Interdisciplinary Foundations for Multiple-task Human Performance Modeling in OMAR

Stephen Deutsch (sdeutsch@bbn.com)
BBN Technologies
10 Moulton Street
Cambridge, MA 02138

Abstract

The Operator Model Architecture (OMAR) provides a computational framework in which to develop human performance models that generate reasonable multiple-task behaviors. An interdisciplinary foundation that reached beyond the experimental psychology and artificial intelligence literatures was considered essential to the construction of successful models. Brain imaging and clinical studies suggest that tasks are accomplished through the coordinated execution of function-specific perceptual, cognitive and motor capabilities. These studies together with philosophically grounded cautions, further suggest that the mediation of task contention be accomplished in a framework that does not require an executive that manages task execution. The computational framework for building models sensitive to these considerations is described. Examples from a commercial air traffic control domain are used to illustrate OMAR modeling capabilities.

1. Introduction

The Operator Model Architecture (OMAR) provides a simulation environment in which to model human operators, the workplaces at which they operate and the entities of the larger world that are reflected in their workplaces. An important goal has been to provide human performance models with sufficient fidelity to usefully explore and develop operating procedures for complex environments. Much of the research has focused on the commercial air traffic control environment with aircrews and air traffic controllers as the principal players. Each of the players typically has several tasks in process and interruptions are commonplace. To address the fidelity requirement, the OMAR operator models must exhibit reasonable multiple-task behaviors.

The modeling of multiple-task behaviors has been explored extensively in EPIC (Meyer & Kieras, 1997), and SOAR (Newell, 1990; Laird, Newell, & Rosenbloom, 1987) has also been adapted to model multiple-task behaviors. In particular, Meyer and Kieras (1997) report considerable success in developing a production rule-based model of the psychological refractory period (PRP) procedure. The basic components of their model are a cognitive processor comprised a production rule interpreter with inputs from long-term and production memory, and a working memory, with auditory and visual processor inputs that interact with the production rule interpreter. The model relies heavily on a centralized, synchronous production rule framework. A production rule-based executive process administers the task scheduling strategy for regulating the execution of competing tasks. The implementation is just one of a theoretically infinite number of computational frameworks that might give rise to human-like multiple-task behaviors. In building the OMAR framework, particular attention has been paid to developing multiple-task behaviors from an assembly of concurrently operating functional centers absent an executive or central controller.

The motivation for this approach to human performance modeling, derived from a selective a reading in several disciplines, is outlined in the Section 2. Section 3 provides background on the aircrew/ATC domain and describes implementation of the models of aircrew in-person conversations and their interruptions by ATC directives. Section 4 provides a description of the computational elements for constructing OMAR human performance models.

2. Interdisciplinary Foundations for Modeling Multiple-task Behaviors in OMAR

The process of building a human performance model capable of emulating the operators of a complex system is a somewhat speculative endeavor at best. Drawing on the research from a number of disciplines, a modest goal of this undertaking has been to put in place a neuropsychologically motivated distributed processing framework from which to explore some of the human performance issues, related principally to multiple-task behaviors, that impact the operation of complex systems. The modeling environment developed is symbolic, but does not preclude the inclusion of connectionist components.

Over the years, experimental psychologists have conducted extensive experiments providing a wealth of interpreted data, philosophical discussion dates back through the millennia, and more recently, cognitive neuroscience and clinical studies have provided electroencephalography (EEG), magnetoencephalography (MEG), positron emission tomography (PET) and functional magnetic resonance imaging (fMRI) of the brain at work (Posner, 1993; Raichle, 1994) that identify the locus of specific perceptual, cognitive and motor functionalities. Many readings of this literature are possible and it should not be surprising that the computational architecture for models developed in the OMAR framework differs from that of EPIC and SOAR in fundamental ways: (1) stimuli impinge directly on, activate, and propagate through long term procedural memory—the knowledge of how to do things (see Figure 1); (2) tasks, skilled cognitively-driven behaviors, are accomplished through the coordinated actions of function-specific procedures representing the contributions of specific brain areas; (3) to the extent that the resulting behaviors may be considered intelligent, that intelligence is the product of the pattern matching implicit in the changing sensitivities of the network of procedures as stimuli evoke responses at network nodes; (4) task contention outcome, rather than being determined by a central executive, is mediated on a pair-wise basis among contending tasks. The foundations for these choices are discussed in this section.

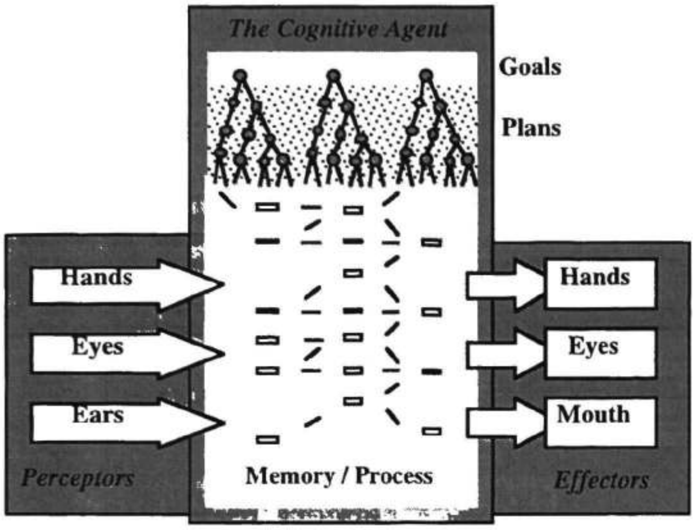

Figure 1: The OMAR Human Performance Model

The now routine accounts of the early PET studies and the more recent fMRI studies portray the execution of each experiment as being the product of a small number of brain centers—small areas of activity at widely dispersed major brain centers. Posner, Peterson, Fox, and Raichle (1988) draw on the evidence of a series of their PET experiments to suggest that "the mental operations that form the basis of cognitive analysis are localized in the human brain." To further support their assertion of the localization of cognitive function, they cite studies of patients with lesions and their related deficits. Based on these studies, the basic architectural framework seems reasonable well established. Tasks, made up of perceptual, cognitive and motor components, appear to be accomplished through the collective actions of small specialized areas of activity that take place in each of several widely dispersed brain centers.

On a closely related but more conjectural plane, Edelman (1987) discusses the psychological functions of "development, perception (in particular, perceptual categorization), memory, and learning" and how they relate to the brain. Edelman (1989) extends his analysis to consider "perceptual experience—the interaction of memory with the present awareness of the individual animal," that is, perceptual awareness and conscious experience. He describes neural maps as the ordered arrangement and activity of large groups of neurons as distinct from single-neuron connections. They are highly and individually variant in their intrinsic connectivity. Changes in the behavior of the network are the result of changes within particular *populations* of synapses. "These structures provide the basis for the formation of large numbers of degenerate neuronal groups in different repertoires linked in ways that permit reentrant signaling" (Edelman, 1987, p. 240) where, in *degenerate* systems, functional elements in a repertoire may perform more than one function and a function may be performed by more than one element (Edelman, 1987, p. 57). Reentry is a basic mechanism suitable for synchronizing the neuronal activity across the mappings at diverse hierarchical levels. *Global mappings* have a dynamic structure that reaches across reentrant local maps and unmapped regions of the brain to account for the flow from perception to action. Motor activity, an essential input to perceptual categorization, closes the loop.

Moscovitch (1994), while ascribing central-systems functions to the prefrontal cortex, describes it as "a large, heterogeneous structure consisting of a number of distinct areas, each with its own projections to and from other brain regions and each having presumably different functions (Pandya & Barnes, 1987)." He goes on to state that "... the functions of other smaller regions can also be distinguished one from another (Goldman-Rakic, 1987; Petrides, 1989)."

Taken together, Moscovitch, Posner et al. and Edelman present a picture of the execution of a task as the coordinated activities of small, specialized local sites operating at several remotely located brain centers. In Edelman's terms, reentrant signals link the components within the local sites, while global mappings connect the activities of the broadly dispersed major centers. The OMAR models attempt to emulate this basic computational framework. That the smallest operating units are large groups of neurons is taken as license to build the models at a symbolic level.

Edelman, referencing Bartlett (1932), goes on to present a view of memory as *process*. For him, memory is the "*ability* to categorize or generalize associatively" (Edelman's italics, 1987, p. 241). Categorization occurs at the level of a global map and is degenerate. Edelman is well aware of the distinctions between declarative and procedural memory, but he is also quick to point out that these distinctions may be less than generally assumed. He suggests that there may be a procedural base supporting declarative memory.

In Edelman's view of memory as process, perception, categorization, generalization, and memory are closely linked. "Memory is a form of recategorization based upon current input; as such, it is transformational rather than replicative" (Edelman, 1987, p. 265). Memory is an active process of classification leading to recategorization and, thus, a partitioning of the world that is presented as one "without labels." Storage, to the extent that it exists, is one of procedures for mapping inputs to responses; hence, full representations of objects are neither stored nor required: *"It is the complex of capacities to carry out a particular set of procedures (or acts) leading to recategorization that is recollected"* (Edelman's italics 1987, p. 267). This view contrasts sharply with that of memory cast as data residing in a data base where the content is passive. In such schemes, something operates on memory as data, reinforcing some of it and degrading other parts of it. In the models developed here, memory is an integral part of the processes that employ it.

The scope or very presence of executive controllers is an important issue. Dennett (1991) expresses considerable concern over such homunculi-based theories. Centering his discussion around the metaphor of the Cartesian Theater where "everything" comes together, he suggests that the theater provides catchall for awkward elements leading to the failure to address difficult underlying questions. Dennett offers a Multiple Drafts model of consciousness in which "all varieties of perception—indeed all varieties of thought or mental activity—are accomplished in the brain by parallel, multiple-track processes of interpretation and elaboration of sensory inputs." He speaks in terms of an on-going process of "editorial revision." Dennett reinforces parallel processing as essential to modeling task execution and reminds us to be firm in our disavowal of homuncular concepts.

Cognitive modeling systems based on production rules (e.g., EPIC and SOAR) take a different stand on the issue of control. Production rule interpreters operate as executives with broad administrative responsibilities. Rule conditions may have oversight of one or more active tasks and memory stores, while rule actions may initiate, interrupt or terminate tasks and execute operations on memory or other capabilities central to the functioning of a model.

Following Dennett's admonition, the models developed in OMAR do not employ an executive or controlling process. The position explored here interprets the brain as home to a broad range of specific and spatially separated perceptual, cognitive and motor functional capabilities and attempts to model selected component parts as a dynamically configured network of computational elements operating in parallel and at times contending with one another in producing human multiple-task behaviors.

3. An Aircrew/Air-Traffic-Control Scenario

While scenarios in the commercial air traffic control domain can be developed to an arbitrary level of complexity, even the simplest scenarios can make multiple-task demands on aircrews and air traffic controllers (ATC).

3.1 Aircrew/Air-Traffic-Control Communication

Verbal communication, frequently the point of convergence for task contention in the air traffic control environment, takes place in three modes: in-person conversation between aircrew members, party-line radio communication between aircrews and the ATC managing their airspace, and telephone communication between ATCs in adjacent sectors. At the discretion of the aircraft's captain, either the captain or the first officer may undertake the task of handling ATC communication. The aircrew member not handling the ATC communication will monitor all ATC communications expect for occasional periods when communication with, for example, a company dispatcher is required. The party-line nature of radio communication means that ATC communication with each aircraft is heard by the aircrews of all aircraft under control of that ATC. Hence, an ATC will identify the designated aircraft call sign as the first segment of an utterance.

Conversation on the flight deck between the captain and first officer is the more typical person-to-person conversation of everyday life, but it is subject to interruption by ATC communication. The interruptions may take the form of directives addressed to their aircraft or to another aircraft under control of the ATC. In the interests of clarity and efficiency, most of the aircrew/ATC communications are highly stylized exchanges initiated with a directive or a question and completed by an acknowledgment of the directive or a response to the question. Established policy plays an important role in these exchanges. Verbal transactions between aircrew members must be suspended for ATC-initiated communication, even when the communication is directed to another aircraft. An aircrew member wishing to initiate a communication with an ATC must wait for the completion of an on-going transaction before initiating the communication. Typical directives to an aircraft might involve changes in heading, altitude and airspeed. The crew member handling the communication will acknowledge the communication and monitor the execution of the directive by the other crew member. Policy dictates cross checking—each crew member's expectations of exactly what the other crew member will do must be confirmed or the exception addressed. The crew must remember to resume their intra-crew transaction on completion of the ATC interruption. The domain is a fertile one in which to examine multiple-task behaviors.

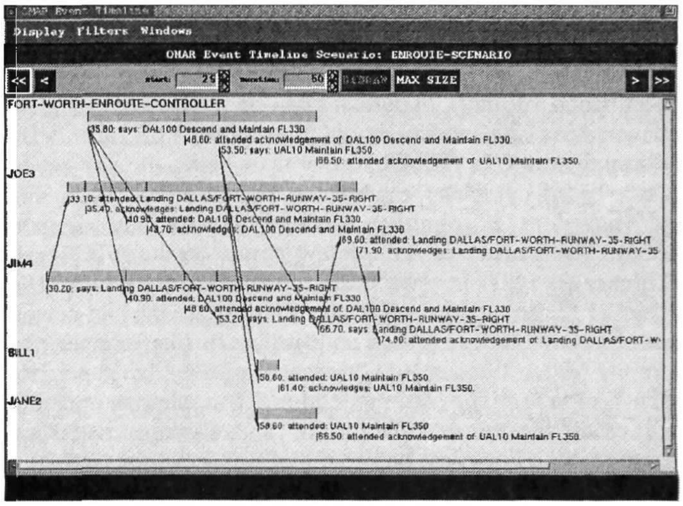

```
35.80: says: DAL100 Descend and Maintain FL330
48.60: attended acknowledgement of DAL100 Descend and Maintain FL330
53.50: says: UAL10 Maintain FL350
66.50: attended acknowledgement of UAL10 Maintain FL350
33.10: attended: Landing DALLAS/FORT-WORTH-RUNWAY-35-RIGHT
35.40: acknowledges: Landing DALLAS/FORT-WORTH-RUNWAY-35-RIGHT
40.90: attended: DAL100 Descend and Maintain FL330
43.70: acknowledges: DAL100 Descend and Maintain FL330
69.60: attended: Landing DALLAS/FORT-WORTH-RUNWAY-35-RIGHT
71.90: acknowledges: Landing DALLAS/FORT-WORTH-RUNWAY-35
30.20: says: Landing DALLAS/FORT-WORTH-RUNWAY-35-RIGHT
40.90: attended: DAL100 Descend and Maintain FL330
48.60: attended acknowledgement of DAL100 Descend and Maintain FL330
53.20: says: Landing DALLAS/FORT-WORTH-RUNWAY-35-RIGHT
66.70: says: Landing DALLAS/FORT-WORTH-RUNWAY-35-RIGHT
74.80: attended acknowledgement of Landing DALLAS/FORT-W
58.60: attended: UAL10 Maintain FL350
61.40: acknowledges: UAL10 Maintain FL350.
58.60: attended: UAL10 Maintain FL350
66.50: attended acknowledgement of UAL10 Maintain FL350
```

Figure 2: Aircrew Conversation Timeline

Figure 2 provides a timeline of a aircrew conversation interrupted once by ATC directive that they must attend to and then by an ATC directive for another aircraft causing them to further delay their conversation. Jim, the captain of flight DAL100, has just initiated a conversation with his first officer Joe when they are interrupted by a communication from the ATC. Jim acknowledges the ATC directive and Joe, having initiated the flight level change (not shown in the figure), resumes the in-person conversation, but it is immediately interrupted by another ATC communication, this time directed to Jane and Bill's flight, UAL10. Jim must pause once more before he once again resumes the interrupted communication with his first officer.

3.2 Modeling Task Contention and its Resolution

Established policy plays an important role in determining aircrew response to communicative acts: in-person communication is deferred in response to the onset of ATC radio communication; cross-checking dictates overlapping responsibilities with ATC communication managed by one crew member, while ATC directives are acted on by the other crew member; expectations must be satisfied and those that are not meet must be called out to secure safe aircraft operation; initiation of a party-line communication must await the completion of ongoing transactions. In SOAR or EPIC, each of these "decision" events might be viewed as the appropriate subject of an executive process and implemented as a rule set. In these tick-based simulation environments, each decision might be revisited numerous times before it is resolved and the concurrent nature of the ongoing tasks might dictate that several separate rule sets be evaluated at each tick.

The OMAR simulator is an event-based simulator to accommodate the particular and varied time steps at which each of several concurrent processes can be expected to operate. An aircrew member may initiate the action required by the change-altitude portion of an ATC directive (perhaps by setting the new altitude on the mode control panel (MCP)), while continuing to attend to subsequent speed and heading directives. These activities go on concurrently, each implemented as tasks with appropriate time frames. Established policy dictates that an in-person aircrew conversation be deferred at the onset of an ATC communication. In OMAR, rather than being the subject of a rule-based decision, established policy-driven behaviors are viewed as a cognitive form of automaticity (Logan, 1988). The priority of the aircrew "listen to the ATC" task is higher than the aircrew "in-person conversation" task. The onset of "listen to the ATC" task interrupts the "in-person conversation" task based on its priority. In like manner, the aircrew "listen to other ATC transaction" has higher priority than "initiate ATC communication." An aircrew member will wait for the completion of an on-going party-line transaction to complete before initiating a new transaction.

Policy-based decisions are viewed, not as the product of a centralized executive process, but rather as the outcome of contention among the particular subset of tasks activated and competing to execute in response to events initiated externally or internally. Events, be they externally or internally initiated, impinge, not on short-term memory, but on activated long term memories in the form of schemata with established policy-based priorities. In acting on an the initial directive of an ATC directive while attending to subsequent directives there may be no contention, but when contention is present, as in initiating a party-line communication, policy-based priorities mediate action. Given that several dispersed functional components may contribute to each of the contending tasks, when the contention is resolved, the component functions must act in accordance with the resolution.

3.3 Modeling Three Functional Components of Listening

The implementation of the listening tasks is representative of task modeling in OMAR. To elaborate on the implementation, it is necessary to examine how tasks are constructed from goals, and their plans and procedures (see Figure 1), and how competition between tasks is mediated. The aircrew members each have goals to manage in-person (*handle-voice-communication*) and radio communication (*handle-atc-communication* or *manage-atc-communication* depending on whether the crew member is responsible for managing or simply monitoring ATC communication). Each of these goals is implemented as a plan made up of subgoals and procedures. The goals and subgoals express the proactive agenda of the agent in addressing anticipated contingencies, while the procedures express the actions to be taken to accomplish each goal. These particular goals are distinct to the extent that the protocols for conducting in-person and radio communication are distinct. The communication goals are activated with the procedures *listen-for-voice-message* and *listen-for-radio-message* in a wait-state. The goals and their plans instantiate the cognitive capability to conduct an in-person or radio conversation using the appropriate protocol for each communication media. Their procedures are in a wait-state pending the onset of a communication in their particular media. These goals, subgoals and procedures form the attended cognitive component of the "listening" complex of tasks. Additional goals and procedures stand ready to assess and act on the content of the communication, for example, setting a new target altitude using the MCP.

As currently implemented, the listening tasks have two additional components. The listening task complex is activated by a verbal communication. A separate procedure for processing the auditory input, initiated through a separate goal, awaits the onset of an auditory communication. Shortly after the auditory procedure is activated, it in turn activates a speech understanding procedure to develop the propositional form of the communication that the attended cognitive task will respond to. In the simulation, the communication content is conveyed as an object and the auditory and speech understanding processes are time-consuming process stubs.

The development of the listening task posits three distinct functional areas of processing. Separate goal and subgoal trees set up each of the functional capabilities. The onset of the auditory communication initiates the processing with the activities of the three functional areas coordinated through a series of messages, or *signals*, as they are defined in OMAR. The functional areas and signals are a symbolic analogue of Edelman's (1987) reentrant nets. The procedural bias in the modeling approach is taken a step further. Motivated by Edelman's (1987) process view of memory and reinforced by his references to Bartlett (1932), short-term memory, rather than being treated as a faculty in its own right, is modeled as a set of distinct capabilities (Martin, 1993; Schneider & Detweiler, 1987) distributed among a family of modality-specific functional procedures. Auditory memory for a verbal communication is a component of the auditory process, while the propositional memory is a component of the language understanding process. Their persistence, different for each modality, is envisioned as, but not yet implemented as, a product of the persistence of their enclosing procedures.

The model makes explicit a proposal for how component functionalities might be coordinated during task execution and how task contention might be mediated. Given a task, postulated to be the product of contributions from several dispersed functional capabilities, the event of the execution and interruption of that task has been explored in a manner that does not require an executive controller. The implementation suggests how an attended listening task might interact with auditory and language understanding components of the task and how actions based on the communication's content might be coordinated.

4. OMAR Support for Multiple-task Human Performance Modeling

OMAR's strengths as a human performance modeling environment lie in its representation languages and their graphical editors and browsers, its simulator and its post-run analysis tools. The principal representation languages are the Simple Frame Language (SFL) and the Simulation Core (SCORE) language, a language for specifying the concurrent execution of goals and procedures. SFL is a direct descendent of the KL-ONE (Brachman & Schmolze, 1985) family of frame languages, while SCORE is a descendent of ACTORS (Agha, 1986), a model of concurrent computation in distributed systems. A rule based language provides the capability to develop rule packets as models of decision making. This section focuses on the aspects of the languages that support the development the models of human multiple-task behaviors.

4.1 Concurrent Task Execution and Mediating Task Contention

Language constructs in SCORE provide the building blocks for modeling task execution. Tasks are expressed as goals to be accomplished by means of plans that are made up of subgoals and procedures. Concurrent execution essential to the multiple-task capability in the models is specified using *race* and *join* forms in the language. A *race* form completes when the first of its enclosed forms completes. A *join* form completes when all of its enclosed forms complete. A *spawn* form is available to initiate an independent execution thread.

The contention between tasks is a more complex concern. At least three levels of contention can be envisioned. Attended thoughtful deliberation can lead to the selection of one course of action over another. This class of deliberation processes that might be explicitly modeled on a case by case basis is not addressed here. The concern in the current effort has been with the simpler cases of policy-driven decisions as described in the aircrew scenarios above and the still simpler contention based on access to particular, identifiable resources. The contention between tasks can occur high in the goal tree as in the contention between "listen to ATC" and "in-person conversation" procedures or near the leaves of the tree, as in contention between tasks for access to the dominant hand for a skilled manual operation.

All SCORE procedures are SFL concepts and each may be classified as a procedure that contends with another particular procedure (as in the case of "listen to ATC" and "in-person conversation") or with other instances of their own class (as in the case of the dominant hand requirement). A new procedure about to run must either establish that it does not contend with a running procedure or that it has sufficient priority to block the execution of the running procedure. If a new procedure has sufficient priority, it begins execution and execution of the contending procedure is halted until execution of the new procedure has completed. At this point, barring intervening events affecting the contending procedures, the original procedures resumes execution. If the priority of the new procedure is not sufficient to block the running procedure, it must wait for the running procedure to terminate. Procedure priorities are computed dynamically and procedure contention is revisited as priorities change. Contention is mediated on a pair-wise basis that does not require management by a third-party controller.

4.2 Pattern Matching in Coordinated Functional Component Execution

As we have seen in the aircrew scenario, goals are employed to establish a network of procedures, each of which assumes a wait-state sensitive to particular externally or internally generated events. The events take the form of *signals* in SCORE. The signals are implemented as lists with the first element of the list defining the signal type and additional elements of the list provide the data required for the signal type. The SCORE form, *signal-event,* takes a list as an argument and generates the signal. A procedure may enqueue on a signal by using the *with-signal* form. Execution of the procedure invoking the *with-signal* form enters a wait-state pending the occurrence of a signal of the designated type. The *with-signal* form may include a test on any of the data elements of the signal that must be satisfied before the signal is accepted for processing. Once a signal is accepted, processing of the enclosing procedure continues. If the signal type is of further interest, the *with-signal* form must be employed again.

A procedure issuing a signal continues operation—it does not wait upon or receive any returned values. The issuing procedure has no knowledge of the other procedures that have enqueued on the signal. There may any number of procedures enqueued on it or none at all. A given procedure may enqueue on a signal once or in each of two or more parallel threads employed to explore different patterns of events that each include this particular event. Signal-based coordination of procedure execution bears some resemblance to a dataflow architecture (Arvind & Culler, 1983). A procedure in a wait-state resumes execution when data meeting its pattern matching requirements arrives.

The capability to enqueue on signals forms the basis for the coordination of the functional capabilities that make up the multiple-task model of human performance. Signals are used both as the representation of external events that trigger the model's human receptors (eyes and ears in the current implementation) and as the representation of the elements of the subsequent internal cascade of events that is produced in developing the coordinated multiple-level response to external events. The network of activation of *with-signal* forms can change rapidly over time to reflect the occurrence of external events and the evolving reconfiguration of procedures that represent the functional capabilities that combine to form tasks governing the response to those events. The changing network of active procedures, each sensitive to particular external or internal events, forms a pattern matcher that incorporates a temporal dimension and determines the behaviors of the model. The proactive component of the behaviors is provided by the goals and subgoals that govern the initial network configuration and activation of event-sensitive procedures. Each of the behaviors in the performance of a task is the result of the activation of a mix of the proactive and reactive components of the task. The signal-driven activation of network nodes that represent functional capabilities provides an emulation of Edelman's reentrant and global maps.

5. Future Work

Simulation studies in the commercial air traffic control domain employing professional aircrews and air traffic controllers have been conducted on a regular basis over the past several years. Access to the data from these studies would provide the basis for an assessment and further refinement of the modeling described here. In refining the human performance model, several areas are of particular interest. The first is to further explore the link between process and memory, and in particular, to model the persistence of memory instances as the residuals of procedure execution. A second area of interest is the increased tempo in performance that people exhibit as workload increases. In part, this will lead to a reexamination and further elaboration of the priority-based mediation of task contention.

6. Acknowledgments

Over the years Marilyn Jager Adams, Glenn Abrett and Jeff Palmucci have made significant contributions to the OMAR human performance modeling effort at BBN. The author also wishes to thank Michael Young of the Air Force Research Laboratory for his continued support for this research effort. The research reported on here was conducted under USAF Armstrong Laboratory Contracts F33615-91-D-0009 and F41624-9T-D-5002.

7. References

Agha, G. A. (1986). *ACTORS: A model of concurrent computation in distributed systems.* Cambridge, MA: MIT Press.

Arvind, & Culler, D. E. (1983). *Why dataflow architectures.* Computational Structures Group Memo 229-1, Laboratory for Computer Science, Massachusetts Institute of Technology.

Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology.* Cambridge, U. K.: Cambridge University Press.

Brachman, R. J., & Schmolze, J. G. (1985). An overview of the KL-ONE knowledge representation system. *Cognitive Science, 9,* 171-216.

Dennett, D. C. (1991). *Consciousness explained.* Boston, MA: Little, Brown and Company.

Edelman, G. M. (1987). *Neural Darwinism: The theory of neuronal group selection.* New York: Basic Books.

Edelman, G. M. (1989). *The remembered present: A biological theory of consciousness.* New York: Basic Books.

Goldman-Rakic, P. S. (1987). Circuitry of primate cortex and regulation of behavior by representational memory. In F. Plum (Ed.), *Handbook of physiology—The nervous system* (Vol. 5). Bethesda, MD: American Physiological Society.

Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). SOAR: An architecture for general intelligence. *Artificial Intelligence, 33,* 1-64.

Logan, G. D. (1988). Automaticity, resources, and memory: Theoretical controversies and practical implications. *Human Factors, 30,* 583-598.

Martin, R. C. (1993). Short-term memory and sentence processing: Evidence from neuropsychology. *Memory & Cognition, 21,* 176-183.

Meyer, D. E., & Kieras, D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance: Part 1. Basic mechanisms. *Psychological Review, 104,* 3-65.

Moscovitch, M. (1994). Memory and working with memory: Evaluation of a component process model and comparisons with other models. In D. L. Schacter & E. Tulving (Eds.), *Memory systems.* Cambridge, MA: MIT Press.

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

Pandya, D., & Barnes, C. L. (1987). Architecture and connections of the frontal lobe. In E. Perecman (Ed.), *The frontal lobes revisited.* New York: IRBN Press.

Petrides, M. (1989). Frontal lobes and memory. In F. Boller & J. Grafman (Eds.), *Handbook of neuropsychology* (Vol. 3). North Holland: Elsevier Science Publishers B. V.

Posner, M. I. (1993). Seeing the mind. *Science, 262,* 673-674.

Posner, M. I., Peterson, S. E., Fox, P. T., & Raichle, M. E. (1988). Localization of cognitive operations in the human brain. *Science, 240,* 1627-1631.

Raichle, M. E. (1994). Visualizing the mind. *Scientific American, 270,* 58-64.

Schneider, W. & Detweiler, M. (1987). A connectionist/control architecture for working memory. In G. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 21). New York: Academic Press.