



Kent Academic Repository

Barnard, Philip and Bowman, Howard (2004) *Rendering Information Processing Models of Cognition and Affect Computationally Explicit: Distributed Executive Control and the Deployment of Attention*. Cognitive Science Quarterly, 3 (3). pp. 297-328. ISSN 1466-6553.

Downloaded from

<https://kar.kent.ac.uk/14180/> The University of Kent's Academic Repository KAR

The version of record is available from

This document version

UNSPECIFIED

DOI for this version

Licence for this version

UNSPECIFIED

Additional information

Versions of research works

Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

Author Accepted Manuscripts

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in *Title of Journal*, Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

Enquiries

If you have questions about this document contact ResearchSupport@kent.ac.uk. Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our [Take Down policy](https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies) (available from <https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies>).

Rendering information processing models of cognition and affect computationally explicit: distributed executive control and the deployment of attention

Philip J. Barnard

MRC Cognition and Brain Sciences Unit

Howard Bowman

*University of Kent at Canterbury*¹

In this paper we illustrate the potential of process algebra to implement modular mental architectures of wide scope in which control is distributed rather than centralised. Drawing on the Interacting Cognitive Subsystems (ICS) mental architecture, we present an implemented model of the attentional blink effect. The model relies on process exchanges between propositional meaning and a more abstract, implicational level of meaning, at which affect is represented and experienced. We also discuss how the proposed mechanism of buffer movement can, in the context of the ICS architecture, be extended to account for effects of emotional stimuli and brain damage.

Keywords: Distributed systems, Process Algebra, Executive Control, Attentional Blink, Emotion, Interacting Cognitive Subsystems.

Introduction

Production system architectures and connectionism have dominated research on computational realisation of theories of cognition. However, in spite of great technical progress in sub-symbolic, symbolic and hybrid methods, much influential theory still lacks computational realisation and interactions between processes, subsystems or mental modules are still often represented as box and arrow diagrams. These recruit the information processing metaphor but present theoretical constructs in only informal abstract

¹ Address for correspondence: Computing Laboratory, University of Kent at Canterbury, Canterbury, Kent, UK, CT1 2DE (Email: H.Bowman@ukc.ac.uk, Telephone: +44 1227 823815, Fax: +44 1227 762811).

style. Much criticised by the computational modelling community for lack of specificity and their inherent ambiguity, such box and arrow models nonetheless enable hypotheses to be framed in sufficient detail for guiding experimentation in circumstances where theorists are either unable, or reluctant, to commit to the kind of detailed assumptions required to run formal computational simulations. In many domains of enquiry, such as emotional influences on cognition, it is widely acknowledged that the theoretical picture must be complex, (e.g. Leventhal, 1979; Teasdale & Barnard, 1993). Multiple influences are at play which require many “modules.” Consequently, a huge number of potentially unwarranted assumptions would be needed to simulate such models with conventional techniques, let alone to accommodate the kind of individual variation that lies at the heart of our understanding of the effects of mood state on cognition and disorders such as anxiety or depression.

The development of parallel computing and large scale networking of computers, typified by the world wide web, has presented computer scientists with similar kinds of technical problems in modelling interactions amongst components of complex distributed systems. They responded by developing new mathematical formalisms and tools for modelling systems of interacting modules. These fall within the category of *formal specification techniques*. We will argue that such specification formalisms operate at a level of abstraction comparable to the box models of psychological theory. Cooper (1995) has made related arguments. Such formalisms and tools offer the potential for modelling complex mental architectures in either abstract formal mathematics (Duke, Barnard et al. 1998) or to “run” specifications rather like a conventional simulation (Bowman & Faconti 1999), but without imposing particular implementation details.

The idea that abstract specification of cognitive models can enhance our understanding of the deeper properties of cognitive theories, and support the development of tools for their implementation, has also emerged in the context of production system methodologies (Cooper, Fox et al. 1996). In this paper, we give a new perspective on this general orientation by exploiting the potential of process algebra (a formal specification technique developed by computer scientists) to build a computationally explicit model of human attention in cognitive-affective settings.

Humans pay attention to information that matters, as a result of the cognitive task they are required to perform (Duncan 2000), as a function of that information’s personal salience (Moray 1959) or as some function of their motivational and emotional state. Anxious people preferentially pay attention to external threat (MacLeod, Mathews et al. 1986). Over more extended periods, depressed people focus their internal attention on negative self-

related thoughts (Beck 1976). In all these domains the key questions concern the dynamic redeployment of attention over time. Empirical paradigms like the psychological refractory period (Pashler and Johnston 1998), task shifting (Allport, Styles et al. 1994) or the attentional blink (Raymond, Shapiro et al. 1992), all illustrate subtle restrictions on our ability to re-deploy attention. We will describe the latter of these in some depth in the subsection *The Attentional Blink* on page 7.

The computational model presented here focuses on the attentional blink because it has recently been shown to be influenced by emotional factors. Our model development has been guided by a particular box and arrow mental architecture, Interacting Cognitive Subsystems, or ICS (Barnard 1985; Barnard 1999), which has been developed to address the effects of affect on cognition (Teasdale & Barnard 1993). This architecture assumes that executive control is distributed, central aspects of which emerge from interactions between two subsystems. These process qualitatively distinct types of meaning. One is classically “rational,” being based upon propositional representation. The other involves a yet more abstract encoding, implicational meaning, at which affect is represented and experienced. ICS refers not to a central executive but to a central engine of mentation to reflect the distributed nature of executive control (Teasdale & Barnard 1993).

This paper has three aims. The first is to illustrate the use of process algebra to render box and arrow models of psychological theory computationally explicit. The second is to present a specific, implemented, model of attentional mechanisms based upon the distributed control hypothesis specifically applied to the processing of meaning. The third is to discuss how such modelling could in principle be extended to the effects on attention of emotional meanings.

Computational Issues: Concurrency and Distributed Control

Following Fodor (1983), among others, numerous theories now assume that many mental modules, or their neural substrates, are processing information at the same time. Consistent with this position, any realistic computational model of the mind must, at some level, be *concurrent*. The control of concurrent processing can be centralised or distributed, emerging from interaction amongst modules whose behaviour *evolves independently*. Historically, centralised control played a key role in the development of information processing psychology (Broadbent 1958) and subsequently computational models (Newell and Simon 1972). Such a notion, and Von Neumann-like architectures, fitted well with a restricted capacity to attend to information, to memorise it, or to deal with more than one task at a time. However, a substantial body of neuropsychological evidence now indicates that brain damage selectively impairs particular processing capabilities, while leaving oth-

ers intact. Even under circumstances where executive functions are impaired, complex behaviours can still be co-ordinated (Shallice 1988). This suggests that computational models should embrace some form of *distributed executive control* in which there is no single locus of control.

Architectures based on the distributed control hypothesis are becoming increasingly common, including those that address interactions between motivation, affect and cognition. For example, the component process theory of emotion (Scherer 2000) explores interactions among five subsystems: cognitive, autonomic, motor, motivational and a monitor subsystem. Interactions among these subsystems determine behaviour and dynamic systems theory provides a perspective on self-organising patterns of control. Carver & Scheier (1998) also appeal to concepts of self-regulation, but emphasise the hierarchical organisation of control mechanisms. The systems-level theory of Bond (1999), which has been implemented, models motivational and social aspects of primate behaviour. These emerge as a function of computationally realised interactions among modules with no locus of central control. The ICS architecture contrasts with these other approaches by assuming that all modules process information according to the same fundamental principles – subsystems differ only in the way the information they process is encoded. In this case, hierarchy is also implicit in the abstraction of higher order regularities in information patterns (Teasdale and Barnard 1993).

Some notion of hierarchy is essential for the sort of “distributed systems” modelling that we are advocating. Not only is it parsimonious to view modules as being composed of modules, but the same arguments concerning decentralised overall control of the mind can be applied within modules. This combination of distributed control and hierarchical decomposition is reflected in many current theories, whether implemented or not. In our case, subsystems in the ICS architecture are themselves systematically decomposed into components with distinct internal functions, such as an array representing the input data to a subsystem, an image record, and processes that transform data from one type of mental representation to another. Likewise, Baddeley (2000) decomposes his Working Memory model into phonological, visuo-spatial and executive components, each of which can be decomposed into storage and processing resources. Even at the level of the brain, coarse divisions can be subdivided. Schneider (1999) presents a hybrid model that reflects such a hierarchy. Thus, it appears natural when adopting a modular theory that the behaviour of individual components emerges from a set of interacting sub-modules, yielding a hierarchical component structure.

The benefits of connectionist approaches have been widely discussed (e.g. Rumelhart, McClelland, & PDP Group, 1986) as have their drawbacks

(e.g. Fodor & Pylyshyn, 1988). The benefits of production system architectures such as SOAR have also been extensively debated (e.g. Newell, 1990), as well as problematic issues that arise with them (e.g. Cooper et al., 1996). With respect to our particular concern with mechanisms of distributed control and hierarchy, both can of course be realised in either connectionist or production system frameworks. However, neither fully satisfy the requirements we regard as important to address when modelling complex concurrent systems with distributed control (Hybrid approaches, which combine symbolic and connectionist techniques, are also relevant, but we postpone their consideration until the final discussion section of this paper).

Concurrency can, of course, be generated by allowing multiple productions to fire on each cycle of a symbolic architecture's operation, as is done in the EPIC model (Meyer and Kieras 1997). However, control remains centralised, being focussed on the working memory embedded in their single cognitive processor and since this centralised structure is not reflected in the brain, when relating symbolic architectures to brain-level implementation, some sort of re-alignment needs to be assumed. As noted earlier, postulating such structure re-aligning mappings from abstract architecture to brain is not always straightforward (again see Shallice, 1988). Also, a large proportion of the psychological theories available are themselves distributed, being expressed in terms of independently evolving interacting modules. Thus, structure re-aligning maps need to be postulated both when relating "upwards" from production systems architectures to high-level psychological theories and when relating "downwards" to low-level brain models. Bond (1999) has also made similar arguments.

In contrast, while connectionism inherently embraces distribution of control, it does so at a very low level. Almost all the uses of connectionism in cognitive modelling have been *specialised* in nature. Neural networks have proved highly effective for modelling specific cognitive phenomena, such as the Stroop effect (Cohen, Dunbar et al. 1990), word reading (Plaut 1998), serial order recall (Page and Norris 1998) and many others. However, in extrapolating from these specific phenomena to the big architectural picture, they have not done so well. This is in no small part due to the fact that it is very hard to construct large architectural models using connectionist paradigms, which in turn, we would argue is influenced by the inability to describe hierarchical structures. A level of compositionality is obtained in connectionism through the interconnection of layers, each of which can be viewed as a module. However, there is only one level of composition and it is not easy to nest interacting components within interacting components. The component structure of neural networks is essentially flat – the primitive elements of neural networks are neuron-like nodes, not neural networks. A further reason for not using connectionism concerns abstraction.

Modelling based on neural networks is, in certain respects, very low-level in character. In particular, one has to work hard in order to obtain what are “primitive” constructs and data structures in higher-level computational notations preferred, for example, by the symbolic modelling community (see Newell, 1990). For the variety of modelling we are undertaking, the larger architectural picture is essential. In addition, the psychological theories we are considering are typically couched in terms of representations being transmitted between components and such passing of data items is characteristic of symbolic computational paradigms, rather than the connectionist paradigm. Unlike either paradigm, process algebra was explicitly developed to model modular architectures with distributed control at an abstract level. It therefore offers a modelling paradigm that appears particularly well suited for implementing broadly scoped box and arrow theories.

Process Algebra

Process algebra originated in the late 1970’s and early 1980’s and there is now an extensive research literature surrounding them. They have been widely applied to the specification and analysis of communication networks, distributed systems and telecommunications systems. Here we summarise the key elements of process algebra, while referring interested reader to comprehensive texts (e.g. Hoare, 1985; Milner, 1989; Roscoe, 1998; Schneider, 2000). We will be using a technique named LOTOS (Bolognesi and Brinksma 1988). Throughout this paper, we will describe our model mostly in intuitive terms, rather than presenting detailed fragments of specification (although the full specification is available from the second author). Those elements of LOTOS notation that we do use are introduced as they arise in the text. Three aspects of LOTOS specifications impinge upon the arguments we are making here:

1. **Structural Decomposition and Distribution.** We seek a notation for describing systems with distributed control. This is realised in LOTOS since the basic element of structuring is a *process*. *Processes* possess their own local state and are computationally autonomous. The reader familiar with use of the term process in psychology should beware, since the term has a very specific meaning in the process algebra domain. In fact, in the context of this paper the concept is interchangeable with those of a module or subsystem.
2. **Interaction.** Processes evolve independently and asynchronously of one another subject to interaction through message exchange. If a process wishes to communicate with another process it offers to perform *an action* with that other process. If the other process is willing to perform the action, a synchronised message exchange occurs. Message passing inter-

action is different to activation exchange in neural networks. On each update cycle of a neural network, activation is relayed over all links in the net and interactions between neurons are often globally synchronised. In contrast, in process algebra there is no global control of interaction. Processes make autonomous local decisions about when they wish to communicate and whom they wish to communicate with. While the data relayed between neurons is of a single type, a real number, this is not so with process algebra communication. Any data type can be sent from one process to another. The dynamics of communication and interaction that can be set-up in process algebra are thus extremely flexible.

3. **Control and Data.** LOTOS is really a composite of two languages - a control and a data language. The previous two points focussed on the former, a language for describing processes and their interaction. The data language uses a style of data specification called algebraic data types (de Meer, Roth et al. 1992). Operations on data types are defined in an equational style and data expressions are executed through an ordered application of the equations. Each equation application rewrites the data expression towards an irreducible form. Such a declarative style of computation is similar to that employed in functional programming languages, such as Haskell (Thompson 1999) or even Lisp. Thus, the data language is expressively rich and allows the construction of information-rich symbolic representations. Furthermore, the symbol system is combinatorial in the manner advocated by say Fodor & Pylyshyn (1988). Broadly speaking, the data language is as computationally expressive as Lisp, which from a theoretical perspective is justified by the data language being Turing complete. In fact, both the control and the data languages are, on their own, computationally powerful enough to simulate Turing machines. Finally, with regard to the relationship between data and distributed control, in LOTOS, data is fundamentally local - operations on data have an effect, which is local to a particular module. This is consistent with our view that there is no shared memory, that could implicitly act as a focus for centralised control, as it does in production systems, such as EPIC (Meyer and Kieras 1997).

The Attentional Blink

Although several related phenomena predated it (e.g. Broadbent & Broadbent, 1987), the phenomenon, which is robust and thoroughly investigated, was first reported by Raymond et al. (1992). Typically, letters are presented using rapid serial visual presentation (RSVP) at around ten items a second. One letter (T1) is presented in a distinct colour. It is the target whose identity must be reported. A second target (T2) follows after some number of inter-

vening items. For example, the person may have to report whether the letter “X” was among list items that followed T1. Detection of T2 is impaired with a characteristic serial position curve (fig. 1). If T2 occurs immediately after T1, then its presence is accurately detected. Detection then declines and recovers to baseline at around a half second lag.

The empirical literature and alternative theoretical accounts of it have been summarised elsewhere (Shapiro, Arnell et al. 1997; Potter 1999; Shapiro and Luck 1999). These authors conclude that the blink does not appear to be a product of simple perceptual, memory or response output limitations. Various theories have been advanced to account for the influences on the serial position curves (e.g. Chun & Potter, 1995; Duncan, Ward, & Shapiro, 1994; Raymond, Shapiro, & Arnell, 1995). These theories all naturally assume that allocating attention to T1 leaves less attention for T2, but details of their proposed mechanisms, like decay, interference, similarity, and bottlenecks obviously vary. There also remain significant areas of empirical uncertainty - such as conditions under which effects do or do not occur with cross modal presentations (Shapiro, Arnell et al. 1997; Potter 1999).




Figure 1. The basic “Attentional Blink” effect for letter stimuli (adapted from Raymond et al., 1992). Here, baseline represents a person’s ability to report the presence of T2 in the absence of a T1.

As research on the blink and RSVP in general has progressed, it is becoming clear that the allocation of attention is affected by the meaning of items (Maki, Frigen et al. 1997) and their personal salience (Shapiro, Caldwell et al. 1997). Indeed, similar serial position curves are readily obtained when words are used as list items. There is also evidence from electrophysiological recordings that suggests that the meaning of a target is being processed even

when it remains unreported (Shapiro and Luck 1999). Most recently there are reports of specific effects of affective variables. Holmes & Richard (1999) report differences in target detection in the Attentional Blink (AB) paradigm for high and low anxious people. More dramatically, Anderson (2001) has shown that the blink effect is markedly attenuated when the second target is an aversive word. Anderson & Phelps (2001) also report data for patients with amygdala damage. These patients all displayed impaired T2 report comparable to control conditions with affectively neutral material, a standard blink effect. However, a group of patients with unilateral damage to left amygdala showed no attenuated blink effect to aversive words, whereas those with damage to the right amygdala, like the control group, showed reliable attenuation of the blink effect.

Of the theories that have so far been developed, most of those reviewed by Shapiro et al. (1997a) would require considerable modification and extension to address, in any depth, the full range of neuropsychological and affective factors now known to shape the serial position curves. To our knowledge none of them have been implemented in computational terms. Of these models, that proposed by Chun & Potter (1995) is most closely related to the computational model we will propose shortly. Their model assumes two stages of processing. The first stage performs an initial evaluation to determine "categorical" features of the item. This stage is not capacity limited, the identity of items is unavailable, and the representation is open to rapid forgetting. The second stage builds upon and consolidates the results of the first in order to develop a representation of the target, sufficient for subsequent report. This stage is capacity-limited, invokes central conceptual representations and storage, and is only initiated by detection of the target on the first stage.

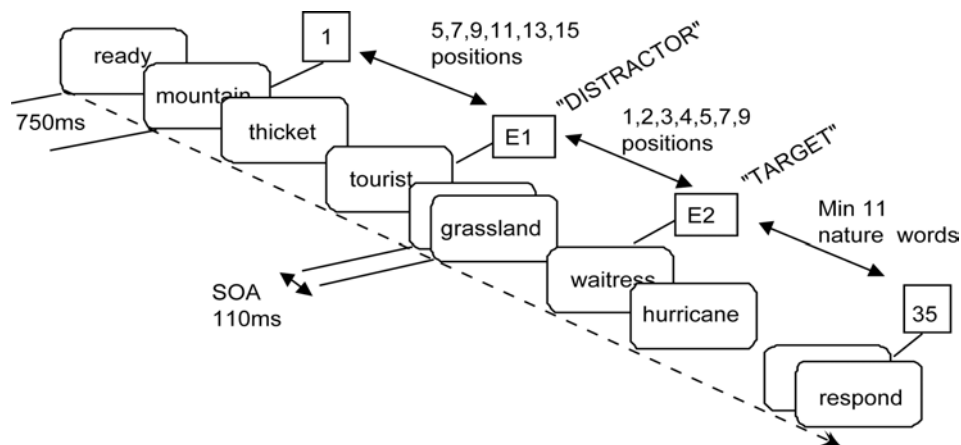


Figure 2: Task schema for the word-based blink; adapted from Barnard, Scott, Taylor, May, & Knightley (2001).

Research on the blink phenomenon has relied on marking the to-be-reported target identity with some form of perceptual feature such as a distinct colour. In order to examine semantic effects in more detail, Barnard et al. (2001) used a variant of the paradigm in which no perceptual features were present to distinguish the target from the background list. Targets were only distinguishable from background items in terms of their meaning. This variant of the paradigm did not rely on dual target report. Rather, participants were simply asked to report a word if it refers to a job or profession for which people get paid, such as “waitress” and these targets were embedded in a list of background words that all belong to the same category. In this case, they were inanimate things or phenomena encountered in natural environments (fig. 2). Participants could report the target word (accurate report), say “Yes” if they were confident a job word had been there but couldn’t say exactly what it was, or say “No” if they did not see a target, and there were, of course, trials on which no target was presented. When the distractors were household items, a different category from both background and target words, there was little influence on target report. However, distractors that referenced a property of a human agent, but not one for which they were paid, like tourist or husband, gave rise to a classic and deep blink effect not unlike that already shown in fig. 1.

Barnard et al. used latent semantic analysis (Landauer and Dumais 1997) to assess similarities between “human” distractors and job targets. Being aware of a target’s presence and being totally unaware were linked with rather different profiles of semantic similarities, and these authors present an argument that this effectively ruled out an explanation based upon positive or negative priming effects (Tipper 1985; Maki, Frigen et al. 1997). Like Chun & Potter (1995), they argued for a two-stage model, but this time recast to focus exclusively on semantic analysis and executive processing. In the first stage, a generic level of semantic representation is monitored and initially used to determine if an incoming item is salient in the context of the specified task. If it is found to be so, then the specific referential meaning of the word is subjected to more detailed semantic scrutiny. In this second stage a word’s meaning is actively evaluated in relation to the required referential properties of the target category. If this reveals a match then the target is reported. The first of these stages is somewhat akin to first taking a “glance” at generic meaning with the second akin to taking a closer “look” at the relationship between the meaning of the distractor and the meaning of the target category. Assuming that semantic processes cannot glance at incoming items while looking at and scrutinising another, a blink would re-

sult. We now proceed to model this effect using mechanisms relying on distributed executive control of semantic representations.

Interacting Cognitive Subsystems

The account offered by Barnard et al. for their semantic effects was shaped by the ICS architecture, and the idea that executive processing involves distributed control of the processing of two types of meaning. This architecture initially decomposes into nine subsystems, each of which process different mental representations (fig. 3). There are three sensory subsystems, four central subsystems, and two effector subsystems. These representational subsystems are supplemented by peripheral somatic and visceral response systems. These realise the bodily effects of emotional reactions that are in turn picked up and represented by the Body State subsystem. All subsystems have an identical internal decomposition. They differ only in terms of the representations input to them, stored in their own local memories and output by them as a consequence of the action of processes that transform inputs into outputs.

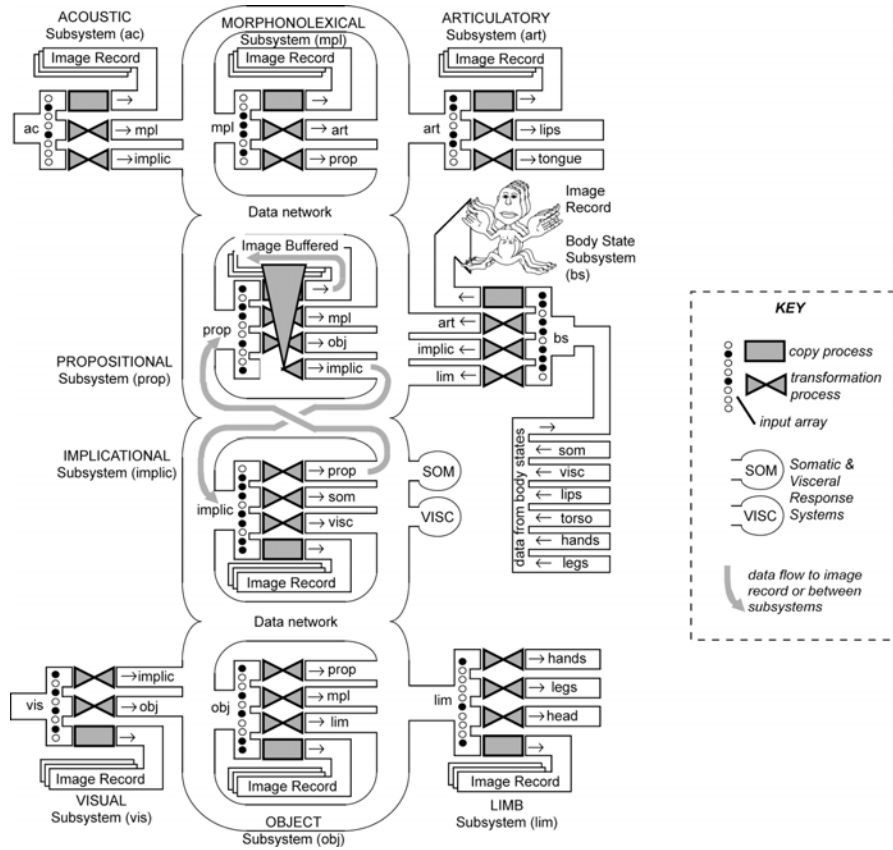


Figure 3: The ICS architecture in outline, showing data flow between two meaning subsystems.

For full description of the ICS architecture, readers are referred to either informal descriptions of its operation (e.g. Barnard, 1999; Teasdale & Barnard, 1993), or more formal models of its operation in complex real-world tasks (Duke, Barnard et al. 1998; Barnard and May 1999). Here we focus only on those aspects that directly constrain the modelling of the blink effect. ICS specifies paths of communication and makes particular assumptions concerning the internal organisation, representations, and dynamics of process operation. As information arrives at an ICS subsystem, it is mapped into an input array. From this point a dedicated process copies the basic units of input into an image record (see key in fig. 3). This is a local memory that preserves a trace or image of recent input as well as preserving a long-term record of information patterns. In parallel with the copy process, each subsystem contains a set of processes that transform inputs to outputs and these

mediate communication between subsystems, either by operating directly on the input array or by re-configuring to use the image record.

Since the image record holds representations in input code, the processes can configure to access recently arrived material, in a mode referred to as buffered processing. In fig. 3, the propositional process communicating with the implicational subsystem is depicted as configured in this buffered mode. In what follows, the recent image is given computational realisation as delay lines. Although not a mechanism currently widely used in cognitive modelling, related mechanisms have been postulated for, and simulated in, neural systems (Abeles 1991). With its explicit hierarchical structure and definition of processes that communicate within and between modules, coupled with its adherence to the distributed control hypothesis, ICS is specified at a more or less ideal level of abstraction for realisation in process algebra. Here we concentrate on the central engine in which propositional and implicational representations are exchanged.

Implicational subsystem

This subsystem will be used to implement the kind of “glance” introduced earlier. As a generic form of encoding, implicational meaning represents the broad “categorical” analysis of items also considered in Chun and Potter’s first stage of processing. This subsystem builds and uses implicational representations. For a lexical task, one of its feeds involves rapid and direct recoding of visual form (via Vis → Implic, fig. 3). We shall refer later to implicationally salient items, as those that “pass the implicational subsystem test”. Since implicational meaning takes as inputs the immediate, and rather unrefined, products of processing visual, auditory and body state patterns, it provides a platform from which, not only the rapid consequences of preliminary semantic processing can be modelled, but also those associated with personal salience and affect.

Propositional subsystem

This subsystem will be used to implement the more detailed look. It builds upon the implicational representation generated from the glance, and it also takes input from the longer processing configurations, not implemented in the model described below, but important for later discussion, that arrive at this level via the object and/or morphonolexical subsystems that interpret visually presented words. It builds a referentially specific, propositional, representation, which is sufficient for report. We will describe items that “pass the propositional test” as propositionally salient. The existence of a propositionally explicit representation is required in ICS to compute the word form to be passed through an output configuration for generating a

response. In ICS, implicational representations cannot directly pass data to the subsystems that mediate lexical output.

Buffering

Our modelling of emergent attentional mechanisms focuses on buffer movement. Processing activity and its distributed control is subject to specific constraints. In ICS, a process can only be in one of three modes: direct, buffered, or record access (not discussed here), it can only process one coherent stream of data at any instant, and only one process in a wider configuration can be buffered at a given time. The buffered mode of processing enables a process to deal with novel combinations of units of representation that have recently arrived. It is also linked to focal awareness of a particular representation (Teasdale and Barnard 1993) and thus implies the allocation of limited attentional resources to a particular process. With a delay line representation, the automatic copying of data to the image record means that all input is represented in temporally extended form.

A Computational Model of the Attentional Blink

Specification Structure

The top-level structure of our model is shown in fig. 4. The heart of the model is the central-engine. In accordance with the wider ICS specification, we assume that input into the central-engine is received from perceptual systems and output from the central-engine is relayed to a response system from which detected targets are reported. Our modelling of these “peripheral” systems is however not detailed. While this figure seems like an informal box-and-arrow diagram like fig. 3, in fact, the specification exactly follows the structure of this diagram. At the top-level of decomposition, our specification contains three processes, one encapsulating the perceptual subsystems, another modelling the central-engine and a third encapsulating the response subsystems. This fits with the observation that we have been making that explicitly having facilities to express distributed control in a modelling notation yields a natural formalisation of psychologist’s box and arrow diagrams. The specification evolves through message exchange based interaction between the three top-level processes. The links over which message exchanges occur are indicated by action names. At the top level these actions are *ext_implic_data* and *prop_ext_data*.

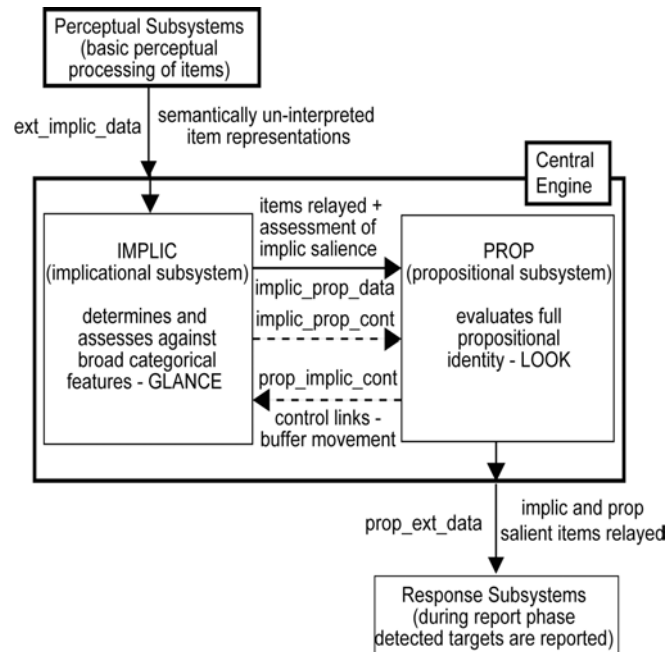


Figure 4: Top-level Structure of Model

The central-engine contains two interacting processes, represented as internal boxes - IMPLIC and PROP. These implement the implicational and propositional systems respectively. The fact that we can structure our specification in this way indicates one of the strengths of process algebra - since arbitrary process hierarchies are available to us, the complex decompositional structure of theories such as ICS can be directly represented. The behaviour of the central-engine emerges through interaction between processes at this level, with the communication channels between IMPLIC and PROP playing a major role. These channels are divided into two types - a data channel (`implic_prop_data`) and control channels (`implic_prop_cont` and `prop_implic_cont`). The former of these is used to relay item representations from the implicational to the propositional system, while the latter two links are used in order to control buffer movement. As an indication of the style of specification arising with LOTOS, the central-engine in our specification would have the following basic format (although, of course, the specification is in fact much more complex than this and in addition, in order to simplify presentation, we have pared down the LOTOS syntax somewhat).

```

process CENT_ENG :=
hide implic_prop_data, implic_prop_cont, prop_implic_cont in

```



```

IMPLIC
|[ ext_implic_data, prop_implic_data, implic_prop_data,
  implic_prop_cont, prop_implic_cont ]|
PROP

```

where

```

process IMPLIC :=
(* The body of the Process would go here. *)
endproc

process PROP :=
(* The body of the Process would go here. *)
endproc

```

endproc

To pick out a few aspects of this specification, notice that there is a separate definition for each process (e.g. the IMPLIC and PROP definitions). Also the syntax “hide x_1, \dots, x_n in Y ” states that the synchronisation actions x_1, \dots, x_n are local and thus, not available to the context in which a process (here CENT_ENG) is placed. The expression:

```

IMPLIC
|[ ext_implic_data, prop_implic_data, implic_prop_data,
  implic_prop_cont, prop_implic_cont ]|
PROP

```

denotes that the implicational and propositional systems execute independently in parallel subject to interaction via the actions: `ext_implic_data`, `prop_implic_data`, `implic_prop_data`, `implic_prop_cont` and `prop_implic_cont`. Thus, the, so called, parallel composition operator,

$$P \mid [Y_1, \dots, Y_n] \mid Q$$

yields distributed control, it allows two processes (here P and Q) to execute independently of one another. However, it states that the two processes can interact by exchanging messages over the action links Y_1, \dots, Y_n .

Data Representations

The entities in the RSVP task we are modelling are words and our data representations need to reflect this. Firstly, note that there are three types of words in the paradigm - background, target and distractor. Secondly, at this level of theoretical abstraction, all we need from our data representations is to distinguish between the different word types. We use an enumerated type, which states that a word identifier can either be `Back`, `Targ` or `Dist`.

```
Word_id ::= Back | Targ | Dist
```

The purpose of the model is to associate salience assessments with words in the RSVP. As a result, the actual data representation passed through the model is a triple that contains three “slots” - a word identifier (as just introduced), an implicational salience assessment and a propositional salience assessment, i.e.

```
Rep ::= ( Word_id , Sal , Sal )
```

When entering the central engine, the last two slots would be set to `U`, indicating that they are un-interpreted. The `IMPLIC` process strives to place a `T` or an `F` in the second slot (indicating `True` if the word is implicationally salient or `False` if it is not) and `PROP` has a similar role for the third slot.

We could have constructed our model with complete words as the items entering and being passed through it. A new item would then enter the model at RSVP rates of between 90 and 120ms per item (the SOA rate of the experiment). In fact, we use a more fine-grained timing compatible with the assumption that underlying neural mechanisms can represent updates around every 20ms (e.g. see Bond, 1999; Rolls & Stringer, 2001). Thus, a new item enters the system every 20ms. An item can be thought of as a word constituent and a 90-120ms word is comprised of 5 to 6 items. Each item is modelled as a triple in the form of `Rep` above. Our explanation of the blink data will be in terms of the time-course of allocation of attentional resources. Consequently, it is useful to have a fine grain of timing so that we can be more discriminating with regard to this time-course. Also, we wish ultimately to refine the model in terms of the temporal build-up of perceptual, semantic and affective aspects of word representations. We would not be able to do this if every update corresponded to a complete word.

An important aspect of this approach is that we largely abstract away from the process of mapping words to meanings. Word category is built into our representation in that the first slot contains `Back`, `Targ` or `Dist`, which is available when items enter the central engine. This may seem strange when modelling an experiment that concerns assessing whether or not words are in particular categories. However, in common with the majority of the literature on the attentional blink, we explain the phenomenon in terms of how attentional resources are *allocated* and this can be done without detailed modelling of the extraction of meaning from words. In fact, abstracting away from such analysis allows us to concentrate our efforts on modelling the time-course of information processing which is the pivotal issue.

Thus, our items do not directly correspond to either letters, graphemes or even features (which are standard components of word recognition). Although we would accept the involvement of such entities at some level of analysis, we have abstracted from these here. Our delay lines are time constrained and this is reflected in our items, each of which represents a 20ms time slice and, as previously discussed, 5-6 items correspond to a word. One can interpret the first item in such a sequence of 5-6, as something like the amount of information in the first 20ms of processing the word. Thus, the model employs a time sliced decomposition of words rather than an object based or featural decomposition.

We would argue that such an abstract encoding of words is an advantage since we are not imposing any constraints on the nature of the word recognition process. All we are assuming is that there are two salience assessment mechanisms, which when completed are registered in item representations. Thus, we impose no constraints on the mechanism by which salience assessment is achieved. All we need for the model to work is to know that salience assessments are made and that different types of words have different salience outcomes, viz background words, target words etc.




Figure 5: (a) Example delay line; (b) typical delay line state; (c) IMPLIC processing in buffered mode.

Delay-Lines

At one level, we can view the model as implementing a pipeline. New constituent items enter the pipeline via action `ext_implic_data`, items are then fed through IMPLIC and passed into PROP via action `implic_prop_data` and then they reach the end of the pipeline via action `prop_ext_data`. Every 20ms a new item enters the pipeline and all items currently in transit are pushed along one place. The IMPLIC and PROP subsystems perform their corresponding salience assessments as items pass through the pipeline. The key data structure that implements this pipeline metaphor is a delay-line. This is a simple data structure for recording time

constrained serial order. Thus, it is a very natural mechanism to use for capturing the temporal properties of the blink experiment, which is a time constrained order task. A simple delay-line of 4 elements is depicted in fig. 5(a). This records the last 4 time instants of data. In our model, each instant corresponds to a 20ms time-slice. The pipeline employed in our model is, in overall length, considerably longer than 4 units and we will not depict it in full here. However, it is worth representing a typical state of a 12 item portion of the overall delay-line during our blink simulations (fig. 5(b)). As a simple approximation some 6 updates correspond to a 120ms SOA, i.e. to one word, or, in other words, one RSVP word is modelled by a sequence of 6 delay-line items.

Salience assessment, Buffering and Attention

Each central-engine subsystem contains an input array and a main subsystem delay-line, which represent, in simplified form, the “image” of recently arrived data. As previously suggested, every 20ms a new item enters a subsystem’s input array, which, in effect, acts as a mini (3 item) delay-line for the most recently arriving data. As shown in fig. 5c, items in the main subsystem delay-line (in ICS terms, basic units) are constructed by looking across the input array to see how the representation builds-up over time. At this level, updates are made every 60ms corresponding to update of the entire input array (i.e. 3 times 20ms). In this kind of paradigm, each new word has to be assessed for salience in the context of the instructed category. Our theoretical proposal is that in the context of “novel” information patterns in each subsystem salience assessment is performed by this processes of looking across the input array (fig. 5c).

We should note in passing that basic units of implicational representation are effectively composed by summarising the current dynamic state of the input array and placing this in a higher order delay line with a slower rate of change. It also captures, in elemental form, what amounts to a hierarchical structuring of information content. In the broader ICS architecture, constituents arrive from multiple sources and this mechanism allows for multimodal integration into higher order units of representation.

In the model, a process in a buffered state represents the focusing of central attentional resources at that subsystem. In particular, a subsystem can only assess salience of novel, unautomated patterns, if it is buffered. Since, as previously stated, only one subsystem can be buffered at any one instant internal attention can only be focussed at one processing location. Thus, in terms of the current model, only one subsystem can be assessing salience at any instant. Buffering ensures the serial allocation of attentional resources.

In understanding the model, it is important to realise that all constituent items entering the central engine (at action `ext_implic_data`), will pass through the pipeline (coming out at `prop_ext_data`) and they will be output in the same order that they are input. However, due to the seriality of buffering there is no guarantee that the item will have been assessed for salience at both `IMPLIC` and `PROP`. If a subsystem is not able to perform a salience assessment on an item, because it is not buffered, then the corresponding slot in the item representation is left as `U` (i.e. un-interpreted). In addition, a word will only be reported if both its implicational and propositional salience have been assessed and both have been found to hold, i.e. both the second and third slots in the item representation are `T`.

The algorithm that controls buffer movement is thus central to realising the attentional blink. Two principles control buffer movement:

1. The subsystem that is buffered decides when the buffer moves and actively passes it on to the other subsystem (by sending a signal down a control link, see Figure 4).
2. A subsystem decides to pass on the buffer by observing the representations it is processing.

The exact algorithm can be explained as follows:

Initially IMPLIC is buffered

*If IMPLIC is buffered and it detects an implicationally salient item
then the buffer is passed to PROP*

*If PROP is buffered and it detects an implicationally un-interpreted item
then the buffer is passed to IMPLIC*

There is a delay between the time at which a subsystem decides that the buffer should move and when it actually does. This is justified on the grounds that redirection of attentional resources is computationally demanding and thus, time constrained. In other words, the system becomes “locked into” a particular processing configuration, which causes inertia when the buffer needs to move. This inertia will provide the explicit underlying mechanism for the delayed onset of the attentional blink that, according to Shapiro et al. (1997a), is only accounted for in a rather post-hoc manner by most existing theories.

Summary of how the model blinks

1. Targets are missed when an earlier distractor is found to be implicationally salient, causing the buffer to move from `IMPLIC` to `PROP`. While the system is buffered at `PROP`, implicational salience cannot be assessed. Consequently, the implicational salience of a target would fail to be assessed and hence, during the reporting phase, the system would fail to

register the overall salience of the word. In other words, the model would have blinked!

2. The delayed onset of the blink arises because (as just mentioned) there is a delay between implicational salience assessment of an item and that item passing out of `IMPLIC`. Thus, closely following items will be at early stages in the pipeline before the distractor item has fully passed out of it and the buffer moves. This ensures that targets immediately following a distractor are likely to be processed for implicational salience.
3. Recovery from the blink arises because if there is sufficient separation between the distractor and the target, `PROP` level assessment of salience of the distractor will have completed before the target enters `IMPLIC`. Consequently, the buffer will have returned to `IMPLIC` in time to assess the implicational salience of the target.

Results

The key parameters that need to be set when running simulations are the buffer movement delays and the salience assessment thresholds for both subsystems. In standard fashion, we add noise into our simulations by sampling our parameters randomly from distributions. This is done in order to give variability between simulation runs and thus to reflect individual differences. For the simulations reported here buffer movement delays were set as follows:

1. *Implicational Buffer Movement.* The delay between an implicationally salient item entering `IMPLIC` and the buffer moving is set to either 180 or 240ms. These correspond to 9 and 12 update cycles, and act to represent individual differences in redirecting attentional resources from `IMPLIC`. We sample randomly from these possibilities with equal probability.
2. *Propositional Buffer Movement.* The delay between an implicationally uninterpreted item entering `PROP` and the buffer moving is sampled from the following six possibilities: 60, 120, 180, 240, 300 and 360ms. Again we sample randomly, however, here the probability mass associated with each possibility is determined according to a normal distribution. A probability mass of 0.341 is associated with 180 and 240; 0.136 with 120 and 240; and 0.023 with 60 and 360.

Broadly speaking, the speed of implicational buffer movement controls the shape of the blink onset. Smaller implicational buffer movement values reduce lag 1 performance. This is because the longer the gap between implicational salience detection of the distractor and the buffer moving to `PROP`, the more the chance that following target items are completely processed for

implicational salience. Values of 180 and 240ms ensure that lag 1 performance is not impaired. In addition, because the distribution of values is small, just two values, there is not a great deal of variability in assessment outcome at early serial positions; this ensures that blink onset is steep.

In contrast, the propositional buffer movement value plays a significant role in controlling the shape of the blink offset, although it is not the only factor. As previously suggested, recovery from the blink arises because the gap between distractor and target is long enough that *PROP* salience assessment of the distractor has been completed and the buffer has moved back to *IMPLIC* by the time that the target starts being processed. Consequently, longer *PROP* buffer movement times tend to slow blink recovery. Furthermore, the fact that propositional buffer movement values range over a sample set of 6 values (compared to the 2 for *IMPLIC*) contributes to obtaining a slow blink offset. In psychological terms this implies that there is greater variability in the time course of propositional evaluation than implicational evaluation. This fits with the perspective that, over the range of possible distractor and target items, there is a greater diversity of propositional salience levels, which in turn yields greater variability in the processing required to assess propositional salience.

Salience assessment thresholds also have to be set. Assessing an item for salience yields a value. The smaller the value, the more salient the item. The salience assessment threshold determines how small this value has to be for an item to be adjudged to be salient.

1. *IMPLIC Threshold*. The salience assessment threshold used at *IMPLIC* is sampled from three possible values, each has equal probability of being selected (one third). For all the three possible threshold values, background words will be judged not to be implicationally salient. For the smallest threshold value, neither distractor nor target words will be judged to be salient. For the intermediate value, target words will be salient and distractor words will not be salient, while for the largest value, both target and distractor words will be salient.
2. *PROP Threshold*. The *PROP* salience assessment threshold is not randomly sampled. It is set to a value that ensures that background and distractor words are not propositionally salient and target words are.

These threshold settings enable us to capture psychologically plausible ideas concerning salience in our task setting. Rather like wide and narrow sensory filters, there is considerably less precision in implicational salience assessment than in propositional salience assessment. For example, *PROP* will never mistakenly interpret a distractor word as propositionally salient. In contrast, not only can *IMPLIC* interpret distractor words as implication-

ally salient, it can also assess target words as not implicationally salient, although this is a low probability outcome. This could be argued to fit in with the theory that *IMPLIC* is implementing a “superficial glance”, while *PROP* is implementing a “detailed look”. In other words, *IMPLIC* can more easily be fooled than *PROP*.

Figure 5 reproduces a subset of behavioural data from the Barnard et al paradigm. In these data, no blink effect was obtained for the control condition in which the control distractors were words that were members of the background list category. When the distractors were in the category human and thus semantically related to the job/occupation target words, a substantial blink occurs. The model fit to these curves is shown. The fit to the data is very close across the first few serial positions, but the simulation of report levels recovers for human distractors faster than the behavioural data. However, it will be recalled that the Barnard et al. paradigm made use of *three* response categories - full and accurate report of the identity of the job word, reporting that they saw a job but were not sure what it was, and no they didn’t see a job. In fact, in the Barnard et al data, failure of report levels to recover to control levels over the later serial positions could be attributed to reports of target presence but lack of identity. In our existing model, we made no attempt to implement the stage of deciding exactly which response to give. Indeed, other than a steeper onset, the shape of the blink function in the simulation is very similar to that observed in the original report of a letter-based blink (fig. 1) in a dual target setting where the only requirement is to report awareness of T2.

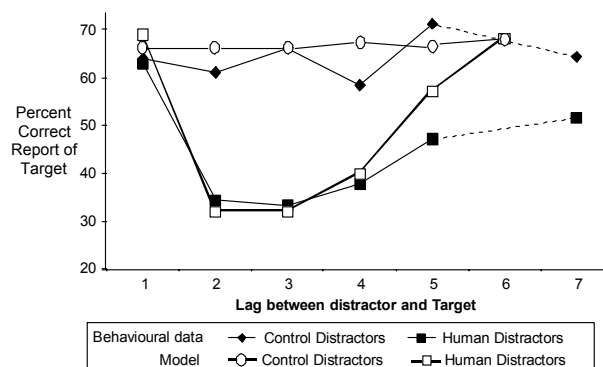


Figure 5: Simulation results from 3000 trials using the parameter settings given above.

Barnard et al. also found that accurate report of the job targets was related to their semantic similarity to distractors. At least qualitatively, our

model will give this effect. If we make the IMPLIC salience assessment threshold more generous, then more distractors will be assessed as IMPLIC salient and consequently, more targets will be missed, i.e. the blink will become deeper. A larger IMPLIC threshold could model a smaller distance between distractor and target in semantic space. Indicating that more distractor assessment mistakes will be made in IMPLIC.

Discussion

Our first aim was to illustrate the potential of process algebra to implement modular mental architectures of wide scope. Those aspects of psychological theory captured are the idea that executive control is distributed, illustrated here by interactions between ICS subsystems that represent two types of meaning. Each of these subsystems only had access to its own internal state. Executive control emerged as a function of their interaction which itself involved an exchange of data and control information. The internal structure of each subsystem was represented by component processes operating on the content of delay lines that represented an image of those basic units that have recently arrived. For the purposes of presentation, discussion of both mental architecture and the actual model have necessarily been brief. The full specification and how its behaviour depends on the parameters are rather more intricate than we have been able to elaborate here. An extended presentation of the computational issues, as well as a more detailed comparison with other psychological models, and a discussion of the extent to which the buffering and pipeline mechanisms fully capture processing modes in ICS is provided in Bowman & Barnard (2001). The full specification is available from Bowman.

From a computational perspective, the core mechanisms involve concurrent processing of different types of representation. They also realise the idea that information can be processed within a subsystem in a number of different modes, with a state of buffered processing being associated with focal awareness and the consequent allocation of limited attentional resources. Effectively, the buffer acts as a movable locus of control within distributed processing activity. Our account of the blink phenomena rests upon the ideas that buffer movement from one subsystem, which represents generic meaning, to another, which represents more specific meaning, takes time, and that the state of buffering is intimately connected to “what” representations get processed in real time. The particular trajectory of representational exchange that determines “what” gets processed is also only a function of the representation that is locally available. Here those representations merely marked class, salience and state of interpretation. The process algebra notation enabled us to capture the hierarchical composition of process-

ing resources, the emergence of executive control from distributed activity and it enabled us to deal with explicit process action at a level of abstraction comparable to the original psychological reasoning. The model runs and produces a blink effect without requiring more detailed assumptions about perceptual, lexical or semantic representations.

In the introduction we discussed why we have preferred process algebra to either connectionism or production systems. However, the status of hybrid techniques in this debate is also worth dwelling on. Such techniques combine symbolic and connectionist ideas. One element of, which is indeed adding a form of, distributed control to symbolic approaches. This might, for example, allow activation levels to be associated with concepts and symbolic computations. Although these techniques have made an important contribution to topics such as representation of semantic concepts, planning and analogical reasoning, they do not really fulfil our modelling requirements. For example, in many of these techniques, especially those that have their roots in semantic networks, the unit of distributed control is much more fine grain and computationally limited than we seek. For example, the units of distribution in (Hendler 1989) are at the level of concepts and do not have threads of computational control associated with them. In contrast, each ICS subsystem (our unit of distribution) is computationally rich and this needs to be reflected in the modelling paradigm chosen.

Some of the more recent agent-based hybrid techniques come closer to our requirements. For example, in the DUAL architecture (Kokinov 1994; Petrov and Kokinov 1999) the units of distribution – agents – are computationally more rich and it would indeed be interesting to attempt to implement an architecture such as ICS in a framework such as DUAL. However, we believe such an implementation would still be limited in respect of our requirements. In particular, we have advocated abstract specification of cognitive behaviour without the imposition of implementation assumptions. Formal specification techniques, such as process algebra, are exactly targeted at obtaining such a non-prescriptive style of description. In contrast, hybrid agent-based models, such as DUAL, are tied to programming-level implementation in languages such as Lisp. It is also unclear whether the communication between distributed units (either of the marker passing or activation exchange variety) is as rich as it is in process algebra. However, a detailed comparison of the techniques would be required in order to answer this question.

Our second objective was to outline an implementation of a specific model of the attentional blink based upon the distributed control hypothesis. Throughout the history of attention research, from Broadbent (1958) onwards, much research has been directed at positioning “the serial bottle-

neck" and its attributes, in a chain of processing through sensory, perceptual and central mechanisms. Over time, concern has increasingly drifted towards post-perceptual factors. From a psychological perspective, we have effectively moved this argument to an extreme point. The running model shows that the basic blink phenomenon is open to explanation in terms of entirely central, or executive processing activity determined by meaning. Although the wider ICS architecture does allow for earlier effects of stream selection in terms of what is passed from visual processing through to the centre, the serial position curves for the blink are determined entirely by semantic stages. Teasdale and Barnard (1993) also equate buffered processing in ICS with focal awareness of the content of a particular type of mental representation at the subsystem where buffering resides. Hence, buffering of the implicational image would be directly linked to momentary generic awareness of the arrival of a salient item, while buffering of the propositional image would be associated with focal awareness of the referential identity of an item, and that this current propositional information affords report.

Bowman and Barnard (2001) provide more detailed comparisons with other models, and in particular, that advanced by Chun and Potter (1995). Among the key comparisons, our model does not rely on interference or similarity effects at perceptual levels of encoding, decay of stored material, or a simple filter-like "bottleneck", since in our architecture data can pass from one stage to another differing only in its representational attributes. Central to these differences is the handling of concurrency. The serial pipeline of processing is a central metaphor in our approach. Nonetheless, different items are being processed simultaneously at different stages in that pipeline (e.g. at *IMPLIC* and *PROP*). We are aware of no other running models against which the present one can be compared in detail, and without such implementation ambiguities arise. For example, we could directly compare the Chun and Potter first stage with our "glance" mechanism and our "look" with their capacity limited, second stage. However, we could also regard their first stage as more equivalent to the process in our model that passes data to the central engine (*Vis* → *Implic*), and consider our "look and glance" to be more of an elaboration of their second stage in which material is consolidated for report.

Our account of the blink phenomenon in terms of central mechanisms alone may appear to fit quite uncomfortably in this empirical domain. The vast bulk of the research has focused upon influences such as lexical priming (Shapiro, Driver et al. 1997), perceptual masking, or interference effects (Raymond, Shapiro et al. 1992; Chun and Potter 1995; Shapiro, Arnell et al. 1997). It is, for example, clear from these latter studies that the T1+1 and T2+1 items play an important role in obtaining the blink. It has been argued

that they act as masks for the T1 and T2 items, and that they prevent consolidated processing of the targets. For example, Chun and Potter reduced the strength of the blink by placing a blank at the T1+1 position. Furthermore, they varied the ease with which the T1+1 item could be discriminated (by type) from T1 and found that the blink deepened with increased difficulty in discriminating T1+1 from T1.

Qualitatively such an effect could occur in our semantic subsystems by adapting the implicational buffer movement mechanism without recourse to peripheral mechanisms. Currently, *IMPLIC* passes the buffer to *PROP* at the point at which implicationally salient items start to be output from *IMPLIC*. However, we could assume that the buffer moves earlier if the level of implicational salience of items on the input array (i.e. at the input end of the *IMPLIC* part of the pipeline) falls *below* a certain level. Then the buffer would move more quickly from *IMPLIC* to *PROP* if the T1+1 item could easily be distinguished from T1, the extreme case being when T1+1 is a blank. Now if the buffer moves more quickly to *PROP*, it will also return more quickly to *IMPLIC*. As a result, the blink will be shorter and is also likely to be shallower. Such an adaptation is psychologically really rather plausible. It would enable processing resources to be re-directed under low demand as well as high demand. This would be the case if a blank, or easily discriminated items, appear on the *IMPLIC* input array. Although unimplemented, such arguments suggest that a range of effects may be open to explanation in terms of mechanisms evaluating semantic representations of current input in relation to semantic representation of the current task demands.

Naturally, the basic blink effect is open to a number of alternative explanations (e.g. see Shapiro, Arnell et al., 1997), each of which could potentially be given computational realisation, although we are currently aware of no published implementations. Our reproduction of human data gives our existing model the status of a *candidate* explanation of the basic attentional blink effect. As with all such models, it serves first to illustrate that an informal explanation of the blink (see the section *Summary of how the model blinks*) can be realised computationally, and developed to frame predictions that can be tested in further empirical research. In this context, an important aspect of our work is that we have developed our model in terms of a macro-theoretic architecture, ICS, that has already been extensively developed to address the influence of emotion on normal and dysfunctional cognition (e.g. see Barnard, 1999; Teasdale & Barnard, 1993). Our current implementation of the distributed control hypothesis thus provides a basic platform which can be readily extended to address not only how attention is captured by emotionally, as well as cognitively, salient material, but also some reported effects of brain damage.

For example, our choice of a delay-line representation for the image record component in an ICS subsystem is motivated in part by the requirement to enable a process to be able to see over data that arrives via different configural routes through a distributed processing mechanism (see Fig 3). Word meanings derived from longer routes that involve sequential parsing through object and morphonolexical subsystems, will not necessarily arrive in the propositional image of the ICS architecture at exactly the same moment as perhaps coarser grain information arriving by different routes. Some information extracted directly from sensory patterns like facial expression or tone of voice may arrive earlier, while the build up of bodily experiences and the feedback this generates could even arrive somewhat later. Delay lines provide a simple mechanism whereby a locally represented state can bring temporally offset information into a unified structure. Our explanation of the blink effect relies on buffer movement that is in part determined by the concurrent states of both propositional and implicational meaning. The implicational level of representation represents generic meanings that can be affectively charged on receipt of appropriate combinations of constituents arriving from the sensory subsystems and the propositional subsystem. The representations are built concurrently in the ICS architecture, but out of different inputs with different lags. Reference back to fig. 3 will show that four subsystems send constituent data to the implicational subsystem: the acoustic, visual, body state and propositional subsystems. In effect, the implicational subsystem builds schematic models of meaning in the context of particular sensory patterns. It is at this level, for example, that tone of voice, facial expression and activated or lowered body states can influence our “understanding”.

Such an approach provides a relatively straightforward way of addressing other influences on the blink effect while retaining the same core mechanisms. The onset of the blink in our model is effectively determined by the time that the buffering mechanism dwells at the implicational level of representation, while the recovery of report depends on the dwell time of buffering at the propositional level of representation. This allows relatively straightforward extensions to deal with emotional influences on the shape of the report function as well as the effects of neurological damage. For example, the build up of affect depends upon the current state of implicational representations. The presence of a “threat marker” at that level of representation can be assumed to influence implicational salience assessment, and this provides a means whereby aversive words occurring as a second target (T2) can break through the blink, i.e. can interrupt T1 processing. In ICS, information can be processed in direct mode rather than buffered mode if the transformation is fully automated. So, one’s own name or even threat material for anxious individuals, can be assigned salience without requiring

buffering and this would enable data to be passed from implic to prop with its implicational salience parameter set. Were this to be implemented in the model, our own names or threat words could break through the blink. Both types of effect have been reported, the former by Shapiro, Caldwell, et al., (1997) and the latter by Anderson & Phelps (2001). Of potentially far greater importance are interactions. So, for example, Anderson & Phelps (2001) also found that the attenuated blink, with aversive T2s, does not occur in several cases of left amygdala damage. If we were to assume that salience is influenced by feedback from body state representations, via the implicational-body state loop, and that the amygdala is among those brain structures that mediate body state representation of attributes related to fear, then any disruption to that circuitry would prevent the body state consequences of threat being fed back to the implicational subsystem. This is a more or less direct analogue of the somatic marker hypothesis proposed by Damasio, Tranel, and Damasio (1991). Importantly, the prevention of body state feedback does not affect our core blink mechanism of buffer movement based upon a semantic “glance” and a more detailed “look.” Accordingly the model allows a normal blink to occur while enabling us to accommodate the absence of an attenuated effect with amygdala damage.

Another key finding from neuropsychological studies of the blink effect is that patients with right parietal brain damage, giving rise to visual neglect, also display a blink effect, but one that shows very shallow recovery to baseline levels of accurate report (Husain, Shapiro et al. 1997). Indeed, it can take around 1500ms for report to fully recover to baseline levels. It seems as if such patients have a massively protracted “look” stage of processing. We can assume that a full account of what goes on in our look stage involves not only implicational input to propositional representations but later arriving inputs from more extended processing of visuo-spatial (via the object subsystem) and lexical form (via the morphonolexical subsystem in ICS). Consequently, brain damage impairing the capabilities of either of these subsystems would then require more extensive exchanges between these subsystems and the propositional subsystem before a stable propositional representation emerges. Accordingly, buffering would dwell in the propositional subsystem, or at object and MPL subsystems, for a considerably extended duration prior to returning to implic and hence lead to a shallower recovery curve. As with the arguments concerning feedback of body state information, the basic mechanisms of buffer movement dependent on salience assessment would be retained, the only thing that alters concerns the time taken to build either the implicational or propositional representations required to trigger buffer movement.

These particular ideas are naturally speculative, since the details are as yet unimplemented. They nonetheless provide concrete illustrations of how

the distributed control of two qualitatively distinct types of semantic representation, each preserved in functionally dissociated delay line memories, can yield explanations of related phenomena without a requirement to incorporate additional processing resources or mechanisms. More widely, the approach itself adds a new tactic for macro-theoretic level modelling of the behaviour of complex mental architectures to supplement the use of connectionist, symbolic or hybrid modelling technologies. Barnard, May, Duke, and Duce (2000) present a more detailed discussion of the use of computer science based methods for modelling macro-theoretic, as opposed to micro-theoretic, issues in the governance of behaviour. Our use of process algebra in this paper has provided the means for demonstrating how interactions between two subsystems with decentralised or distributed control can give rise to behavioural data. It does so at a level of abstraction that limits the number of micro-theoretic commitments, which go beyond what the psychological theory was seeking to implement. We have effectively implemented a core model of executive control of attention, and extensions to it require us to implement, at an equally abstract level of specification, interacting loops between meaning subsystems and other subsystems such as body-states or object and lexical levels of representation. Process algebra provides us with a notation that makes such modelling tractable. Indeed, we believe that specification languages of this class could well open up many domains of cognitive-affective enquiry, that have either been closed to or only so far exposed to computational modelling of restricted scope.

Acknowledgements

We would like to thank the reviewers of this paper who made valuable suggestions, which have greatly improved the presentation of our ideas.

References

- Abeles, M. (1991). *Corticonics*. Cambridge, UK, Cambridge University Press.
- Allport, D. A., E. A. Styles, et al. (1994). Shifting Intentional set: Exploring the dynamic control of tasks. *Attention & Performance XV: Conscious and non-conscious information processing*. C. Umiltà and M. Moscovitch. Cambridge Mass, MIT Press: 421-452.
- Anderson, A. K. (2001). "Affective influences on attentional dynamics during perceptual encoding: investigations with the attentional blink." (*Under Review*).
- Anderson, A. K. and E. A. Phelps (2001). "Lesions of the human amygdala impair enhanced perception of emotionally salient events." *Nature* **411**: 306-309.
- Baddeley, A. D. (2000). "The Episodic Buffer: a New Component of Working Memory?" *Trends in Cognitive Sciences* **4**(11): 417-423.

- Barnard, P. (1985). *Interacting Cognitive Subsystems: A Psycholinguistic Approach to Short-term Memory*. Progress in the Psychology of Language, Vol. 2. A. Ellis. Hove, UK, Lawrence Erlbaum Associates Ltd: 197-258.
- Barnard, P. and J. May (1999). "Representing cognitive activity in complex tasks." Human-Computer Interaction **14**: 93-158.
- Barnard, P., J. May, et al. (2000). "Systems, interactions and macrotheory." ACM Transactions on Human-Computer Interaction **7**(2): 222-262.
- Barnard, P. and C. Ramponi (2001). "Effects of state and trait anxiety in a semantic variant of the attentional blink paradigm." (*in Prep*).
- Barnard, P. J. (1999). *Interacting Cognitive Subsystems: Modelling Working Memory Phenomena with a Multi-processor Architecture*. Models of Working Memory. A. Miyake and P. Shah, Cambridge University Press.
- Barnard, P. J., S. Scott, et al. (2001). "Paying Attention to Meaning." Under Review.
- Beck, A. T. (1976). Cognitive Therapy and the Emotional Disorders. New York, International Universities Press.
- Bolognesi, T. and E. Brinksma (1988). "Introduction to the ISO Specification Language LOTOS." Computer Networks and ISDN Systems **14**(1): 25-39.
- Bond, A. H. (1999). "Describing Behavioural States Using a System Model of the Primate Brain." American Journal of Primatology **49**: 315-338.
- Bowman, H. and P. J. Barnard (2001). *Computational Modelling of Distributed Executive Control*, Computing Laboratory, University of Kent at Canterbury.
- Bowman, H. and G. Faconti (1999). "Analysing Cognitive Behaviour using LOTOS and Mexitl." Formal Aspects of Computing **11**: 132-159.
- Broadbent, D. E. (1958). Perception and Communication. London, Pergamon Press.
- Broadbent, D. E. and M. H. P. Broadbent (1987). "From Detection to Identification: Response to Multiple Targets in Rapid Serial Visual Presentation." Perception and Psychophysics **42**: 105-113.
- Carver, C. S. and M. F. Scheier (1998). On the Self-Regulation of Behaviour. Cambridge, Cambridge University Press.
- Chun, M. M. and M. C. Potter (1995). "A Two-Stage Model for Multiple Target Detection in Rapid Serial Visual Presentation." Journal of Experimental Psychology: Human Perception and Performance **21**(1): 109-127.
- Cohen, J. D., K. Dunbar, et al. (1990). "On the Control of Automatic Processes: A Parallel Distributed Processing Account of the Stroop Effect." Psychological Review **97**(3): 332-361.
- Cooper, R. (1995). Towards an Object-Oriented Language for Cognitive Modelling. 17th Annual Conference of the Cognitive Science Society, Pittsburgh, Pennsylvania, Lawrence Erlbaum Associates.
- Cooper, R., J. Fox, et al. (1996). "A Systematic Methodology for Cognitive Modelling." Artificial Intelligence **85**: 3-44.
- Damasio, A. R., D. Tranel, et al. (1991). Somatic markers and the guidance of behavior: theory and preliminary testing. Frontal lobe function and dysfunction. H. S. Levin, H. M. Eisenberg and A. L. Bemton. New York, Oxford University Press: 217-219.
- de Meer, J., R. Roth, et al. (1992). "Introduction to Algebraic Specification Based on the Language ACT-ONE." Computer Networks and ISDN Systems **23**: 363-392.
- Duke, D. J., P. J. Barnard, et al. (1998). "Syndetic Modelling." Human-Computer Interaction **13**(4): 337-393.

- Duncan, J. (2000). Visual attention in mind and brain. Brain mechanisms of perception, learning and memory: Essays for Gabriel Horn. J. J. Bolhuis. Oxford, Oxford University Press: 49-68.
- Duncan, J., R. Ward, et al. (1994). "Direct Measurement of Attentional Dwell Time in Human Vision." Nature **369**: 313-315.
- Fodor, J. A. (1983). Modularity of Mind: An Essay on Faculty Psychology. Cambridge, Massachusetts, MIT Press.
- Fodor, J. A. and Z. W. Pylyshyn (1988). "Connectionism and Cognitive Architecture: A Critical Analysis." Cognition **28**: 3-71.
- Hendler, J. A. (1989). "Marker-passing over Microfeatures: Towards a Hybrid Symbolic/ Connectionist Model." Cognitive Science **13**: 79-106.
- Hoare, C. A. R. (1985). Communicating Sequential Processes. London, Prentice-Hall.
- Holmes, A. and A. Richard (1999). Attentional bias to threat related material in anxiety: a resource allocation or a practice effect (poster presentation at September Conference). Cognitive Section of the British Psychological Society.
- Husain, M., K. L. Shapiro, et al. (1997). "Temporal dynamics of visual attention reveal a non-spatial abnormality in spatial neglect." Nature **385**: 154-156.
- Kokinov, N. B. (1994). The DUAL Cognitive Architecture: A Hybrid Multi-Agent Approach. ECAI, 94, 11th European Conference on Artificial Intelligence, John Wiley.
- Landauer, T. K. and S. T. Dumais (1997). "A Solution to Plato's Problem: The Latent Semantic Analysis Theory of the Acquisition, Induction and Representation of Knowledge." Psychological Review **104**: 211-240.
- Leventhal, H. (1979). A perceptual-motor processing model of emotion. Perception of Emotion in Self and Others. P. Pilner, K. Blankstein and I. M. Spigel. New York, Plenum. **5**: 1-46.
- MacLeod, C., A. Mathews, et al. (1986). "Attentional bias in emotional disorders." Journal of Abnormal Psychology **95**: 15-20.
- Maki, W. S., K. Frigen, et al. (1997). "Associative Priming by Targets and Distractors During Rapid Serial Presentation." Journal of Experimental Psychology: Human Perception and Performance **23**: 1014-1034.
- Meyer, D. E. and D. E. Kieras (1997). "A computational theory of executive cognitive processes and multiple task performance: Part 1. Basic mechanisms." Psychological Review **104**: 3-65.
- Milner, R. (1989). Communication and Concurrency, Prentice-Hall.
- Moray, N. (1959). "Attention in dichotic listening: Affective cues and the influence of instruction." Quarterly Journal of Experimental Psychology **11**: 56-60.
- Newell, A. (1990). Unified Theories of Cognition. Cambridge, Massachusetts, Harvard University Press.
- Newell, A. and H. A. Simon (1972). Human Problem Solving. Englewood Cliffs, N.J., Prentice-Hall.
- Page, M. P. A. and D. Norris (1998). "The Primacy Model: A New Model of Immediate Serial Recall." Psychological Review **105**(4): 761-781.
- Pashler, H. and J. C. Johnston (1998). Attentional Limitations in Dual-task Performance. Attention. H. Pashler. Hove, UK, Psychology Press: 155-189.
- Petrov, A. A. and N. B. Kokinov (1999). Processing Symbols at Variable Speed in DUAL: Connectionist Activation as Power Supply. 16th International Joint Conference on Artificial Intelligence.

- Plaut, D. C. (1998). "A Connectionist Approach to Word Reading and Acquired Dyslexia: Extension to Sequential Processing." Cognitive Science, Special Issue on Connectionist Models of Human Language Processing: Progress and Prospects **23**: 543-568.
- Potter, M. C. (1999). Understanding sentences and scenes: the role of conceptual short term memory. Fleeting Memories. V. Coltheart. Cambridge Mass, Bradford Books: 13-46.
- Raymond, J. E., K. L. Shapiro, et al. (1992). "Temporary Suppression of Visual Processing in an RSVP Task: An Attentional Blink." Journal of Experimental Psychology: Human Perception and Performance **18**(3): 849-860.
- Raymond, J. E., K. L. Shapiro, et al. (1995). "Similarity Determines the Attentional Blink." Journal of Experimental Psychology: Human Perception and Performance **21**: 653-662.
- Rolls, E. T. and S. M. Stringer (2001). "A Model of the Interaction between Mood and Memory." Network: Computation in Neural Systems **12**: 89-109.
- Roscoe, A. W. (1998). The Theory and Practice of Concurrency, Prentice-Hall.
- Rumelhart, D. E., J. L. McClelland, et al. (1986). Parallel Distributed Processing, Explorations in the Microstructure of Cognition. Volume 1: Foundations and Volume 2: Psychological and Biological Models. Cambridge, Massachusetts, A Bradford Book, The MIT Press.
- Scherer, K. R. (2000). Emotions as Episodes of Subsystem Synchronisation Driven by Non-linear Appraisal Processes. Emotion, Development and Self-Organisation. M. D. Lewis and I. Granic, Cambridge University Press: 70-99.
- Schneider, S. (2000). Concurrent and Real-time Systems, The CSP Approach. Chichester, Wiley.
- Schneider, W. (1999). Working Memory in a Multilevel Hybrid Connectionist Control Architecture (CAP2). Models of Working Memory, Mechanisms of Active Maintenance and Executive Control. A. Miyake and P. Shah. Cambridge, Cambridge University Press: 340-374.
- Shallice, T. (1988). From neuropsychology to mental structure. Cambridge UK:, Cambridge University Press.
- Shapiro, K. L., K. M. Arnell, et al. (1997). "The Attentional Blink." Trends in Cognitive Sciences **1**(8): 291-296.
- Shapiro, K. L., J. I. Caldwell, et al. (1997). "Personal names and the attentional blink: The cocktail party revisited." Journal of Experimental Psychology: Human Perception and Performance **23**: 504-514.
- Shapiro, K. L., J. Driver, et al. (1997). "Priming from the attentional blink: A failure to extract visual tokens but not visual types." Psychological Science **8**: 95-100.
- Shapiro, K. L. and S. J. Luck (1999). The attentional blink: A front-end mechanism for fleeting memories. Fleeting memories, Cognition of Brief Visual Stimuli. V. Coltheart. Boston, Massachusetts, A Bradford Book, MIT Press.
- Teasdale, J. D. (1999). "Emotional processing, three modes of mind, and the prevention of relapse in depression." Behaviour Research and Therapy **37**: S53-S77.
- Teasdale, J. D. and P. J. Barnard (1993). Affect, Cognition and Change: re-modelling depressive thought. Hove, Lawrence Erlbaum Associates.
- Thompson, S. J. (1999). The Craft of Functional Programming. Harlow, Essex, UK, Addison Wesley Longman Ltd.

Tipper, S. P. (1985). "The Negative Priming Effect: Inhibitory Effects of Ignored Primes." Quarterly Journal of Experimental Psychology **37A**: 571-590.