenomena.

Covert Movements of Spatial Attention. People are capable of choosing particular regions of a visual display and processing stimuli in these regions more intensively than other stimuli. Furthermore, they are capable of rapidly altering which stimuli are attended and which are unattended without moving the eyes (Helmholtz, 1924, p. 455; James, 1890). Two lines of evidence suggest that these covert movements of visual attention are not subject to the central bottleneck constraint.

Pashler (1991), for example, presented a tone stimulus followed after a variable SOA by a masked visual stimulus consisting of an array of eight letters and a bar probe. Subjects made a speeded response to the tone, then indicated (at their leisure) the identity of the letter located next to the bar probe. If subjects were unable to allocate covert attention to the probed item until after they had selected a response to the tone then, at short SOAs, attentional allocation would not have occurred until well after the display had been masked. This delay in the allocation of attention should have catastrophic effects on performance, as demonstrated by a control condition in which the <u>presentation</u> of the bar probe was delayed until after the display had been masked. Contrary to this prediction, letter task accuracy was very high at all SOAs, indicating that spatial attention was in fact allocated to the probed item well before tone-task response selection had been completed. Hence, Pashler concluded that movements of spatial attention are not limited by the central bottleneck.

Johnston, McCann, and Remington (1995) arrived at a similar conclusion regarding movements of spatial attention. These authors used locus-of-slack logic to show that the processing stage(s) affected by a spatial cueing manipulation come <u>at or before</u> stimulus identification, whereas the processing stage(s) that cause interference in the PRP paradigm come <u>after</u> stimulus identification. Thus it appears that the locus of the spatial attention effect is earlier than the locus of the central bottleneck.

Repetition Blindness. When the same target is presented twice in an RSVP stream, subjects often miss the second presentation of the target (e.g. Kanwisher, 1987). This effect, known as repetition blindness (or "**RB**"), occurs even when the RSVP stream consists of words and failure to detect a repeated word makes a sentence ungrammatical. Some have attributed RB to a failure of token individuation; in other words, subjects fail to establish a separate token for the repeated target (e.g. Kanwisher, 1987). This view can be incorporated into the UCB model by assuming that token individuation belongs to the set of demanding central operations; hence, token individuation of the second instance of a target is postponed until token individuation of the first instance of that target has been completed.

On the other hand, a variety of evidence suggests that failures in retrieval may play a key role in RB (Armstrong & Mewhort, 1995; Fagot & Pashler, 1995; Whittlesea, Dorken, & Podrouzek, 1995; Whittlesea & Wai, 1997). Fagot and Pashler (1995), for instance, argued that repeated items are perceived just as clearly (or unclearly) as non-repeated items, but that subjects are biased against reporting the repeated items. In favor of this account, Fagot and Pashler showed that RB effects are greatly reduced or eliminated in designs that minimized the subject's memory load (but see also Hochhaus & Johnston, 1996). Further, they were able to show that RB is sensitive to the type of retrieval cue presented after the RSVP stream has ended, directly implicating retrieval failure rather than perceptual failure as the cause of RB. These data argue against a central bottleneck account of the RB effect. In addition, Chun (1997) has shown that (a) AB and RB have different time courses, (b) RB is found in cases where AB is absent, and (c) AB is found in cases where RB is absent. Given the evidence in this chapter that the AB effect is due in part to a central bottleneck, Chun's results would appear to directly indicate that RB is not due to a central bottleneck.

Summary

This chapter examined the possibility that the AB effect (interference between two unspeeded RSVP tasks) and the PRP effect (interference between two speeded-response tasks) are attributable to the same central bottleneck. If this unified central bottleneck model is correct, then substantial dual-task interference should occur in hybrid AB/PRP designs where Task 1 is an unspeeded RSVP task and Task 2 is a speeded-response task, and vice versa. Experiments 1 and 2 confirmed this prediction. Furthermore, Experiments 3 and 4 showed that prolongation of the Task 1 central stage caused a deeper and longer-lasting interference effect on an unspeeded RSVP Task 2. This "Task-1 carryover" effect directly supports the assertion that central stages on Task 1 postpone the critical operations needed to form a reportable percept of the Task 2 RSVP target. Although the AB effect appears to be due in part to a central bottleneck, it appears to be due to specifically visual interference as well. Experiment 5 demonstrated substantially greater impairment of an RSVP Task 2 when the Task 1 stimulus was visual than when the Task 1 stimulus was auditory, even though both Task 1 judgments required similar amounts of central processing. Whereas the present chapter showed that a central bottleneck model can account (at least in part) for interference in two very different types of dual-task designs (AB and PRP), the central bottleneck appears not to provide a satisfactory account of certain other dual-task phenomena, such as repetition blindness. Furthermore, the central bottleneck does not appear to play a major role in covert movement of spatial attention across the visual field.

References

Armstrong, I. T. & Mewhort, D. J. (1995). Repetition deficit in rapid-serial-visual-presentation displays: Encoding failure or retrieval failure? *Journal of Experimental Psychology: Human Perception and Performance,* **21**, 1044-1052.

Arnell, K. M. & Duncan, J. (1999). Specific and general sources of dual-task cost in stimulus identification and response selection. Manuscript submitted for publication.

Broadbent, D. E., & Broadbent, M. H. P. (1987). From detection to identification: Response to multiple targets in rapid serial visual presentation. *Perception & Psychophysics,* **42**, 105-113.

Carrier & Pashler (1995). Attentional limitations in memory retrieval. *Journal of Experimental Psychology: Learning, Memory & Cognition,* **21**, 1339-1348.

Chun, M. M. (1997). Types and tokens in visual processing: A double dissociation between the attentional blink and repetition blindness. *Journal of Experimental Psychology: Human*

Perception and Performance, **23**, 738-755.

Chun, M. M. & Potter, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance*, **21**, 109-127.

Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General*, **113**, 501-517.

Fagot, C. & Pashler, H. E. (1992). Making two responses to a single object: Implications for the central attentional bottleneck. *Journal of Experimental Psychology: Human Perception and Performance*, **18**, 1058-1079.

Fagot, C. & Pashler, H. E. (1995).Repetition blindness: Perception or memory failure? *Journal of Experimental Psychology: Human Perception and Performance*, **21**, 275-292.

Gottsdanker, R. (1980). The ubiquitous role of preparation. In G. E. Stelmach and J. Requin (Eds.), *Tutorials in motor behavior* (pp. 315-371). Amsterdam: North-Holland.

Greenwald, A. (1972). On doing two things at once: Time sharing as a function of ideomotor compatibility. *Journal of Experimental Psychology*, **94**, 52-57.

Greenwald, A. & Shulman, H. G. (1973). On doing two things at once: II. Elimination of the psychological refractory period effect. *Journal of Experimental Psychology*, **11**, 72-89.

Helmholtz, H. (1924). *Helmholtz's Treatise on Physiological Optics*, translated from the 3d German edition. Edited by James P. C. Southall. Rochester, NY: The Optical Society of America.

Hochhaus, L., & Johnston, J. C. (1996). Perceptual repetition blindness effects. *Journal of Experimental Psychology: Human Perception and Performance*, **22**, 355-366.

James, W. (1890/1950). *The Principles of Psychology*, Vol 1. New York: Dover.

Johnston, J. C. & Delgado, D. F. (1993). Bypassing the single-channel bottleneck in dual-task performance. Paper presented at the annual meeting of the Psychonomic Society, Washington, D.C.

Johnston, J. C., McCann, R. S., & Remington, R. W. (1995). Chronometric evidence for two types of attention. *Psychological Science*, **5**, 365-369.

Johnston, J. C., Ruthruff, E., & McCann, R. S. (1999). Chronometric evidence for a late selection locus to the attentional blink. Unpublished manuscript.

Jolicoeur, P. (1998). Modulation of the attentional blink by on-line response selection: Evidence from speeded and unspeeded Task$_1$ decisions. *Memory & Cognition*, **26**, 1014-1032.

Jolicoeur, P. (1999a). Concurrent response selection demands modulate the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance*, in press.

Jolicoeur, P. (1999b). Restricted attentional capacity between sensory modalities. *Psychonomic Bulletin & Review*, **6**, 87-92.

Jolicoeur, P., & Dell'Acqua (1998). The demonstration of short-term consolidation. *Cognitive Psychology*, **36**, 138-202.

Jolicoeur, P., & Dell'Acqua (1999). Attentional and structural constraints on visual encoding. *Psychological Research*, in press.

Kanwisher (1987). Repetition blindness: Type recognition without token individuation. *Cognition*, **27**, 117-143.

Luck, S. J.,Vogel, E. K., & Shapiro, K. L. (1996). Word meanings can be accessed but not reported during the attentional blink. *Nature*, **383**, 616-618.

Maki, W. S., Frigen, K., & Paulson, K. (1997). Associative priming by targets and distractors during rapid serial visual presentation: Does word meaning survive the attentional blink? *Journal of Experimental Psychology: Human Perception and Performance*, **23**, 1014-1034.

McCann, R. S., & Johnston, J. C. (1992). Locus of the single-channel bottleneck in dual-task interference. *Journal of Experimental Psychology: Human Perception and Performance*, **18**, 471-484.

Pashler, H. E. (1984). Processing stages in overlapping tasks: Evidence for a central bottleneck. *Journal of Experimental Psychology: Human Perception and Performance*, **10**, 358-377.

Pashler, H. E. (1989). Dissociation and dependencies between speed and accuracy: Evidence for a two-component theory of divided attention in simple tasks. *Cognitive Psychology*, **21**, 469-514.

Pashler, H. E. (1991). Shifting visual attention and selecting motor responses: Distinct attentional mechanisms. *Journal of Experimental Psychology: Human Perception and Performance*, **17**, 1023-1040.

Pashler, H. E., Carrier, M., & Hoffman, J. (1993). Saccadic eye movements and dual-task interference. *Quarterly Journal of Experimental Psychology*, **46A**, 51-82.

Pashler, H. E. & Johnston, J. C. (1989). Chronometric evidence for central postponement in temporally overlapping tasks. *Quarterly Journal of Experimental Psychology, Section A: Human Experimental Psychology*, **41**, 19-45.

Pashler, H. E. & Johnston, J.C. (1998). Attentional limitations in dual-task performance. In H. Pashler (Ed.) *Attention* (pp. 155-189). Hove, UK: Psychology Press.

Raymond, J. E., Shapiro, K., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance*, **18**, 849-860.

Ruthruff, E., Miller, J. O., & Lachmann, T. (1995). Does mental rotation require central mechanisms? *Journal of Experimental Psychology: Human Perception and Performance*, **21**, 552-570.

Ruthruff, E., Johnston, J. C., & Van Selst, M. A. (in press). Why practice reduces the dual-task interference. *Journal of Experimental Psychology: Human Perception and Performance*.

Shapiro, K., Raymond, J. E., & Arnell, K. M. (1994). Attention to visual pattern information produces the attentional blink in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance*, **20**, 357-371.

Shapiro, K. L., Caldwell, J., & Sorensen, R. E. (1997). Personal names and the attentional blink: A visual “cocktail party” effect. *Journal of Experimental Psychology: Human Perception and Performance*, **23**, 504-514.

Shapiro, K., Driver, J., Ward, R. & Sorensen, R. E. (1997). Priming from the attentional blink: A failure to extract visual tokens not visual types. *Psychological Science*, **8**, 95-100.

Van Selst, M. A., Ruthruff, E., & Johnston, J. C. (in press). Can practice eliminate the Psychological Refractory Period effect? *Journal of Experimental Psychology: Human Perception and Performance*.

Vince, M. A. (1948). The intermittency of control movements and the psychological refractory period. *British Journal of Psychology*, **38**, 149-157.

Vogel, E. K., Luck, S. J., & Shapiro, K. L. (1998). Electrophysiological evidence for a postperceptual locus of suppression during the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance*, **24**, 1656-1674.

Welford, A. T. (1952). The "psychological refractory period" and the timing of high-speed performance – A review and a theory. *British Journal of Psychology*, **43**, 2-19.

Whittlesea, B. W. A., Dorken, M. D., & Podrouzek, K. W. (1995). Repeated events in rapid lists: Part I. Encoding and representation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **21**, 1670-1688.

Whittlesea, B. W. A., Wai, K. H. (1997). Reverse "repetition blindness" and release from "repetition blindness": Constructive variations on the "repetition blindness" effect. *Psychological Research*, **60**, 173-182.

Author Notes

This work was supported by a National Research Council postdoctoral fellowship to Eric Ruthruff and by grants from the National Institute of Mental Health (1-R01-MH45584) and the National Science Foundation (SBR #9729778). We thank James C. Johnston, Robert McCann, Roger Remington, Valerie Lewis, Beth Siegel, Julie Creagh, and Mark Van Selst for useful discussions.

TASK 1

		Speeded-Response	**Unspeeded RSVP**
TASK 2	**Speeded-Response**	PRP	Hybrid AB/PRP
	Unspeeded RSVP	Hybrid PRP/AB	AB

Table 1: Dual-task designs created by pairing speeded-response and/or unspeeded RSVP tasks. PRP = Psychological Refractory Period design; AB = Attentional Blink design.

		LAG							
		1	**2**	**3**	**4**	**5**	**6**	**7**	**8**
Exp. 1	**UV**	-----	97.4	88.6	88.6	89.3	87.7	90.6	87.0
Exp. 2	**SA**	95.1	97.0	95.1	94.8	95.3	96.0	93.8	95.3
Exp. 3	**SV**	75.8	79.5	78.9	81.3	81.9	81.3	79.7	81.8
	UV	76.5	79.2	80.9	80.5	79.5	80.0	80.6	80.1
Exp. 4	**SV-C**	93.5	95.2	95.0	93.5	92.9	94.4	95.0	90.2
	SV-I	90.2	89.9	90.2	88.1	89.0	91.4	88.7	89.3
Exp. 5	**SV**	93.5	93.8	92.9	94.4	95.0	93.8	95.9	94.7
	SA	89.9	90.5	92.6	93.2	91.4	89.9	90.5	89.3

Table 2 : Task 1 percent correct as a function of lag for each experiment. S = Speeded; U = Unspeeded; V = Visual; A = Auditory; C = Compatible; I = Incompatible.

Figure Captions

Figure 1: (a) Typical data from a Psychological Refractory Period (PRP) experiment as a function of the stimulus onset asynchrony (SOA). (b) Typical data from an Attentional Blink (AB) experiment as a function of the lag between the first and second targets.

Figure 2: The central bottleneck model. Processing on each task has been divided into three super-stages (A, B, and C) each of which might be further divisible into substages. Stage B is the bottleneck stage, meaning that while stage B on Task 1 (B1) is underway, stage B on Task 2 (B2) is postponed.

Figure 3: Task 2 response time in Experiment 1 as a function of lag.

Figure 4: Task 2 percent correct and Task 1 response time in Experiment 2 as a function of lag. Data from trials in which the Task 1 stimulus was present are represented by squares. Data from trials in which the Task 1 stimulus was absent are represented by triangles.

Figure 5: Task 2 percent correct and Task 1 response time in Experiment 3 as a function of lag. The speeded Task 1 data are represented by squares. The unspeeded Task 1 data are represented by diamonds.

Figure 6: Task 2 percent correct and Task 1 response time in Experiment 4 as a function of lag. Data from the compatible Task 1 condition are represented by squares. Data from the incompatible Task 1 condition are represented by circles.

Figure 7: Task 2 percent correct and Task 1 response time in Experiment 5 as a function of lag. Data from the letter Task 1 condition are represented by squares. Data from the tone Task 1 condition are represented by the 'X' symbols.

---

[1] The PRP effect sometimes fails to occur for speeded tasks with extremely high stimulus-response compatibility, such as fixating a visual stimuli (Pashler, Carrier, & Hoffman, 1993) or zero-order stick-tracking (Johnston & Delgado, 1993; see also Greenwald, 1972; Greenwald and Shulman, 1973). In addition, the PRP effect can in some cases be greatly reduced with extensive task practice (Ruthruff, Johnston, & Van Selst, in press; Van Selst, Ruthruff, & Johnston, in press).

[2] Dual-task slowing at short SOAs might also be due in part to reduced task preparation or perceptual interference.

[3] It is conceivable that subjects do select an immediate on-line response to the Task 1 stimulus, even though response selection could in principle be deferred until after the RSVP stream has ended. However, it seems unlikely that Task-1 response selection is responsible for the AB effect. Johnston et al. (1999) found robust AB effects even in a condition where the response could not be selected on-line because the response mapping was not revealed until well after the RSVP stream had ended.

[4] One obvious way to test the early bottleneck model would be to find a Task 1 judgment that does not contain this early stage and then see if such a Task 1 interferes with an unspeeded RSVP Task 2. Unfortunately this approach is not feasible. We cannot specify what stage causes the putative early bottleneck, with sacrificing generality; hence we cannot find a task that lacks an unspecified stage (if one even exists).