Birkbeck ePrints: an open access repository of the research output of Birkbeck College

http://eprints.bbk.ac.uk

Baughman, Frank D. and Cooper, Richard P. (2007). Inhibition and young children's performance on the Tower of London task. *Cognitive Systems Research* **8** (3) 216-226.

Citation for this version:
Baughman, Frank D. and Cooper, Richard P. (2007). Inhibition and young children's performance on the Tower of London task. *London: Birkbeck ePrints.* Available at: http://eprints.bbk.ac.uk/636

Citation for the publisher's version:
Baughman, Frank D. and Cooper, Richard P. (2007). Inhibition and young children's performance on the Tower of London task. *Cognitive Systems Research* **8** (3) 216-226.

Inhibition and Young Children's Performance on the Tower of London Task

Frank D. Baughman (f.baughman@psychology.bbk.ac.uk)

Richard P. Cooper (r.cooper@bbk.ac.uk)

School of Psychology, Birkbeck, University of London,
London WC1E 7HX, UK

Abstract: Young children, when performing problem solving tasks, show a tendency to break task rules and produce incomplete solutions. We propose that this tendency can be explained by understanding problem solving within the context of the development of "executive functions" – general cognitive control functions which serve to regulate the operation of the cognitive system. This proposal is supported by the construction of two computational models that simulate separately the performance of 3-4 year old and 5-6 year old children on the Tower of London planning task. We seek in particular to capture the emerging role of inhibition in the older group. The basic framework within which the models are developed is derived from Fox and Das' domino model (2000) and Norman and Shallice's (1986) theory of willed and automatic action. Two strategies and a simple perceptual bias are implemented within the models and comparisons between model and child performance reveal a good fit for the key dependent measures (number of rule breaks and percentage of incomplete solutions) of the two groups.

1   Introduction

Since the seminal work of Newell and Simon (1972), many researchers have studied problem solving using concepts such as problem-spaces, states and operators or moves. From this perspective, problem solving involves selecting a series of moves which allows one to move from a representation of the current state of a problem to a desired or goal state. The emphasis in this work (as exemplified by studies of tasks ranging from the Hobbits and Orcs problem (Thomas, 1974) to the Tower of Hanoi (Simon, 1975)) has generally been to investigate the processes by which moves are proposed or evaluated.

A different literature views problem solving as one of many cognitive skills that are under the control of the "central executive". Within this approach, the primary concern is with understanding how the human cognitive system is organized and how it operates at a mechanistic level in providing high-level cognitive functions, such as reasoning and problem solving, within the context of general computational control functions. Although accounts of the executive (see e.g., Zelazo, Reznick & Frye, 1997; Shallice, 2002) and the extent to which control is achieved by a single, unitary process or by many, diverse processes are contended (Miyake, Friedman, Emerson, Witzki, Howerter & Wager, 2000), a consensus does exist with respect to the notion that certain control functions appear central to human cognition. These so-called *executive functions* are commonly viewed as a set of "general purpose control mechanisms that modulate the operation of various cognitive subprocesses and thereby regulate the dynamics of cognition" (Miyake et al., 2000, p. 50) and are widely assumed to form the basis of our ability to perform complex tasks, including reasoning, planning and problem solving.

A substantial body of research has focused on the role of three specific possible executive functions: (1) inhibition of prepotent responses ('*inhibition*'), (2) shifting of mental sets ('*shifting*') and (3) updating of working memory ('*updating*'). These mechanisms have formed the basis of a variety of theoretical accounts, often drawing on data from the neurosciences to help elucidate hypothesized processes (e.g., Aron, Robbins & Poldrack, 2004; Garavan, Ross, Murphy, Roche & Stein, 2002; Sylvester, Wager, Lacey, Hernandez, Nichols, Smith & Jonides, 2003). Thus, Sylvester et al., (2003) demonstrated converging evidence for separable and differential levels of involvement of executive function mechanisms across tasks through fMRI studies of attention switching and response

interference. Yet within the literature terms such as "mental flexibility", "inhibition", "mental set shifting", "planning", "problem solving" and "categorisation" (see e.g., Bull, Espy & Seen, 2004; Miyake et al., 2000) are frequently used to describe the same or similar processes. In order to properly evaluate theories founded on such general control mechanisms one must be explicit in detailing what the theorized functions are and how they may be implemented in mechanistic terms. Computational modeling offers one way in which accounts of hypothesized mechanisms may be assessed. It thus offers the distinct advantage of allowing us to understand which attributes of a particular process may have important effects in producing behavior and which do not.

Our particular concern in this paper is how the putative development of executive functions may impact upon problem solving performance. The Tower of London task (Shallice, 1982) is a problem solving task that has frequently been employed in studies of executive functioning (e.g., Zook, et al., 2004), and a number of developmental accounts of performance on the Tower of London exist (see e.g., Blair, Zelazo, & Greenberg, 2005) that are broadly supportive of a complex and dynamic interaction of mechanisms underlying high-level cognitive processes. For this reason our focus here is on modeling the behavior of young children on the Tower of London task at two different points in development. Additionally, Houdé (2000) has argued that a key mechanism of cognitive maturation is the development of task-general inhibitory mechanisms. Consistent with these views on the development of executive function mechanisms we adopt a position that assumes an association between age and an increased ability to inhibit a prepotent response.

In the Tower of London task, subjects are presented with an apparatus that consists of three colored balls and a board with three pegs of different lengths. The length of each peg constrains the number of balls that can be placed on it to 1, 2, and 3 balls, respectively. Subjects are then presented with a picture of the goal state and are asked to move the balls, one at a time, to match the goal using the picture as a reference. Two simple Tower of London tasks are shown in Figure 1. Each task can be solved in exactly three moves, but more complex tasks, requiring more moves, may easily be generated by manipulating the starting and goal states.
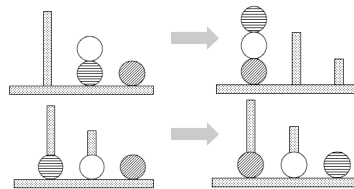

Figure 1. Two Tower of London tasks, each of which may be solved in three moves

Within studies of problem solving on the Tower of London, efforts have focused primarily on the ability to generate and manage subgoals. This has been assessed through analyses of the number of correct solutions on problems of different goal configurations (e.g., flat-ending vs. tower-ending, Klahr & Robinson, 1981) and performance on problems of different complexity (i.e., problems in which the minimum number of moves required to reach the solution varies) between children of different ages. However, a key characteristic of young children's performance on problem solving tasks is their tendency to break task rules and produce partial or incomplete solutions (Waldau, 1999). In the context of the Tower of London task, no attempts have thus far been made to account for young children's tendency towards (1) *rule breaks* (e.g., holding two balls at the same time), or (2) solutions that are *partially complete* (i.e., solutions within only one or two balls in their correct place).

In line with the aim of this special issue, we demonstrate the advantage of applying cognitive modeling to cognitive theory on the Tower of London task. Focusing specifically on the occurrence of rule breaks and partial completion on this task, this paper details two computational models that provides an explicit account of how these behaviors may arise from differences in the ability to inhibit responses and thus offers a bridge between traditional approaches to problem solving and accounts based on theories of executive functions and data from the neurosciences. In order to ground our account of inhibition on a solid foundation, we phrase our models within a framework derived from

two sources: (1) the Domino model of Fox and Das (2000) and (2) Norman and Shallice's theory of willed and automatic action (1986).

2   Executive Functions and Cognitive Architecture

2.1   Experimental evidence for Executive Functions

A number of tasks, such as the Wisconsin Card Sorting task, the Stroop task and the Tower of Hanoi task, that are held to load heavily on different executive functions have been developed. Thus, the behavioral effects of the executive function of inhibitory control are typically considered demonstrable by superior performance on tasks where an automatic or dominant response should be suppressed. This may be the successful inhibition of the tendency to process the semantics rather than the actual color of word items on the Stroop test. Conversely, a deficiency in the ability to inhibit is implied by poorer performance.

Theorists arguing both sides of the unity versus diversity debate have used executive tasks to explicate the role of executive functions (see e.g., Duncan & Owen, 2000; Miyake et al., 2000). Thus, the claim that these tasks measure specific dissociable cognitive functions is derived from observed dissociations of performance on them (e.g., Miyake, et al., 2000). However, progress in the isolation and study of executive functions has been hindered by the multiple and somewhat "arbitrary and post-hoc" interpretations of resulting data (Miyake et al., 2000, p. 53; see also Bull et al., 2004).

With the objective of clarifying some of these issues, Miyake et al. (2000) detail a study in which they find support for separable mechanisms of executive functions. Consistent with many previous studies and of special relevance to this paper were their findings that inhibition was more strongly associated with performance on the Tower of Hanoi task than a range of other executive function tasks. Miyake et al. offer a plausible interpretation of these findings, reasoning that in the Tower of Hanoi one is influenced by the tendency to move *towards* greater perceptual similarity rather than move away (see also, Simon, 1975). This interpretation fits with numerous other studies in which moves that take the configuration of the current state of a problem away from the goal state are described as counter-intuitive, or undesirable, whilst in fact they are necessary for task completion (Gilhooly, 2002). Such moves may also be needed to solve certain Tower of London problems.

2.2   A framework of behavioral control (SAS and CS)

Most accounts of executive functions isolate particular putative functions but fail to provide an integrative computational framework within which those functions may operate. One plausible framework within which various postulated executive functions might be implemented is Norman and Shallice's (1986) theory of willed and automatic action. This is perhaps one of the best-known frameworks to embody the diversity view of cognition. Divisible into two distinct but significantly related processes that operate according to specific parameters, it comprises an automatic or reactive system – contention scheduling (CS) – that is held to control behavior in routine situations, and a controlled or deliberative system – the supervisory attentional system (SAS) – that is held to control behavior in non-routine situations. Briefly, the CS organizes routine behaviors in the form of schemas and is characterized by low-level, predominantly autonomous processes that control everyday actions. The SAS imposes a heavy top-down influence on behavior by way of generating goals, creating schemas for CS to carry out and monitoring behavior. Problem solving and behavior in general is thus held to be the product of the influences of these two interrelated systems, with the SAS more involved in novel tasks but the CS taking over when tasks become familiar. The depth and breadth of behavior the CS-SAS theory is intended to account for makes it a suitable starting point for modeling the Tower of London.

2.3   The Domino model

In Norman and Shallice's original description, the SAS was specified mainly in terms of its functions (i.e., how and when it was held to influence the operation of CS). The framework was therefore not sufficiently fleshed out to allow the construction of a fully mechanistic account of problem solving on

a task such as the Tower of London. However, in more recent work, Shallice and colleagues (e.g., Shallice & Burgess, 1996; Burgess, 2000; Shallice, 2002) have attempted to fractionate or decompose the SAS into component subprocesses, such as generating and setting intentions and monitoring behavior against current goals. On the basis of this, Glasspool and Cooper (2002: see also Shallice, 2002, Glasspool, 2005) have shown that the CS-SAS framework is compatible with a computationally more explicit general cognitive architecture, Fox and Das' (2000) Domino model (see Figure 2).
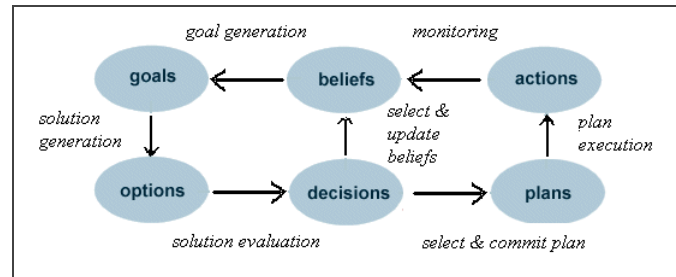

Figure 2. The generalized Domino model of Fox and Das (2000)

The Domino model represents a highly organized system for decomposing elements of a problem. While it is not assumed to be a model of human cognition, it has been used extensively in AI work on expert systems (see Fox & Das, 2000), and the principal elements of the model may be related to those involved in a GOMS-style analysis of problem solving (Card, Moran & Newell, 1983). Thus, processing within the domino involves generating goals (or subgoals) based on beliefs, generating possible solutions to those goals, evaluation and selecting from those solutions, etc.

2.4   Linking the Domino and the CS-SAS framework

The majority of the Domino processes flesh out the possible operation of the SAS. CS fits in to the picture only in the right-most processes associated with actions and plan execution. CS is held to consist of a hierarchically structured network of schemas in which processes of interactive activation work to select highly active schemas that then trigger basic actions (Cooper & Shallice, 2000). SAS, in the form of the other Domono processes, provides a top-down bias to schemas within CS (which correspond to actions and routinized sequences of actions), thus encouraging the performance of routine behaviors. SAS (and the Domino) only has indirect control of action, however, as the interactive activation processes of CS may lead to schemas being selected even without top-down excitation.

3   Problem solving and the Tower of London: Developmental Findings

The Tower of London task has become a popular tool for measuring problem solving abilities of children and adults with neurological impairments. Both the Tower of London and the Tower of Hanoi (on which the Tower of London is based) have been held to load heavily on executive functions and, in the case of the Tower of London in particular, on inhibition (Bull et al., 2004; Miyake et al, 2000).

Previous developmental research on the ToL suggests two strategies that young children may use in problem solving. These form specific components within the computational models developed here. The strategies are (1) an *immediate-hit* strategy (the tendency to place a ball in its target position immediately if the target position is free and the target ball is free to move) and (2) a *one-move look-ahead* strategy (the tendency to plan moves up to one-move away). Bull et al., (2004) and Goel, Pullara and Grafman (2001) describe these strategies in more detail and provide evidence for their use in solving Tower of London problems.

The target behavioral data for the present work come from a study in which children's performance on two different types of Tower of London problem, tower-ending and flat-ending problems, was compared (Waldau, 1999). In this study, two groups of children (3-4 year olds and 5-6 year olds) completed six problems (three problems of each type). There were 17 children in each group. Our concern here is not the children's behavior on the specific problem types. Rather it is the degree to

which children in the different age groups produced incomplete solutions or broke task rules in coming to their solutions. Table 1 therefore shows the key dependent variables for each age group (collapsed over the six problems). In this table, "configuration" refers to the percentage of problems in which children's solutions matched the target solution, regardless of the colors of the balls (e.g., having one ball on each peg in the case of a flat-ending problem). "Colors" refers to the percentage of problems in which balls were also arranged correctly by color. As can be seen from the table, children in both groups were generally good at matching the correct configuration, but on about one-third of occasions children in the younger group did so while neglecting the colors of the balls. The between group difference was highly significant ($t(32) = -3.37$, $p = 0.002$, two-tailed).

Table 1: Developmental data for the Tower of London

	3-4 yr olds	5-6 yr olds	t(df = 32)	p
Configuration (%)	95.83	96.05	−0.42	0.677
Colors (%)	66.32	94.74	−3.37	0.002
Rule breaks (%)	52.08	21.92	2.67	0.012
Avg. no. moves	11.26	9.79	2.65	0.012

Table 1 also shows that young children broke the task rules on over half of the trials, while the older group broke the rules on less than a quarter of trials. The between group difference was highly significant ($t(32) = 2.67$, $p = 0.012$, two-tailed). At the same time, children in the older group required significantly fewer moves than those in the younger group to produce their solutions.

4 Modeling the Tower of London

Computational approaches using cognitive architectures such as Soar (Newell, 1990) and ACT-R (e.g., Anderson & Lebiere, 1998) have contributed greatly to the study of a wide range of cognitive behavior. However, their use has also drawn criticisms on a number of theoretical and technical fronts (see e.g., Cooper & Shallice, 1995 and Altman & Trafton, 1999, respectively). In this paper, we seek to limit the number of underlying theoretical assumptions in modeling the Tower of London task, and therefore rather than use a cognitive architecture with all its architectural assumptions, we adopt the COGENT modeling environment (Cooper & Fox, 1998; Cooper, 2002) as a platform in which to implement the relevant aspects of the Domino model.

4.1 The COGENT modeling environment

COGENT is a visual environment for cognitive modeling that attempts to provide the modeler with maximal freedom in developing their model while imposing minimal assumptions on the model. The environment builds upon the box and arrow notation popular within information processing psychology, while at the same time addressing key limitations of that notation by underpinning it with concepts from object-oriented programming and design. Thus, a COGENT model consists of a collection of computational "objects", with each object being an instance of a specified class. Standard classes include buffers, rule-based processes and networks. In addition, each class has a set of properties. For any instance of an object, the values these properties fully define the computational behavior of the instance. Thus, a box within a box and arrow diagram may be specified as a short-term limited capacity buffer by specifying that the box is an instance of the buffer class (or a subclass of the buffer class) and then by specifying appropriate values for the buffer's capacity and decay parameters.

Three COGENT classes are particularly relevant to the models described here. *Buffers* are components that allow information to be stored, either temporarily or permanently, and either with or without capacity limitations. *Rule-based processes* contain sets of condition-action rules. The conditions of such rules may match elements in buffers, or carry out more complex logical operations (such as matching an element within a list), while the actions may modify buffers or send messages to other components. Rules may also have triggering patterns associated with them. Such rules will only fire when the process that contains them receives a message that matches their trigger. By default, all rules within a rule-based process will operate in parallel. That is, if multiple rules match on a single

processing cycle, or if a rule can be matched in multiple ways on a given processing cycle, then all instantiations of those rules will be fired. Finally, *interactive activation networks* are components that contain a set of labeled nodes, each with an activation value. Nodes may be excited or inhibited by messages sent to a network, or by processes internal to the network such as self activation and lateral inhibition.

Processing within COGENT is based on a simple blackboard model, in which all components (i.e., all boxes, or equivalently, all specified instances of the various classes) operate in parallel, reading from the blackboard at the beginning of the processing cycle and writing to it at the end. Thus, in a model with several rule-based processes, those processes will all operate in parallel. They will also operate in parallel with buffers (whose contents may be decaying over time) and with interactive activation networks (whose nodes will be increasing or decreasing in activation over time).

As should be clear from the above, COGENT allows the inclusion of symbolic and connectionist mechanisms within the same model. The mechanisms are just specified via boxes of different classes, with the standard blackboard processing model which allows communication between all components within a heterogeneous model. These features make COGENT an attractive environment within which to develop the mix of automatic low-level processes assumed within the CS and the higher-level processes within the SAS. Furthermore, in contrast to Soar and ACT-R, and as highlighted above, COGENT does not specify any particular theoretical architecture within which models must be placed. It therefore supports the inclusion of only aspects of theory that are deemed especially relevant.

4.2   General principles of the model

The COGENT model consists of two over-arching components: a subject model and an environment with which the subject model interacts. The role of the environment is to present the subject model with tasks (the six tasks used in the original Waldau (1999) study) and record and collate the subject model's behavior. Programmatic details of the environment are not considered further here, as it is the subject model that is of prime importance.

4.2.1   The Subject
The arrangement of sub-processes adopted for the subject model (for both younger and older children) is displayed in figure 3. This shows the influence of the domino model and Norman and Shallice's CS/SAS theory. Within the figure, round-edged oblongs represent buffers that store information during processing, while hexagonal boxes represent processes that operate on buffers or transform information.

The basic operation of the model is as follows. At the beginning of a problem **Current State** and **Desired State** (parts of the environment) are initialized with representations of the problem. The **Subject** model derives representations of the **Current State** and **Desired State** through **Perception of World**. This process extracts simple properties of the task and maintains representations of the **Current State** and **Desired State** in **Working Memory**. Problems are recognized as problems by **Monitoring & Goal Generation** if the **Current State** and **Desired State** do not exactly match. When there is a discrepancy between these two states, a message is produced by **Monitoring & Goal Generation** and sent to **Goals.** This triggers the use of existing strategies aimed at reducing the difference. **Strategies** delivers representations of the immediate-hits and one-move look ahead strategies to **Candidate Strategies**, which are then analyzed by **Evaluate Solutions**.
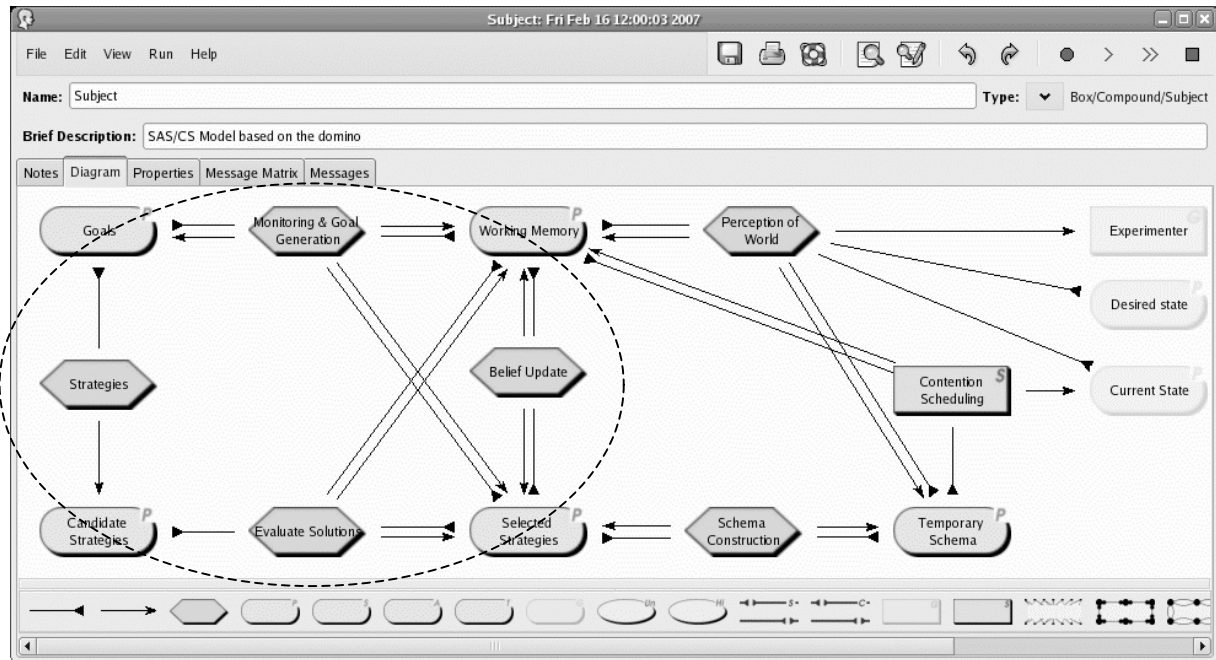
Figure 3. The processes within Subject

4.2.2 Evaluate Solutions

Evaluate Solutions provides intensive processing of information represented within **Working Memory** and to a lesser extent **Selected Strategies**. The primary objective for this process is to evaluate the outcome of proposed solutions, or moves. **Evaluate Solutions** is responsible for identifying immediate-hits (see Figure 4) and look ahead moves (see Figure 5). In the event where none exist, **Evaluate Solutions** starts a process whereby possible moves are proposed to **Working Memory**. **Evaluate Solutions** calculates what the resultant state would be if that possible move was actioned. If an immediate-hit is possible given a resultant state, the possible move is initiated.[1] If a possible move does not yield an outcome whereby an immediate hit is possible then it is temporarily black-listed in **Working Memory** and another possible move is explored. Once a decision has been made to move a ball, automatic processes within **Contention Scheduling** take over.

```
IF the strategy "Immediate Hit" exists AND
      Ball X is Free to Move AND
      Ball X is not currently Held AND
      The Position for Ball X is Free, THEN
      Initiate the Action Pick Up Ball X to go to Position Y
ENDIF


IF the strategy "Immediate Hit" exists AND
      Ball X is Free to Move AND
      Ball X is currently Held AND
      The Position for Ball X is Free, THEN
      Initiate the Action Put Down Ball X to go to Position Y
ENDIF
```

Figure 4. Pseudo code showing rules pertaining to the Immediate-hit strategy

The left-hand side of the model (highlighted by the dashed ellipse, Figure 3) is thus given over to decision-making as strategies are proposed and evaluated. For each selected strategy action schemas

---

[1] This tendency to abandon full searches of the problem space for other beneficial moves is consistent with the literature (Gilhooly, 2002).

are created (via **Schema Construction**) and fed into **Temporary Schemas**, serving to excite elements within **Contention Scheduling**.

```
(1)
IF the strategy "Look ahead One" exists AND
    no Immediate Hits are available, Then
    Identify any available Ball X
    Identify Possible Position (1) for Ball X
    In Working Memory create representation of state if Ball X moved
to Position (1)
ENDIF
(2)
IF an Immediate Hit is available given Working Memory state, THEN
    Initiate Action move Ball X to Position (1)
ENDIF
IF No Immediate Hit available,
    Identify Possible Position (2) for Ball X
    In Working Memory create representation of state if Ball X moved
to Position (2)
ENDIF
(3)
IF an Immediate Hit is available given Working Memory state, THEN
    Initiate Action move Ball X to Position (2)
ENDIF
(4)
IF No Immediate Hit Available, Then
    Identify any available Ball Y
    Identify Possible Position (1) for Ball Y
    In Working Memory create representation of state if Ball Y moved
to Position (1)
ENDIF
…
[Repeat steps 2-3]
```

Figure 5. Pseudo code showing rules pertaining to the one-move look-ahead strategy

4.2.3   Contention Scheduling
Within **Shared Schema Hierarchy** (see figure 6), two basic types of action schemas ('pick up a ball' and 'place ball on a peg') exist that may be applied to six specific actions (*pick up red*, *pick up blue*, *pick up green*, *put down on left peg*, *put down on centre peg*, *put down on right peg*).

In the absence of any strategy (perceptual, or otherwise) schemas corresponding to all possible moves receive excitation from the representation of the current state. Specifically, individual schemas that correspond to picking up a color ball receive activation if that ball is free to move and schemas for putting a ball down on a peg receive activation if the ball is held and there is space on a peg for a ball. A move is selected for action and carried out through **Act** (see figure 7) if the activation of its schema passes a threshold value of 0.75. Processes of lateral inhibition and self excitation ensure that, even in the absence of strong top-down excitation, one or more schemas will become active and be proposed for action. Though weak inhibition operates between competing schemas it does not prevent two schemas becoming selected for action at the same time. Thus, the basic CS mechanism is prone to rule breaks that will occur if two schemas exceed the selection threshold simultaneously.

The behavior of the CS in the absence of top-down input is random – moves are selected and performed subject to environmental constraints but with no concern for the achieving the goal state. Whilst correct solutions are eventually reached, they may take several hundred moves. This basic operation of CS is modulated, however, by one low-level perceptual strategy and two higher-level processes within the SAS.
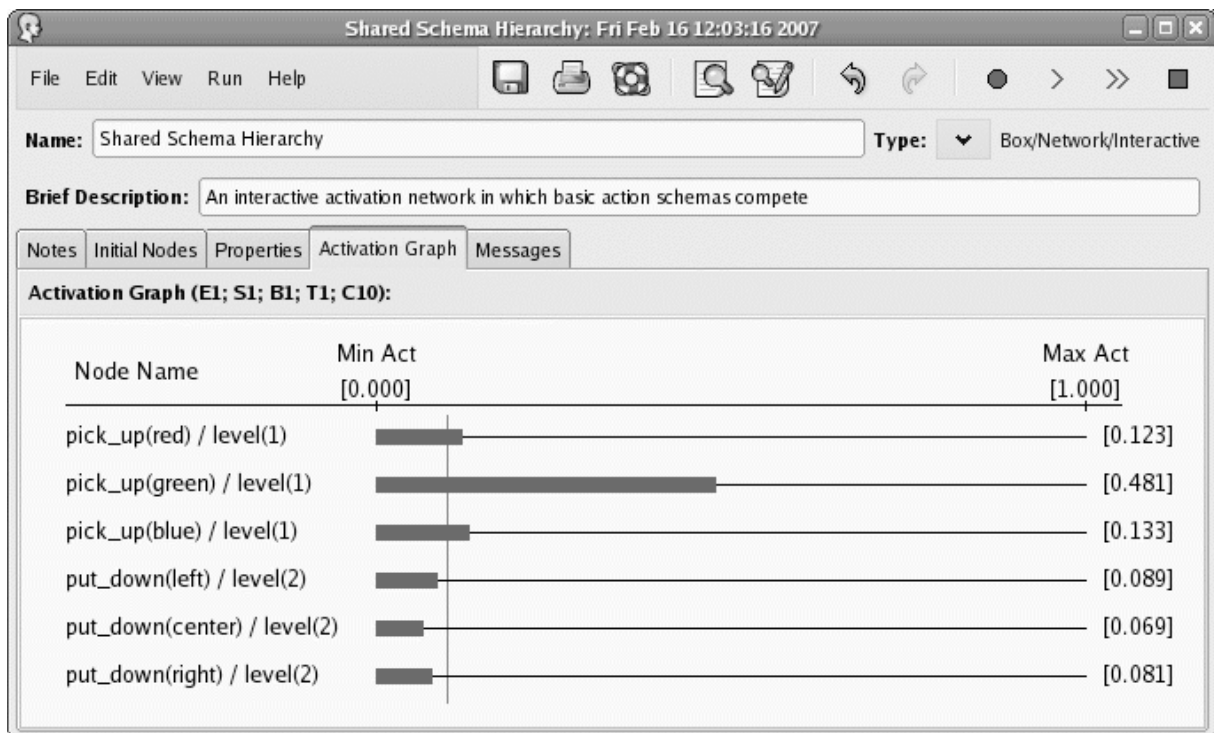
Figure 6. The interactive-activation network of Contention Scheduling

Within **Contention Scheduling** a direct perceptual bias, to move balls to match the configuration of the goal state, provides an additional early influence in proposing moves. The level of excitation that schemas receive through this perceptual bias is dependent on the level of similarity shared between the configuration of the **Current State** and **Desired State**. That is, moves leading to one configurally correct ball get 0.22 excitation, moves leading to two get 0.44, and moves leading to three get 0.66. This low-level bias organizes behavior in such a way that moves are no longer random and the total number of moves required to solve a problem is greatly reduced, but any ball that is placed in its correct color position is purely coincidental. This bias to move a ball to a location that matches the overall configuration of the goal state, but not necessarily the correct color position mirrors findings from experimental data and serves to allow basic responses to take place in more complex situations where existing strategies do not appear suited.

Two further strategies that originate from the SAS provide more powerful analysis of the problem space. These are the (1) immediate-hit and (2) one-move look-ahead strategies. The immediate-hit strategy influences moves that relate to the immediate placement of a ball in its correct position if the ball is free and the target position is free. The one-move look-ahead strategy provides a representation of the resultant state of possible moves and assesses them with respect to whether they provide opportunities for an immediate-hit.

Figure 7 illustrates the processes within **Contention Scheduling** in which action schemas (such as 'pick up green ball' and 'put down on centre peg') are ultimately produced. Automatic processes within **Trigger Schemas** operate on (1) balls that can be moved and (2) pegs that have space, by reading from the **Current State**. Once a move is made, the contents of **Current State** and **Working Memory** are updated to reflect the new positions of balls and the process of determining the next possible move begins.
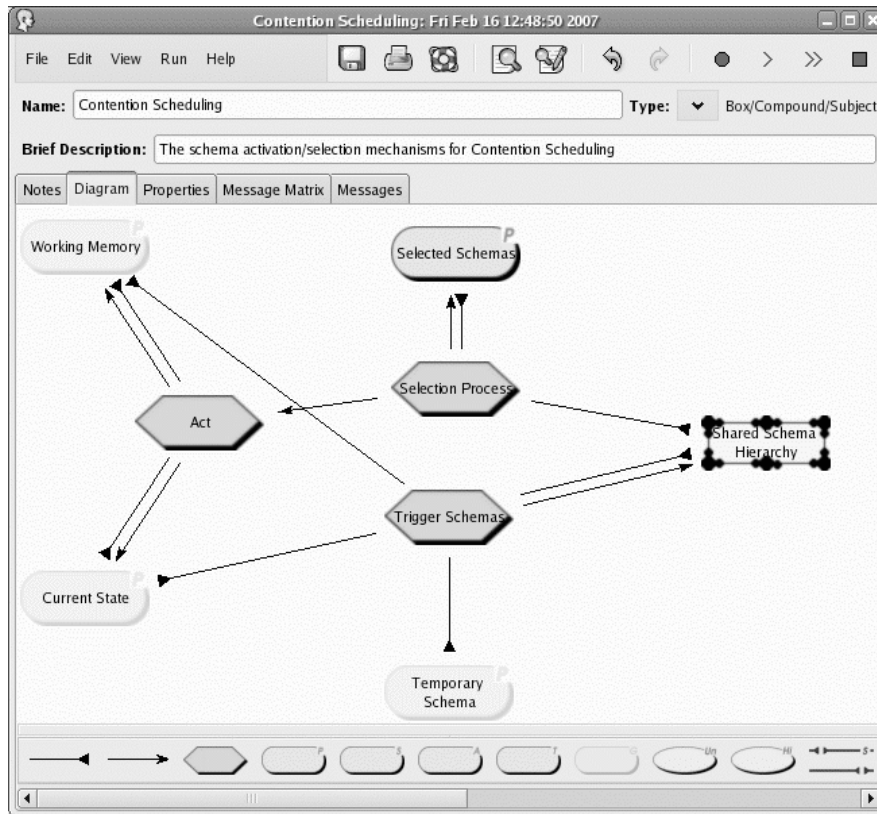
Figure 7. The contention scheduling process

4.3 The Younger Child model

The performance of the Younger Child model is influenced by the existence of two strategies and a simple perceptual bias to match the current configuration to the Goal State. The model was run 17 times on the six problems from the empirical study. The fit of the dependent measures with the experimental data appears good for the 3-4 year olds on all criterion measurements (see Table 2), though statistical comparisons reveal significant differences on rule breaks and average number of moves.

Table 2. Summary of behavior of younger children and model 1

	3-4 yr olds	Model 1	t(df = 32)	p
Configuration (%)	95.83	100.00	−1.33	0.193
Colors (%)	**66.32**	**65.74**	**0.35**	**0.729**
Rule breaks (%)	**52.08**	**44.11**	**2.35**	**0.025**
Avg. no. moves	11.26	6.8	4.95	0.001

Two interrelated processes may account for both the simulated number of rule-breaks and the proportion of balls in their correct positions. Firstly, in the absence of strong lateral inhibition between nodes, activation of two nodes e.g., 'pick-up red' and 'pick-up blue' can reach threshold and result in both actions being taken. As a ball is picked up, it may reveal another ball to which a second strategy applies. Hence, the schema to pickup the second ball can reach threshold before the first ball has been placed. This pattern of behavior is consistent with observations of rule breaks by children.

Secondly, partial completion (or, the mixture of lower number of correct colors and high number of correct configurations) is explained as the result of a combination of effects of the immediate-hit strategy and the direct, configural bias within **Contention Scheduling**. Both of these influences are concerned with immediate perceptual properties of the **Current State**. In the case of the former, the strategy is concerned purely with placing a ball in its target position and does not process the placement of other balls. In the latter, only configural properties are processed. In the course of

problem solving the perceptual features of a problem are present *before* the results of processing of the various strategies have been carried out (i.e., these more intensive processes take longer to return proposed moves). **Contention Scheduling** (containing the bias for configural similarity) is dependent on perceptual information only and so has an early advantage at influencing the selection of schemas. Thus, the direct influence of **Contention Scheduling** goes unchecked and impacts on the total number of correctly placed balls. The lack of co-ordination between **Contention Scheduling** and supervisory processes suggests the need for greater monitoring and control to *inhibit* the influence of simple and direct perceptual biases that originate from **Contention Scheduling**.

4.4   The Older Child model

In the Older Child model a mechanism of inhibition is introduced in an attempt to limit the occurrence of rule breaks and simulate the performance of the 5-6 year olds. Operationally, the second model extends upon the first by interrupting the combined influence of immediate-hits and configural bias, thereby enabling a greater degree of influence from higher-level strategies.

The difference between the Younger Child model and Older Child model is embodied as a single rule within **Monitoring & Goal Generation** (see Figure 8 below). If its condition is met a more detailed examination of the positions of other balls in the current state is triggered. If this reveals the ball under the target position for the immediate-hit is not in place, the strategy is terminated and a new move considered.

```
IF the Action to Move Ball X to Position 1 has been Initiated AND
    the Placement for Ball X is Above any Ball Y, THEN
    Halt the Action to Move Ball X to Position 1 AND
    check Placement for Ball Y is correct
ENDIF
```

Figure 8. Inhibiting the immediate-hit strategy if the ball under is not in place

A comparison of the behavior of this model (run over 17 attempts at each of the six problems) and older children is given in Table 3. Overall, it appears that inhibition holds considerable weight on the overall behavior of the Older Child model. The proportion of rule breaks is reduced and the overall number of balls in the correct color position is higher than in the Younger Child model. T-tests on these measures reveal that the difference between model and human performance does not differ significantly for configuration or color, while the difference for rule breaks is marginally significant (with the model producing more rule breaks than the human participants). The average number of moves to completion also differs significantly between the model and human data. This discrepancy is discussed below. We argue that while it is an essential component of older children's performance, the final outcome on each task is reliant on a number of other processes. This interpretations fit well both with diversity accounts and studies emphasizing a strong involvement of inhibition on tasks of executive functions, including the Tower of London (Miyake et al., 2000).

Table 3. Summary of behavior of older children and model 2

	5-6 yr olds	Model 2	*t*(df=32)	*p*
Configuration (%)	96.05	100.00	−1.32	0.196
Colors (%)	**94.74**	**95.05**	**−0.71**	**0.483**
Rule breaks (%)	**21.92**	**30.47**	**−2.05**	**0.049**
Avg. no. moves	9.79	7.17	2.70	0.011

The Older Child model adds an important feature that serves to inhibit actions based on simple and direct perceptual biases. These biases are suppressed via a rule that triggers a deeper search of the problem state. The effects of this mechanism within the **Monitoring & Goal Generation** process is in

reducing the chances of the model being 'led astray' by superficial characteristics of the problem and increasing the proportion of balls being placed in their correct color position.

5   General discussion

The models presented here integrate a number of aspects from the work by Fox and Das (2000) and Norman and Shallice (1986) to achieve a framework capable of testing a candidate role of inhibition on the Tower of London. Overall, these models demonstrate a good fit of the key dependent measures of task completion and rule breaks for 3-4 and 5-6 year olds.

The effect of the simple perceptual bias enables moves to be made towards the configuration of a goal. Determining which ball to move and which peg to place it at is a result of competition within an interactive-activation based network; with the level of configural similarity governing the amount of excitation competing nodes receive. If configural similarity is high and either no strategies apply, or more extensive processing is required to apply a strategy, this bias has a stronger possibility of influencing a move to configuration, increasing the chances of only a partially complete solution.

In our analysis we included a number of additional measures including average number of moves made on problems. Of some interest is the difference on this measure between Model 1 (6.8) and Model 2 (7.17) compared to the younger children (11.26) and the older children (9.79). Here, our models do not match the patterns observed in the child data but appear to show a slight cross-over interaction. The causes for the differences between the younger and older children may be attributable to their initial lack of apparent willingness to move balls away from the current state when the configuration matched closely that of the goal state. This may have been further compacted in instances where one ball was in its correct color position and where the goal to match the desired configuration was strongest. Though delayed within the children's behavior, backtracking was a feature of many children's behavior as it was within both models. However, whereas the models took immediate steps to backtrack and investigate other strategies, the children appeared to delay. Further work is necessary to clarify this account.

The models described here accurately simulate the performance of the younger and older children on the other dependent measures. The close fit of our models for the data on rule breaks is in our view a distinctive feature of this work. Not only do the models mirror the shift in performance on the proportion of balls in their correct position between younger and older children, but they also simulate the reduction of rule breaks. Both these results are a consequence of one view of the possible role that inhibition may play in problem solving on the Tower of London. This view established in the Older Child model was built on conceptualizations offered by Miyake et al (2000) and the implementation given to inhibition within this model accounts for a shift in performance, from one resembling the behavior of 3-4 year olds to one resembling the behavior of 5-6 year olds.

Furthermore, rather than indicating the need for one process to control overall functioning, these models demonstrate behaviors that are the result of a range of interacting processes. Although in the second model, the role given to inhibition is instrumental in accounting for specific differences between the Younger Child and Older Child models, influences of both strategies and the perceptual bias converge to affect performance. Thus, these models strongly favor diversity views of executive functions.

The work presented here is consistent with the view that younger children's poorer performance on the Tower of London is a product of their failure to inhibit simpler strategies. In contrast to the view that younger and older children possess qualitatively different cognitive strategies these models demonstrate that a lack of ability to inhibit may mask the existence of more complex skills.

The account offered here is a functional one. In this paper we have demonstrated that a computational implementation of inhibition can explain differences in performance between younger and older children. However, it remains to be explained what drives the development of mechanisms that underlie cognitive development.

References

Anderson, J. R., & Lebiere, C. (1998*). Atomic Components of Thought.* Mahwah, NJ: Erlbaum.

Altmann, E.M. & Trafton, J.G. (1999). Memory for goals: an architectural perspective. *Proceedings of the twenty-first annual meeting of the Cognitive Science Society,* 19–24. Hillsdale, NJ: Erlbaum.

Aron, A.R., Robbins, T.W., & Poldrack, R.A. (2004). Inhibition and the right inferior frontal cortex. *Trends in Cognitive Sciences*, 8, 170-177.

Blair, C., Zelazo, P.D., & Greenberg. M.T. (2005). The measurement of executive function in early childhood. *Developmental Neuropsychology, 28(2),* 561–571.

Bull, R., Espy, K. A., & Senn, T. E. (2004). A comparison of performance on the Towers of London and Hanoi in young children. *Journal of Child Psychology and Psychiatry, 45(4)*, 743–754.

Burgess, P.W. (2000). Real-world multitasking from a cognitive neuroscience perspective. In S. Monsell & J. Driver (eds.) *Control of Cognitive Processes*. Attention and Performance, XVIII, pp. 465–472. MIT Press, Cambridge, MA.

Card, S., Moran, T.P., & Newell, A. (1983). *The Psychology of Human Computer Interaction.* Lawrence Erlbaum Associates, Hillsdale, NJ.

Cooper, R.P. & Fox, J. (1998). COGENT: A visual design environment for cognitive modelling. *Behaviour Research Methods, Instruments, & Computers, 30*, 553–564.

Cooper, R.P. & Shallice, T. (1995). Soar and the case for Unified Theories of Cognition. *Cognition, 55,* 115–149.

Cooper, R.P., & Shallice, T. (2000). Contention Scheduling and the control of routine activities. *Cognitive Neuropsychology, 17*, 297–338.

Cooper, R.P. (2002). *Modelling High-Level Cognitive Processes*. Lawrence Erlbaum Associates, Mahwah, NJ.

Duncan, J. & Owen, A.M. (2000). Dissociative methods in the study of frontal lobe function. In S. Monsell & J. Driver (eds.) *Control of Cognitive Processes*. Attention and Performance, XVIII, pp. 567–576. MIT Press, Cambridge, MA.

Fox, J. & Das, S. (2000). *Safe and Sound: Artificial Intelligence in Hazardous Applications*. MIT press, Cambridge, MA.

Garavan, H., Ross, T.J., Murphy, K., Roche, R.A.P., & Stein, E.A. (2002). Dissociable executive functions in the dynamic control of behavior: inhibition, error detection, and correction. *NeuroImage*, 17, 1820-1829.

Gilhooly, K. J. (2002). *Thinking: Directed, Undirected and Creative*. Academic Press.

Glasspool D.W. (2005). The integration and control of behaviour: Insights from neuroscience and AI. In D. N. Davis (Ed.) *Visions of Mind: Architectures for Cognition and Affect* (pp. 208–234). Hershey, PA: IDEA Group.

Glasspool, D.W. & Cooper, R.P. (2002). Executive processes. In Cooper, R.P. (ed.) *Modelling High-Level Cognitive Processes* (pp. 313–362). Lawrence Erlbaum Associates, Mahwah, NJ.

Goel, V., Pullara, S. D., & Grafman, J. (2001). A computational model of frontal lobe dysfunction: Working memory and the Tower of Hanoi task. *Cognitive Sciences, 25,* 287–313.

Houdé, O. (2000). Inhibition and cognitive development: object, number, categorization, and reasoning, *Cognitive Development*, 15, 63–73.

Klahr, D. & Robinson, M. (1981). Formal assessment of problem-solving and planning processes in preschool children. *Cognitive Psychology*, 13, 113–148.

Miyake, A., Friedman, N.P., Emerson, M.J., Witzki, A.H., Howerter, A., & Wager, T.D. (2000). The unity and diversity of executive functions and their contributions to complex frontal lobe tasks: A latent variable analysis. *Cognitive Psychology, 41,* 49–100.

Newell, A. (1990). *Unified Theories of Cognition*. Cambridge MA: Harvard University Press.

Newell, A. & Simon, H.A. (1972). *Human Problem Solving*. Prentice-Hall. Englewood Cliffs, NJ.

Norman, D.A., & Shallice, T. (1986). Attention to action: willed and automatic control of behaviour. In R. Davidson, G. Schwartz, & D. Shapiro (Eds.), *Consciousness and Self Regulation* (Vol. 4). Plenum: New York.

Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society London B 298*, 199–209.

Shallice, T. (2002). Fractionation of the supervisory system. In Stuss, D.T. & Knight, R.T. (eds.) *Principles of Frontal Lobe Function*, pp. 261–277. Oxford University Press, Oxford, UK.

Shallice, T., & Burgess, P.W. (1996). The domain of supervisory processes and temporal organisation of behaviour. *Philosophical Transactions of the Royal Society of London, B351,* 1405–1412.

Simon, H.A. (1975). The functional equivalence of problem solving skills. *Cognitive Psychology*, *8*, 86–97.

Sylvester, C.Y.C., Wager, T. D., Lacey, S. C., Hernandez, L., Nichols, T. E., Smith, E. E., & Jonides, J. (2003). Switching attention and resolving interferences: fMRI measures of executive functions. *Neuropsychologia, 41*, 357–370.

Thomas, J.C. (1974). An analysis of behavior in the hobits-orcs problem. *Cognitive Psychology*, *6*, 257–269.

Waldau, R. (1999). A developmental study of problem solving in children aged 3–6: Development of planning strategies. *Unpublished undergraduate thesis,* University of London.

Zelazo, P.D., Carter, A., Reznick, J.S., & Frye, D. (1997). Early development of executive function: A problem-solving framework. *Review of General Psychology, 1,* 198–226.

Zook, N.A., Davalos, D.B., DeLosh, E.L., and Davis, H.P. (2004). Working memory, inhibition, and fluid intelligence as predictors of performance on Tower of Hanoi and London tasks. *Brain and Cognition*, *56*, 286-292.