

1.2 Task switching

What happens when subjects change the goal of cognitive control? If the basic characteristics of control hold in this situation, then the chains of communication that were previously being stabilized will now provide an error signal, and some form of negative feedback will replace those chains with the new target. This change of cognitive control goal and subsequent changes in cognition is called task switching.

This cybernetics-based approach to cognitive control pays off when reading task-switching literature. Effects on reaction time and accuracy due to switching tasks in a variety of experimental setups may be due to a variety of influences. How much time does the change in goal take? Does the change always take place at all? Given that a new goal is encoded, how long does it take for the negative feedback to realize the target chain of communication? What is necessary for that realization to be completed, in the specific situation in which it takes place?

1.2.1 Proactive interference 1: task set inertia

Allport et al. [3] studied seven situations in which subjects had to switch between different tasks. Their initial expectation was to find increases on reaction time when subjects had to switch to more difficult tasks, or had to perform a more complex switch. After all, switching was conceived to be an example of executive control, so such difficulty manipulations should have such an impact on switching. Allport et al. describe how, over the course of their experiments, their basic assumptions about task switching were reversed.

In their first experiment, stimuli were a group of one to nine numbers, all of which had the same value (again, one to nine). This provided two stimulus dimensions - group size and numerical value - with which subjects were required to perform two types of judgement: whether the relevant number was odd or even, or whether it was more or less than five. Stimuli were presented in lists, with only one stimulus visible at a time, and the time to complete the list was used as the dependent variable. Responses were given vocally. Three kinds of switch condition were presented in which subjects had to a) switch between relevant stimulus dimension, keeping the judgement type the same; b) switch between judgement types, while the same stimulus dimension remained relevant; and c) switching between both judgement type and stimulus dimension. In these conditions, the task alternated each trial. The reaction times of these conditions were compared with the average completion time of the tasks being switched between, when those tasks were performed in isolation. The results showed that odd - even judgements were slower than more - less judgments, and that judgements based on group size were slower than those based on numerical value. For all types of switching, switch conditions had higher completion times than the baseline conditions. This difference is the cost of switching between tasks in this experimental context. The size of these switch costs was not different for the different types of switching - stimulus dimension, judgement and both stimulus and judgement switches.

In the second experiment, task difficulty was manipulated. Now, two number groups were presented adjacently. The left and right values and group sizes all differed from one another. The two judgment tasks were 1) to declare whether the left or right number, based on the relevant stimulus dimension, was greatest, and 2) to name the greatest number. Difficulty was manipulated by the difference in numbers on the left and right stimuli. Near stimuli differed by one to three, far stimuli by three to five units. It is unclear why this overlap in distances was used in preference to simply not using a distance of three, as the overlap would apparently reduce differences between

the near and far groups. Regardless, the manipulation was effective in that difficulty caused a significant increase in list-completion reaction time. Switch costs, calculated between the shift stimulus dimension and shift response switch conditions and their baselines, were also found. There was no interaction of switch costs with difficulty.

In the third experiment, stimuli were either a number group or a Stroop word. The tasks involved naming the value or the group size for the number groups, and reading the word or naming the color for the Stroop words. Uniform lists with only one stimulus type were presented on which subjects alternated between tasks or performed each task in isolation. Also, mixed lists were presented with alternated number group and Stroop word stimuli. These lists were performed using either the dominant (value and word) or nondominant (group size and color) tasks. Results showed smaller switch costs for mixed than for uniform lists (230 ms against 440 ms per item). No interaction of dominance with switch cost was found on either the uniform or mixed lists.

These first three experiments had been expected to show effects on switch costs based on some form of limited capacity attributed to the supervisory attentional system. No such effects were found - due to multiple stimulus dimensions to be switched between, task difficulty or task dominance - that suggest that switch costs are sensitive to increasing demands on a unitary, limited-capacity central executive. Three further experiments were performed to explore the contrasting idea that control is distributed amongst subsystems.

The fourth experiment was aimed at explaining the persistence of switch costs in mixed lists, which was in contrast to the expectation, based on Spector and Biederman [217], that switch costs should not occur when the stimulus unambiguously cues the response. In the mixed lists of the third experiment, stimuli were actually associated with both the dominant and nondominant task, even though only one of those was "officially" relevant in a specific condition. In this experiment, the order in which subjects were introduced to the tasks associated with the stimuli was manipulated. In the first block, subjects completed mixed lists using only either the dominant or the nondominant tasks. Baseline lists were only performed for the tasks to be used in the mixed list. In the second block the tasks were switched, and in the third block, switched back. So, in the third block, the only difference compared to the first block was that the stimuli had been contaminated by the other tasks. Each block contained eight runs of three lists (two baselines and the mixed list). The results were as follows. Dominant tasks were faster than nondominant tasks. Strong switch costs were found on the first one or two runs of the second two blocks, but not the first block. Switch costs on subsequent runs remained small but significant in blocks two and three, but were zero in block one. This evidence of proactive interference led to the task set inertia hypothesis: switch costs are due to competing stimulus-response mappings associated with the same stimuli used in performing a different task on previous trials. This interference persisted for a time of the order of one minute.

The task set inertia hypothesis predicted that increasing the intertrial interval, allowing subjects more time to prepare for the new task, will not reduce switch costs. In the experiment to test this hypothesis, the interval between a response and the next stimulus was manipulated. Stimuli were incongruent Stroop words or neutral stimuli: color words displayed in black and colored rows of x's. The stimuli in the lists were now presented in pairs. After the first response, a delay of 20, 550 or 1100 ms was placed before the next stimulus was presented. The delays were kept constant within lists and blocked over eight lists, the order of the blocks being counterbalanced over subjects. Two instructions were given for completing the alternating lists, differing on the order of tasks: either word reading, then color naming, or the reverse. Alternating lists were presented using either the incongruent Stroop stimuli or the neutral stimuli. The dependent variable was

the reaction time to the second stimulus. On incongruent lists, stimulus type did not cue the task. In these lists, switch costs were found only for the word reading task. These costs seemed to decrease as the response-stimulus interval was increased, but this was nonsignificant. The costs remained present (around 140 ms) even at the longest interval of 1100 ms. For neutral stimuli, switch costs were also present, but they were not different for the color naming and word reading tasks. Differences between reaction times for incongruent and neutral stimuli - Stroop interference - were also studied. On baseline lists, a large Stroop effect was found for color naming, as expected, but also a reverse Stroop effect: word reading was also influenced by the incongruent color. On alternating lists, the reverse Stroop effect was larger, converging to the normal (color naming) Stroop effect at the longest response-stimulus interval. The important points of this experiment were, first, that increasing response-stimulus interval did not remove or even significantly reduce switch costs. Second, no switch costs were found for the color naming task, even at short intervals. If this reflects a complete endogenous switch, how could that take so much longer for the word reading task? Finally, word reading was influenced by a different task associated with the Stroop stimuli, even when that task was performed on a different list, as shown by the unexpected reverse Stroop effect on baseline lists. That is, subjects did not seem (able?) to disengage the irrelevant task from the Stroop stimuli. Such effects could be explained by task set inertia. The more difficult task (color naming) requires strong biasing for its action schema to be associated with a stimulus feature, which causes more proactive interference (e.g. due to inhibition of irrelevant but dominant tasks), so that the easier task (word reading) suffers higher switch costs.

Two final experiments studied switching in a different context, namely rapid sequential visual presentation. In these tasks, sequences of words were presented one after another. The words were taken from three semantic categories. One or two of the words in the list were from a target category; subject had to detect and subsequently, after the list presentation, report these words. The sequence could change position on screen, and this spatial shift cued a switch in target category. In the first of these experiments, accuracy following a switch decreased and gradually, over the course of seven stimuli or, equivalently, 1100 ms, returned to the pre-switch baseline. The final experiment was aimed at understanding whether it was time or the number of stimulus presentations that caused the decay of switch costs. Now, a fast and a slow sequence were presented. The switch cost decay was shown to be slower for the slow sequence when plotted against time, but was exactly the same in the two lists when plotted against the number of presentations. Spatial shifts without a semantic shift were also included, and were shown to have much smaller effects on accuracy, with the effect gone after one or two items.

Allport et al. conclude that this evidence points to a dependence on external stimuli of task set control. Of an endogenous / autonomous, unitary, limited-capacity central executive concept, they conclude

As we have seen, the stimulus-independent autonomy of such control operations must be seriously questioned. If the metaphor of a limited-capacity central executive can, in other respects, provide so little guidance - still less any testable predictions - for experiments of this kind, it may be time to look for a better metaphor.

1.2.2 Task set reconfiguration

Rogers and Monsell [197] established a number of task variations that influence switch costs. They varied trial-type - switch or non-switch - within blocks using an alternating runs design. In this

design, tasks A and B are presented in an order of the form AABBAABB... The first presentation of a task in a run thus requires a switch from the other task, while the second presentation does not. AABB is of course not the only possible scheme, and Rogers and Monsell also used a AAAABBBB scheme, to test a prediction they derived from the proactive interference hypothesis on switch costs. The use of alternating runs - more generally, of any within-block presentations of switch and non-switch trials - avoided the problems in comparing an alternating block with a pure block as in [103], [217] and [3]. In such studies, only the alternation block involves keeping two tasks instead of one task in mind and this may have influenced differences between the alternation and pure blocks. The task was cued exogenously (which was in principle redundant) by the location in which stimuli were presented. The screen was divided into four quadrants (in the AABB) design or eight slices (in the AAAABBBB) design. Stimuli were presented in a clockwise circuit, and tasks were assigned to locations so that the desired task scheme had to be performed. The time between a response and the next stimulus was manipulated experimentally. The tasks used were a vowel - consonant letter task, and an even - odd digit task. Trials involved the presentation of two stimuli, the left or right (randomly over trials) of which belonging to the relevant task, close together. Crosstalk and no-crosstalk conditions were created based on what the irrelevant stimulus was: a member of the set belonging to the other task, or a non-alphanumeric character. In the crosstalk condition, the stimulus could be congruent or incongruent depending on whether the associated response was the same as the correct response, based on the relevant character.

Using a variety of these alternating runs experiments, the following results were found for crosstalk and preparation interval. Crosstalk increased reaction time for both switch and non-switch trials, but more for switch trials. Within the crosstalk condition, incongruent characters slowed reaction time for both switch and non-switch trials, but more for switch trials. Switch costs declined as the response-stimulus interval was increased, but only if that interval was kept constant within a block. Switch costs at long (1200 ms) response-stimulus intervals were still present (around 40 ms), even in a no-crosstalk condition. These switch costs were termed residual switch costs. Three final effects were found that further constrained switch costs as they occurred in Rogers & Monsell's spatially-cued task-sequence paradigm. In the AAAABBBB scheme, only the first trial of a run of tasks showed a switch cost, and the subsequent three trials did not differ from each other. Adding a warning signal 500 ms before stimulus presentation reduced reaction time and increased error rate but did not affect switch costs. Only non-switch trials benefitted from response repetition.

The results led to the following main conclusions. First, a residual switch cost exists that endogenous control cannot remove. Switching between tasks has an exogenous component that requires the presentation of a stimulus to complete. Second, task sets - or action schemas [167] - can be cued by stimuli associated with that task, as shown by the effect of crosstalk stimuli. Third, proactive interference was expected to cause a gradual decline in reaction time after a switch. However, no such pattern was found. A single trial was sufficient to achieve an apparently asymptotic non-switch condition. Some perhaps apparently more basic effects were found: a) switch costs did decrease with increased response-stimulus interval, under the right conditions, b) switch costs were not sensitive to a manipulation (the warning signal) that influenced the general state of preparedness, and c) switching disrupted the positive proactive effect of response repetition. Such effects are in line with a general idea that switch costs reflect - at least in part and under certain conditions - some kind of "executive" reconfiguration prior to stimulus presentation.

Meiran [149] specified two conditions that must be fulfilled for switch costs to reflect an executive

process, in the sense of Norman and Shallice's model [167], of reconfiguration. He described a task switch as a combination of control processes and lower-level processes. Control processes shift the task and subsequently reconfigure the lower level processes which, once configured, automatically perform it. An executive reconfiguration process must 1) be specifically related to shifting the task, and 2) be proactive, preceding actual performance. So, for switch costs to reflect such a process, they must decay with preparation time, as this would reflect proactivity, and this decay must be specific to switch trials. Meiran held the inter-trial interval constant while manipulating preparation time by presenting an explicit task cue at different distances in time from the imperative stimulus (the stimulus demanding a response). Thus, proactive interference was assumed to be held constant over different preparation intervals, in contrast to the Rogers and Monsell study, in which preparation time and decay time were confounded. Meiran found that switch costs did decrease at a longer cue-stimulus interval, although it turned out that instructions and / or training were essential in finding the effect. The effects of practice were also studied, from the perspective that practice would strengthen the cue-task association (in Norman and Shallice's terms, automatize the schema selection process). It was found that practice reduced switch costs only for the short cue-stimulus interval, that is, when switch costs were interpreted to be due to incomplete reconfiguration. The residual switch costs, at cue-stimulus intervals that were so long (more than twice the mean reaction time) that reconfiguration could, in terms of duration, have been completed, remained even after practice.

An important contrast between the studies by Meiran and Rogers and Monsell is that Meiran found effects of preparation interval when that was varied within-block. Meiran attributed this to the potency of his cues in comparison with the spatial locations. Meiran further noted that effects of the cue-stimulus interval were found for both switch and non-switch trials, and switch costs were still found at long cue-stimulus intervals, which were argued not to be due to reconfiguration. It was the specific interaction of switch - non-switch and (non-confounded) preparation interval on performance that suggested that task switching involves an executive process involving some form of reconfiguration that is only necessary for switch trials.

1.2.3 Modelling versus interpretations

In their AAAABBBB scheme, Rogers and Monsell [197] showed that reaction time following a switch does not gradually decline, but shows a peak for the first trial and subsequently a level reaction time for the following three non-switch trials. But does such a result prove that an endogenous switch with an exogenous component is responsible for switch costs, as opposed to proactive interference? A modelling study by Gibert and Shallice, using a parallel distributed processing model, suggested that this is not the case [70]. They extended an existing model for Stroop interference [38] so that the state of the model at the end of a trial influenced the state at the start of the next trial. The essence of the model was that task-demand units biased mutually inhibitory "word" and "color" responses. These responses were activated by their respective stimulus dimensions. Differences in the weights between the word and color stimulus-dimensions and their responses made color naming require biasing from the task-demand units to win the response competition. In the extended model, a proportion of the activation of the task-demand units (which was specified as desired by the modeller to determine which task should be performed) persisted. So, on the trial following a color-naming trial, the color-naming task-demand unit would still have a high activation due to its required role on the previous trial. In contrast, the word-naming task-demand unit was not needed to bias the already favorable stimulus - response connections. Switching from a

color-naming to a word-reading trial therefore resulted in more interference than the reverse switch, since the color-to-word switch involves higher persistent activation of the incorrect task-demand unit. By the time a (correct) response can be given, however, the correct task-demand unit has regained dominance. Subsequent trials - as in the AAAABBBB scheme - are thereby protected from proactive interference, despite that the switch costs on the first trial are due to proactive interference. Furthermore, preparation was shown to reduce switch costs, preparation involving simply an advance activation of the task-demand units (as could occur either following a response [197] or a cue [149]). The essential point is that a computational model can be formulated in which precisely the same top-down process prepare for switch and hold trials, and important patterns found by Rogers and Monsell and Meiran on switch costs follow purely from the current state of the network, as determined by interference from the previous trial. It was emphasized by Gilbert and Shallice that their point was not that no endogenous switching process exists - it clearly does, and in their model simply exists as a change in which top-down unit gets activated - but that such a process may not determine switch costs.

1.2.4 Residual switch costs and the mixture model

While the endogenous part of switch costs was being suggested to actually reflect an exogenous process, the exogenous part was being suggested, by the failure-to-engage hypothesis, to actually reflect an endogenous process, or more precisely, its intermittent failure. The failure-to-engage hypothesis [47] is an explanation of residual switch costs that is based on reaction time distribution rather than average reaction time. The residual switch costs found when subjects are provided with the opportunity to prepare for a switch in task prior to stimulus onset have been found to have a distribution that consists of a mixture of reaction times drawn from two other distributions: those when subjects were already prepared and those when subjects had not had a chance to prepare [47, 163]. Residual switch costs may not, therefore, be due to an effect present on every trial when a switch must be performed, but to the proportions of prepared and unprepared trials. According to the failure-to-engage hypothesis, the intention to engage in advance preparation (note that this refers specifically to the context of switch trials) can or can not be activated by an internal or external cue. This activation was argued to be effortful and controlled, as shown by results that suggested that subjects engage in advance preparation less often when they were confronted with long blocks of trials. This effect on the proportion of prepared trials was found even if subjects only were aware that they might have to perform such blocks, but knew that they were at the moment performing a short block. The mixture model of the reaction-time distribution for residual switch costs can be used to estimate trial-by-trial preparation odds based on reaction time. In one of the experiments described in this thesis, these odds were used to define conditions for EEG analyses. Further details will be given in section 2.1.

1.2.5 Proactive interference 2: the retrieval hypothesis and backwards inhibition

After the publication of Rogers and Monsells task set reconfiguration paper, Allport and Wylie [4] again used Stroop stimuli to study task switching, but now in an alternating runs paradigm as in [197]. Three types of stimuli were used: colored words, words in black and colored rows of x's. Colored words were always incongruent: the color designated by the word and the color of its letters never matched. The tasks were to name the word or name the color. Three conditions were used.

In an all-neutral conditions, only black words and colored x's wer presented. This removed any incongruity from the stimuli, and also removed crosstalk: each stimulus was only relevant for one of the tasks. In a color-neutral condition, colored words were presented instead of black words for the word-naming task. Color naming was still done using colored x's. So in the color-neutral condition, only words were incongruent while color stimuli only allowed one task. The stimuli for tasks are still unique: no crosstalk occurs. In the final all-stroop condition, all stimuli were incongruent color words. Crosstalk in that case is complete. Only one response-stimulus interval of one second was used.

Results showed switch costs and reverse Stroop effects (slowing of word-reading due to an incongruent color) that were due to proactive interference from previous tasks. First, reaction times for word reading in the all-neutral and color-neutral (that is, word reading incongruent) conditions were not different, but word reading was slowed in the all-stroop relative to the color-neutral condition. That is, the reversed Stroop effect occurred only when the color-naming task involved the same incongruent, colored words stimuli as the word-reading task. Second, the switch costs for word reading were higher in the all-stroop task than the color-neutral condition. So, the effect of switching to the word-reading task, as opposed to repeating it, was influenced by the task being switched from: if the other task involves the same, incongruent stimuli, then switch costs are higher. Third, switch costs for color naming did not increase from the color-neutral to the all-stroop condition, despite a change from neutral to incongruent Stroop stimuli for that task. The word reading task, however, was the same in these conditions, just as color namings switch costs. Together, the results indicate that switch costs in this situation are due to the characteristics of the task being switched from, not the task being switched to.

Two further experiments were run which enabled the time course of proactive interference to be studied. Now, conditions were presented continually in mini-blocks of 6 cycles, a cycle being a run of 4 tasks in an AABB scheme. In the first of these experiments, only the color-neutral and all-stroop conditions were used. Conditions were cued by high or low tones at the beginning of each cycle in a mini-block. The results were as follows. Word-reading responses became faster over the course of the color-neutral mini-block, but became slower over the course of the all-stroop mini-block. Switch costs for word reading did not change over cycles in either the color-neutral or all-stroop mini-blocks. The responses for color naming also became faster during the color-neutral mini-block, but now an interaction occurred: the overall reaction time improvement was due to decreasing switch costs. The same pattern occurred in the all-stroop block: reaction times for repeat trials remained the same, while reaction times for switch trials decreased. Now, in contrast to Rogers and Monsell's results, a gradual effect of interference from a previous task was found: the presence of the color-naming task in all-Stroop mini-blocks decreasingly influenced word reading in color-neutral mini-blocks. In the discussion of these effects, Allport and Wylie noted that the original task set inertia hypothesis could not explain why there were switch costs at all on word-reading reaction times within the color-neutral mini-block. Performing the color-naming task in these mini-blocks did not require biasing the word-reading task-set, so why should there be switch costs?

The retrieval hypothesis was formulated to account for these within-mini-block switch costs: subjects learn specific stimulus-response associations as they perform the tasks, and these associations are triggered automatically when the stimuli they were based on are presented again. This idea was applied to the results of the mini-block experiment. During color-naming in all-Stroop mini-blocks, associations were learned between colors and a naming response, and between words

and a null response. Subsequent word-reading would then suffer interference because the stimuli trigger the associations relevant to the color-naming response. After a change to a color-neutral mini-block, the colored words are only associated with the word-reading response. It was proposed that the color-naming stimulus-response associations would then become less associated with the colored words (Stroop) stimuli, while the word-reading stimulus-response associations were still being used and learned. This could explain the improvement in reaction time for word reading in the color-neutral condition. The retrieval hypothesis also explains the gradual increase in reaction time for word reading on switch trials in all-Stroop mini-blocks. In all-Stroop mini-blocks, colored words become associated with the color naming task, after a color-neutral mini-block in which colored words were unambiguously related to the reading response, and the color naming task was associated with other stimuli. The colored-words to color-naming association is (re-)learned during the all-Stroop mini-block, so that interference with word reading increases due to the increasingly learned retrieval of the color-naming task on the presentation of a colored stimulus.

A third experiment was performed to test the retrieval hypothesis. Now, a new condition was used instead of color-neutral, namely word-neutral. In the word-neutral condition, color-words are presented in an outlined font. Subjects now always encounter Stroop stimuli for the color-naming task, but alternate between neutral and Stroop stimuli for the word-reading task. They thus have more opportunity to learn the colored-word to color-naming task than in the previous experiment, so that the interfering effects of retrieval should be greater. This was the case: switch costs on the word-reading task in the all-Stroop mini-blocks were greater now that the Stroop stimuli were likely to be more strongly associated with the color-naming task.

Two points were brought forward in this paper. First, repeat trials in the alternating runs paradigm are not fully prepared, as shown by comparisons to pure blocks and by gradual effects in mini-blocks. Second, switch costs can be due to characteristics of the task being switched from, as opposed to the task being switched to. Switch costs, it was concluded, might then not be a good way to index the amount of executive control involved in switching between tasks.

A different kind of proactive interference is the backwards inhibition of task sets [142, 143]. Consider a sequence of three tasks, taken from the set of tasks A, B and C. It was shown that reaction times are slower for the third trial in the sequence A-B-A than the sequence C-B-A. That is, when a switch trial involves switching back to a task performed two trials ago, reaction times are higher than when the trial being switched to was performed more than two trials ago. This backwards inhibition effect was not sensitive to preparation interval, and was not modulated by the repetition of stimulus values or responses, or by negative priming conditions. Thus, the effect seemed to concern the abstract task set as opposed to its perceptual or motor components. Backwards inhibition was sensitive to the top-down control of the task set. This was shown by comparing subjects who received a cue informing them of the stimulus dimension (color, orientation and movement) with subjects who received no cue, but had to detect the deviant stimulus dimension from the display. For both groups, the task was to respond to indicate which of four displayed stimulus was deviant. Thus, both groups performed sequences of tasks including the A-B-A and C-B-A sequences, but only the cued subjects showed backwards inhibition. Residual switch costs at the long (650 ms) cue-stimulus interval were found to be much higher (127 ms versus 26 ms) when the switch was to the recent (lag-2) task set, than when the switch was to a less recent task set. Finally, backwards inhibition was also found when sequences of four tasks were cued at once, so that cues had to be retrieved from memory instead of being presented externally. The existence of backwards inhibition was related to the need to balance flexibility and stability.

Backwards inhibition was proposed to play a part in that balance by allowing us to move away from an initially highly activated (or stable, or dominant) task set, without a kind of weapons race of activations of subsequent task sets.

1.2.6 Task switching and task memory

Altmann [5] manipulated the interval between task-cue and imperative stimulus (100 or 900 ms), and looked at the effect of this manipulation when it was done between or within subjects. A further variable was how many trials were cued. In one-trial experiments, each trial was cued, and the cue remained available during the trial. In multiple-trial experiments, a single cue was followed by a run of, on average, seven trials during which the cue was only available in memory. The tasks involved judging the width or height of a rectangle. The results showed, first, decreasing switch costs over cue-stimulus interval, but only when the interval was varied within subjects. In the multiple-trial experiments, reaction time increased over the course of a run. Within the between-subject interval manipulations, a significant main effect of interval was only found for the multiple-trial experiment. Altmann emphasizes two results. First, whether switch costs decrease as preparation interval increases depends on the exposure of the cognitive system to multiple intervals. So, the preparation effect is limited to certain situations, and models of a switching process should predict this situational dependence. Second, the reaction time of the first trial in a run was a) higher than the single-trial reaction times, and b) sensitive to interval in both the within- and between-subject design. This effect was attributed to memory as opposed to a switch process. Prior to a run, the activation of the task in memory was conceived to be more important than when cues were present continually, and the effect of cue-stimulus interval suggests that this activation process is time-consuming. In these experiments, it would then seem that task activation is a more robust explanation for phenomena than task switching. Further, a methodological problem was claimed to exist in the alternating runs paradigm, as switch costs would always be based on comparisons between the first and second trials in a run, confounding switching with the (re-)activation of the memory of the task on the first trial of the run.

Logan [134] also studied task switching in the context of the activation of task-memories, using the task span procedure. In this procedure, subjects had to memorize a list of one to ten tasks, and then either actually perform the sequence of tasks (performance sequences) or simply indicate what the sequence was (recall sequences). A further condition was performing a sequence of only one task. Stimuli were numerals and number words from one to nine, and the tasks were judgements of parity (even or odd), magnitude (greater or less than five) and form (number or word). The dependent variables were accuracy (chance of correctly completing the sequence) plotted against list length, accuracy against time to retrieval ($t = 0$ being the start of the test sequence, which began 1 sec after end of the presentation of the sequence for memorization) and reaction time against position in the sequence. Accuracy was shown to decrease with list length, with the performance and recall sequences strongly overlapping, thus suggesting that there was no trade-off between sequence storage and task performance in working memory. The performance and recall accuracy decreases were no longer similar when plotted against time, suggesting that the sequence information was lost as a function of item-interference as opposed to decay-time. This conclusion was strengthened by manipulating response-stimulus interval, which resulted in overlapping accuracy curves against list length (over both interval and sequence-type conditions) but not against time. Reaction times showed increases at list positions where subjects were expected to retrieve a chunk of information from long-term memory. The time needed for task-set reconfiguration was estimated by the dif-

ference between the reaction times for a) the perform condition (which involve, in Logan's terms, retrieval of task goals, their translation into an effective representation that specifies the task rules, and the subsequent communication of these control settings to subordinate processes) and b) the sum of the recall and single-task sequences. Except at the first position in the list, 480 ms of increased reaction time was unaccounted for by the recall + single-task reaction times. Note that this measure of switch-related discrepancy controls for artefacts of task activation or run position [5]. In further contrast to the previous study, Logan found response-stimulus interval effects that suggested that the effects of interval duration are partially influenced by task switching, and not only memory activation. For one of the sequence sets used, the effect of interval duration was greater for the performance than the memory sequences. Subjects were presented with one of two sets of sequences: either of length 2, 4, 6 and 8, or of length 2, 3, 6 and 9. Sequences were created by concatenating subsets of permutations of the three tasks. The goal of these sequence sets was to determine the expected moments of chunk-recall, but in terms of understanding task switching, this interaction with response-stimulus interval seems at least as interesting. Unfortunately, the effect was not further discussed. It was only the 2n-length sequences that showed an effect of interval on switch costs, not the 3n-length sequences. Due to the concatenation procedure, repeat trials would often occur (1/3 of the time) at the expected chunk-retrieval times for the 3n-length sequences. For the 2n-sequences, chunk-retrieval would often definitely coincide with a task switch (at positions three and five), and only on position seven would there be a 1/3 chance of a repeat. Could this coincidence of retrieval and switching have caused the response-stimulus effect on the switch discrepancy for the 2n-sequences alone?

Retrieval of task cues from memory was compared with externally cued tasks in a study by Koch using an alternating runs design [118]. In this study, a short and long response-stimulus duration was varied between blocks, and subjects were divided into a group who had to perform the sequence (AABB...) from memory, and a group who also were presented with redundant cues. Only the group with external cues showed a reduction of switch costs with increasing preparation interval. One consequence of this is that decay of proactive interference could not have been the cause of the reduced switch costs, as that should be the same for both groups. Koch decomposes advance preparation for task switching to account for the cue x switch cost interaction. First, a task "goal" must be selected, and second, the task-specific stimulus-responses rules must be retrieved. Since the no-cue group did not have low accuracy scores, Koch suggests that the task selection process is not problematic with internal cues. However, the second, specific rule-retrieval process was considered to be possibly sensitive to the salience of cues. This division of preparatory work seems to fit well with Logan's results, in which performing tasks involving a switch required more time than the sum of performing the task and retrieving the task identifier.

Mayr and Kliegl [144] present a somewhat similar decomposition of task set reconfiguration. First, a cue must initiate the retrieval of task rules from long-term memory. This is termed the retrieval stage. Only one task set was proposed to be able to be active at once. Second, in the application stage, the task rules are automatically applied to a stimulus (either presented or imagined). The question posed by Mayr and Kliegl was whether switch costs result from changes in the task, or changes in the retrieval path (from cue to task rules). Their study used an information-reduction paradigm, in which a single task could be cued by two different cues. Thus, cues could be made to switch independently from tasks, providing two pure cost variables: a cue-switch cost (cue switch - no-switch) and a task-switch cost (task switch - cue switch). The results were as follows. First, switching cues while repeating tasks was found to result in cue-switch costs. Thus, in other

studies in which task switches and cue switches were confounded, part of the switch cost may have been due to long-term memory retrieval, as opposed to reconfiguration. Second, response-repetition benefits were found following cue changes but not task changes. Changing the cue thus does not seem to disturb the task set configuration in the way that changing the task does. Third, practice only reduced the cue switch cost, but not the task-switch cost. Fourth, when response-cue and cue-stimulus intervals were manipulated (so that preparation and decay periods could be manipulated independently, as argued for by Meiran), the only interval-related cost reduction was a reduction of the cue-switch cost due to an increased cue-stimulus interval. So the task-switch cost was not reduced either by increased decay or preparation interval. Finally, backwards inhibition was found to occur even when the cues changed, so that backwards inhibition is linked to the application but not the retrieval stage.

1.2.7 Conclusion

From a control perspective, task switching would be expected to involve a change in goal and a subsequent process of error reduction. The broad literature concerning the various kinds of switch costs can be understood in these terms. The studies by de Jong, Logan, Koch and Mayr and coworkers provide information on goal changes. Goals appear to be decomposable into an identifier component and an effective-goal component, and so goal changes also may show effects specific to one or the other components, or their interaction. The identifier, serving as an internal cue to the goal-memory, plays a part in Logan's recall sequences and Koch's internally cued switched. The effective component, that results in a steering of neural chains of communication towards the goal state, seems to be what was manipulated by cue switching in Mayr's studies. The negative reinforcement part of the control perspective is very similarly to Mayr and Kliegl's automatic application of task rules. Mayr and Kliegl do not seem to make the identifier - effective goal distinction of Logan and Koch. De Jong's mixture model and failure-to-engage hypothesis fits well with the identifier - effective goal decomposition, suggesting that the effective goal can be retrieved but may not be, even if the identifier (the "ineffective" or "unengaged" intention) is in fact retrieved. People may want to switch in an abstract, unspecified way, but the intention may not be concretely or specifically realized so as to change the chains of communication, in other words, to reconfigure. The state of "wanting to switch" may conceivably be present in the brain, while for instance the effective goal "map 'X' to left index finger, map 'O' to right index finger", such that negative feedback could occur, is not at all present. Notably, the actual stimulus, of which the representation is part of this goal memory, seems quite likely to help activate an effective goal, which is perhaps the kind of exogenous process hypothesized to exist by Rogers and Monsell. A final point is that, if goal retrieval is a controlled process, it should actually be expected to be fallible [47] and sensitive to context [5], motivation [47] and strategy [149]. That is, in contrast to claims that the lack of preparatory switching under certain conditions is problematic for a general controlled-processing conceptualization of task switching [5], it would seem to be precisely what such an idea would predict.

The various kinds of proactive interference - task set inertia, the retrieval hypothesis and backwards inhibition - can be seen as affecting the second part of control, namely the negative feedback phase. It is actually this phase which may most deserve the term reconfiguration, as argued for by Mayr based on its wiping-out of the response-repetition benefit. This is a notable departure in conceptualization and terminology from earlier studies, in which reconfiguration was assumed to be the proactive process. It seems likely, however, that the proactive process is task retrieval (both

task identifier and effective task goal), while reconfiguration is a subsequent automatic process of negative feedback. This automatic reconfiguration process is however the aspect of control specific to task switching as opposed to task preparation in general. Failures to retrieve an effective task goal, proactive interference and a simple lack of sufficient preparation time could all preclude or interfere with reconfiguration.

All in all, it would seem that a cooperative stance between results and hypotheses would be reasonable in studying task switching. Unfortunately, so far many studies seem to be geared towards verifying generalized claims concerning task switching or switch costs - that costs are or are not due to proactive interference or advance reconfiguration or cue retrieval for instance. This seems likely to have suppressed potentially interesting cross-studies, for instance, concerning the effects of proactive interference on the odds of the retrieval of an effective task goal. A second problem is a lack of data. Not in the sense of a lack of paradigms, effects, interpretations and contradictions, but in terms of data dimensions, or data types. Behavioral data has the inherent limitation that it reflects the end product of all the inferred systems and processes that are in no way directly measured, or even measurable without some kind of more precise, physical specification. Physiological data can to some extent imply theoretical consequences on mediating processes and structures, as described in subsequent paragraphs.