Psychonomic Bulletin & Review
2001, 8 (1), 73-80

Processing bottlenecks in dual-task performance: Structural limitation or strategic postponement?

ERIC RUTHRUFF
NASA Ames Research Center, Moffett Field, California

and

HAROLD E. PASHLER and ALWIN KLAASSEN
University of California, San Diego, California

Recent evidence indicates that a central bottleneck causes much of the slowing that occurs when two tasks are performed at the same time. This bottleneck might reflect a structural limitation inherent in the cognitive architecture. Alternatively, the bottleneck might reflect strategic (i.e., voluntary) postponement, induced by instructions to emphasize one task over the other. To distinguish structural limitations from strategic postponement, we examine a new paradigm in which subjects are told to place equal emphasis on both tasks and to emit both responses at about the same time. An experiment using this paradigm demonstrated patterns of interference that cannot easily be attributed to strategic postponement, preparation effects, or conflicts in response production. The data conform closely to the predictions of structural central bottleneck models.

When people try to carry out two separate choice-response tasks at nearly the same time, responses to one or both tasks are usually delayed. The slowing often shows up on the second task and becomes greater as the time between stimulus onsets (the stimulus onset asynchrony, or SOA) is reduced. This effect is known as the *psychological refractory period* (PRP) effect. Welford (1952), who noted that the PRP effect occurs even when two tasks do not obviously share input or output systems, was the first to suggest that the effect is due to a central bottleneck. According to this hypothesis, central operations (response selection and perhaps other operations) can take place on only one task at a time.

The central bottleneck hypothesis makes several distinctive predictions regarding how manipulations of Task 1 and Task 2 stages should affect reaction time (RT) to Task 2 (see Pashler, 1984; Schweickert, 1978; Schweickert & Townsend, 1989). It predicts that the effects of manipulating Task 2 central stages should combine additively with the dual-task slowing that arises when the SOA is reduced.[1] In contrast, it predicts underadditive interactions between manipulations of prebottleneck Task 2 stages and SOA. These and other bottleneck model predictions have been repeatedly confirmed (e.g., McCann & Johnston, 1992; Pashler, 1984; Pashler & Johnston, 1989; see Pashler, 1994a, for a review).

Although it seems clear that a central bottleneck can occur in PRP experiments, one can still question whether the bottleneck is "structural" (i.e., a basic limitation that arises due to cognitive/neural architecture). In many PRP studies, subjects are instructed to emphasize the speed of Task 1 responses. If they take this to mean "Never make the Task 2 response before the Task 1 response," then central postponement might arise as a strategic (i.e., voluntary) effort to prevent response reversals. Thus, "interference" would occur, but not as a consequence of any fundamental limitation in cognitive architecture (e.g., Koch, 1995; Meyer & Kieras, 1997a, 1997b; Ruthruff, Miller, & Lachmann, 1995).

It has also been noted that much of the dual-task slowing attributed to the central bottleneck might instead be due to other factors. Gottsdanker (1980) and Koch (1995), for instance, pointed out that subjects must prepare for two tasks at short SOAs but only one task at long SOAs. Thus, slowing of RT at short SOAs could be due to reduced task preparation.[2] Also, dual-task interference might be due in part to structural but not central interference—for example, conflicts in response initiation or production (De Jong, 1993; Keele, 1973).

Although several authors have suggested the possibility that there is no structural central interference between tasks, there have been few attempts to address this issue empirically. The present paper addresses this gap. For this purpose, we propose a new dual-task paradigm designed to determine whether there is interference between tasks that cannot be attributed to strategic post-

During the course of this research, E.R. received support from National Research Service Award 1F31-MH10922 and from a National Research Council Postdoctoral Fellowship. The research was also supported by Grant 1-R29-MH45584 from the National Institute of Mental Health to H.E.P. The authors thank Veronica Dark, James Hoffman, James C. Johnston, Steven Keele, Ray Klein, Valerie Lewis, Robert McCann, David E. Meyer, Jeff Miller, Roger Remington, Richard Schweickert, Saul Sternberg, Mark Van Selst, and Rolf Ulrich for helpful comments. Correspondence should be addressed to E. Ruthruff, NASA Ames Research Center, MS 262-4, Moffett Field, CA 94035 (e-mail: eruthruff@mail.arc.nasa.gov).

ponement, changes in preparatory state, and peripheral interference. When interference is found, the paradigm allows us to assess whether the interference occurs in central processing.

A NEW PARADIGM FOR ANALYZING DUAL-TASK SLOWING

In planning the experiment described here, the goal was to eliminate noncentral sources of interference (e.g., peripheral interference, voluntary postponement, preparation changes) as much as possible and then to examine whatever dual-task slowing remains. To eliminate peripheral interference, we present subjects with two tasks that do not share input or output modalities. To strongly discourage strategic postponement in dual-task blocks, we tell subjects that both tasks are equally important, and, therefore, they should not delay one task in favor of the other. To further eliminate any incentive for voluntary postponement, we urge subjects to emit both responses at approximately the same time (i.e., to "group" their responses). Beginning with the earliest PRP research, it has been noted that subjects readily group responses (Borger, 1963), and indeed the traditional practice of emphasizing Task 1 responding was adopted to prevent grouping. In the present design, by contrast, response grouping is explicitly encouraged. To further encourage subjects to process both tasks at the same time, we present the stimuli simultaneously on every dual-task trial (rather than using a variable SOA, as in the typical PRP paradigm).[3]

To examine the difference between single- and dual-task performance without contamination from differences in task preparation (Koch, 1995), we use a nontraditional single-task condition. Instead of presenting the same task on every trial within a single-task block, we randomly intermix stimuli from the two tasks. Because subjects do not know which task will be required, they should prepare both tasks insofar as possible, just as in dual-task blocks.

For reasons to be explained later, the tasks are chosen so that one generally takes longer to perform than the other. In the experiment below, the easier task required the subject verbally to report the number of tones presented (one or two); the more difficult task was to manually indicate whether an upside-down letter was a normal or mirror image (i.e., a "mental rotation" task; see Cooper & Shepard, 1973). We refer to these as the *tone task* and the *letter task*, respectively.

Two tests are performed on the data to determine whether there is a structural central bottleneck. The first test involves a comparison of single- and dual-task RT. As discussed in the next section, if there is no central interference, then only a relatively small amount of dual-task slowing should occur. Structural central bottleneck models, on the other hand, predict several hundred milliseconds of slowing. If slowing does occur, a second test will determine whether the effects of prolonging central operations on the easier of the two tasks "carry over" onto dual-task RT. Structural central bottleneck models predict full carryover, whereas models with no interference (or interference only in postcentral processes, such as response initiation) predict relatively little carryover.

Predictions for Single-Task RT Versus Dual-Task RT

Processing diagrams for the single- and dual-task conditions are shown in Figures 1A–1C.

Model with no interference. Suppose there were no central interference in the dual-task condition. Presumably, subjects would comply with grouping instructions by withholding production of the response selected first on each trial until the other response has also been selected, as shown in Figure 1B. With peripheral interference, preparation effects, and voluntary postponement

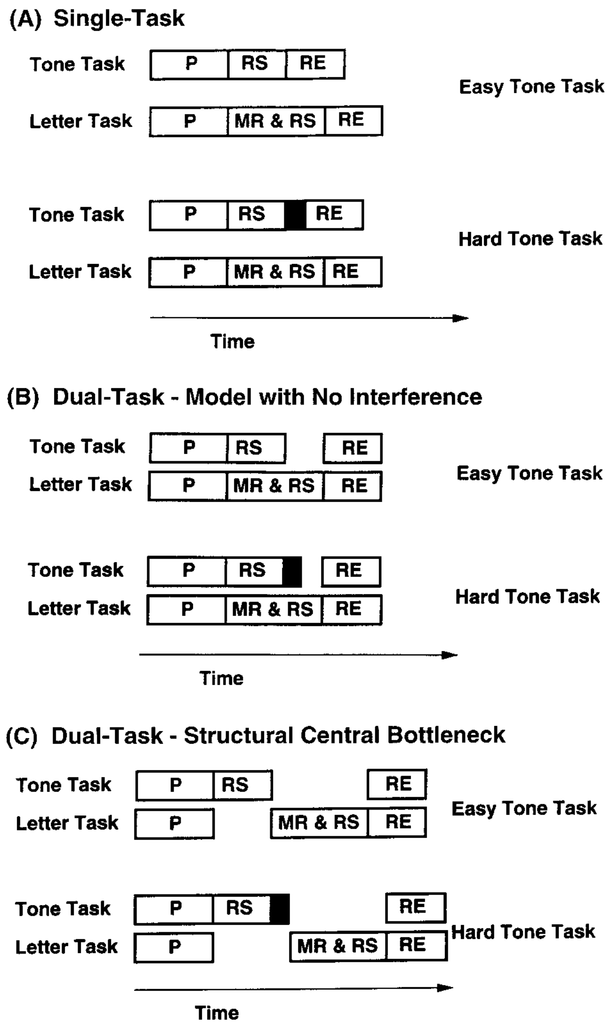

Figure 1. Processing time diagrams: (A) single-task condition; (B) dual-task condition according to a model with no central interference; (C) dual-task condition according to a structural central bottleneck model. P = perception; MR = mental rotation; RS = response selection; RE = response execution.

out of the picture, the tone and letter tasks should be processed with little interference (i.e., about as quickly as in single-task blocks). Why then would there be any substantial dual-task slowing?

One possibility is that it takes more time to produce a grouped response than a single response. It seems unlikely that grouping costs are very large, however, because subjects often prefer to group responses when not instructed otherwise (e.g., Pashler, 1994b). In addition, Fagot and Pashler (1992) found that subjects could produce both a vocal and a manual response to a visual attribute about as fast as they could produce either response alone (the difference averaged less than 20 msec).[4] Furthermore, Pashler and Johnston (1989) found that subjects encouraged to group responses performed the tasks about as quickly as those who did not group responses. Thus, while grouping might increase RT somewhat, the increase appears to be an order of magnitude smaller than typical dual-task interference effects (350+ msec). Later, we present a control experiment that directly supports this assertion. Also note that our main analyses use the faster of the two grouped responses on each trial, thus eliminating grouping costs due to any refractoriness in response production.

A second reason for dual-task slowing to occur, even if there is no central interference, is that grouped response production depends on the completion of the slowest of the two response selections on every trial. The average finishing time of the slowest of two stochastic processes will be equal to or slower than the finishing time of either one alone. Fortunately, the use of tasks differing substantially in finishing time greatly reduces the amount of dual-task "slowing" predicted by this statistical factor. Nevertheless, to estimate the slowing attributable to this statistical factor, we use the single-task RT distributions of the tone and letter tasks to calculate how long it would take to finish both tasks when performed simultaneously, without interference.

Let $F_T(t)$ and $F_L(t)$ represent the (observable) cumulative distribution functions (CDFs) of RTs for the tone and letter tasks in single-task blocks. Let $F_D(t)$ represent the predicted dual-task CDF (for the first of the two grouped responses), based on a model with no interference. As shown in the Appendix,

$$F_D(t) \geq F_T(t) + F_L(t) - 1.$$

The result is expressed as an inequality because it depends on the (unknown) correlation between the finishing times of the two tasks. In any case, the dual-task CDF must be greater than $F_T + F_L - 1$ at each percentile (ignoring sampling error).[5] If the observed dual-task CDF lies significantly below this limit at any of the percentiles (i.e., demonstrates more slowing than can be explained by statistical factors alone), then the data indicate that there is some source of interference between tasks.

Structural central bottleneck models. According to structural central bottleneck models, on the other hand, central operations in the tone and letter tasks must proceed serially. In the letter task, there is evidence that central (i.e., bottleneck-prone) operations include not only response selection but also "mental rotation" of the stimulus letters (Ruthruff et al., 1995). Figure 1C shows the sequence of processing stages in the two tasks predicted by a structural central bottleneck model, for the case in which tone-task central operations are performed first (the processing order of the two tasks should make little difference). Dual-task RTs should be roughly equal to letter-task RTs when performed alone (i.e., in single-task blocks) plus the time spent waiting for access to the bottleneck. This waiting time is roughly equal to the duration of tone-task central operations (plus any task-switching time). Thus, on the basis of previous PRP data, we should observe 300 msec or more of dual-task slowing.

Predictions for the Effects of Tone-Task Difficulty

To determine whether the observed dual-task interference, if any, is due to a central bottleneck, we manipulated the duration of central processing on the easier of the two tasks (in this case, the tone task). Specifically, we varied the compatibility of tone-task response selection. Naturally, this manipulation will affect tone-task RT in the single-task condition. The critical question is whether these effects will carry over onto dual-task RT.

Model with no interference. Because the letter task takes more time to complete, dual-task RT will usually depend on the time to complete the letter task, not the time to complete the tone task (see Figure 1B). Therefore, tone-task difficulty should have a smaller effect on dual-task RT than it does on single-task RT (i.e., there should be less than full carryover onto dual-task RT). This is often referred to as *absorption* of a factor effect into cognitive "slack." Note that temporal overlap of response selections can lead to absorption even if there are later conflicts in response initiation or production.

Structural central bottleneck models. According to structural central bottleneck models, central operations required by the tone and letter tasks must proceed serially (see Figure 1C). Therefore, the tone-task difficulty effects should carry over fully onto dual-task RT.

METHOD

Subjects

Eight undergraduates at the University of California, San Diego, participated in three daily sessions for partial course credit or $7 per hour.

Stimuli

Tone stimuli consisted of one 800-Hz tone for 17 msec or two such tones separated by 50 msec. Letter stimuli (F, R, J, g), subtending 1.4° horizontally × 1.9° vertically, were white against a black background.

Procedure

The subjects were to count the number of tones (one or two). In the easy version, the subjects said this number aloud; in the hard version, they said the opposite number aloud (i.e., "two" if they heard one tone, "one" if they heard two tones).[6] The subjects also indicated whether letter stimuli were normal or mirror images by pressing the "j" and "k" keys, respectively. This task is similar to "mental rotation" tasks (e.g., Cooper & Shepard, 1973; Ruthruff et al., 1995), except our stimuli were always upside-down.

In dual-task blocks, the tone and the letter always appeared simultaneously. The subjects were instructed to emphasize both tasks equally and to emit both responses at about the same time. In single-task blocks, either a tone or a letter appeared (chosen at random). The subjects were asked to respond quickly while maintaining high accuracy.

The subjects completed three daily sessions, the first of which was considered practice. Each session consisted of four blocks of 100 trials (which included 20 warm-up trials). Blocks alternated between single and dual tasks. Tone-task difficulty changed after the second block. Block type order was counterbalanced across subjects. The subjects received RT and accuracy feedback at the midpoint and end of each block.

Each trial began with the presentation of a fixation cross for 800 msec, then a blank field for 300 msec, followed by the tone and/or the letter stimuli. After all required responses had been made, accuracy feedback was displayed for 1 sec. The next trial began 1 sec later.

RESULTS

Trials with an RT less than 150 msec or greater than 4 sec (< 1% of all trials) were rerun later in the block. The subjects were reasonably successful at grouping (mean absolute interresponse interval [IRI] = 64 msec); the subjects with small IRIs showed the same results as those with larger IRIs. Unless indicated otherwise, the analyses of variance (ANOVAs) included the following factors: letter/tone task, single/dual task, easy/hard tone task, and session.

Main Effects

Figure 2 shows mean RTs and error rates[7] for the tone and letter tasks. An ANOVA on mean RT from single-task blocks revealed that the subjects responded faster to the easy version than to the hard version of the tone task [$F(1,7) = 165$, $p < .001$]. In addition, an ANOVA on single- and dual-task RT indicated that the subjects responded to each task more slowly in dual-task conditions than in single-task conditions [tone task, $F(1,7) = 36.8$, $p < .001$; letter task, $F(1,7) = 36.7$, $p < .001$]. Furthermore, there was a main effect of practice [$F(1,7) = 14.9$, $p < .01$].

Predicted Versus Actual Dual-Task RT

Figure 3 shows the observed CDFs of dual-task RTs (defined as the time to complete the first dual-task response) obtained by averaging ("Vincentizing"; see Vincent, 1912) across subjects; also shown are the "slowest" dual-task CDFs that can be predicted by a model with no interference (see Appendix). The observed dual-task CDF lies significantly below the predicted CDF [$F(1,7) > 18$, $p < .01$] at almost all percentiles. In other words, observed dual-task RTs tended to be slower than predicted by models with no interference. Not only were the violations significant, but, in several cases, they exceeded 350 msec. Thus, the data demonstrate considerable dual-

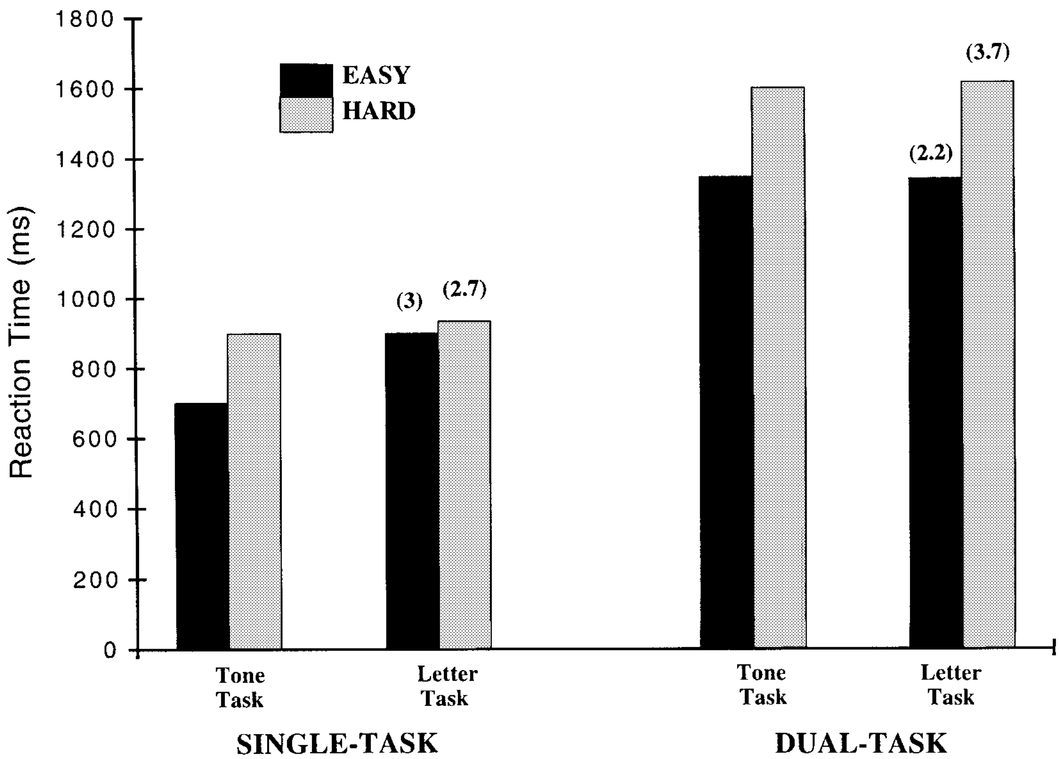

Figure 2. Mean reaction time as a function of experimental condition. Easy tone task: black bars. Hard tone task: gray bars. Mean letter-task error rates are shown in parentheses.

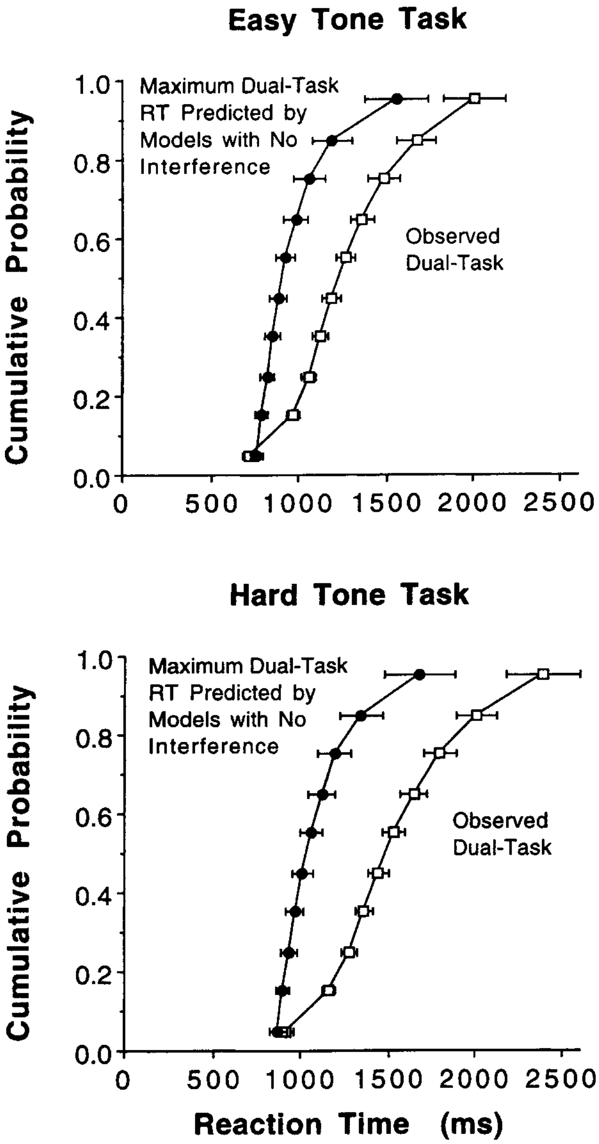

Figure 3. Observed cumulative distribution function (CDF) of dual-task RTs, along with the slowest dual-task CDF that can be predicted by models with no interference. Standard error bars, derived from the CDF type (actual vs. predicted) × subject interactions, are shown for each percentile.

task slowing that cannot readily be attributed to voluntary postponement, changes in preparatory state, peripheral interference, or statistical factors. The observed dual-task slowing, however, is fully consistent with structural central bottleneck models and with previous data from traditional PRP experiments.

Effects of Tone-Task Difficulty

The tone-task difficulty effect was 200 msec in single-task blocks. Models with no interference (or with interference only in response initiation/production) predict that this effect should not carry over fully onto dual-task RT (i.e., due to "absorption into cognitive slack").[8] Contrary to this prediction, tone-task difficulty, if anything, had a somewhat larger effect in the dual-task condition (260 msec) than in the single-task condition (200 msec).[9] These results are consistent with structural central bottleneck models.

CONTROL EXPERIMENTS

Grouping Costs

Can the substantial dual-task slowing observed here be attributed to response grouping? We argued above that this is unlikely, on the basis of previous studies of response grouping. To provide a more precise estimate, we ran a control experiment ($N = 6$) in which we manipulated whether or not the subjects grouped responses. The tone always preceded the letter by 1.5 sec, effectively eliminating central interference. Response grouping slowed letter-task responses by only 21 msec (SE = 35 msec). This confirms that the costs of holding onto the tone-task response and grouping it together with the letter-task response are at least an order of magnitude too small to account for the dual-task slowing observed in the main experiment.[10] Note that, even if all our observed dual-task slowing were due to grouping costs (a claim rendered unlikely by this control condition), one should still not expect tone-task difficulty to carry over fully onto dual-task RT as it did.

Replication Without Instructions to Group

To further test the conclusion that the interference observed in the main experiment was not attributable to grouping costs, we replicated the main experiment ($N = 8$) without instructions to group responses. However, the subjects were explicitly and repeatedly told that both tasks were important and that their primary goal was to complete the tasks in as short a time as possible, just as in the main experiment. Thus, there still should have been little motivation to voluntarily postpone central operations on one task or the other. The subjects adopted a range of processing strategies (mean absolute IRI = 201 msec). Some typically responded to the letter task well before the tone task, some responded in the opposite order, and others grouped their responses (see Pashler, 1994b, for an investigation of dual-task processing strategies adopted under equal task emphasis). Despite the complications introduced by this heterogeneity of strategies, we were able to assess dual-task interference by partitioning the dual-task data according to whether the subjects responded to the tone task or the letter task first on that trial. When the subjects responded first to the tone, mean letter-task RT was 374 msec slower than that observed in single-task conditions. When the subjects responded first to the letter, mean tone-task RT was 411 msec slower than that observed in single-task conditions.[11] Thus, even though the subjects were not asked to group responses or to prioritize one task over the other, we still observed very substantial dual-task interference.

Mental Rotation with Concurrent Articulation

We ran a final control experiment ($N = 6$) to evaluate the possibility that mental rotation requires subvocalization of letter identities and thus interferes with selection and/or production of the vocal response to the tone. These control subjects performed the letter task by itself or while rapidly and constantly articulating the syllable "bah" (roughly three times per second). If subvocalization is necessary to mentally rotate letters, then the articulation condition should have severely disrupted performance. Instead, concurrent articulation slowed mean letter-task RT by only 36 msec ($SE = 27$ msec).[12]

DISCUSSION

The results revealed large amounts of dual-task slowing—similar in size to that found using traditional dual-task paradigms—despite concerted efforts to eliminate noncentral sources of interference (e.g., voluntary postponement and peripheral interference). In fact, there was no sign that these efforts reduced dual-task slowing at all, relative to previous experiments.

The present results cannot be reconciled in any straightforward way with the absence of central interference. First, that view would seem to predict little or no interference in our paradigm, yet the observed interference was very large. Second, if the central stages on the two tasks proceeded concurrently, as shown in Figure 1B, then the effects of prolonging the central stage of the easier task (the tone task) should have been absorbed into cognitive slack. However, we found no evidence for any such absorption.

The Executive Process/Interactive Control Model

Meyer and Kieras (1997a, 1997b) have attempted to explain PRP effects using a complex quantitative model in which there is no structural central interference. On this view, dual-task slowing arises from strategic postponement, delays in motor response programming and production (first suggested by Keele, 1973), and/or peripheral interference (e.g., when both tasks require visual processing). They refer to their model as the *executive process/interactive control* (EPIC) architecture.

We explicitly designed the present paradigm to avoid the noncentral sources of dual-task interference proposed by EPIC. Therefore, the fact that substantial dual-task interference remained is certainly a mark against EPIC. Put another way, we looked in the most obvious place for evidence supporting EPIC's core assertion (concurrent central processing) but failed to find any.

Nevertheless, EPIC is a complex model, and, with appropriate modification, it might account for the present data. How could this be done? First, one could argue that the subjects were capable of performing central operations on both tasks concurrently but elected to postpone central operations as a matter of strategy. This possibility is unappealing, however, given that our paradigm was explicitly designed to promote simultaneous processing on both tasks. Second, one could argue that the observed dual-task slowing was due to the overhead associated with grouping responses (because Meyer & Kieras, 1997a, 1997b, did not directly address the issue of response grouping, we cannot say exactly what causes this overhead). This possibility seems implausible, given previous research suggesting that grouping costs are at least an order of magnitude too small to account for our results (which probably helps explain why subjects under speed pressure often spontaneously settle into a response-grouping strategy; e.g., Pashler, 1994b). In addition, a control experiment in which the subjects were not asked to group responses produced very large interference effects. Third, one might postulate that our experiments engendered competition for peripheral, rather than central, resources. As far as we can tell, none of the peripheral conflicts enumerated by Meyer and Kieras would apply to our tasks. Input conflicts (competition for auditory or visual modules) seem unlikely, given that stimuli were clearly presented (with no masking stimuli) in distinct input modalities. Response conflicts should not have occurred because responses were made using distinct output modalities. Furthermore, the effects of our tone-task central stage manipulation were not absorbed into cognitive slack, suggesting that interference occurred at a central or earlier stage.

Although the present data pose serious problems for EPIC, there appear as yet to be no clearly specified constraints on when different sources of interference postulated by EPIC (e.g., voluntary bottlenecks, peripheral interference) can be invoked. Consequently, we would not claim that these data (or any other potential data for that matter) conclusively refute the EPIC model. (See Coltheart & Coltheart, 1972, and Roberts & Pashler, 2000, for a discussion of theory testing with very flexible models.) However, if EPIC were further developed, so that it clearly stated the conditions under which the different sources of interference do or do not occur, then the present paradigm might well provide a means of definitively testing the model.

The Central Bottleneck Hypothesis

The present results closely match the predictions of the central bottleneck hypothesis, which says that central operations (e.g., response selection) can proceed on only one task at a time. In the present design, we made no attempt to reduce central interference, so the central bottleneck hypothesis predicted large interference effects, similar to those observed in traditional PRP paradigms. This prediction was confirmed. The central bottleneck hypothesis also correctly predicted full carryover of tone-task difficulty effects onto dual-task RT. Because our paradigm strongly discouraged voluntary postponement, it appears that the central bottleneck reflects a structural limitation inherent in the cognitive architecture (although conceivably able to be circumvented by degrees of practice beyond those examined here).

These data offer perhaps the most definitive evidence to date for the existence of a structural, rather than a voluntary, central bottleneck. However, other data sets—

some old and neglected, some recently published—provide additional evidence. Kalsbeek and Sykes (1967), for instance, asked the subjects to perform a continuous choice-RT task (pressing a footpedal in response to tones) at rates ranging up to 100% of maximum capacity. When the subjects performed concurrent decision-making tasks (e.g., maze-tracing, handwritten composition, or another choice-RT task), severe interference occurred even after extensive practice. Gladstones, Regan, and Lee (1989) confirmed this finding, combining various serial choice-RT tasks not evidently similar in either input or output modality. Of particular interest is the finding of Kalsbeek and Sykes that tone–footpedal task caused semantic deterioration and "primitivization" of handwritten composition. It is hard to see how such interference could be attributed to conflicts in motor production. In addition, several dual-task studies have found evidence for a central bottleneck even when both tasks were emphasized equally (Carrier & Pashler, 1995; Ruthruff et al., 1995).

SUMMARY

We present data from a new dual-task paradigm, designed to determine whether there is a structural central bottleneck. The results demonstrated considerable dual-task interference that cannot easily be attributed to voluntary postponement, preparation effects, or conflicts in response production. We also observed full carryover of tone-task difficulty effects onto dual-task RT. Both of these results conform well to the predictions of structural central bottleneck models. Further work is needed to see whether these results generalize to other task combinations and whether central interference persists after extended practice.

REFERENCES

Borger, R. (1963). The refractory period and serial choice-reactions. *Quarterly Journal of Experimental Psychology*, **15**, 1-12.

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

Coltheart, M., & Coltheart, V. (1972). On Rumelhart's model of visual information-processing. *Canadian Journal of Psychology*, **26**, 292-295.

Cooper, L., & Shepard, R. N. (1973). Chronometric studies of the rotation of mental images. In W. G. Chase (Ed.), *Visual information processing* (pp. 75-176). New York: Academic Press.

De Jong, R. (1993). Multiple bottlenecks in overlapping task performance. *Journal of Experimental Psychology: Human Perception & Performance*, **19**, 965-989.

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

Gladstones, W., Regan, M., & Lee, R. (1989). Division of attention: The single-channel hypothesis revisited. *Quarterly Journal of Experimental Psychology*, **41A**, 1-17.

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

Kalsbeek, J., & Sykes, R. (1967). Objective measurement of mental load. *Acta Psychologica*, **27**, 253-261.

Keele, S. W. (1973). *Attention and human performance*. Pacific Palisades, CA: Goodyear.

Koch, R. (1995, Month?). *Hick's law and the psychological refractory period*. Paper presented at the Royal Dutch Academy of Sciences, Amsterdam.

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

Meyer, D. E., & Kieras, D. E. (1997a). A computational theory of human multiple-task performance: The EPIC information-processing architecture and strategic response deferment model. *Psychological Review*, **104**, 1-65.

Meyer, D. E., & Kieras, D. E. (1997b). A computational theory of human multiple-task performance: Part 2. Accounts of psychological refractory phenomena. *Psychological Review*, **107**, 749-791.

Miller, J. O. (1982). Divided attention: Evidence for coactivation with redundant signals. *Cognitive Psychology*, **14**, 247-279.

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

Pashler, H. (1994a). Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin*, **16**, 220-244.

Pashler, H. (1994b). Graded capacity-sharing in dual-task interference? *Journal of Experimental Psychology: Human Perception & Performance*, **20**, 330-342.

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

Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory-testing. *Psychological Review*, **107**, 358-367.

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

Schweickert, R. (1978). A critical path generalization of the additive factor method: Analysis of a Stroop task. *Journal of Mathematical Psychology*, **18**, 105-139.

Schweickert, R., & Townsend, J. T. (1989). A trichotomy: Interactions of factors prolonging sequential and concurrent mental processes in stochastic discrete mental (PERT) networks. *Journal of Mathematical Psychology*, **33**, 328-347.

Sternberg, S. (1969). The discovery of processing stages: Extensions of Donder's method. *Acta Psychologica*, **30**, 276-315.

Vincent, S. B. (1912). The function of the viborissae in the behavior of the white rat. *Behavioral Monographs*, **1** (No. 5).

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.

NOTES

1. This assumes that each task engenders a single chain of serial processing stages, as shown in Figure 1 (see Sternberg, 1969). See Schweickert and Townsend (1989) for a discussion of more complicated processing networks.

2. Bottleneck proponents acknowledge that preparatory changes occur but argue that they account for only a portion of dual-task slowing (e.g., Pashler & Johnston, 1989).

3. When the stimuli are consistently presented simultaneously (SOA = 0 msec), subjects might learn to conjoin the two tasks into one. Because our goal was to study dual-task interference, this would defeat the main purpose of our study. Fortunately, however, we found little evidence that the subjects conjoined the tasks.

4. Two redundant responses were made to the same stimulus attribute. A single operation presumably selected both responses. It is also easy to verify that one can add a constant response element (e.g., saying "boo") in synchrony with a choice-RT response with negligible effects on RT.

5. Miller (1982) used very similar logic to test race models of performance on redundant-target trials in a visual search paradigm.

6. The differences in RT between the one- and two-tone conditions, if any, should average out across the experimental factors of interest.

7. A spot check of recorded tone-task responses revealed high accuracy rates for all subjects (> 97%).

8. To estimate how much smaller the tone-task difficulty effect should be, we simulated a model with no interference. We paired every single-task RT from the letter task with every single-task RT from the tone task and then calculated the maximum of each pair. The resulting mean RT was 90 msec greater in the hard tone-task condition than in the easy condition. This is significantly smaller than the observed value of 260 msec [$F(1,7) = 24$, $p < .01$].

9. Although not statistically significant, the increase was reasonably large. Therefore it is worth considering what might have caused it. First, the subjects might have prepared less for the letter task when paired with the difficult tone task. In fact, letter-task responses in single-task blocks were 34 msec slower when paired with the difficult tone task. Also, task-switching costs might have been greater following the hard tone task.

10. This control experiment would not detect slowing of the letter task, if any, due to encoding of the tone-task response into memory. On the other hand, any such costs would arguably represent a type of central interference and, thus, should not be predicted by the very models we are trying to test.

11. These interference effects, although quite large, appear to be somewhat smaller than those observed in the main experiment. This difference may be attributable to subject differences or to a speed–accuracy tradeoff: The subjects made 7.1% letter-task errors and had a mean single-task RT of 720 msec in the control experiment but made 3% errors and had a mean single-task RT of 858 msec in the main experiment.

12. The result is perhaps unsurprising given that both R and rotated-R have the same name; hence, the letter name does not indicate which response should be made.

APPENDIX
Dual-Task RT Predicted by Models With No Interference

The purpose of this Appendix is to determine the relationship between single-task and dual-task RT distributions predicted by a model in which there is no interference between tasks. The main assumptions are that (1) once begun, all stages proceed just as fast in the dual-task condition as they do in the single-task condition (i.e., without interference), and (2) response initiation and execution begin only once response selection has been completed on both tasks (see Figure 1B). The latter assumption allows for the grouping of responses. For convenience, we derive predictions for the first (i.e., fastest) of the two grouped responses on each dual-task trial.

Let the random variables RT_T and RT_L represent the completion times of the tone and letter tasks in the single-task condition, and let RT_D represent the completion time of the first response in the dual-task condition. The total RT can be broken down into two additive components: (1) the stages up to and including response selection, which we call RS, and (2) the response stages (initiation, execution, etc.) that follow response selection, which we call RE. Thus, $RT_T = RS_T + RE_T$, and $RT_L = RS_L + RE_L$.

$$RT_D = \max(RS_T, RS_L) + \min(RE_T, RE_L).$$

This equation can be conditioned on the relative values of RE_T and RE_L.

$$RT_D = \begin{cases} \max\left(RS_T + RE_T,\ RS_L + RE_T\right), & \text{if } RE_T \le RE_L \\ \max\left(RS_T + RE_L,\ RS_L + RE_L\right), & \text{if } RE_T > RE_L \end{cases} \qquad \text{(A1)}$$

Note that when $RE_T \le RE_L$, it must be the case that $RS_L + RE_T \le RS_L + RE_L$, and when $RE_T > RE_L$, it must be the case that $RS_T + RE_L \le RS_T + RE_T$. Equation A1 can therefore be rewritten as

$$RT_D \le \max(RT_T, RT_L).$$

The probability that RT_D is greater than t milliseconds (i.e., the "survivor function") is

$$p(RT_D > t) \le p(RT_T > t) + p(RT_L > t) - p(RT_T > t \text{ and } RT_L > t), \quad \text{for all } t. \qquad \text{(A2)}$$

The first two terms on the right side of Equation A2 can be estimated from single-task blocks. The third term, meanwhile, depends on the correlation between RT_T and RT_L, which is unknown. It is a probability, however, so it can never be less than zero. Therefore, it must be the case that

$$p(RT_D > t) \le p(RT_T > t) + p(RT_L > t).$$

Letting F_D, F_L, and F_T denote the cumulative distribution functions [i.e., $p(RT \le t)$] of RT_D, RT_L, and RT_T, respectively, this equation reduces to

$$F_D(t) \ge F_T(t) + F_L(t) - 1.$$

(Manuscript received October 21, 1998;
revision accepted for publication January 3, 2000.)