another. As Broadbent writes, "different arms of the cross can of course hold information simultaneously, and this information can be relevant to two or more tasks. So long as tasks do not require transfer of information across the middle of the cross at the same time, there is no reason they should interfere."

This feature of the Maltese cross opens it to some of the same criticisms that have been aimed at the Y-model. The Y-model was developed from a series of studies that examined subjects' performance in divided attention tasks with little or no practice. Yet performance can change quite dramatically with practice. These changes make the bottleneck quite elusive. Tasks that interfere at one point no longer interfere at other points.

Broadbent's discussion of the Maltese cross never directly addresses this issue. In fact, he never mentions practice effects in the context of the Maltese cross, and as with the Y-model, the vast majority of experiments that he invokes in support of his model involve unpracticed subjects. He does cite two studies that involved extensive practice, but discusses them as if practice were not a consideration.

Consider his discussion of the work of Schneider and Shiffrin (1977; Shiffrin & Schneider, 1977), in which it was shown that the number of irrelevant nontarget items had no effect on scanning time when the targets come from a well-practiced set. Schneider and Shiffrin argued that with practice, processing of the irrelevant items had become parallel. Broadbent rebuts that "in Maltese cross terms, there is no evidence that information about [the irrelevant nontargets] has ever left the sensory store and crossed the centre of the cross." The changes that Schneider and Shiffrin observed with practice are not mentioned. This neglect causes problems. Presumably, the irrelevant nontargets should stay in sensory store whether or not they are from a highly practiced set. Broadbent simply does not state directly or indirectly why practice has isolated the nontargets in sensory store.

The same neglect occurs when Broadbent discusses the work of Hirst, Spelke, Reeves, Caharack, and Neisser (1980), in which subjects learned to read while writing from dictation. He argues that subjects could be processing the dictated words during saccades. If this had been the case, why would the task of reading and writing have been hard to learn? It took subjects almost a half a year to learn the task. Again, Broadbent does not offer a detailed consideration about what might have been going on in these six months.

Although I doubt that Broadbent is right that the bottleneck is a fixed structural property, I will readily agree that the issue is empirical. If the debate is to take place in an experimental context, however, then Broadbent must seriously explain the mechanisms by which people learn to do two things at once. Limitations on processing cannot be definitively derived from studies of unpracticed subjects.

Broadbent's Maltese cross memory model: Something old, something new, something borrowed, something missing

Elizabeth F. Loftus, Geoffrey R. Loftus, and Earl B. Hunt
Psychology Department, University of Washington, Seattle, Wash. 98195

The pipeline or "series of stages" framework that Broadbent critiques has been taking shape since the mid-1960s, and has served as a useful vehicle for describing and discussing an enormous body of research on human information processing. We have been loyal fans of this framework, using it as the basis of our textbook, *Human Memory* (Loftus & Loftus 1976), and in our undergraduate and graduate level courses. The stage framework viewed people as constantly taking in information from the environment and then storing, manipulating, and recoding portions of this information in a succession of stages. The key tasks involved in the scientific investigation of information processing were seen as (1) identifying the stages themselves, and (2) investigating what types of information processing characterize each stage.

Broadbent now proposes a new framework, one that resembles the earlier one in many ways but incorporates significant differences. The Maltese cross retains the notions of sensory store, long-term store, and others that were central to prior conceptions. However, it is different in certain key ways, principally in the inclusion of a motor output store. Rarely have past models included motor memory as a component, although an earlier model of Sperling (1967) is a notable exception. The Maltese cross is also claimed to be different from its ancestors in that there is no assumed bias in the directionality in information flow. Information flows from the processing system to the other components, and from those components back to the processing system. There is no favored direction. Is this a reasonable assumption? Logically it would seem that certain directions of information flow do occur more commonly than others, and so there should be favored directions. But while we await the resolution of this issue it is worth noting that even in the earlier frameworks the flow of information was quite flexible. Information flowed from short-term memory to long-term memory (information storage) and back again from long- to short-term memory (information retrieval).

Broadbent's Maltese cross model is surprisingly close to an organization of memory that has been proposed for artificial intelligence programs that deal with the real world, especially with speech comprehension. These models are generally referred to as "blackboard" models because of the terminology used in the Hearsay comprehension project where they were introduced (Erman & Lesser 1975; Hayes-Roth & Hayes-Roth 1979). The programs receive knowledge about the world from a variety of sources. Information is "written" on a (conceptual) blackboard, using a variety of codes. The system's long-term memory consists of a set of *productions*. Each production is a pattern-action pair, where "pattern" denotes some condition, in some code, that could be written on the blackboard, and "action" denotes some action that the system could take, including the action of writing another note, in another code, on the blackboard. The mapping of the blackboard organization onto Broadbent's Maltese cross is fairly obvious (after all, a blackboard is a desk rotated 90 degrees), with one apparent exception. Broadbent seems to assign certain parts of the Maltese cross to particular sensory systems. But does he really? We do not think so, because he permits the transfer of information into a "sensory" store from another store. Thus it would be more correct to say that each of the arms of Broadbent's Maltese cross is actually a message that is written in a particular type of code. The code is one that is accessible to the sensory system involved, which can't write in any other code, and is one of the several codes that the central processing system can use.

If it is accepted that Broadbent's stores are actually codes, then it appears to us that the Maltese cross proposal is a variant of blackboard models. What we find interesting is that Broadbent justifies his model by reference to the data obtained in laboratory studies of human information processing. Blackboard models were originally proposed in a quite different setting – as the best way to mimic the impressive human ability to understand human speech. Specialists in artificial intelligence have since proposed them for models of planning in (computer) problem solving and, to a lesser extent, as the appropriate way to organize knowledge in expert systems. It seems that there is a convergence in concepts although the situations that the theorists are examining are quite different.

The analogies of a man sitting at his desk or of a stockbroker reading a blackboard bring the models to life. They also allow us to see what is missing. The biggest problem is that there is no model of the executive at his desk! The blackboard model avoids this by placing all higher order information into the pattern recognition procedure. But then that has to be specified. There are other problems as well. Both analogies suggest that the pieces of information in this model are discrete – one piece gets put on top of another piece. But how, in this system, can we account for the distortions that are so common in memory? In many instances distortions arise because of a confusion between two items; however, in some instances distortions result from a combination of elements from different sources. For example, in one study (Loftus 1979), subjects saw slides of people going about routine activities. One person in the slides was seen reading a book with a green cover. Later, a leading question suggested that the book had actually been blue. When subjects were finally asked to select the remembered color on a color wheel their choices tended to be a compromise between what they actually saw and what they were told later on their questionnaire. How would "the man at the desk" make these sorts of errors? He could not simply file a letter in the wrong file folder. An analogous error might be one in which the man at the desk recalled a letter received by Mary Smith as one that had been signed by Jane Doe.

In any event, Broadbent's Maltese cross focuses its emphasis toward the shorter-term memories and away from long-term memories and the well-known and interesting distortions of the latter. In fact, Broadbent has rather little to say about long-term memory, except to stress its associative nature. Other recent models take the opposite tack (e.g. CHARM; Eich 1982). They can successfully account for a number of long-term memory phenomena. We see no reason why one or more of these long-term memory models could not be successfully combined with the Maltese cross, perhaps by some operation like "convolution" (Eich 1982), to provide a more complete explanation of the human information-processing system.

ACKNOWLEDGMENTS
The writing of this commentary was supported by National Science Foundation grants to E. Loftus and G. Loftus, and by an Office of Naval Research contract to E. Hunt.

The homunculus as bureaucrat

Alan K. Mackworth
Department of Computer Science, University of British Columbia, Vancouver, B.C., Canada V6T 1W5

The Maltese cross metaphor offered as a replacement for the outmoded telephone exchange and pipeline metaphors might better be called "the homunculus as bureaucrat." The image of a gnomish, desk-bound homunculus is indeed appealing. As Broadbent points out, introducing a homunculus does not lead to an infinite regress if he is confined to executing a set of well-defined rules. The interpreter for that set of rules may well be a second rule-governed homunculus defining the virtual machine for the first; however, that recursion eventually bottoms out on the real machinery of the brain. This view of the structure of intelligent systems is the communal metaphor of artificial intelligence. The primary scientific task now is the determination of the actual computational architecture of those levels.

There are many universal computing devices; they are all equivalent in theoretical power but differ markedly in their expressive power for building cognitive systems and in their suitability for implementing on neural hardware. The two poles of the cognitive task and the real machine exist in a space of possible architectures, creating a tension within it. That space can and must be spanned by a hierarchy of machines, but workers tend to focus their efforts near one pole or the other. Those clustered around the task pole favor production systems (Newell & Simon 1972) or schemata (Havens & Mackworth 1983) while those nearer the real machine pole favor relaxation-based cooperative computation (Marr 1982) and connectionist schemes (Hinton & Anderson 1981). The research strategies of theoretical and experimental psychologists must be designed to identify the most appropriate architecture for the top-level cognitive machine and to specify a necessary and sufficient set of primitives for that machine. Broadbent's target article should be judged by its success in defining and achieving that goal.

Although Broadbent tries to disarm his potential critics by modestly presenting a simplistic theory, in broad brush strokes, some of the details that are present can be questioned. The system architecture is problematic. Broadbent favors production systems but does not pay sufficient attention to their structure. His proposal is in part modeled on the von Neumann model of a stored program computer with separated processing system and memory. In that model the instructions are stored in memory with the memory location of the next instruction stored in the processor's program counter. Since location-addressable memory is assumed, retrieving the next instruction from memory is a single fetch operation. However, Post production systems require a different architecture. Broadbent is not clear on the contents of his long-term associative memory. He says it "retains a running total of the number of times Event A and Event B have occurred together." But later he says, "Although the rules on which it acts are held in long-term memory, it is left to the central control system to recognise the pattern of conditions appropriate to each of the possible rules." Hence the central processing unit must permanently store all the condition predicates in order to carry out the function of the program counter in the stored program model and retrieve the applicable rule(s) from memory. Since the set of rules has no apparent internal structure, the processor must evaluate a very large number of predicates in a very short time to decide what to do next. Moreover, Broadbent states categorically that the learned equivalence of "2" and "two" is *not* stored in the associative long-term memory, but in the processor itself. All this suggests that the serial processor and separate memory architecture proposed is simply inappropriate for the cognitive machine.

Broadbent provides an excellent summary of some of the discontent with the pipeline or linear stage model, but one is still left with a feeling that the new model is a hybrid of the telephone exchange and pipeline models dressed up in new computational clothes. The confusion over the nature and location of "memory" seems to originate from the associationist exchange model. Most of the experimental data provided derive from a school of psychology firmly embedded in the pipeline model. Many of those experiments ask questions that are simply not meaningful in a computational model of perception and cognition. For example, the "logogen" pipeline theory of word recognition is offered as the model for "the pattern recognition system," an otherwise unexplained subhomunculus living inside the central processing system. The notion of a sensory store of "relatively raw input" paradoxically indexed, for sound, by speaker ("the nature of the voice") is another example. One can also discern an attempt to smuggle in the pipeline theory of attention ("the selection of one part of sensory store for further transmission"). Perhaps, as Neisser (1976, p. 84) says, "when perception is treated as something we do rather than as something thrust upon us, no internal mechanisms of selection are required at all."

However appealing the bureaucrat at his desk may be, the metaphor still conveys too passive an image of the perceiver. His filing cabinet stores statistical correlations between events, his many in-baskets fill up with paper, independently of his own purposes, plans, and actions, which he must then attend to or ignore. Requirements for an active, schema-based cycle of perception theory have been outlined elsewhere for psychology (Neisser 1976) and artificial intelligence (Mackworth 1978). The psychology of mind must undergo a thoroughgoing revision more radical than we see here if it is to adopt the computational model.