Chapter in press.

The Attentional Blink Bottleneck

PIERRE JOLICŒUR
University of Waterloo

ROBERTO DELL'ACQUA
University of Padova

and

JACQUELYN CREBOLDER
Defence and Civil Institute of Environmental Medicine

Evidence suggesting that the attentional blink (AB) phenomenon is caused by a central processing bottleneck is reviewed. A bottleneck model of the psychological refractory period paradigm is used to motivate four major predictions concerning the patterns of results expected in dual-task experiments. This model, which was designed to explain results from experiments in which both tasks are speeded is modified to make predictions for experiments in which one task, or the other, is not speeded. This extention of the bottleneck model allows us to test four major predictions in the context of the AB paradigm. The bulk of the evidence is consistent with these predictions, and so we conclude that the AB phenomenon is caused by a processing bottleneck. The evidence also provides some contraints on the possible locus of the bottleneck in the information processing stream.

The AB phenomenon can be succinctly defined as an increase in the difficulty of reporting a second (masked) target that follows (after a short delay) a first target that required immediate processing. In this chapter we review evidence supporting the view that the attentional blink (AB) phenomenon results from a bottleneck in the information processing stream required to perform tasks designed to reveal the AB phenomenon.

The chapter is organized into six sections. In the first four we describe a major prediction of bottleneck models of dual-task interference (Pashler & Johnston, 1989), followed by relevant evidence from AB and related paradigms. In the fifth section we review evidence pertinent to the issue of the locus of the bottleneck in the AB phenomenon. The last section presents some conclusions.

Carry-over to Task$_2$ of pre-bottleneck and bottleneck Task$_1$ manipulations

Prediction of Postponement Model

Figure 1 shows two sets of stage diagrams that illustrate one of the major predictions of one class of bottleneck models of dual-task interference. These diagrams are meant to capture task interactions in experiments designed to investigate the psychological refractory period (PRP). In these experiments, both tasks are speeded. In this class of models it is assumed that one stage of processing constitutes a processing bottleneck. This stage is labeled B in Figure 1. This stage can only process information from one information processing stream at a time. If the bottleneck stage is occupied by the processing required for one task, then it is unavailable for the processing in another task. Stages before (labeled A in the figure) or after (labeled C) the bottleneck can go on in parallel with other stages of processing in a concurrent information processing stream. For example, as shown in the figure, pre-bottleneck processing required for Task$_2$ (A_2) can overlap pre-bottleneck (A_1) and bottleneck (B_1) processing required for Task$_1$.

Processing at the bottleneck stage in Task$_2$ (B_2), however, cannot begin while the bottleneck is busy with Task$_1$. A period of waiting in the information processing stream for Task$_2$ will result if pre-bottleneck processing

---

We thank Marg Ingleton for technical support. This work was supported by research grants from NSERC and HFSP. Send correspondance to: Pierre Jolicœur, Department of Psychology, University of Waterloo, Waterloo, Ontario, Canada, N2L 3G1. Phone: 519-885-1211 x2142; FAX: 519-746-8631; email: pjolicoe@cgl.uwaterloo.ca; web: http://www.cgl.uwaterloo.ca/~pjolicoe

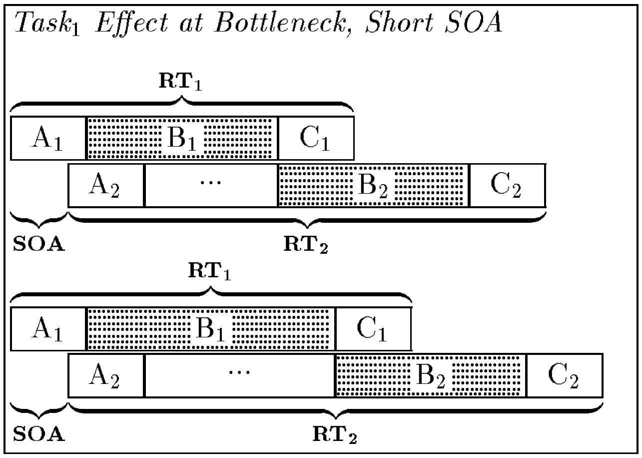

Figure 1. Stage diagrams showing the predicted task interactions in the PRP paradigm for a manipulation in Task$_1$ that affects the bottleneck stage. At short SOA the onset of processing at the bottleneck stage in Task$_2$ is delayed by processing at the bottleneck in Task$_1$. This delay is longer when Task$_1$ processing at the bottleneck is longer.

finishes before the bottleneck is available. The probability and duration of such a period of waiting, called slack in scheduling theory, increases as the stimulus onset asynchrony (SOA) between the target stimulus for Task$_1$ and the target stimulus for Task$_2$ is reduced. We call the first target T$_1$ and the second target T$_2$.

As can be seen in Figure 1 across the top and bottom pair of stage diagrams, a longer duration of processing at stage B$_1$ will increase reaction times in Task$_2$. We refer to this dual-task interaction as a carry-over from Task$_1$ to Task$_2$. Clearly, the model predicts that this carry-over will only occur at short SOAs. Furthermore, carry-over would also be expected if the duration of pre-bottleneck processing in Task$_1$ was varied.

This class of models can be applied to the AB paradigm in the following way. Rather than measuring RT in Task$_2$, experimenters investigating the AB effect measure accuracy in Task$_2$ to a briefly-presented and masked target stimulus (T$_2$). Figure 2 illustrates how the model in Figure 1 can be extended to the AB paradigm. We suppose that the representation of T$_2$ decays (or is degraded by the mask) during the period of waiting. This is represented in the figure by the falling sequence of dots during the period of slack. We also assume that a longer period of waiting allows more decay to take place before the representation of T$_2$ can be stabilized by processing at the bottleneck stage (Arnell & Jolicœur, 1999; Chun & Potter, 1995; Dell'Acqua & Jolicœur, 2000; Jolicœur, 1998, 1999bc; Jolicœur & Dell'Acqua, 1998, 1999; Ross & Jolicœur, 1999). With these assumptions, it is clear that a carry-over from a manipulation affecting the duration of bottleneck processing in Task$_1$ to accuracy in Task$_2$ is predicted by the model. More waiting, and hence more decay of T$_2$, would occur as the duration of bottleneck or pre-bottleneck processing in Task$_1$ is increased. As in the PRP paradigm, this carry-over is predicted to occur only at short SOAs.

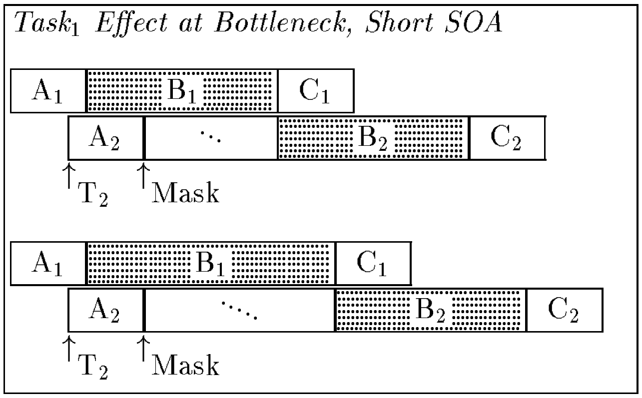

Figure 2. Stage diagrams showing the predicted task interactions in the attentional blink paradigms for a manipulation in Task$_1$ that affects the bottleneck stage. At short SOA the onset of processing at the bottleneck stage in Task$_2$ is delayed by processing at the bottleneck in Task$_1$. This delay is longer when Task$_1$ processing at the bottleneck is longer. The representation of T$_2$ decays or is overwritten during the period of waiting if T$_2$ is masked; more decay results from a longer waiting period.

Tests of the Carry-over Prediction

Several tests have confirmed the Task$_1$ carry-over to Task$_2$ predicted by postponement models in the AB paradigm. We consider first the experiments of Crebolder and Jolicœur (1999). They presented upper-case letters using rapid serial visual presentation (RSVP). T$_1$ was an H, O, or S, presented in red; T$_2$ was an X or a Y, presented in white; the other (filler) letters were all white; the background was black. T$_2$ was presented in every trial, but T$_1$ was presented only on some trials. Task$_1$ was to report the identity of T$_1$, when T$_1$ was presented, or to press the space bar, when T$_1$ was not presented. Task$_2$ was to report the identity of T$_2$.

The manipulation of greatest interest in these experiments concerned the relative frequency of occurrence of the three possible letters used as targets in Task$_1$. For any given subject, one letter was shown 9 times in every 14 T$_1$-present trials, one was shown 4 times, and one was shown once. The assignment of letters to probability conditions was counterbalanced across subjects.

The results from two AB experiments are shown in Figure 3. In both cases, the magnitude of the AB effect was largest for the least frequent T_1, intermediate for the intermediate-frequency T_1, and smallest for the most frequent T_1. The pattern of results was found both when $Task_1$ required a deferred unspeeded response (middle panel), as in the usual AB paradigm, and also when $Task_1$ was performed with a speeded response (top panel). $Task_2$ was always unspeeded. Results from the experiment in which $Task_1$ was speeded confirmed the expectation that RT_1 should increase as the relative frequency of T_1 decreased (Kornblum, 1968). The results shown in the bottom panel are discussed in a subsequent section.

Jolicœur, Dell'Acqua, and Crebolder (2000) found clear-cut evidence for carry-over in a cross-modal AB experiment in which $Task_1$ required either a speeded two-alternative or four-alternative discrimination based on the pitch of a tone presented concurrently with an RSVP stream that contained T_2 (X vs. Y). As can be seen in Figure 4, a larger AB effect was found when the duration of central processing in $Task_1$ was lengthened.

Jolicœur (1999c) also manipulated $Task_1$ variables in AB paradigms. In Experiment 1, $Task_1$ required either a two-alternative discrimination or a four-alternative discrimination, with a deferred unspeeded response. In Experiment 2, the same discriminations were made, but with a speeded response. $Task_2$ was the same X–Y two-alternative discrimination as in the Crebolder and Jolicœur (1999) experiments. An AB effect was found in both experiments. A larger AB effect was found when $Task_1$ required a four-alternative discrimination than when it required a two-alternative discrimination, but only when $Task_1$ was speeded. Equivalent AB effects in the two-alternative discrimination and four-alternative discrimination versions of $Task_1$ were found when $Task_1$ was unspeeded, as reported by Ward, Duncan, and Shapiro (1996) who had also found equivalent interference on $Task_2$ from a two-alternative and a four-alternative $Task_1$ in a modified AB paradigm. Jolicœur (1999c) argued that response selection in $Task_1$ was deferred when the response was deferred, but that it could not be deferred when a speeded response was required in $Task_1$. Interference between $Task_1$ and $Task_2$ is only expected if the capacity-demanding processing required to perform each of these two tasks have the potential to overlap temporally. For these tasks and stimuli, it appears that subjects could defer one of the capacity-demanding operations (response selection) when $Task_1$ was not speeded.

The effects of requiring a speeded response in $Task_1$ on the carry-over effect highlights the importance of ensuring that manipulations in $Task_1$ affect processing taking place concurrently with $Task_2$ processing. It is possible that some previous failures to find carry-over effects were due to an ability to schedule capacity-demanding operations after the time-critical processing in $Task_2$ had run to completion (e.g., Raymond, Shapiro, & Arnell, 1995; Ward et al., 1996, 1997).

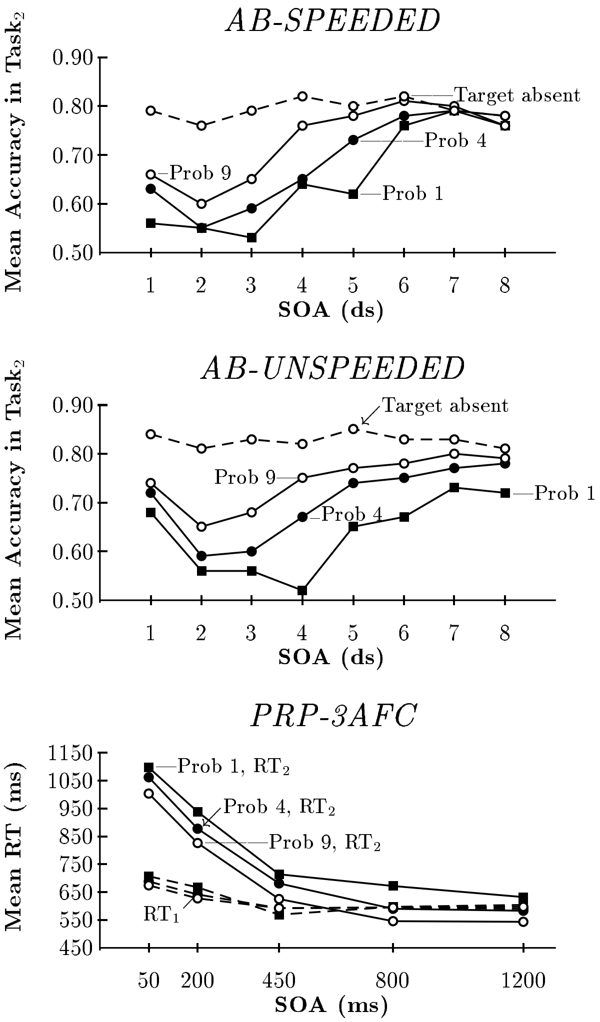

Figure 3. Results from Crebolder and Jolicœur (1999). Top Panel: Mean accuracy in $Task_2$ when $Task_1$ was speeded; for control trials (dashed line) and experimental trials (solid lines) for each level of relative signal frequency of T_1, as a function of T_1–T_2 SOA. Middle Panel: results with $Task_1$ not speeded. Bottom Panel: Mean response time in $Task_1$ and $Task_2$ of a PRP experiment in which relative signal probability of T_2 was manipulated in $Task_2$.

In a third experiment, $Task_1$ was either a simple RT task, or a speeded two-alternative discrimination based on letter identity (H vs. S; Jolicœur, 1999c). A much larger AB effect was caused by the two-alternative discrimination task than by the simple RT task, providing another demonstration of the carry-over of a $Task_1$ manipulation on $Task_2$ performance in the AB paradigm. These results are shown in Figure 5.

Jolicœur (1998) provided additional converging evidence simply by comparing the magnitude of the AB effect across conditions in which $Task_1$ was unspeeded or speeded. At short SOAs, larger deficits were found

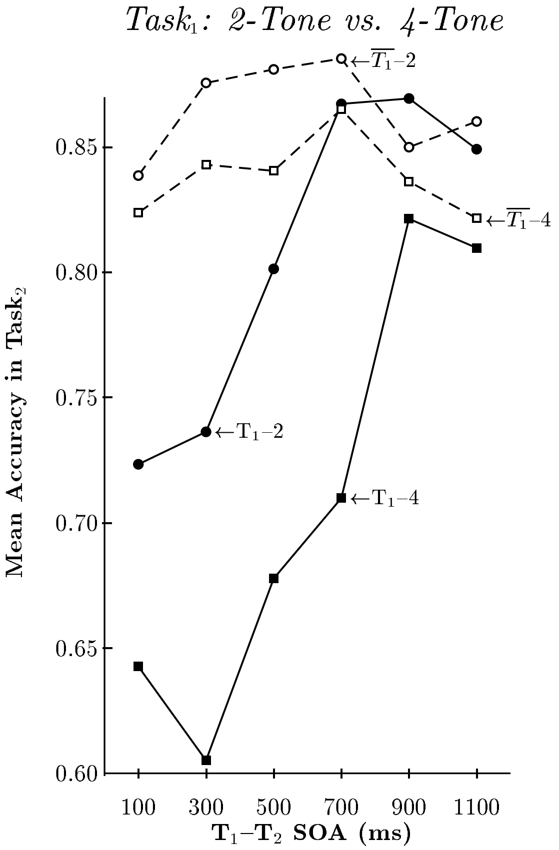

Figure 4. Results from Jolicœur, Dell'Acqua, and Crebolder (1999). Mean accuracy in $Task_2$ for control trials (unfilled symbols, dashed lines) vs. experimental trials (filled symbols, solid lines) as a function of T_1–T_2 SOA, for two levels of $Task_1$ difficulty ($Task_1$ had either two response alternatives (T_1–2) or four response alternatives (T_2–4).

in $Task_2$ when $Task_1$ was speeded than when it was unspeeded. If we assume that a longer period of central processing was required to perform $Task_1$ when it was speeded, then the results provide additional evidence of carry-over.

In summary, several experiments in which a longer period of central processing was required to process T_1 while T_2 was encoded have confirmed the carry-over prediction of postponement models in the context of the AB paradigm (Crebolder & Jolicœur, 1999) as well as in AB paradigms modified by requiring a speeded response in $Task_1$ (Jolicœur, 1998, 1999bc; Jolicœur et al., 2000).

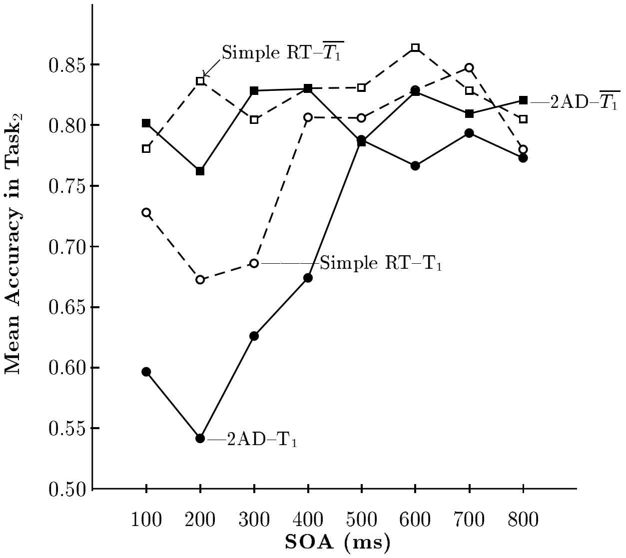

Figure 5. Results from Jolicœur (1999c). Mean accuracy in $Task_2$ for control trials (square symbols) vs. experimental trials (round symbols) as a function of T_1–T_2 SOA, for two levels of $Task_1$ difficulty ($Task_1$ required either a simple RT response, or a speeded two-alternative discrimination (2AD).

Correlation between $Task_1$ and $Task_2$ performance at short SOA

Prediction of Postponement Model

Figure 1 can also be used to understand a major prediction of bottleneck models of dual-task interference. As in previous sections, we begin with an application of bottleneck models to the PRP paradigm. At short SOAs between T_1 and T_2, variations in the duration of processing in the bottleneck stage (or at earlier stages) in $Task_1$ should cause corresponding variations in the duration of processing in $Task_2$. This is illustrated in the figure by supposing that the top and bottom pairs of stage diagrams represent different trials in which the duration of $Task_1$ bottleneck processing was either short (top pair) or long (bottom pair) due to randomly-occurring variance in the duration of processing. On the assumption that a significant proportion of the variance in $Task_1$ processing duration is due to variance in the bottleneck stage (Pashler, 1994), the model leads to the prediction of a positive correlation between $Task_1$ RTs and $Task_2$ RTs. This correlation should be larger at shorter SOAs than at longer SOAs. This prediction has been confirmed in several experiments for the PRP paradigm (Pashler, 1994).

Tests of Correlation Prediction

Jolicœur and his colleagues have examined whether a similar prediction holds for the dual-task interference observed in the AB paradigm. In order to estimate the duration of processing in $Task_1$, the paradigm was modified slightly by requiring an immediate and speeded response in $Task_1$, rather than a deferred and unspeeded response, as had typically been used in previous work. As stated in the foregoing section, we suppose that the representation of T_2 decays or is overwritten by subsequent stimulation (Chun & Potter, 1995; Giesbrecht & Di Lollo, 1998) if it is not consolidated. The model in Figure 2 leads to the prediction that a longer period of processing in Task$_1$ should be associate with a longer period of decay and hence with lower accuracy in Task$_2$.

The most direct way to estimate the duration of processing in Task$_1$ is to require an immediate and speeded response and to measure RT in Task$_1$ (RT$_1$). Jolicœur (1998) was the first to examine the effects of requiring a speeded response in Task$_1$ of the AB paradigm, which made it possible to test the prediction of postponement models that RT$_1$ and Task$_2$ accuracy should be negatively correlated at short SOAs, and that this relationship should attenuate at longer SOAs. Results from Jolicœur (1998, Experiments 2 and 3) are shown in Figure 6. Task$_1$ was to decide whether a red letter, presented in an RSVP stream of white letters, was an H or an S. The response in Task$_1$ was a button press and it was to be made as quickly as possible, while keeping errors to a minimum. Task$_2$ was to decide whether the RSVP stream contained an X or a Y, and the response in Task$_2$ was unspeeded and deferred until the end of the trial. Response times in Task$_1$ were first sorted based on RT$_1$, for each subject and each SOA, and aggregated into four bins defined by RT$_1$ quartiles. For each bin, accuracy in Task$_2$ was computed, and the resulting means are shown in Figure 6. Accuracy in Task$_2$ was highest for trials that were in the first RT$_1$ quartile (labeled Q$_1$) and lowest for trials in the fourth RT$_1$ quartile (Q$_4$), as expected from a postponement process, and this relationship between RT$_1$ and Task$_2$ accuracy was not found at the longest SOA, where we expect the relationship to be weaker.

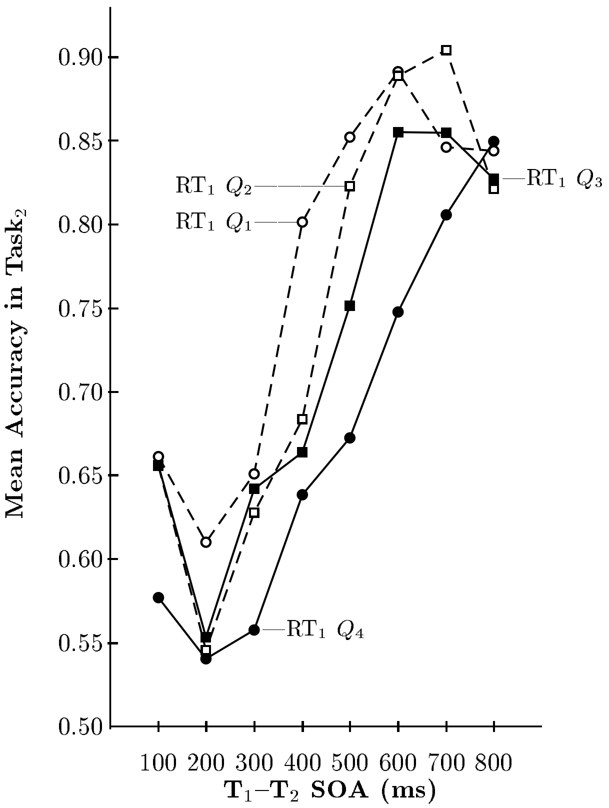

Figure 6. Results from Jolicœur (1998). Mean accuracy in Task$_2$ depending on the speed of the response in Task$_1$ and on T$_1$–T$_2$ SOA. First quartile (Q$_1$) RT$_1$s were the shortest.

This relationship between speed of processing in Task$_1$ and the magnitude of the AB effect has been found in numerous experiments in which a direct measure of processing time in Task$_1$ has been available (Jolicœur, 1998, Experiments 2 and 3; Jolicœur, 1999a, Experiments 1–2; Jolicœur, 1999c, Experiments 2 and 3; Jolicœur et al., 2000, Experiment 1; Ross & Jolicœur, 1999, Experiments 1–5;). A similar relationship was also found in related experiments designed to provide converging evidence with experiments using the AB paradigm (see Dell'Acqua & Jolicœur, 2000; Jolicœur & Dell'Acqua, 1999).

The relationship between RT$_1$ and Task$_2$ accuracy is correlational in nature, and as such it is open to alternative interpretations. However, this correlation was predicted by postponement models of dual-task interference, and the absence of this correlation would have disconfirmed this class of models. The available published studies, however, provide consistent and strong confirmations of the prediction.

Absorption of pre-bottleneck Task$_2$ effect

Prediction of Postponement Model

Next we consider what is probably the most striking prediction of postponement models of dual-task interference: The absorption into slack of a Task$_2$ pre-bottleneck effect. Figure 7 illustrates the key concepts behind this prediction. It is assumed that a variable affecting the duration of stage A$_2$ (a pre-bottleneck stage in Task$_2$ processing) has been manipulated. As shown in the figure, although the duration of stage A$_2$ has been affected by the manipulation, this has no effect on RT$_2$. The additional processing at stage A$_2$ has been absorbed into the period of slack. At long SOAs, however, there is no period of slack (no postponement), and so the effect of the manipulation is observed in the RT$_2$s. Thus, postponement models predict that effects of a variable altering the duration of a pre-bottleneck stage in Task$_2$ should be largest at longer SOAs and should become increasing smaller as SOA is reduced. This interaction is often referred to as an underadditive interaction as SOA is reduced. This prediction of postponement models has been verified several times in the context of the PRP paradigm (e.g., Pashler & Johnston, 1989; Pashler, 1994; Van Selst & Jolicœur, 1997).

Tests of Absorption Into Slack Prediction

Jolicœur (1999d) performed an experiment to de-

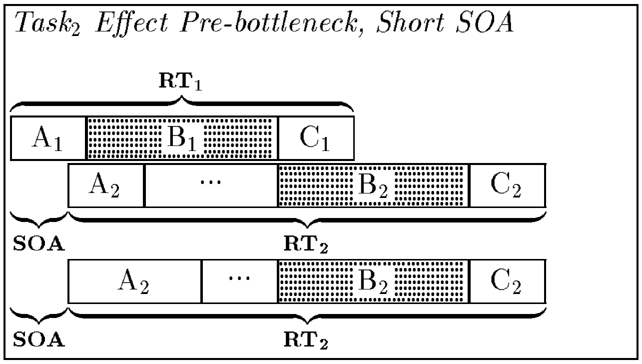

Figure 7. Stage diagrams showing the predicted task interactions in the PRP paradigm for a manipulation in Task$_2$ that affects a pre-bottleneck stage. At short SOA the onset of processing at the bottleneck stage in Task$_2$ is determined by when the bottleneck is freed by Task$_1$ rather than by when Task$_2$ pre-bottleneck processing runs to completion. The Task$_2$ manipulation has no effect on RT$_2$.

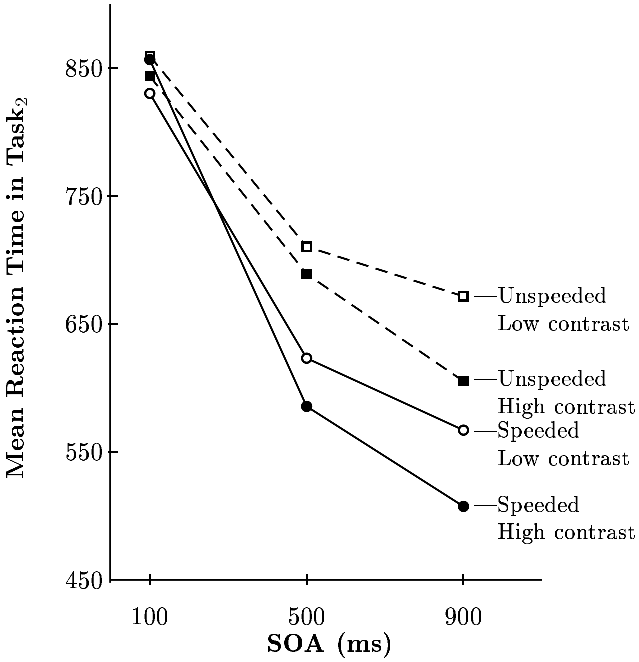

Figure 8. Results from Jolicœur (1999d). Mean reaction time, in milliseconds, for Task$_2$ responses depending on the SOA between T$_1$ and T$_2$, on whether the response in Task$_1$ was speeded or unspeeded, and on the contrast of T$_2$.

termine whether the type of encoding task that triggers an AB effect in the AB paradigm would also trigger a form of dual-task interference that could be characterized as a bottleneck. The logic of the experiment was to begin with a classic AB paradigm, with a typical Task$_1$, but with Task$_2$ modified in such a way to allow for a test of the absorption into slack prediction. The first task, not speeded, required visual encoding, and had been used in previous experiments to investigate the AB phenomenon. T$_1$ was either an H, an O, or an S, presented in red in a stream of white letters shown using RSVP, and was reported after the end of the RSVP stream. Task$_2$ was a speeded discrimination between two possible T$_2$s (X vs. Y). T$_2$ was always the last item in the RSVP stream. In addition, the contrast of T$_2$ was manipulated (two levels: normal contrast vs. low contrast). Response times in Task$_2$ were measured. The most important results concern the pattern of interaction of this contrast manipulation in Task$_2$ with SOA. This manipulation, in the context of the PRP paradigm, has produced underadditive interactions with decreasing SOA in numerous previous experiments (e.g., Van Selst & Jolicœur, 1997).

The results are shown in Figure 8. The contrast effect was strongly underadditive with decreasing SOA. At longer SOAs, low-contrast T$_2$s caused longer RTs than high-contrast T$_2$s. As SOA was reduced the contrast effect decreased and became very small at the shortest SOA. These results provide strong evidence that the interference in the AB paradigm takes the form of a processing bottleneck and that the bottleneck is after a stage of processing affected by stimulus contrast. The additional processing time caused by the low-contrast stimulus was not strongly reflected in RT$_2$ at short SOA presumably because it was absorbed into a period of slack in the processing stream for Task$_2$.

The underadditive interaction between contrast and SOA provides particularly strong support for a bottleneck interpretation of the interference observed in the AB paradigm because other forms of interference, such as capacity sharing, do not easily predict this form of interaction. Figure 8 also shows results from a companion PRP experiment that was in all ways identical to the AB experiment described in the foregoing paragraphs, except that a speeded response was also required in Task$_1$. The two experiments generally produced strikingly similar results, with both sets of results exhibiting clear and similar underadditive interactions between contrast and SOA. Interestingly, RT$_2$s were actually longer when Task$_1$ was not speeded relative to when Task$_1$ was speeded, a result that has been replicated in other similar experiments. We discuss this results further in the last paragraphs of the chapter.

Converging evidence for a bottleneck account of the interference in the AB paradigm was provided by Ruthruff, Johnston, and McCann (1998). T$_1$ was an S or a T presented in an RSVP stream (100 ms/item). A speeded response to T$_2$ was required, which was presented either just above or below an RSVP stream that contained T$_1$ and continued after the presentation of T$_2$. Rather than manipulating the contrast of T$_2$, the letters that could be used as T$_2$ characters were distorted, as illustrated in Figure 9. T$_2$ could be either an A or an H, shown in a normal format as in the left pair of letters in Figure 9, or distorted as in the right pair in Figure 9. Task$_2$ was to decide, as quickly as possible, whether T_2 was an A or an H. As expected, distorted letters produced longer RT_2s than normal letters. The most important finding, however, was that the effect of letter distortion was largest at longer SOAs and that it decreased as SOA was reduced. That is, an underadditive interaction was found between letter distortion and decreasing SOA.

Ruthruff et al. (1998) interpreted their results as evidence for a bottleneck after the stage at which visual patterns are classified into letter categories, on the assumption that the letter distortion effect occurred at a stage where letters are categorized.

Figure 9. Illustration of the normal (left) and distorted (right) As and Hs used by Ruthruff, Johnston, and McCann (1998).

The results of Ruthruff et al. (1998) and of Jolicœur (1999d) both confirm the prediction of absorption into slack made by bottleneck models of dual-task interference in the context of the AB paradigm. The results also provide evidence concerning the locus of the bottleneck, which will be discussed in more detail in a subsequent section.

Additivity of bottleneck and post-bottleneck $Task_2$ effect

Prediction of Postponement Model

Another prediction of bottleneck models is that changes in the duration of processing in $Task_2$ at the bottleneck stage, or after the bottleneck stage, should produce additive effects with SOA on response times in $Task_2$. Unlike effects taking place before the bottleneck (Figure 7), effects at the bottleneck stage in $Task_2$ are expected to be fully reflected in RT_2 even at short SOAs, as shown in Figure 10. Because the effect size of bottleneck or post-bottleneck manipulations is expected to be the same at all SOAs, the model predicts additive effects with SOA.

Tests of Additivity Prediction for Bottleneck and Post-bottleneck Effects

Additivity of $Task_2$ stimulus-response compatibility and SOA in a PRP paradigm was reported by McCann and Johnston (1992). Jolicœur (1999d) examined the additivity prediction in the AB paradigm by manipulating stimulus-response compatibility in a speeded second task that followed a typical unspeeded first task known to produce a robust AB effect in a typical AB paradigm (e.g., Crebolder & Jolicœur, 1999).

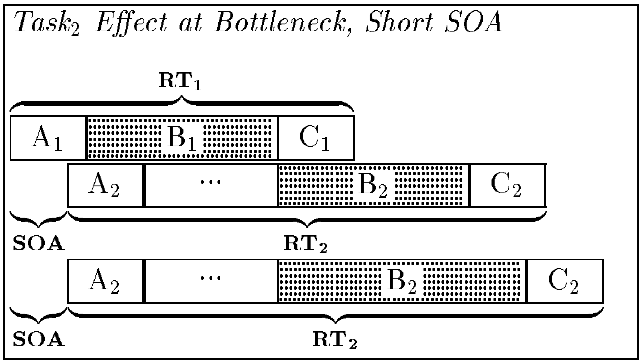

Figure 10. Stage diagrams showing the predicted task interactions in the PRP paradigm for a manipulation in $Task_2$ that affects the bottleneck stage. Even at short SOA RT_2s reflect the full magnitude of the manipulation.

T_1 was an H, O, or S, presented in red in an RSVP stream of white letters. $Task_1$ was to report the identity of T_1, at the end of the trial, without speed pressure. The RSVP stream always ended with T_2, which was an L or an R. $Task_2$ was a speeded two-alternative discrimination based on the identity of T_2. The subjects were instructed to think of the L and R as meaning 'left' and 'right,' respectively. Responses in $Task_2$ were performed with the index and middle fingers of the left hand, using response buttons on the bottom row of a computer keyboard. Every subject performed two blocks of trials. In one block they responded to the L with the left response button and to the R with the right response button. This was the compatible stimulus-response condition. In the other block they responded to the L with the right button and to the R with the left button. This was the incompatible stimulus-response condition. The incompatible condition was expected to produce a longer mean response time than the compatible condition.

If we assume that the stimulus-response compatibility manipulation has an effect at response selection (McCann & Johnston, 1992), and if response selection is at or after a processing bottleneck that causes the AB effect, then the compatibility manipulation should have additive effects with SOA in this paradigm. That is, the difference between the mean RT across the incompatible and compatible conditions should be the same at all SOAs. The contrast of T_2 was also manipulated (low vs. high). This manipulation should produce an underadditive interaction with decreasing SOA, replicating the results shown in Figure 8 for the unspeeded $Task_1$ condition.

The logic of this experiment demanded that subjects perform $Task_1$ (encoding and later reporting a single letter) with high accuracy, otherwise subjects could tradeoff accuracy in $Task_1$ in order to process T_2. For this reason, only subjects who maintained an accuracy of $87\frac{1}{2}\%$ or better in $Task_1$ were included in the final analyses. An accuracy level of 85% was also required in Task$_2$. The results of 14 subjects who met these requirements are shown in Figure 11. As expected, the incompatible stimulus-response condition produced a longer mean RT than the compatible condition. Although there was a tendency for the compatibility effect to become smaller with decreasing SOA, the interaction between SOA and compatibility was not statistically significant ($p > .13$). The contrast effect, on the other hand, was clearly underadditive with decreasing SOA: the contrast effect was largest at the longest SOA and smaller at two shorter SOAs ($p < .012$). There was a large change in the measured contrast effect across the middle and longest SOAs, and no change across the middle and shortest SOAs. The residual contrast effect at the two shorter SOAs was significant ($p < .002$).

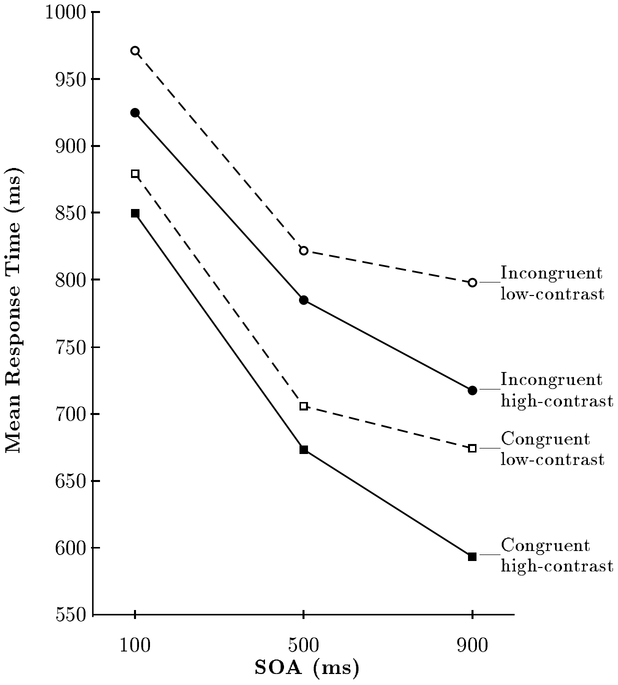

Figure 11. Results from Jolicœur (1999d). Mean reaction time, in milliseconds, for Task$_2$ responses depending on the SOA between T$_1$ and T$_2$, on the contrast of T$_2$, and on the stimulus-response compatibility in Task$_2$.

The stimulus-response compatibility results were not as clear as I would have wished. Although the interaction between compatibility and SOA was not significant, there was a tendency for the compatibility effect to decrease as SOA was reduced; there was a 124 ms congruency effect at 900 ms SOA, a 114 ms effect at 500 ms SOA, and an 84 ms effect at 100 ms SOA. Such a trend, even if not significant, should probably not be ignored (the trend was significant if subjects with lower accuracy in Task$_1$ were included in the analysis). Interpreted in the context of bottleneck models, this trend suggests that some of the stimulus-compatibility effect is at a stage that is before the bottleneck. The large residual effect at the shortest SOA, however, suggests that much of the effect is at or after the bottleneck. Generally, the results are consistent with the view that the AB bottleneck coincides with a stage affected by stimulus-response compatibility, but that the bottleneck may extend to an even later stage.

The absorption into slack of the contrast effect was expected given the results shown in Figure 8 and the results of Ruthruff et al. (1998). This suggests that a large component of the contrast effect was in a stage that was before the AB bottleneck. Interestingly, not all of the contrast effect was absorbed, leaving a significant residual effect of about 36 ms at the two shorter SOAs (also evident in Figure 8 for the unspeeded condition, but the effect there was 15-20 ms). According to postponement models, this would suggest that the contrast manipulation also had an effect in or after the AB bottleneck.

Another line of evidence that appears consistent with the additivity prediction has been provided by Maki, Frigen, & Paulson, (1997). T$_2$ stimuli (words) were preceded by neutral or semantically-related prime words at various point in the RSVP sequence. The results of several experiments revealed additive effects of priming and SOA on accuracy in Task$_2$. Although there are scaling issues when we interpret additivity or interactions for accuracy results because accuracy scores may not be on an interval scale, the results of Maki et al. provide some evidence that appears to confirm the additivity prediction. Given that the variable manipulated involved semantic relationships, which presumably occur fairly late in the information processing stream, the results suggest that the AB bottleneck occurs at a stage at which semantics is extracted, or at a later stage, but not at an earlier stage.

Locus of the bottleneck

Several converging lines of evidence suggest that the locus of dual-task interference in the AB paradigm (i.e., the locus of the bottleneck), is relatively late in the information processing sequence. Luck and his colleagues have used evoked response potential (ERP) techniques to constrain the locus of the effect (e.g., Luck, Vogel, & Shapiro, 1996). They focused on the N400 ERP response, which is generally evoked when a stimulus does not match the present semantic context. Each trial began with the presentation of a word. The purpose of this word was to create a semantic context for the remainder of the trial. Then an RSVP stream was presented. Each frame of the RSVP stream contained several characters. One of the frames, T$_1$, was a repeated string of digits (e.g., 22222). Subjects were asked to report the digit (e.g., 2) shown in T$_1$. At different lags following T$_1$, a word, T$_2$, was also presented. This word could either match or mismatch the semantic context created by the word shown at the beginning of the trial. As expected, at long SOAs, when subjects could report the identity of

T_2, T_2 words that did not match the semantic context of the initial word elicited a robust N400 ERP response relative to words that matched the semantic context. More importantly, at a lag where a large proportion of T_2 words could not be overtly reported, because the behavioural data at this lag exhibited a robust AB effect, the N400 response was just as large as at longer lags. The results suggest that the stimuli were, in some sense, deeply processed, despite the poor performance in the overt responses. The suggestion is that the AB bottleneck is relatively late in the information processing stream, occurring after stimuli like words have activated representations in semantic memory.

Jolicœur and his colleagues (e.g., Arnell & Jolicœur, 1999; Jolicœur, 1998, 1999abc; Jolicœur & Dell'Acqua, 1998, 1999; Ross & Jolicœur, 1999), as had Chun and Potter (1995), proposed that the interference occurs at a stage of processing where perceptually identified representations are transferred to a form of durable storage (Coltheart, 1984). They called the transfer process short-term consolidation (Jolicœur & Dell'Acqua, 1998). According to Chun and Potter (1995) the consolidation of T_2 is postponed by the consolidation of T_1. Jolicœur and his colleagues made the same supposition, but in addition they hypothesized that other processes can interfere with, and be interfered by, short-term consolidation. For example, Jolicœur and Dell'Acqua (1998, 1999) hypothesized that response selection in a speeded pitch discrimination task is postponed by the short-term consolidation process required to encode a visually-presented letter. Indeed, response times to a tone, presented after a visual character to be encoded into memory, increased as the SOA between the character and the tone was reduced, but only if the character had to be remembered (Jolicœur & Dell'Acqua, 1998, Experiment 7).

The results of Ruthruff et al. (1998) are also consistent with a locus of AB interference after letter identification, given that the letter distortion manipulation, which presumably affects the duration of letter identification, produced an underadditive interaction with decreasing SOA. Thus, letter identification had to occur before the bottleneck.

The results of Jolicœur et al. (2000), and Crebolder and Jolicœur (1999), both suggest that some portion of the AB effect may occur at the same stage as the locus of the PRP effect. Recall that Jolicœur et al. (1999c) manipulated the number of stimulus and response alternatives in Task_1 of a speeded AB experiment and found that this manipulation carried over into the accuracy results in Task_2 (Figure 4). Crebolder and Jolicœur (1999) found Task_1 carry-over for a stimulus probability manipulation (Figure 3). These results suggest that effects of the number of S-R alternatives and of relative signal probability occur either at or before the bottleneck producing the AB effect.

Jolicœur et al. (2000) and Crebolder and Jolicœur (1999) also manipulated these variables (probability and number of S-R alternatives) in Task_2 of a standard PRP paradigm, in order to determine the locus of probability and of the number of S-R alternatives relative to the PRP bottleneck. Additive effects of probability and SOA (see Figure 3, bottom panel) and additive effects of number of S-R alternatives and SOA were found, indicating that, these variables had effects that were either at or after the bottleneck stage in the PRP paradigm. If we assume that the stage of processing affected by the manipulations was the same in the AB and PRP paradigms (an assumption that seems especially sound for the effects of S-R numerosity), the results imply that there is a bottleneck in the AB paradigm that is either at the same stage as in the PRP paradigm, or later. Although logically possible, it seems unlikely that the AB bottleneck would be after the PRP bottleneck. Thus, the most likely interpretation of these results is that the two bottlenecks coincide, at least under some conditions.

The results shown in Figure 11, involving the joint effects of SOA and stimulus-response compatibility also suggest a late locus of interference in the AB paradigm. The large residual stimulus-response compatibility effect at the shortest SOA suggests that the bottleneck is at least as late as the stage affected by stimulus-response compatibility. This stage is likely response selection. The trend for a decreasing compatibility effect as SOA was reduced suggests the possibility that the AB bottleneck might extend beyond response selection.

Summary and Conclusions

In this chapter we reviewed findings relevant to the issue of the nature of the dual-task interference observed as the AB phenomenon. The review was organized around four major predictions of bottleneck models of dual-task interference. We began with predictions for response time results in the context of the PRP paradigm (Figures 1 and 7; Pashler & Johnston, 1989) and extended the model to generate predictions for results from the AB paradigm or close variants (Figure 2). For each of the four major predictions:

- carry-over to Task_2 of pre-bottleneck and bottleneck Task_1 manipulations,
- correlation between Task_1 and Task_2 performance at short SOA,
- absorption of pre-bottleneck Task_2 effect,
- additivity of bottleneck and post-bottleneck Task_2 effect,

we found significant converging evidence, often from different labs, confirming the prediction either in the AB paradigm or in a close variant in which Task_2 required a speeded response, but in which Task_1 was unchanged from that used in the usual AB paradigm. We assume that paradigmatic variations in which Task_2 requires a speeded response provides critical evidence for the processing consequences of the encoding operations required to perform Task_1, which are the trigger for the AB phenomenon.

The tests and confirmation of predictions of bottleneck models suggest strongly that the interference in the AB paradigm takes the form of a processing bottleneck. The main alternative to a bottleneck would be some form of capacity sharing. Capacity sharing models, however, have difficulty with two key aspects of the results. First, it is not clear how such models can account for the underadditive interactions with decreasing SOA of contrast and letter distortion. Capacity sharing models predict either overadditivity (Pashler & Johnston, 1989), or additivity (Plourde & Jolicœur, 1999), depending on the details of the timing and sharing assumptions of the model, but not underadditivity with complete or near-complete absorption into slack.

Second, the results from $Task_1$ are generally unaffected by the time at which T_2 is presented relative to T_1 (i.e., by SOA). One might think to argue that there is a channel that has enough capacity for just a little more than one item, such that $Task_1$ processing can be 'protected' from the allocation of some resources to the processing of T_2. This line of argument runs into difficulty, however. Accuracy in $Task_1$ is often well below ceiling, suggesting that all of the capacity in the channel is required to process T_1. Therefore, devoting some resources to the processing of T_2 should take away capacity for T_1, which should produce a drop in accuracy in $Task_1$.

Representative results examining effects of SOA in $Task_1$ and $Task_2$ are shown in Figure 12, in which results of Jolicœur (1998), Experiment 1, are replotted to show $Task_1$ accuracy. T_1 was a red H or an S embedded in an RSVP steam (100 ms/item) of white letters. T_2 was an X or a Y. T_1 was presented on half of the trials (experimental trials) and not shown on the other half (control trials). $Task_1$ was to report whether T_1 was an H or an S, or to report that no red letter had been presented. $Task_2$ was to report the identity of T_2. $Task_1$ and $Task_2$ were unspeeded and responses were deferred to the end of the trial. Clearly, the appearance of T_2 had no effect on mean accuracy in $Task_1$. Furthermore, as can be seen from the error bars, which represent 95% confidence intervals for within-subject designs (Loftus & Masson, 1994), accuracy in $Task_1$ was significantly below 100%.

The paradigmatic AB task can be thought of as a dual-task situation that places high demands on the subjects' abilities to attend, encode, and remember information presented under challenging perceptual conditions. It is not surprising that a complete understanding of the paradigm will likely involve elements from many different levels of analysis and many subsystems. It seems likely, for example, that subjects have some ability to tradeoff accuracy in $Task_1$ for accuracy in $Task_2$, at a relatively macro level, perhaps by preparing more for one task than the other (De Jong & Sweet, 1994; Pashler, 1994). At this level of analysis, it is possible that capacity sharing models might capture significant aspects of the results, although no formal attempts to do so have been proposed to date. We tried to show, for a different level of analysis, at a more micro level, that the processing required to perform $Task_1$ of the AB paradigm is sharply capacity demanding and appears to create a processing bottleneck for other concurrent cognitive operations. This bottleneck is particularly apparent when $Task_2$ is similar to $Task_1$ and presumably requires similar encoding and mnemonic processes. Evidence for a bottleneck can also be found, however, even when $Task_1$ and $Task_2$ are quite different, such as when $Task_1$ requires a speeded response to a tone and $Task_2$ an unspeeded deferred report of a letter embedded in an RSVP stream (Jolicœur, 1998; Jolicœur et al., 2000).

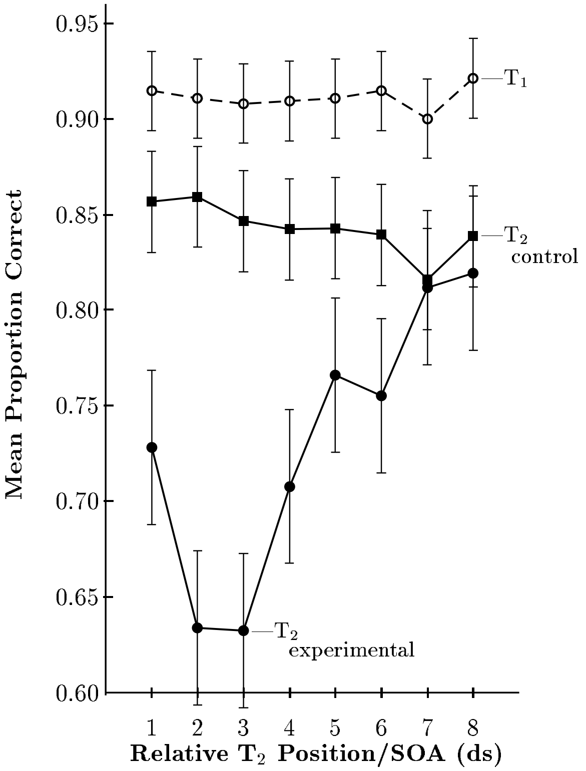

Figure 12. Results from Jolicœur (1998), Experiment 1. Mean accuracy for $Task_2$ control trials (filled squares, solid line), $Task_2$ experimental trials (filled circles, solid line), and $Task_1$ accuracy (unfilled circles, dashed line) as a function of T_1–T_2 SOA. The error bars are 95% confidence intervals for within subject designs.

Some of the evidence that we reviewed suggests that the AB bottleneck is similar to the bottleneck implicated in the PRP paradigm (Pashler, 1994), in terms of the degree of interference caused by the bottleneck and by its general location in the information processing stream (Crebolder & Jolicœur, 1999; Jolicœur & Dell'Acqua, 1998; Jolicœur et al., 2000). For example, the results shown in Figure 8 show that the interference on $Task_2$ caused by encoding T_1 for a deferred unspeeded report (AB paradigm) is quite similar in magnitude to that found when both tasks are speeded (i.e., in the PRP paradigm).

There were also obvious differences, however. First, although similar, the slope of the SOA effect was somewhat larger in the PRP paradigm than in the AB paradigm. Second, at longer SOAs, RT_2s were longer when $Task_1$ was unspeeded relative to when $Task_1$ was speeded. However, at short SOAs, there was no difference in RT_2s across the two $Task_1$ conditions. One interpretation of these results is that, at short SOA, there was equivalent interference on $Task_2$ produced by the two $Task_1$ conditions. As SOA was increased, the need to maintain a durable representation of T_1 for later report in the unspeeded $Task_1$ condition may have created additional interference on $Task_2$. A durable memory for T_1 is not required after the response in the speeded version of $Task_1$. Other interpretations of the results are possible. Our main point is that the different $Task_1$ requirements in the speeded condition versus the unspeeded condition very likely recruit different processing mechanisms (e.g., on-line response selection in the speeded condition, short-term consolidation in the unspeeded condition). Thus, although there are some very interesting similarities between the interference observed in the AB and PRP paradigms, it is clear that there are also some differences. More work will be required to understand fully these similarities and differences, and the mechanisms that underly both of these paradigms. The evidence reviewed in this chapter highlights one important similarity, namely that the interference in both paradigms appears to be caused by a central processing bottleneck.

References

Arnell, K. M., & Jolicœur, P. (1999). The Attentional Blink Across Stimulus Modalities: Evidence for Central Processing Limitations. *Journal of Experimental Psychology: Human Perception and Performance*, 25, 630–648.

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.

Coltheart, M. (1984). Sensory memory: A tutorial review. In H. Bouma and D. G. Bouwhuis (Eds.), *Attention and Performance X, Control of language processes* (pp. 259–285). Hillsdale, NJ: Erlbaum.

Crebolder, J. M., & Jolicœur, P. (1999). *Constraints on the locus/loci of the attentional blink phenomenon.* Unpublished manuscript submitted for publication, University of Waterloo, Waterloo, Ontario, Canada.

Dell'Acqua, R., & Jolicœur, P. (2000). Visual encoding of patterns is subject to dual-task interference. *Memory & Cognition, 28*, 184–191.

De Jong, R., & Sweet, J. B. (1994). Preparatory strategies in overlapping-task performance. *Perception & Psychophysics, 55*, 142–151.

Giesbrecht, B. L., & Di Lollo, V. (1998). Beyond the attentional blink: Visual masking by object substitution. *Journal of Experimental Psychology: Human Perception and Performance, 24*, 1454–1466.

Jolicœur, 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.

Jolicœur, P. (1999a). Restricted attentional capacity between sensory modalities. *Psychonomic Bulletin & Review, 6*, 87–92.

Jolicœur, P. (1999b). Dual-task interference and visual encoding. *Journal of Experimental Psychology: Human Perception and Performance, 25*, 596–616.

Jolicœur, P. (1999c). Concurrent response selection demands modulate the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 25*, 1097–1113.

Jolicœur, P. (1999d). *A processing bottleneck causes the attentional blink.* Unpublished manuscript, University of Waterloo, Waterloo, Ontario, Canada.

Jolicœur, P., & Dell'Acqua, R. (1998) The Demonstration of Short-term consolidation. *Cognitive Psychology*, 36, 138–202.

Jolicœur, P., & Dell'Acqua, R. (1999). Attentional and structural constraints on visual encoding. *Psychological Research, 62*, 154–164.

Jolicœur, P., Dell'Acqua, R., & Crebolder, J. M. (2000). Multitasking performance deficits: Forging some links between the attentional blink and the psychological refractory period. In S. Monsell and J. Driver (Eds.), *Control of cognitive processes: Attention and Performance, XVIII*, 309–330. Cambridge, MA: MIT Press.

Kornblum, S. (1968). Serial-choice reaction time: Inadequacies of the information hypothesis. *Science, 159*, 432–434.

Loftus, G. F., & Masson, M. E. J. (1994). Using confidence intervals in within-subject designs. *Psychonomic Bulletin & Review, 1*, 476–490.

Luck, S. J., Vogel, E. K., & Shapiro, K. L. (1996). Word meaning can be accessed but not reported during the attentional blink. *Nature, 382*, 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. (1994). Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin, 116*, 220–244.

Pashler, H., & Johnston, J. C. (1989). Chronometric evidence for central postponement in temporally overlapping tasks. *Quarterly Journal of Experimental Psychology, 41A*, 19–46.

Plourde, C., & Jolicœur, P. (1999). *Dual-task interference in short-term memory encoding.* Unpublished Manuscript, University of Waterloo, Waterloo, Ontario.

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1995). Similarity determines the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 21*, 653–662.

Ross, N. E., & Jolicœur, P. (1999). Attentional blink for color. *Journal of Experimental Psychology: Human Perception and Performance, 25*, 1483–1494.

Ruthruff, E., Johnston, J. C., & McCann, R. S. (1998). *The attentional blink does not prevent character identification.* Poster presented at the 39th Annual Meeting of the Psychonomic Society, Dallas Texas. Published abstract in *Abstracts of the Psychonomic Society, 3*, 60.

Van Selst, M., & Jolicoeur, P. (1997). Decision and response in dual-task interference. *Cognitive Psychology, 33*, 266–307.

Ward, R., Duncan, J., & Shapiro, K. (1996). The slow time-course of visual attention. *Cognitive Psychology, 30*, 79–109.

Ward, R., Duncan, J., & Shapiro, K. (1997). Effects of similarity, difficulty, and nontarget presentation on the time course of visual attention. *Perception & Psychophysics, 59*, 593–600.