Revisiting Interacting Subsystems Accounts of Cognitive Architecture: The Emergence of Control and Complexity in an Algebra Task

Gareth Miles[1]

[1]University of South Wales

gareth.miles@southwales.ac.uk

Abstract. Symbolic accounts of cognitive architecture most often have a central hub where information is processed (e.g. the production process in ACT-R [1]). An alternative approach is to model cognition as the interaction of multiple largely autonomous subsystems [2, 3]. This latter, Interacting Subsystems, approach is explored in the GLAM-PS cognitive architecture, a theory that operationalizes many of the assumptions of strongly grounded approaches to cognition [4]. The GLAM-PS model of problem solving in algebra is described. Control in the algebra model is passed between three subroutines when solving a problem. These subroutines emerge from the interaction of different subsystems and are not explicitly programmed into the model. By systematically varying two short-term memory parameters it is shown that the model's successful performance of the task depends on the interaction of the contributing modules, and that this interaction demonstrates complexity, with additional memory resources not always improving performance.

Keywords. Cognitive Architecture, Production System, Embodied Cognition

1 The Interacting Subsystems approach to cognition

Cognitive Architecture has been established as a key research area within Cognitive Science following seminal work between 1970 and 1990 [5, 6, 7]. Although a lot of recent work has focused on either the ACT-R Architecture (e.g. [1]) or large-scale neural network models (e.g. [8]), there remain a wide variety of approaches to modeling Cognitive Architecture (e.g. [3], [9]). The purpose of the current paper is to look at cognitive control within a particular subset of these approaches.

In ACT-R and other notable architectures cognitive control is an aspect of cognition that is explicitly modeled, with specialist cognitive modules taking responsibility for the representation of goals and the selection of action. This reflects a consensus view of how the physical brain is specialized within different anatomical areas, notably the identification of the basal ganglia with action selection, and the frontal areas of the brain with the influence of intention [1], [8]. Thus in ACT-R the matching of IF-

THEN production rules is centralized in a module that is mapped on to the basal ganglia, with the representation of goals handled by a separate module mapped on to the anterior cingulate cortex (a frontal area) [1]. In the SPAUN Architecture intention is controlled in neural networks mapped on to the frontal cortex and action selection is mapped on to the basal ganglia [8].

However there are logical objections to this approach that become particularly apparent when one examines the relationship between neural networks and production systems. Both in essence are doing the same thing, the association of an output with a particular configuration of inputs. Whilst there are clearly differences between production systems and neural networks in how areas such as partial matching of configurations, generalization and one trial learning are handled, both can be considered methods of representing configural associations. The logical objections arise because anatomically there are networks of neurons present throughout the brain and it follows that these will be able to compute configural associations. Symbolic approaches to cognition clearly indicate that configural associations are the key underlying process in action selection and cognitive control. Therefore it would seem particularly strange that these key processes are modeled as strongly centralized in leading Cognitive Architectures when configural associations can be computed in many distinct parts of the brain.

An alternative approach is found in Architectures that use distributed interacting subsystems. Barnard's Interacting Cognitive Subsystems (ICS) approach [2] to cognition and emotion theorized how such an approach could model complex tasks. In Barnard's theory there are separate morphonolexical, propositional, object and implicational subsystems, each of which processes and translates symbolic output from the other subsystems. Whilst ICS has proved influential in highlighting the potential of interacting subsystems, the approach was not computationally implemented in full and did not compute configural associations (it's subsystems simply translated one symbol into another). A more recent interacting subsystems approach is 4CAPs [3], an example of a Cognitive Architecture that was directly inspired by knowledge from neuroscience. The emphasis on 4CAPS is on modeling higher cognition, with amodal subsystems modeled including Left and Right Hemisphere Spatial and Executive centres.

The focus however within this paper is on GLAM-PS an interacting subsystems approach to embodied cognition. The idea of emergent control and action selection in a distributed system is particularly relevant to modeling embodied cognition because of the emphasis therein on modal rather than amodal cognitive systems. Modal subsystems are those directly associated with perception and action, in which the grounding of symbols (see [10]) in the external world is clearly indicated. Amodal subsystems are those that are not directly associated with perception or action (e.g. the goal module in ACT-R).

The plan for the paper is as follows, to briefly describe the GLAM-PS cognitive architecture, to demonstrate how cognitive control is modeled in a simple algebra problem solving task, and then finally to demonstrate the emergence of complexity in the algebra model by exploring the effects of small variations in the starting parameters of the GLAM-PS Architecture in the algebra task. Algebra was chosen because it is a paradigmatic task for studying Cognitive Architecture that has often been used by John Anderson to illustrate how ACT-R works (e.g. [1]). In the remainder of the paper ACT-R is used as for comparison purposes as an example of a mature, widely used symbolic Cognitive Architecture.

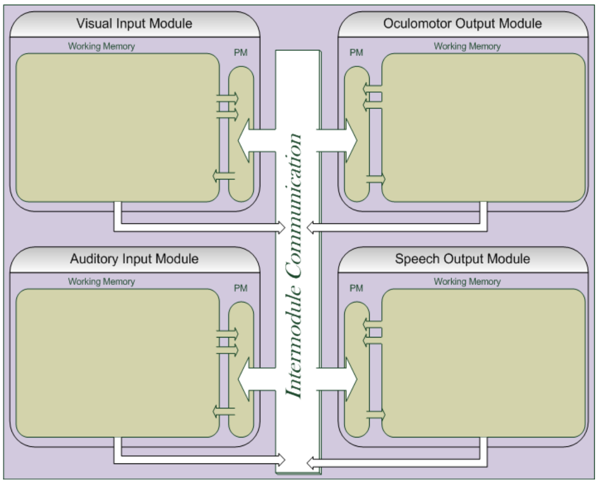

Fig. 1. A simplified view of the GLAM-PS architecture showing the four modules used in the Algebra Model and communication between these modules. PM is the Production Memory.

2 The GLAM-PS Cognitive Architecture

GLAM-PS shares a distributed modular structure with 4CAPS and ICS, however, whilst these Architectures make widespread use of amodal representation, GLAM-PS is intended to explore the implications of a strongly grounded distributed Architecture for cognition (see [4] for a review on Grounded Cognition). Whilst comparisons with ICS are difficult as it was never fully implemented, if we compare GLAM-PS to 4CAPS (arguably the most similar Architecture) it can be seen that GLAM-PS takes an outside-to-inside approach to modeling Cognitive Architecture, wherein peripheral processes dominate cognition. By contrast 4CAPS takes an inside-to-outside approach. The anatomical areas of the brain featured in 4CAPS do not map easily on to the modules described by GLAM-PS, instead the latter features modules that map on to the sensory and motor areas of the cortex. Grounded Cognition [4] suggests much of cognition is driven by these peripheral systems and a major novel contribution of GLAM-PS is to implement these ideas computationally in a symbolic architecture.

A simplified diagrammatic representation of the Architecture is shown in Fig. 1, with the two perception and two action modules used in the algebra task model included (no other modules are used for modeling this task). Both long-term and short-term/working memories are stored and revivified in the modules that originally processed what is being remembered. However, each module influences the behavior of other modules via the mechanism of inter-module communication of the current contents of working memory. In this manner the actions (productions) chosen in a module are based upon a composite view of working memory across all modules. Whilst this mostly acts like a single unified working memory there is a delay associated with inter-module communication. The implication of this is that a given module has an up-to-date view of its own working memory, but a delayed view of working memory in other modules (α is the GLAM-PS global parameter defining this delay in term of production cycles, it is set to 4 in the model reported here).

All long-term memories are stored as productions in GLAM-PS (following early SOAR [8]) using a classic IF-THEN structure. For simplicity and to improve plausibility all productions can only have a single action associated with the THEN side and the IF side is only able to check for the presence or absence of a representation (no programming code is allowed). When actions are represented in the action modules they are not necessarily executed and can be used to reason without action. Actions are only executed once they become 'Super Activated', a process whereby their activation level is raised substantially above the level needed for representation. Only once an Action Execution Threshold (global parameter β) is surpassed will the action be executed. Thus GLAM-PS is able to represent and then reason about actions without necessarily executing them.

Whilst the modules shown in Fig. 1 can be thought of as mapping on to sensory and motor areas of the brain, the processes associated with inter-module communication can be thought of as mapping on to the higher cortical areas (e.g. prefrontal cortex). This is a distinctly different interpretation of cortical function from many existing accounts. Whilst currently GLAM-PS makes no specific claims about how inter-module communication should be mapped on to the brain anatomically, it is nevertheless a potentially interesting future direction.

3 Cognitive Control in the GLAM-PS Algebra Model

The GLAM-PS Algebra Model (GAM) solves simple linear problems of the form $Ax + B = C$, for instance $2x + 4 = 10$ (where the solution is $x = 3$). To solve the problem GLAM-PS, like most human solvers [1], must proceed through three distinct stages or sub-goals, first reading and encoding the problem, then resolving the addend (the B term), before resolving the multiplier (the A term). The cognitive steps used by GLAM-PS are in essence the same as those used by Anderson's ACT-R model of the same task [4], what differs here however is how cognitive control is achieved.

Two types of cognitive control problems occur in GAM, firstly moving between sub-goals and secondly combining actions in such as was as to solve each sub-goal. The latter of these is relatively easy for GLAM-PS as it typically involves a sequence of actions where the result of the preceding action acts as the trigger for the next action in the sequence. In the failed runs reported in section 4 it is rarely the case (< 1%) that failure occurs because of a failure to sequence actions within a sub-goal, instead failures occur because the actions needed to begin a sequence that achieves a sub-goal are not initiated. Hence it is the first type of cognitive control, moving between sub-goals, that GLAM-PS finds difficult (for example beginning the process of resolving the addend once the problem has been encoded).

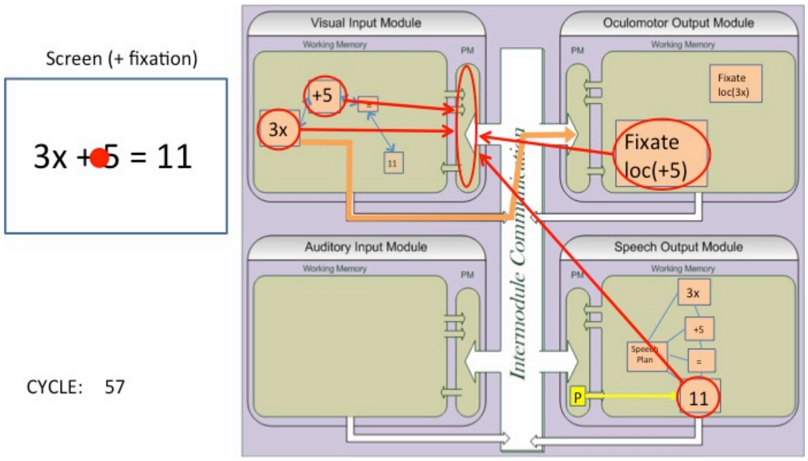

Fig. 2. Visualization of state of the GLAM-PS Algebra Model (GAM) when the control state has been established that begins transition from the Reading sub-goal to Solving the Addend sub-goal when solving $3x + 5 = 11$. Working memory elements (WMEs) are depicted as squares with area proportional to their activation. The WMEs contributing to the control state are circled in red, with their locus of action indicated by arrows pointing to the Visual Input production memory. GAM's current eye fixation is depicted on the left. Cycle indicates the number of production cycles from the beginning of simulation run.

Here we refer to the conditions that need to be satisfied to begin solving a sub-goal as the Control State. Within a distributed cognitive architecture the Control State needed to begin a new sub-goal will often be based on the state of multiple subsystems. If each of these subsystems is largely independent of one another then it can become difficult to achieve the required Control State. This is less of a problem in centralized architectures where a higher degree of control is possible and there is no need to coordinate representations across multiple subsystems. The control state needed to move between reading the algebra problem and solving it is shown in Fig 2.

In Fig 2. the state of GLAM-PS's working memory is visualised after the GAM has read the equation. As well as visual representations of the equation in the Visual Input module, GLAM-PS also has phonological representations of the equation in the Speech output module, the result of having read the equation. The lines between representations indicate structural links. The control state necessary to begin the solving of the equation by unwinding the addend consists of four representations across three different modules, these are the visual representation of the '$3x$' and the '$+5$', the oculomotor representation of the '$+5$' location (indicating attention is focused on the '$+5$') and the phonological representation of the '11' (indicating that the last element of the equation has been read and thus that the equation has been encoded). The production that matches this control state is a Visual Input production that acts by inhibiting the representation of the '$+5$'. Once this representation is inhibited a sequence of actions is initiated that relocates the '$+5$' to after the '11' in the equation (using imagery that is projected into the visual input module), GLAM-PS then changes the sign and then computes their combined value (eleven minus five).

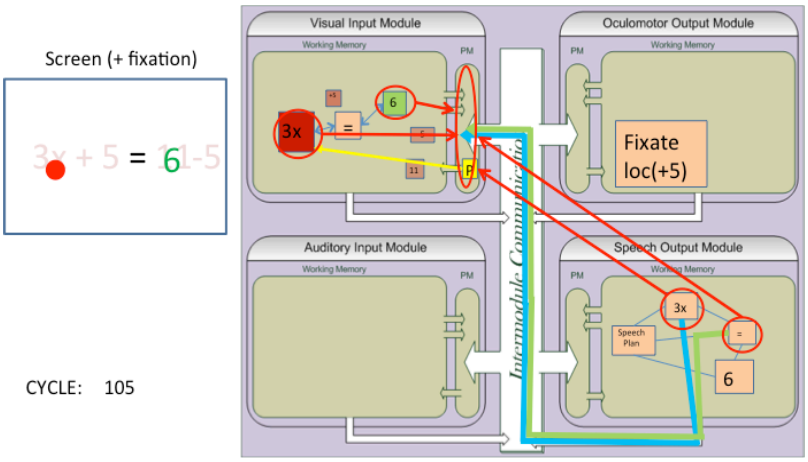

Fig. 3. Visualisation of the state of GAM when the control state has been established that begins the sub-goal of resolving the multiplier (the '3' in '$3x$'). See caption to Fig 2. for key.

The control state that is required to move between resolving the addend and the subsequent sub-goal of resolving the multiplier (the '3' in the example) is shown in Fig 3. Again, the control state is established through the combined presence of four working memory elements, this time across two modules. This consists of visual representations of the '$3x$' and a projected/imagined '6' (the result of the last sub-goal) and adjacent phonological representations of the '$3x$' and the '$=$', together these confirm that the addend has been resolved (the phonological representation is needed to confirm there are no other unresolved terms on the '$3x$' side of the equation). The sequence of actions needed to resolve the multiplier is then initiated by a production in the Visual Input module that inhibits the '$3x$' visual representation, allowing it to be subsequently broken into '3' and 'x' elements using imagery.

A key point is that in both of the transitions illustrated in Fig. 2 and Fig. 3 the control state consists of combinations of perceptual and motor representations, each of these representations is also used for perception or action (respectively), there are no abstract context or goal representations to force a particular cognitive subroutine to take control. This compares to ACT-R and other architectures where sub-goaling is used to ensure that only productions that solve the active sub-goal can be matched and executed, by contrast in GLAM-PS all productions are considered all of the time by the production matching process. Despite this GLAM-PS is able to demonstrate both task sufficiency and subroutine following in an Algebra task that can be considered a classic sub-goaling paradigm. This control is characterised as emergent because of the absence of any explicit control process within the modelling of the task.

In conclusion cognitive control in the GLAM-PS Algebra Model emerges from the interaction of working memory elements in multiple cognitive subsystems. When information from these different subsystems is combined there is sufficient information to indicate what actions the systems has taken previously and what still needs to be achieved. In Taatgen's work on the Minimal Control Principle [11] he indicates that often there will be sufficient information in a system to control action with only minimal need for explicit control representations. Whilst Taatgen clearly imagines that some form of goal representation will remain, in this GLAM-PS model there is no need for explicit goal representation. In short control is totally emergent [12]. Whether some form of goal representation would be needed once a more complex, multi-faceted model is considered is an open question. Certainly sometimes people want to simply read and equation, whilst at other times they need to solve them, though it could be the case that there are always enough clues in the external or internal environment to distinguish the two scenarios and establish an appropriate Control State.

4 The Emergence of Complexity in Interacting Cognitive Subsystems

Symbolic cognitive architectures often behave in a very predictable way, something that is often true of Production System Architectures. Once a set of productions has been 'programmed' into the system then these productions will provide a stable model of performance. This typically reflects the explicit use of goal representation that guides performance toward the achievement of that goal. Failure to achieve the goal would typically be modelled by the forgetting of the goal due to distraction [13]. Sometimes multiple strategies of achieving a set goal might be modelled and it is often the case that random 'noise' parameters will be used to help capture the variation in human performance that is observed from trial to trial in individual participants (e.g. [14]).

Much of the stability seen in established architectures is the result of centralised decision making. For example only one goal can be followed at a time in ACT-R [1] (though see [15]). When an architecture utilising multiple Interacting Subsystems is considered then complexity and instability may well emerge from the unpredictable interaction of the multiple distinct decision cycles in the component subsystems. If information from one module arrives at another module just one decision cycle later in one simulation run as compared to another, then the behaviour of the whole system might change very significantly over the full course of that run.

In order to explore the nature of the interaction of the multiple subsystems used in the GLAM-PS Algebra Model (GAM) a series of 1,170 simulation runs were conducted of the model with systematic variation of two working memory parameters. Note that the model used in the runs was deterministic without any randomised elements.

Working memory in GLAM-PS is module specific, with each module's working memory currently governed by the same global parameters and equations. Each working memory element has an activation varying from 0 to 1. To be matched by a production then a working memory element must have an activation greater than global parameter γ. Each working memory also has a total activation limit, global parameter δ. If the creation or change in activation of a working memory element takes the total activation within a module's working memory above δ, then the activation associated with all other elements in that module's working memory is adjusted so that total activation is equal to δ.

To explore the impact of small changes in working memory availability on GAM the parameter γ was systematically varied from .01 to .39 in increments of .01, this was combined the systematic variation of δ from 1.0 to 3.9 in increments of .1. On each simulation run the total number of cycles taken to solve the algebra equation $3x + 5 = 11$ was measured. The results of these simulation runs are displayed graphically in Fig 4.

The first aspect to consider of the results of these simulations runs is the vulnerability of the GAM model to failure. As δ dips below 3.0 and as γ increases it becomes increasingly more likely that GLAM-PS will not be able to solve the equation. An examination of failed runs clearly indicates that almost all (>99%) result from the failure to establish a control state that allows transition between one sub-goal and the next. According to the GAM model establishing that one sub-goal has been completed and then finding a suitable way to begin the next is difficult and prone to failure if working memory is compromised (e.g. by distraction). This broadly fits in with what has been observed in human participants, who typically take more time to complete steps of a problem that involve starting a new sub-goal [13].

The second aspect we see in the simulation runs is the emergence of complexity. One might reasonably expect that as each module's total working memory capacity, δ, increases then the likelihood of solving the equation would also increase. This is broadly the case, but there are many exceptions to this shown in Fig 4. Similarly as the production matching process becomes increasingly strict, matching fewer working memory elements (as γ increases), one would expect failures to become more likely, but again this is not always the case.

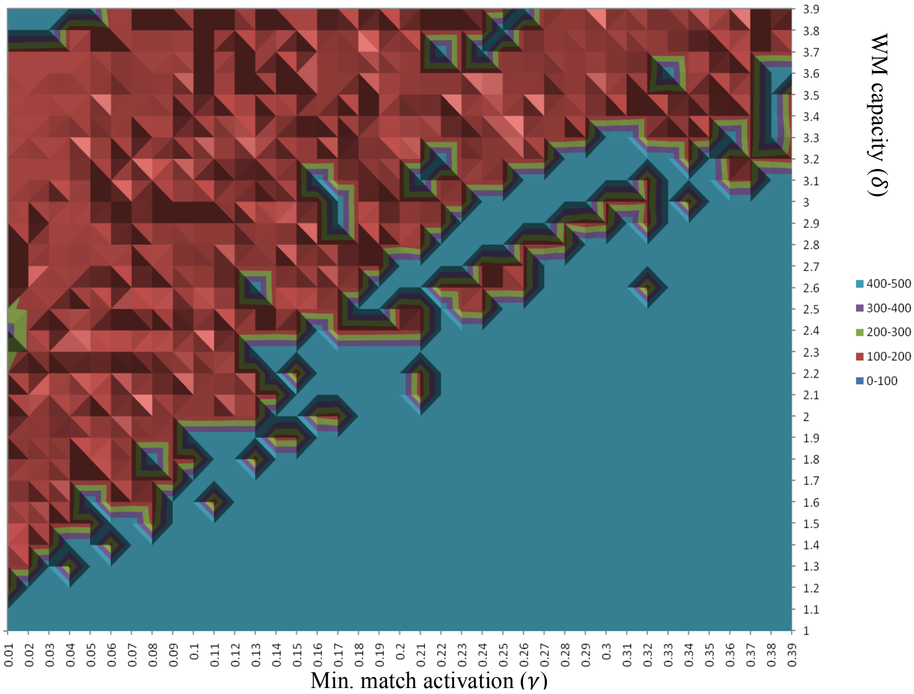

Fig. 4. A graphical display of the number of decision cycles needed to complete the equation $3x + 5 = 11$ by the GLAM-PS Algebra Model when working memory capacity (δ) and the activation needed to match productions (γ) were systematically varied. Light blue indicates the model did not solve the equation. The data is displayed in partial 3D and lit from the x-axis.

Indeed if one examines Fig 4, the parameters determining failure and success appear to influence these outcomes in a non-linear manner. If one considers the point where δ = 2.3 and γ = .13 then GAM fails, yet if we were to either increase or decrease either parameter by a fraction then GAM succeeds. Instead of a smooth curve or a straight line defining the regions where we see success versus where we see failure, what is shown in Fig 4. has more similarity to a geographical coastline. Even where there are successes the number of cycles taken to succeed varies unpredictably, the smoothness of the area in the top left (around δ = 3.5, γ = .03; though note the failures at $\delta > 3.7$, $\gamma < .04$) can be contrasted with the peaks and troughs found in other areas where successes prevail (e.g. around δ = 3, γ = .1, the default parameter settings). The pattern observed in Fig 4. reflects the chaotic nature of the interaction of the multiple subsystems in the GLAM-PS Algebra Model. In short, complexity emerges from Interacting Subsystems.

5 Conclusion

The GLAM-PS Algebra Model demonstrates how both cognitive control and complexity emerges from the Interaction of Multiple Subsystems in Cognitive Architectures that adopt an Interacting Cognitive Subsystems approach [2]. The model is notable for not using any explicit goal representation, instead showing how control is based on Control States in working memory. Each of these Control States contain sufficient information about what the system has done previously and about what the system needs to do, to enable the initiation of purposeful, self-perpetuating sequences of behaviour. The simulation runs reported, exploring working memory parameter space, demonstrate how the model is vulnerable to failure when working memory is reduced or compromised, and how the interaction of cognitive subsystems is chaotic and somewhat unpredictable in nature.

References

1. Anderson, J. R.: How Can the Human Mind Occur in the Physical Universe? Oxford University Press, New York (2007)
2. Barnard, P. J.: Interacting Cognitive Subsystems: Modelling Working Memory Phenomena Within a Multiprocessor Architecture. In A. Miyake & P. Shah (Eds.) Models of Working Memory: Mechanisms of Active Maintenance and Executive Control, pp. 298-339, Cambridge University Press, Cambridge, UK (1999)
3. Just, M. A., Varma, S.: The Organization of Thinking: What Functional Brain Imaging Reveals About the Neuroarchitecture of Complex Cognition. Cognitive, Affective, & Behavioural Neuroscience 7, 153-191 (2007)
4. Barsalou, L. W.: Grounded Cognition. Annual Review of Psychology 59, 617-645 (2008)
5. Newell, A., Simon, H. A.: Human Problem Solving. Prentice-Hall, Englewood Cliffs, NJ (1972)
6. Anderson, J. R.: The Architecture of Cognition. Harvard University Press, Cambridge, MA (1983)
7. Newell, A.: Unified Theories of Cognition. Harvard University Press, Cambridge, MA (1990)
8. Eliasmith, C.: How to Build a Brain: A Neural Architecture for Biological Cognition. Oxford University Press, New York (2013)

9. Sun, R.: Duality of the Mind: A Bottom-up Approach Toward Cognition.: Lawrence Erlbaum Associates, Mahwah, NJ (2002)
10. Harnad, S.: The Symbol Grounding Problem. Physica D 42, 335-346 (1990)
11. Taatgen, N. A.: The Minimal Control Principle. In W. Gray (Ed.), Integrated Models of Cognitive Systems. Oxford University Press, New York (2007)
12. Cooper, R. P.: Cognitive Control: Componential or Emergent? Topics in Cognitive Science 2, 593-613 (2010)
13. Altmann, E. M., Trafton, J. G.: Memory for Goals: An Activation-Based Model. Cognitive Science 26, 39-83 (2002)
14. Lovett, M. C., Daily, L. Z., Reder, L. M.: Modeling Individual Differences in Working Memory Performance: A Source Activation Account in ACT-R. Cognitive Science 25, 315-353 (2001)
15. Salvucci, D. D., Taatgen, N. A.: Threaded Cognition: An Integrated Theory of Concurrent Multitasking. Psychological Review 115, 101-130 (2008)