Systems, Interactions and Macrotheory

PHILIP BARNARD[1], JON MAY[2], DAVID DUKE[3], and DAVID DUCE[4]

[1]MRC Cognition and Brain Sciences Unit; [2]University of Sheffield; [3] University of York; [4]CCLRC Rutherford Appleton Laboratory.

REVISION RE-SUBMITTED TO
TRANSACTIONS ON COMPUTER HUMAN INTERACTION
FEBRUARY 2000

Systems, Interactions and Macrotheory

PHILIP BARNARD[1], JON MAY[2], DAVID DUKE[3], and DAVID DUCE[4]

[1]MRC Cognition and Brain Sciences Unit; [2]University of Sheffield; [3] University of York; [4]CCLRC Rutherford Appleton Laboratory.

A significant proportion of early HCI research was guided by one very clear vision: that the existing theory base in psychology and cognitive science could be developed to yield engineering tools for use in the inter-disciplinary context of HCI design. While interface technologies and heuristic methods for behavioral evaluation have rapidly advanced in both capability and breadth of application, progress towards deeper theory has been modest. A case is presented for developing new forms of theory based around generic 'systems of interactors'. An overlapping, layered structure of macro- and microtheories could then serve not only their traditional explanatory and predictive roles, but also to bind together contributions from the different disciplines. It is also argued that novel routes to formalizing and applying such theories provide a host, not only of interesting, but also tractable problems for future basic research in HCI.

KEYWORDS

Categories and Subject Descriptors: H.1.2 [**Models and Principles**]: User/Machine Systems – *Human-Factors*; H.4.1 [**Models and Principles**]: System and Information Theory.

General Terms: Cognitive Models, Computing System Models, Models of Interaction.

1. Theory development in a boundless domain

In less than a quarter of a century information technologies and their users have diversified into an extraordinary range of socio-technical ecosystems. Scientific studies of HCI ran alongside each advance in application and interface design. A virtual relay-baton of R&D interest passed from line editors and programming applications, through command and WIMP interfaces for word processors and information retrieval systems, all the way to the current range of favorites. These now include: intelligent agents, awareness servers in CSCW, virtual environments, synthetic battlefields, mobile and wearable computers, and, of course, games, the world wide web, and embodied conversational characters capable of exhibiting emotion at the interface. Few would disagree with the proposition that the study of HCI is now effectively a boundless domain.

At the outset, many shared the vision of Card, Moran & Newell (1983) that step-by-step task analysis could be combined with theories of the human information processing mechanism and human knowledge representation. The product would be engineering methods to support design decision-making based upon sound theory subjected to empirical validation. As applications and interfaces diversified, the limitations of rather simple theoretical assumptions about a prototypical user's cognitive mechanism, and of the engineering methods they gave rise to, became all-too-readily apparent. Richer ways of thinking about users, tasks, systems, contexts of use, and design processes were needed.

On the user side, our tooling diversified into situated action, ethnomethodology, distributed cognition, activity theory and co-ordination science (e.g. see Suchman, 1986; Malone & Crowston, 1994). Theory-based evaluations of interface usability were often considered to be of limited value (e.g. Landauer, 1987) and most human factors work continued rely on, or improve, heuristic methods of empirical evaluation (e.g. Nielsen 1993). Design itself was also recognized as a complex process involving many trade-offs. These could beneficially be addressed and resolved by developing new methodologies such as scenario-based design (Carroll & Rosson, 1992), or design rationale (MacLean, Young, Bellotti & Moran, 1991) rather than extending theory-based evaluations.

As the scope of interaction design widened to include multimodal communication, multiple users, and worlds of interaction having no natural counterpart, there was little in the way of a well-developed body of theory ready for direct application. It was, of course, recognized from the outset that the development of theory lagged developments in interface design (Newell & Card, 1985). It was also widely acknowledged that our theories suffered from numerous other deficiencies. They were of restricted scope, applying to local features of interface design. It proved hard to re-use them in novel contexts or to scale up theoretical exercises from laboratory examples to real design settings. In spite of their many limitations, there is little doubt that our theories have evolved to address a broader range of issues. Some have even been better tailored to meet the everyday demands of practical application (e.g. Rudisill, Lewis, Polson & McKay, 1996). Nonetheless, as we move into the new millennium, theory has yet to mature to a point where it could be truly regarded as on a course that is likely to yield large scale benefits on a tangible time-scale.

At this juncture, we could abandon serious attempts to maintain theory development as a key element within the wider HCI enterprise. There is an obvious problem with this strategy. In the absence of a good body of formal theory, practitioners will undoubtedly invent their own informal, or folk theories, to help them represent and think about the problems and issues that are important to them in their context. The practice of HCI could become like that of psychoanalysis, with one school of thought communicating among themselves about a given set of issues in very different terms from those adhering to another school of thought.

Although there is universal agreement that HCI is inter-disciplinary, the vast bulk of theory has been developed along resolutely X-centric lines (System-centered; Application-centered; User-centered; Task-centered; Team-centered; Worksystem-centered; 'Virtual Reality'-centered, or whatever). Continuing with this strategy into the next millennium will inevitably lead to an increasing number of theories of different form and content. We would land up with a range of theories dealing with different facets of individual user performance, of the behavior of groups, and of larger organizations. The problem is that our theories would be unlikely to 'fit together' in a coherent way to resolve the conceptual jigsaw puzzles that exist in real design spaces (Bellotti, Blandford, Duke, MacLean, May & Nigay, 1996; Blandford & Duke, 1997). The vital connective tissue that binds together different topics within a level of analysis and which binds one level of analysis to another has, at best, been under-explored (Malone & Crowston, 1994). At worst, it has been fundamentally obscured by the co-existence of multiple systems of scientific semantics all being used within the same general problem space.

In this paper, our immediate aim is to stimulate further debate on alternative routes to theory development and its integration. Our discipline will not be best served in the new millennium either by abandoning theory or by the unconstrained development of more and more unconnected local theories at different levels and in different domains. After all, Human Computer Interaction theories are not about humans, or computers, but are about their interaction: they should also deal with the interactions between teams of humans, and networks of computers. They should deal with *interactions* between all these entities. They should directly address the problem of linking the different ways of modeling the properties and behaviors of these different entities. We argue here that greater integration within a boundless domain such as HCI is a tractable proposition for research in the new millennium. Integration can be facilitated through the development of generic representations of 'systems of interactors.'

2. Systems of Interactors, Macrotheory, Microtheory and Layered Explanation

Many of the general arguments for layered systems analysis are well known and have been widely discussed. In systems engineering, Chapanis (1996) uses hierarchical diagrams to illustrate how humans, hardware and software are grouped together in sub-subsystems that are embedded in subsystems that make up a complete system, be it a team or wider organization. Likewise, in cognitive science, Newell describes a system level as 'a collection of components that are linked together in some arrangement and that interact, thus producing behavior at that system level. In a system with multiple levels, the components at one level are realized by systems at the next level down and so on for each successive level (p117, 1990).

Newell organized systems of relevance to human behavior into a series of time bands running from those appropriate for the operation of neurons all the way up to evolutionary time bands. Within each band he also introduced layering by time. His cognitive band distinguished deliberate acts (100ms) from operations (1sec) and unit tasks (10sec). Newell's arguments about systems were focused upon human cognition. They were part of a wider argument that emphasized the need to move away from specific models of different phenomena, towards unified theories of cognition that could furnish those phenomena with a common explanation. Also operating within cognitive science, Marr (1982) stratified theories along a dimension of abstraction. He described a computational level of theory, which specifies the essential logic of what needs to be computed for a task to be carried out, and distinguished it from algorithmic and hardware levels of theory. The algorithmic level specifies the implementation of the computational theory while the hardware level refers to the realization of the algorithm in some physical system.

We draw upon components of all these ideas in a generalized form. Rather than taking users, computers, or teams as specific point of departure for theory development, we start by defining all of these entities as 'interactors'. The use of this term originates in computer science (Duke & Harrison, 1993, 1994, 1995a,b). Like the concept of an attractor in mathematics, the concept captures the idea of something that interacts with something else. This term has a number of advantages. First, by being generically, as opposed to specifically X-centric, it enables us to refer to things that interact without carrying the implicit semantic overheads that come with terms such as computers, users or teams. Second, an interactor is something that is composed of other interactors, and as such is a relative rather than absolute construct. Third, an interactor is something whose behavior can, in principle, be mathematically described in terms of the properties of the lower order interactors of which it is composed. Finally, any interactor is an entity that behaves over time.

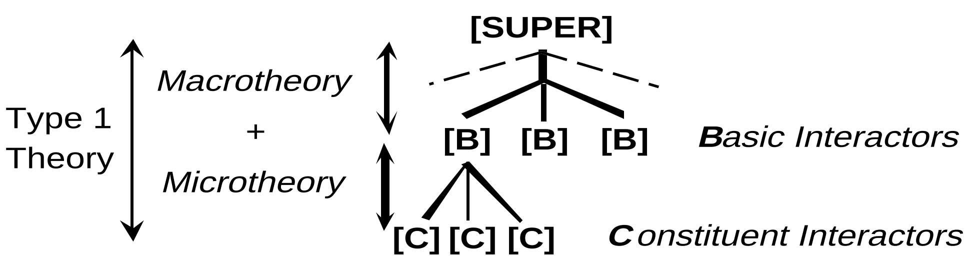

Figure 1. Macrotheory and Microtheory applied to systems of interactors.

To model the behavior of an interactor over time, we need to understand how its behavior is constrained. The behavior of any interactor will be determined in part by constraints originating in its own constituents and in part by the behavior of the other interactors with which it engages as a part of some superordinate system. This implies that a complete theory of the behavior of a system requires two distinct components; a body of microtheory and a body of macrotheory. Figure 1 illustrates a hypothetical system of basic interactors – [B]s.

These are the basic units that interact, and their behavior is constrained by their constituent [C]s, and by the superordinate organization of the system they make up [SUPER].

In the case of human computer interaction, a system of basic interactors might minimally be composed of a user, a computer and other things used in the task context, such a printed document. The behavior of the user (a [B]) will be constrained by the properties of components ([C]s) within their mental architecture (e.g. perceptual mechanisms, decision mechanisms and mechanisms for the control of action etc.) as well as by human biomechanics. Likewise, the behavior of the computer system (a different [B]) will be constrained by its components (I/O devices, processor properties, operating system and application characteristics, etc.). The 'behavior' of the document (a third [B]) would also be constrained by factors like the flexibility of its physical components and how they were bound together. For this system of interactors, at least three types of micro-theory are required: a model of the psychological system; a model of the computational system; and a model of a physical system. However as components of a system these three are not independent – the behavior of each depends upon the others. The relevant macrotheory for this system of interactors would specify how their conjoint behavior is constrained. Only the combination of micro- and macrotheory that would provide a complete theory of this particular system of interactors. A coherent theory made up of an inter-related body of micro- and macrotheory will, for the purposes of a later contrast, be referred to as a 'Type 1' theory.

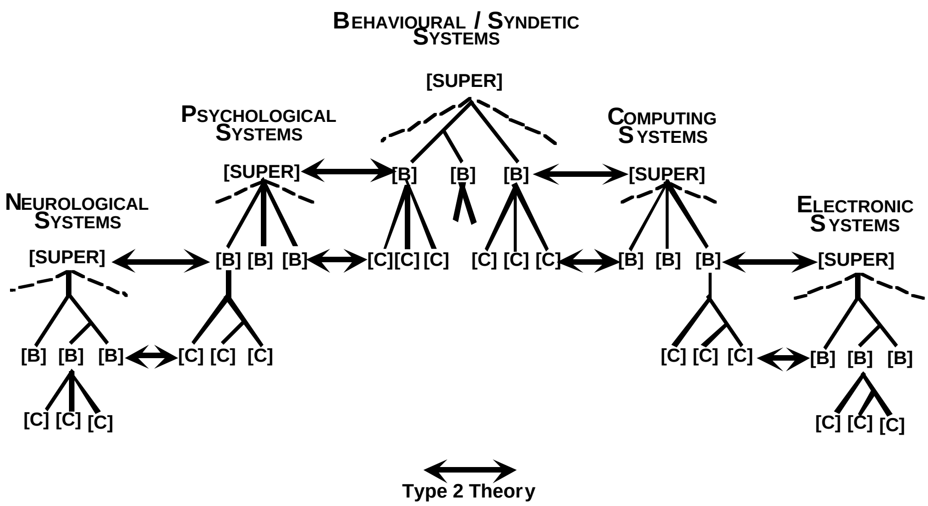

Figure 2. Systems of Interactors at different levels of explanation organized by Type 1 and Type 2 theories.

For those interested in modeling the operation of a psychological system, or of a computing system, the workings of their components require more detailed specification. At the apex of Figure 2 is a system composed of a user, a computer and some other interactor, which might again be documentation. The topmost level is labeled a behavioral/syndetic system. The term syndetic (Duke, Barnard, Duce & May, 1998), is derived from the Greek term *syndesis* meaning to bind together. We use the term syndetic to refer to the specific case of behavioral systems that are composed of interactors of fundamentally different types. To the sides of the apex are shown a psychological system of interactors and a computing system of interactors, and to the side of these are shown their respective refinements into neurological and electronic systems. Those with an interest in the functioning of organizations are unlikely to be interested in such minutiae. They might wish to see an extension of this form of diagram 'upwards' rather than 'downwards,' with higher order systems in which the basic interactors are teams, whose constituents are the type of system at the apex of Figure 2.

Systems are distinguished in terms of the focus of scientific attention. Each system in Figure 2 consists of entities that behave. In the neurological system the things that behave are neurons or glands that release hormones. In the psychological system the things that behave are processes that construct or change mental representations. In a behavioral/syndetic system, the things that behave are humans and technological artifacts. Unlike a strict hierarchical decomposition of successive systems, Figure 2 overlaps hierarchies at different levels. This enables us to highlight two characteristics. First, when we focus our attention on the behavior of a system, we adopt a frame of reference appropriate to the entities that make it up. We must consider both the organization of the superordinating system and the subordinate constituents of the entities. A complete Type 1 theory is composed of macrotheory and microtheory. Newell's arguments were applied to the unification of theories *within* a cognitive layer. We take this form of argument and generalize it to *any* system level. Second, when we adopt a frame of reference for model building or theory development, we do so in terms of the scientific semantics of a particular class of theory. When we move the focus of scientific attention from one level of system to another we use theories with different form and content, as identified by Marr (1982) and others. This is shown in Figure 2 by introducing the notion of a Type 2 theory, which is the mapping from the macrotheory of one level of explanation into the microtheory of another and vice versa.

A Type 2 theory specifies the transformations that occur in the semantics of theories when we move from one layer of systems analysis to another. Figure 2 suggests that these transformations have two components: a mapping from the superordinate composition of one system into the basic units of the higher system; and a mapping from the basic units of the lower system into the constituents of the higher system. When we move up a level we discard the microtheory of the lower level, and when moving down we add microtheory to support more detailed explanation or implementation. This is marked in Figure 2 by the offsets to the right and left of the central behavioral/syndetic system. In moving to either a human or computer system, a basic unit of the behavioral/syndetic system becomes the superordinate entity of our new theoretical domain. Its constituent interactors become the basic interactors of our new theory, and we need a new microtheory to add in the new constituent structures, which were not specifically represented in the theoretical analysis of the behavioral/syndetic system.

The transformations that occur are not just those of the adding or discarding detail. In the regions where macrotheory at one level overlaps with microtheory at a higher level, the two representations differ. One form of transformation is selection. A particular theoretical project may be concerned with abstracting or refining only those components that are relevant to a specific modeling objective. A second is recombination. The syndetic system incorporates basic interactors of different types and must draw on a microtheories of each. Other forms of transformation are qualitative. Macrotheory at one level is not the same as microtheory in an overlapped layer. As Marr (1982) noted, there may be concepts in one level of explanation that have no direct realization in an adjacent level. A key contrast at one level may, for example, be an emergent property of a lower order system. Alternatively, a given level of system may be based upon the same lower order foundations but use the higher order constituents in different ways. In Cantonese, pitch contour may completely change the meaning of a word. Since pitch contour and stress interrelate, the communicative functions of these resources are restricted and another resource must be used to fulfil some of the functions of pitch contour in English and French (e.g. see Barnard & Marcel, 1984). This makes little difference to the macrotheoretic architecture of eastern and western brains. However, it would need to be accommodated in alternative microtheories of language comprehension at a psychological level of explanation. Similarly, in HCI, there can be many alternative ways of combining user and computer capabilities to realize interactions that have particular behavioral properties. These need to be accommodated in modeling higher order systems.

In the years since Newell & Card's (1985) discussion of the problems faced by theory development in HCI, there can be little doubt that significant advances have been made in the development of unified cognitive architectures and placing modeling on firm computational foundations. SOAR (Newell, 1990), ACT-R (Anderson, 1993) and EPIC (e.g. Meyer & Kieras, 1997) have all achieved significant application successes in this domain. However, the progress has been uneven. The advances have been concentrated on particular attributes of tasks. These reflect the purchase that AI architectures, inspired by Newell & Simon (1972), provide on the acquisition and use of knowledge in task execution. Progress has been rather more modest on other topics such as modeling human understanding and use of dynamic graphical interfaces, multimodal perception, and emotion. Our ability to model these aspects of a user's psychological system continues to lag the leading edge of interface design by a substantial margin. In these areas, our existing microtheories require considerable development and forms of macrotheory need to be specified into which they too can be more readily integrated. Only then might we expect a mature Type 1 theory of our complete psychological system to emerge. Existing attempts at unified theories of cognition remain only partial macrotheories of the complete psychological system.

The gulfs also remain wide between different levels of analysis. We are not very good at establishing coherent theoretical connections of Type 2. Someone whose primary interest lies in the overall design of a human-computer work system is unlikely to be interested in the fine grain details of the limitations on human spatial working memory or in the limitations of graphics algorithms used to render a particular image. They wish to know only about how such limitations effect the performance of the work system itself and how they tradeoff with other constraints in their design space (Norman, 1983). When they seek answers from other disciplines they will be frustrated. They will find a range of competing psychological theories. Each of these will be formulated in a different terminology from each other and address different ranges of issues and outcomes. For the purpose of theoretical advance the theorists may even have drawn attention to areas where their model makes contrasting predictions from the competition, rather than highlighting their areas of agreement. All of this can make it extremely difficult to deliver an interdisciplinary synthesis at one level that is based upon principled reasoning grounded in other levels (Bellotti, et al 1996; Blandford & Duke, 1997).

Within the overlapped hierarchy of Figure 2 macrotheory provides the connective tissue that binds together those microtheories of the entities that make up the system of interactors. It also provides the key level of abstraction that should enable us to carry over relevant knowledge, in a systematic way, from the science base of one level of system analysis to the next level up. Our current theory base may well be getting better at modeling and predicting the behavior of humans and computers in specific task contexts. However, it will remain of limited utility until and unless we develop true macrotheories that can meet the challenge of providing the connective tissue for both Type 1 and Type 2 theory.

As with principal components analysis in statistics, vector analysis of physical forces, or Fourier transforms of complex waveforms, the problem of developing better theory needs to be broken down into clearly defined parts. As a boundless domain, HCI needs Type 1 microtheories of interactors and macrotheories of their interaction. It needs such theories at different levels of abstraction that extend from the fundamentals of the behavior of users and computers all the way up to that the co-ordination of people and technologies in large scale organizations (Malone & Crowston, 1994). If HCI as a whole is to maintain some overall unity and coherence, it will also have to nurture Type 2 theories. They are needed to support effective communication between those whose focus of attention is at different levels. Their development is vital to enable knowledge in all the relevant disciplines to be brought to bear systematically on the solution of design problems involving the use of computers by individual users, by groups, and in organizations.

3. Macrotheory and Interaction

Our fundamental conjecture is that macrotheories at all levels of system decomposition can be represented within a general modeling framework. The objective of macrotheory is to capture the interdependencies between interactors at any level of system decomposition. It is intended to provide the scaffolding both for the elaboration of microtheory at that same level and to support moves from one level of systems analysis to another. The key claim is that the behavior of any system of interactors will be determined by four distinct classes of constraint:

System behavior = Fn (*Configuration* of Interactors;
The interactors' individual *Capabilities*;
The *Requirements* that must be met to use those capabilities; and
The regime of *Dynamic Control and Co-ordination* of the interactors)

This four component framework was first introduced as a basis for developing an explicitly formulated body of macrotheory concerning the behavior of the human mental architecture (Barnard, 1987). Here we represent that framework in a form that can be generalized to all systems of interactors.

The *Configuration* defines the identity of the basic interactors that make up a system of concern and specifies their potential for engagement with each other. Their engagement might have physical or informational properties. For example, a system of three interactors might be configured so that they can all communicate with each other directly, or the channels might be more constrained with Interactor 1 being able to communicate with Interactor 3 only indirectly via Interactor 2.

The *Capability* of an interactor is defined as the transformations in information states or of physical state that it can accomplish. As basic interactors within a cognitive architecture, mental processes can be defined as having the generic capability to change the form or content of a mental representation. The capability of individual mental processes would then have a more detailed specification, such as the capability to transform a phonological representation of the 'surface structure' of an utterance into a representation of its underlying propositional meaning. The generic capability of an interactor composed of a human and a technological device might be that of document preparation, with a repertoire of more specific capabilities for the human and the software.

The *Requirements* that must be met for an interactor to realize a specific capability are essentially the states that they need to function – be they physical or information states. The mental process for language understanding may require a clear incoming phonological representation in a language it has learned.

Systems behave over time and the fourth component of the framework, the regime of *Dynamic Control and Co-ordination*, is intended to summarize properties of system activity on a temporal dimension. If we take a time slice of activity within a system, there will be some dynamic properties that characterize the overall state of the interaction. So, for example, a system may be in a state where the pattern of information or physical exchange among interactors is stable over time. With an engaging novel in a comfortable chair, the reader may well be entirely absorbed in this single activity for significant periods. The relevant mental processes may be maintained in essentially the same configuration. When reading the same novel in a railway station while waiting for a train, the activity of reading might well be punctuated by frequent rapid glances up to a train indicator board. Such glances involve only a momentary reconfiguration in the pattern of engagement among interactors within the psychological system and between these and interactors in the physical setting. Were the reader to suspend reading in order to use their mobile phone to call home, then the reconfiguration would last for a more extended period. The first form of reconfiguration might best be described as the co-ordination of a dominant and subsidiary activity through transitory oscillations. The second form of reconfiguration might better described as a more extended interleaving of two activities.

The ways in which activities are synchronized and controlled are also included in this fourth class of constraint on system operation. In some systems, wider control may be an emergent property of synchronous exchanges between interactors. In other systems, particularly military or managerial ones, some interactors have the explicit capability to direct or control the activities of others. At a macrotheoretic level, we still need to capture any states of activity where the effective *locus of control* lies within a set of interactors, and how the pattern of control changes over time.

For example, it is well known that most people can drive a car and hold a conversation at the same time. The moment-to-moment dynamic control of the mental processes required to drive the car may reside mostly within a peripheral perceptuo-motor configuration, while central thought processes are primarily engaged in controlling an independent auditory-verbal configuration. All that might be required under these circumstances is the occasional oscillation in which central processes are momentarily re-deployed to monitor some aspect of the driving task. If a small child were to move from the sidewalk onto the road, then the complete mental system might reconfigure. The peripheral configuration might now be brought directly under the control of central processes. The central control of driving would be interleaved and conversation consequently cease (Barnard, 1999).

The same description can be applied to the workings of a team. In the context of an open-plan call center a team might be taking product orders over the phone while entering them into a computer. At this level individuals, computers and phones are the basic interactors. The supervisor might be working on a primary task, periodically glancing at activity in the center. Were they to notice a problem they might interrupt their current task to go and interleave an activity of helping someone. In this interaction, the supervisor is the locus of control for the team. On a fine grained time scale the activity of the supervisor's psychological architecture, a constituent interactor at this level of system analysis, would be described as oscillating between a primary and secondary task during the first segment of activity. In the second segment the supervisor interleaves an activity. When viewed from the coarse grain of workflow over the course of a week, an interleaved activity at the lower level may be re-represented as an oscillation in the configuration of the workforce.

In the driving example, the analysis focuses on what specific mental processes are doing. In the second example the focus is on what individuals in a team are doing. The four component framework enables us to explore the possibility of macrotheoretical principles that relate configurations, capability and requirements to the behavior of a system over time. For example, the pattern of dynamic control and co-ordination of a system of interactors may alter in a principled manner when particular attributes of configurations, capability or requirements apply. The complexity of dynamic control and co-ordination may rise when capabilities are sub-optimal, when requirements are not met, or when a configuration is depleted. At the level of a psychological system, a novice driver may have few of the skills required to co-ordinate perception and motor aspects of driving without thinking about the driving task. As they become more experienced, the proportion of fully automated skills would rise and with it the proportion of time that central processes could be configured to sustain uninterrupted conversations with a passenger. If the call center team were composed of a high proportion of individuals with less than optimal capability, then workflow might involve a high rate of oscillations in the configuration of the team.

The idea that the behavior of systems can be captured in terms of systematic relationships among four generic classes of constraint does not mean that all systems are governed in the exactly same way. Systems that are configured in different ways and whose activities are subject to different regimes of dynamic control and co-ordination would be expected to exhibit different behaviors. The four component framework is intended to provide a basis for developing macrotheories that capture both the similarities and the differences in behavior of different systems of interactors. The content of the macrotheory is not absolute. It must be bound to relevant microtheories of the individual interactors of which it is composed.

The behavior of any system of interactors evolves over time, and that behavior can usefully be thought of as a trajectory through a set of possible states. The behavior trajectory of a system of interactors, be they cognitive, computational, syndetic, or organizational can itself be decomposed. Figure 3 depicts a trajectory of continuous interaction divided into segments. These each approximate a state of activity among the interactors. There is a transition from segment to another when there is *significant* change in configuration, capabilities, requirements or the pattern of dynamic control and co-ordination of the basic interactors within the system. One segment captures the properties of system behavior in the very short term (VST). A phase of activity is sequence of related short term (ST) transitions among related

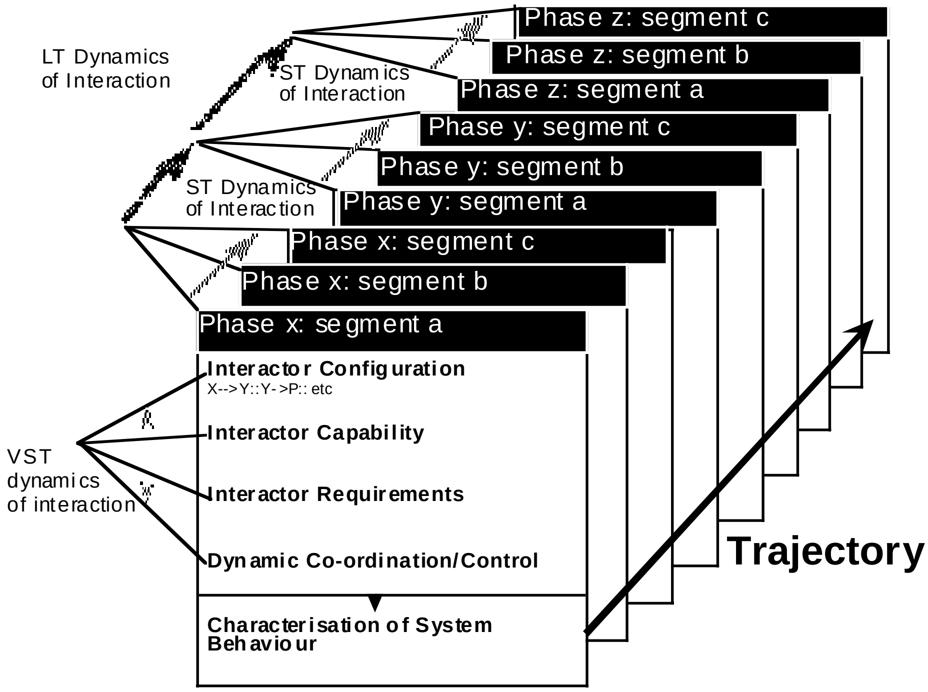

Figure 3. *An outline characterization of a behavior trajectory subject to four 'generic' classes of constraint (Modified for the current perspective from Barnard, Wilson & MacLean, 1987).*

segments. A transition from one phase to another would typically be associated with longer term (LT) changes in the properties of systems. A number of distinct phases may contribute to a trajectory. Returning to Newell's (1990) time bands for human action, in the modeling of a psychological system of interactors 'segments' can be represented by very short term activities lasting a few hundred milliseconds. A 'phase' might be represented by periods of up to 10s (Barnard & May, 1999). Changes in psychological capability brought about by knowledge acquisition and learning would naturally encompass far longer time scales. Exactly what falls within the scope of a segment, phase and trajectory as well as the time scales for very short term, short term and long term transitions would be bound to the definition of basic interactors, and sensitive to the activities modeled.

Thinking about a system's behavior as a trajectory governed by systematically structured sets of constraints is quite different from the more usual forms of step-by-step task analysis conducted in HCI and traditional human factors research. As with simulations of artificial life, the way in which system behavior evolves over time is an interaction of constraints. Each segment or phase of interaction has a point of departure and an outcome. The outcome of one phase is the point of departure for the next, and each will be defined in terms of the attributes of configurations, capabilities, requirements and dynamic control and co-ordination that apply at that point. All of these are variable. So, for example, the capabilities of interactors can change – as they do when an individual or an organization learns. Similarly, new interactors may be introduced, requirements may change and the regime of dynamic control and co-ordination for a given system may change – for example when a business undergoes re-organization or when politicians change the rules of engagement for their military forces.

The whole trajectory can be thought of as analogous to a sentence in natural language, with the segments being analogous to words and the phases analogous to clauses. Just as rules apply to sentences formation in a language, so a class of interaction is governed by a collection of high level rules. Just as language enables an infinite set of sentences to arise from its vocabulary and rules of combination, so there can be an infinite set of behavior trajectories for some systems. Just as different grammars and vocabularies apply to different languages, so different rules and segmental analyses might be applied to different classes of interaction.

4. Capturing Significant Variation in Interaction Trajectories

Thus far, we have introduced the general idea that the problem of theory development for an area as boundless as HCI can be stated as two requirements. One requirement is to develop Type 1 theories of system behavior, composed of a unified body of macrotheory and microtheory. The other requirement was to develop Type 2 theories that map from one level of analysis to another. Type 2 theories also have two components, a mapping from the superordinate organization of one system to the basic units of another, and a mapping from basic units at one level into the constituents of another. We have marked out the development of macrotheories of system behavior as a particularly important and underdeveloped area. We then went on to lay out a specific framework for representing macrotheory. Up to this point, the arguments have necessarily been general. In this section we provide concrete illustrations of what we mean by a behavior trajectory for a system of interactors.

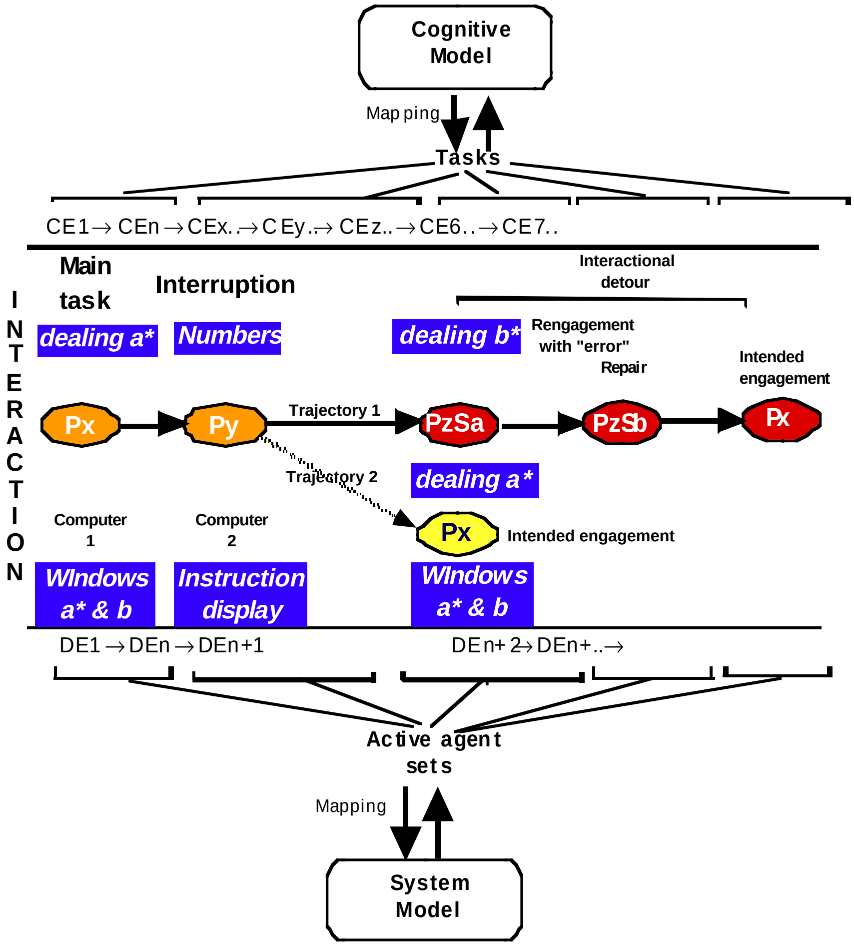

Figure 4. *Two trajectories for interactions with a 'selected' and 'unselected window' (Modified from Barnard & Harrison, 1992).*

When interactions are inefficient or 'go wrong', traditional forms of analysis try to eliminate specific causes of error by re-designing the system, re-designing the task, by re-training users, or by changing the allocation of function between users and technologies. An analysis based around systems of interactors, and the trajectories of their conjoint behavior can help us to think about what is going on in new ways. Couched in these terms, errors represent a special case of what can be referred to as detours in an interaction trajectory (Blandford, Harrison & Barnard, 1995). Once a computer user makes an error, they typically have to take a number of additional steps to recover.

A well-known example of this is the unselected window scenario. A computer user, who is interleaving activities conducted in different windows, may start typing only to find that the products appear in a window they are not looking at, or that one window changes from being inactive to active. Trajectories for this situation are shown in Figure 4. This is a syndetic system composed of a user and two computers. It is assumed that the microtheory for the trajectory of a user's mental activity be captured in a cognitive model. This model constrains the sequence of mental events that cause outcomes in segments of mental activity. These are represented by the linear sequence CE1…CEn linked to the cognitive model. Exactly the same description applies to the model of the devices, in this case computers. These are represented in the lower part of the figure by the sequence DE1…DEn etc. Alternative trajectories for the interaction are shown in the center as a series of phases (Px, Py, Pz), one of which is decomposed further into segments (PzSa, PzSb). Each of these represents not a user state or a device state but a state of interaction, or engagement between the interactors within the system.

The trajectory shows a case where a user is conducting a main task on one computer, in this case share dealing, and is periodically interrupted by requirements to deal with a separate task on a different computer (in this case a numbers task). The share dealing task has two windows and requires frequent movement from one to another. The figure shows a phase of interaction (Px) on the main task followed by an interruption (Py). When returning to the dealing task, interaction may be re-engaged with the active window on trajectory 2, or it may be re-engaged with an inactive window (PzSa) requiring a segment of trajectory for repair (PzSb). The trajectory description is not a combination 'user does this, system does that.' It is a representation of a state of engagement between the two. An attribute of dynamic control and co-ordination would mark the fact that in PzSa that the activity of the two interactors within their engagement was not coordinated.

In many other types of interaction it is clear that people do not just stop one activity and start another. There tend to be transitions that are explicitly marked as distinct phases of interaction. The most obvious cases occur in human conversation. We do not simply start up a conversation with someone, we go through an orientation phase of saying hello, exchanging pleasantries, and only then get on with the real task once common ground is established. A similar transitional activity occurs as conversations are closed down and the various phases appear to be clearly governed by conventions or principles (Clark, 1996).

From this generic principle, we reasoned that the situation with unselected windows could be improved, not by redesigning the way in which an active window was marked, but by re-designing the interaction trajectories to introduce two kinds of transitional phases. These are illustrated in Figure 5. One transitional phase is labeled 'Possible Disengagement' and the other 'Transitional Resumption' (Barnard & Harrison, 1992).

In most forms of interaction of this type, the user generally holds the initiative (Blandford, Harrison & Barnard, 1995). They are the locus of control within the system of interactors. When the user stops doing anything, the conjoint state of the interaction is unclear – the user may be looking at something on the VDU, they may be interleaving an alternative activity, or they might have gone to lunch. In this case, we designed a system that responded to potential disengagement by taking the initiative. After a period of zero input, the computer changed a property of the currently active window – the window border gradually started to 'fizz,' pixels in the border went off and on. If, at any point in the gentle increase in this attribute, the user did anything with the mouse or keyboard the window border returned immediately to its passive state. If the user continued to do nothing the border reached the maximum extent of its 'fizzing' capability (steady state disengagement). At any point when the user re-engages with this system (transitional resumption), the properties of the active window attract the user's attention. As soon as the user carries out an action, the border returns to its more passive indication that it is in an active state. Work with an experimental system demonstrated that it led to substantially fewer unselected window errors than occurred with systems that did not mark the transitions in this way (Lee, 1992).

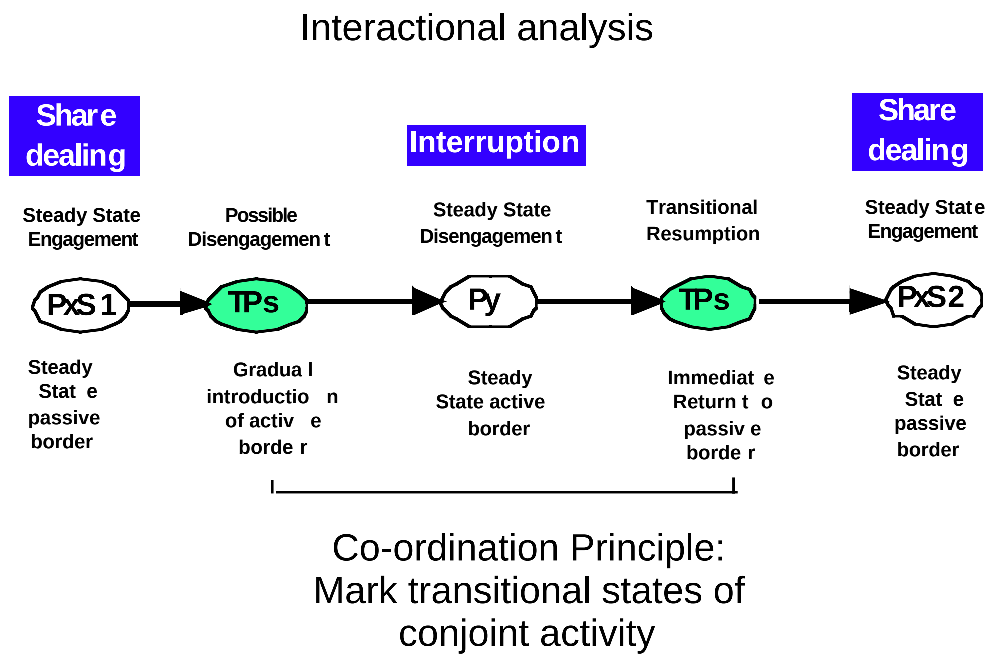

Figure 5. A trajectory designed to include overt marking of transitional phases.

What is interesting about Lee's experimental results is not the reduction in errors, but that design changed the overall pattern of behavior across the whole interaction trajectory. Although there was a small reduction in the occurrence of unselected window errors when the user returned to the main task following an interruption, the greatest reduction occurred during continuous segments of engagement on the main task. When working on the main task such errors occurred in the natural course of the user moving from one window to another. During such phases conditions virtually never arose where the user was inactive for long enough for the border to become 'fizzy' – this was designed to happen only when the subsidiary task was interleaved.

The ultimate explanation of this result lies in a more detailed microtheory of user cognition. Users presumably developed a generic schema for interaction which led to an evaluation of border state on each move among windows, whatever the exact context. However, the general pattern is best understood by reference to more abstract macrotheoretic principle concerning interaction in the higher order system. The dynamic control and co-ordination of the interaction depends on a number of inter-related attributes that can be represented in the four component framework. These link a principle to mark transitional phases in a trajectory with system and user capabilities as well as their coordination:

- A property of system capability - dynamic and passive attributes on window border.
- A property of user capability - the schema for monitoring window state
- A property of dynamic control and coordination - where the system acts as the locus of control in one phase and the user in others.

A change to any of these properties may change the likelihood of detours occurring. Indeed, the same reduction in errors does not occur when active window 'fizzes' all the time and is not sensitive to the transitional phases (Lee, 1992). What is new here is that thinking about conjoint states of interaction, and the abstract attributes that underlie them, quite naturally led to the idea of a solution. That solution would not have been obvious by considering a task, a user model and a system model in isolation. It required a model of the interaction and its trajectory.

We have discussed a number of general properties of interaction trajectories elsewhere in the literature (Barnard & Harrison, 1989, 1992; Harrison & Barnard, 1993; Blandford, Harrison & Barnard, 1995). An important aspect of the approach is to see if properties of trajectories can help us to understand how different phenomena are related. Detours such as that in the unselected window scenario occur in many contexts. A classic series of studies by Carroll and his colleagues (e.g. Carroll & Carrithers, 1984) showed that the availability of large sets of functionality provided an environment in which novice users could go down all sorts of confusing and inessential paths. A 'training wheels' system was designed that explicitly blocked access to a significant range of capability. This was shown to help users learn more effectively. The concept of potential captures the idea that some systems, like use of automatic teller machines, have a very small state space for interaction, while others such as UNIX, have an enormous state space (Barnard & Harrison, 1989). The acquisition of user capability may be best served by designing interaction trajectories of limited potential. At the other extreme, advanced users may benefit from the most direct and immediate route to that same functionality, however infrequently used.

This form of analysis does not represent or distinguish concrete properties of good or bad trajectories. It is necessary to understand the range of trajectories that may be possible, and their relationship to the design context. Complex trajectories may be bad where the requirements include a concern for efficiency and speed. However, both theorists and designers may have other concerns. In the context of safety critical systems, more complex and involved trajectories may be warranted. In the context of computer games, trajectories might be designed to achieve a positively motivating balance between success and failure, as well as challenging skill development. It may be appropriate to guide design by including as a requirement something like: 'for a range of users with varying capabilities (skill levels) the normative segments of interaction trajectory for play should mostly lie within a range of 5-12 exchanges' (Barnard & Harrison, 1992). Below this range, interactions may be frustrating. Above it, the interactions may be boring.

Just as there can be no absolute gold standard for the complexity of trajectory properties, so it is unlikely that there will be any simple recipes for deciding when it is appropriate to add functionality to a system. Successive periods of technology development have often led researchers to ask direct empirical questions about the consequences of adding functionality. Adding video-channels to communication links may deliver benefits only in specific circumstances (Veinott, Olson, Olson & Fu, 1997). The addition to an interface of faces that show emotion, or agents with a human-like persona, can have subtle influences on properties of an interaction, perhaps affecting user satisfaction more than traditional measures of the efficiency of task performance (Walker, Sproull & Subramani, 1994; van Mulken, André & Müller, 1998). As with the unselected window problem, empirical variation in these kinds of setting is unlikely to yield up all of its secrets to step-by-step forms of task analysis. We suggest that modeling the abstract properties of trajectories will be a productive basis for understanding the significant variation evident in the results from user testing in HCI.

Our discussion of the unselected window scenario focused on the properties of behavior trajectories for a syndetic system involving one user and two computers. For the purposes of illustration it assumed an unspecified cognitive model of the user and an equally unspecified model of the computer system. The proposed schema for theory development requires that the analysis of such scenarios can be mapped, via Type 2 theory into well-specified macromodels of the user's psychological systems and into macromodels of the behavior of the computer systems.

5. Towards Macrotheory for Psychological Systems of Interactors

In his argument for integration, Newell (1990) called for the development of unified theories of cognition. As noted in section 2, much of the subsequent work that followed was based in an AI tradition of simulation of a range of phenomena within a single mental architecture. Our approach to integration has followed a different course. Rather than simulating user cognition,

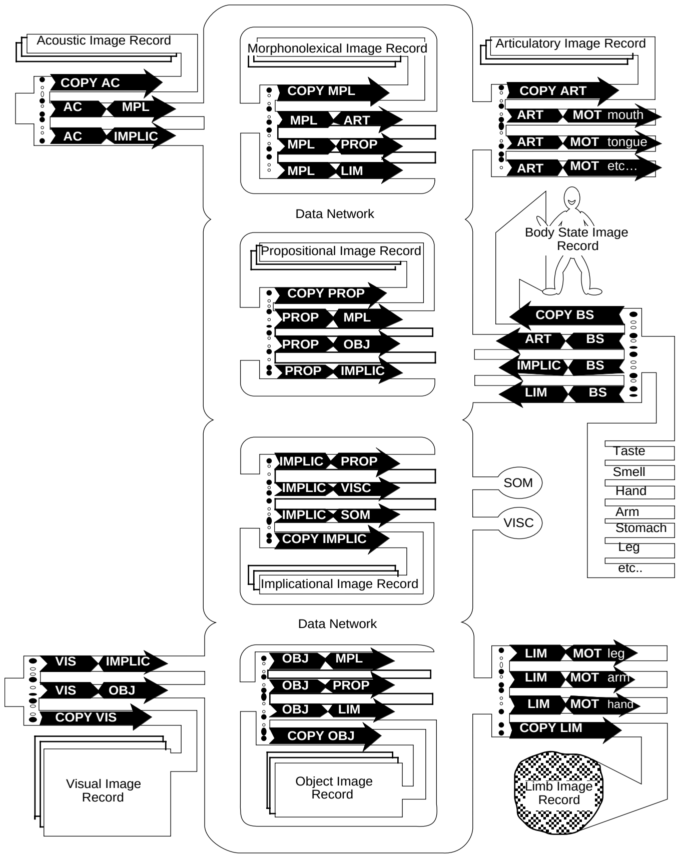

Figure 6. Interacting Cognitive Subsystems (after Teasdale & Barnard, 1993).

we specify a mental architecture and its principles of operation, and then use rules or mathematics to infer properties of its behavior across a range of conditions. This approach is like the modeling of economic systems, where a set of equations are used to infer what is likely to happen across an economy were taxes to be reduced or increased.

The justification of the particular architecture we use, and its approach to modeling, is beyond the scope of the present paper. The architecture was originally developed to integrate accounts of multimodal aspects of human short-term memory with those of language processing, attention and central executive functioning (e.g. Barnard, 1985; 1999). The architecture has been applied to the decomposition of HCI tasks (Barnard, 1987), to multimodal aspects of HCI performance (Barnard & May, 1995), to the understanding of dynamic graphical displays (May & Barnard, 1995b). It has also been applied to the effects of emotion on human cognition and to clinical disorders such as depression (Teasdale & Barnard, 1993).

The basic theory starts from a non-controversial position. It assumes that our mental architecture is a distributed system in which processes, specialized to handle different facets of mental life, operate concurrently. It specifically assumes that our mental architecture is composed of nine subsystems. More controversially the theory assumes that the operation of all these subsystems is governed by a common set of underlying principles. Equally controversial is the assumption that all types of mental representation have common principles of construction – differing only in the form that information is actually encoded.

All of these subsystems are composed of processes. These are the basic interactors within the wider psychological system. The basic interactors are of three types – processes which transform information from one form of mental representation to another; processes which construct representations over time; and a record process which has the capability to regenerate

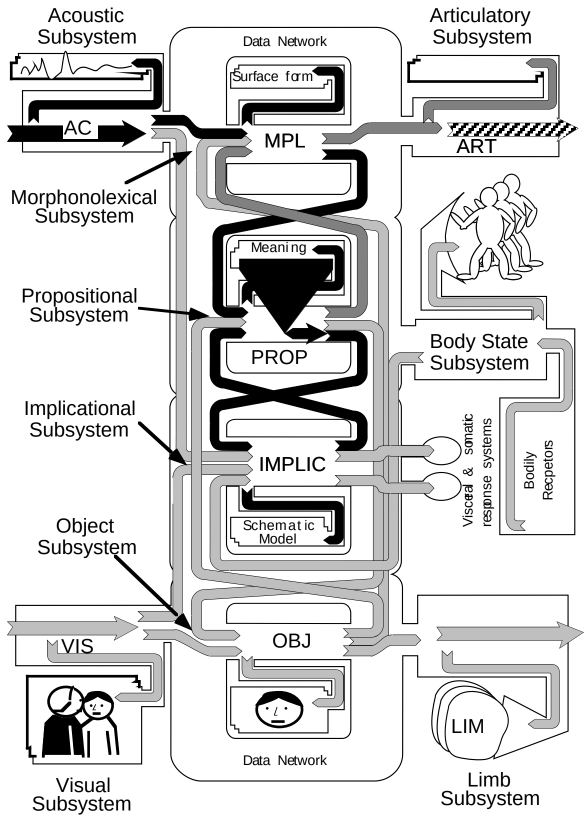

Figure 7. Some of the information flows afforded by the processes shown in Figure 9.

representations of past inputs – a memory system in more conventional terms. Subsystems interact when a process in one subsystem generates an output for use by another subsystem, be it based upon information flowing into a subsystem in real time, or be it based on information regenerated by the record process. Within this schema, mental activity (the behavior of this system of interactors) is composed of interactions between subsystems, hence this theory is called Interacting Cognitive Subsystems, or ICS (Barnard, 1985).

The architecture is shown in Figure 6. Three sensory subsystems deal with information derived either from distal sources (ACoustic and VISual shown on the left) or from Bodily States (BS – shown center right). Two subsystems handle the co-ordination of actions in the world through skeletal movement (the LIMb subsystem) and verbal communication (the ARTiculatory Subsystem). Four subsystems handle higher order abstractions of information. They are (a) the MorphPhonoLexical subsystem, which represents auditory verbal abstractions; (b) the OBJect subsystem (visuo-spatial representation); (c) the PROPositional subsystem which encode semantic representations in a form that is referential and relationally specific; and (d) the IMPLICational subsystem which represents more generic semantic relationships. This subsystem encodes schematic models of the broader existential state of the complete system, in a body, in a wider sensory environment. The content of these models encompasses higher order regularities derived *over* the products of processing within propositional, acoustic, visual and body state subsystems. It is at this level that emotion is experienced and this subsystem handles the wider existential regularities that constrain the content of thought as realized in Propositions, as well as SOMatic and VISceral responses.

The use of upper case marks the abbreviated forms of reference to the different forms of mental representation. These abbreviations index the processes shown in Figure 6. For example, the Acoustic subsystem (upper left) contains a number of processes. One (COPY) transfers an incoming auditory waveform to the memory process (the Acoustic Image Record). A second process (AC→MPL) transforms the incoming speech waveform into a higher order representation of its content in a different code. A third process (AC→IMPLIC) maps other attributes of the acoustic waveform, such as 'tone of voice,' directly into more abstract semantic code. The three processes operate concurrently as indicated by their parallel arrangement within the diagram.

The processes shown in Figure 6 specify the inter-connectivity of the overall mental architecture. They capture information flow between the interactors (Figure 7) and this can be used to define the configurations that play a significant role in different tasks. Some processes can talk to each other directly, others can only do so via an intermediary process; sensory processes can only send to the central ones, effector subsystems typically receive flows from the central ones, and the central subsystem can exchange information one to another in well-defined patterns. The capability of the individual interactors is constrained by a small set of general information processing principles. For example, it is assumed that a given process can only transform a single coherent stream of data at any one point in time.

The patterns of information flow are intimately connected to the nature of the different types of mental representation. The transformations from one mental code to another exactly mirrors the Type 2 theories outlined for scientific abstraction and refinement (Figure 2). By transforming information, mental process are mapping from one system of representation to another. In the transformation from a sensory subsystem to central ones, detailed information about sensory properties is discarded, higher order basic units of representation are formed, and the added value of the higher order mental representation is the new superordinate organization it delivers (Figure 8). In recoding information received from a central subsystem, an effector subsystem would accomplish the opposite of the transformation indicated in the figure. For example, the ARTiculatory representation would be used to compute and co-ordinate the motor instructions for the various musculatures controlling lips, tongue, mouth and breath during speech output.

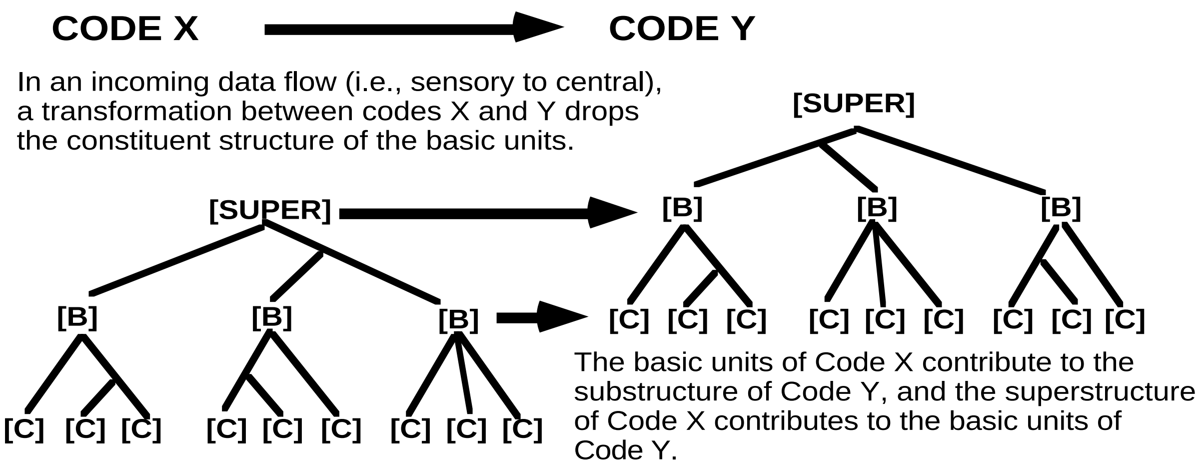

Figure 8. The representational shift accomplished by a process transforming incoming sensory representations (from May, Scott & Barnard, 1995).

The general argument implies that systems of mental representations abstract over the variance in the information received by a subsystem and that, through learning, the processes that transform mental representations must embody regularities in input that are directly mapped onto consequent regularities in output. The set of input-output relationships is the computational function of that process, and represents its capability. Where information is received from different types of sources (e.g. sensory and semantic), the associated system of mental representation will, as with our syndetic theory at the apex of Figure 2, model the inter-dependencies over these inputs. This is the basis for the theoretical treatment of multimodal integration (Barnard & May, 1995).

A relatively small set of processing principles (Barnard, 1985) is assumed to constrain the transformation of mental representations (their capability). All nine forms of mental representation are hierarchically organized with a superordinate structure, basic units, and constituents. They are of similar form but different content. These representations are the input to subsystems over time. They are also the structures that are preserved in their image records. This enables us to formulate equally general principles for requirements at input that must be met to use capability. This renders the problem of developing macrotheory for psychological systems tractable, not only in terms of its verbal formulation and heuristic rules, but also in terms of its formal expression in mathematics, as we seek to show later.

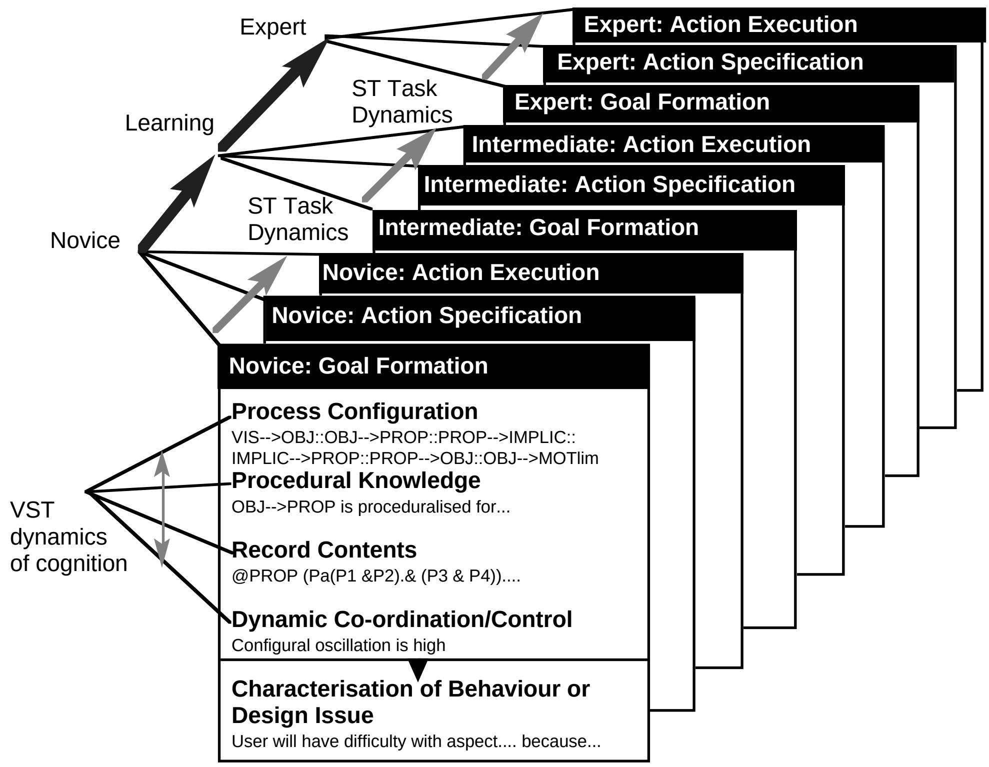

Figure 9 A family of cognitive task models (From Barnard & May, 1993)

In a series of theoretical exercises, first reported at CHI'87, (Barnard Wilson & MacLean, (1987), we have sought to model the behavior of the ICS architecture for a range of problems in HCI in a four component framework (Figure 9), generalized in this paper to all systems of interactors (Figure 3). In this approach, we use the theory to define detailed attribute spaces for configurations, attributes of procedural knowledge (Capability), record content (Requirements) and attributes of the dynamic control and co-ordination of wider processing activity. We approximate over generic phases of mental activity (goal formation, action specification, and execution) and for different parts of a learning trajectory (novice intermediate and expert). The foundations and detail of actual techniques are reported elsewhere (e.g. Barnard, 1987, 1991; Barnard & May, 1993, 1995; May & Barnard, 1995a), their practical application within interface design has been detailed in handbooks and tutorials directed at practitioners and made available on the world-wide-web (May, Scott & Barnard, 1995; May & Barnard, 1997). The overall strategy is more extensively covered in a recent issue of the journal HCI (Barnard & May, 1999).

To do real modeling work, the four attribute spaces need detailed specification. Figure 10 reproduces the most recent attribute space used to represent the capability of mental processes - their 'proceduralized' knowledge. The left hand side of this figure shows a set of generic attributes for a process that transforms input into outputs in a different mental code (a mapping). The right hand side approximates capability of in terms of a wider repertoire of mappings that stand in different relationships to the task setting being modeled. Since the two are closely related, a very similar attribute space for is used to represent record content (requirements). With the attribute space for record content definitions of the relevant repertoires distinguish sets of how-to-do-it-knowledge (common task records and experiential task records) from how-it-works-knowledge (entity property records) and sets of knowledge available in the 'here and now' from current input (active task records). The attribute space for configurations defines the types of process. That for dynamic control makes use of many of the constructs we have already mentioned in passing (e.g. oscillation, interleaving, extent of record access, complexity of process interchange, and locus of control).

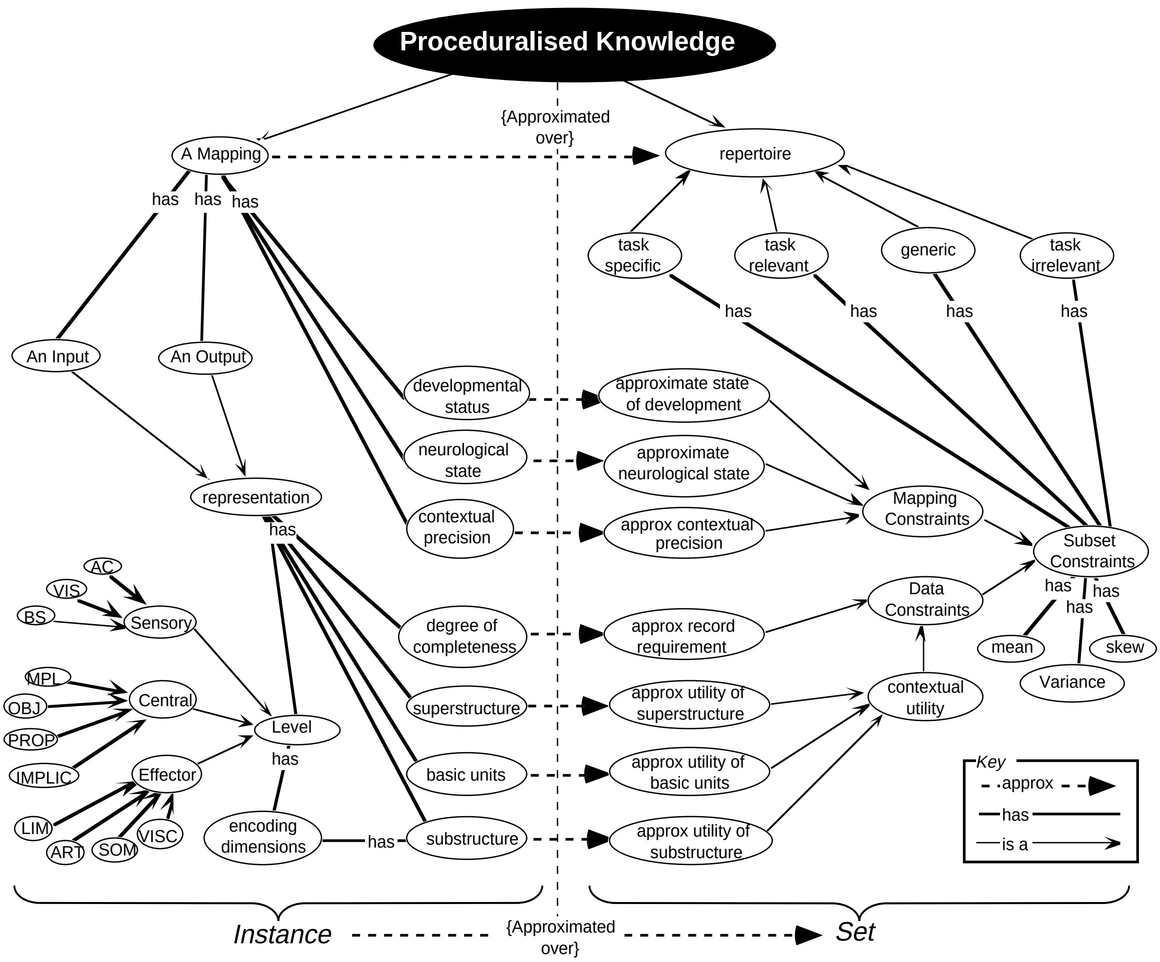

Figure 10. The core attribute space for procedural knowledge (From Barnard & May, 1999)

The four attribute spaces enabled us to build approximate models of mental activity underlying task performance. The fact that all subsystems and mental representations are governed by common principles meant that rules of inheritance could be used. For example, the principle that any one process can only transform a single data stream at a given point in time has consequences. Where a task places two demands on any process, those demands can only be dealt with by some form of oscillation in, or interleaving of, the use of that capability over time. This necessarily increases the complexity of dynamic control and co-ordination across the system. Different generic phases of mental activity, such as goal formation, action specification and action execution would each involve significant variation in how configurations, procedural knowledge (capability) and record content (requirements) are put to use. These, and their consequences for wider system control and coordination, would all be the subject of change over the course of learning. As a user acquires experience, the task relevant content of memory records would increase and more and more capabilities would become automated within specific processes.

In order to provide the rules with some grounding in evidence, the results of a variety of empirical studies were summarized. These were drawn from the general psychological literature and from specific experiments with command languages (Barnard, Grudin & MacLean, 1989; Barnard, MacLean & Hammond, 1984); menu systems (e.g. Hammond & Barnard, 1982); and graphical interfaces (e.g. Green & Barnard, 1990; May, Tweedie & Barnard, 1993). On the basis of this evidence we formulated a set of rules for tying specific properties of tasks and interfaces to the attribute spaces, to generic principles and form these to properties of user behavior. The resulting rules were then embodied in an expert system. To support generalization and extension, the core modeling rules were organized into one set of knowledge bases. These were separated from the two other major components of the system. One was a self-contained set of rules for collecting information about tasks, users and interfaces. The other set of rules was used to generate predictions or design advice.

The approach is neither as well developed nor as extensively tested as the current AI architectures referenced in section 2, but our work illustrates that simulation is not the only route to integration. The rules of the expert system embodied macrotheory and, where we had it for specific task environments, were supported by microtheoretic detail. It therefore conforms to our requirement for the composition of Type 1 theories of the behavior of a system. The inferences do not rely upon a highly detailed simulation like those of the AI tradition within cognitive science. The reasoning is nonetheless explicit and it is based upon theoretical principles whose effects can be traced.

The extent to which principles generalize can also be tested and validated empirically. For example, during the earliest phases of learning, we assume that the ‘how-to-do-it knowledge’ for performing tasks is held in the image record of the Propositional subsystem. The knowledge can be assigned attributes such as the level of item and ordering uncertainty that must be resolved to generate the surface structure of a command sequence from its underlying semantic representation (Barnard et al., 1987). Likewise, current states of a graphical interface are modeled as the set of representations recently copied into the image record of the Object subsystem. In this case an attribute will be set that captures the level of spatial uncertainty associated with the location of an item in a graphical array. Higher levels of uncertainty associated with any one of these attributes are assumed to increase the complexity of processing exchanges, and may also change the locus of control within the mental configuration. These assumptions are testable with laboratory methods (e.g. May, Barnard & Blandford, 1993).

The expert system was used to draw inferences about interaction based upon a model of the human cognitive architecture and the different types of mental representation processed within it. This approach is undeniably unbalanced as a means of representing interactions in a higher order system composed of user and computers. While the representation of the cognitive architecture is reasonably formal and relatively rich, the representation of the computer is both informal and highly impoverished. In terms of our overlapped layers of Figure 2, it may well form a basis for a Type 1 model of a psychological system. Via a Type 2 mapping it may also form one basis for a microtheory of the human component of a higher order system. However, the development of a Type 1 theory of a syndetic or behavioral system requires a more balanced approach. It requires the specification of microtheories of all the relevant interactors at an equivalent level of richness and it requires a macrotheory of their interaction.

6. Realising coherent type 1 theories of interaction grounded in theories of cognitive and computational systems

To achieve a Type 1 theory of the behavior for a system of interactors involving both people and computers, it would be helpful to represent both in a common language. The earlier discussion of the unselected window scenario introduced the general idea of an interaction trajectory, but did so in informal terms. It was not a formally specified model or theory of interaction. Within computer science, it has been acknowledged that the behavior of computational systems needs to be modeled at an abstract level. When designing a system it is, for example, important to establish that it cannot enter deadlock. To model computational systems, the formal methods community have sought to evolve a body of mathematics that enables them to represent abstract properties of systems. Their fundamental concern is to use those models to reason about the behavior of a computer system before it is refined and built. The use of mathematics to describe software and hardware systems has been widely explored. Current work includes its application to human-computer interaction (e.g. Dix, 1991; Duke & Harrison, 1993; Paternó & Palanque, 1997, Harrison & Torres, 1997).

In software development there is a need to carry out rigorous checks on models. The development of such models uses software tools, and this requires the mathematics for software specification to be defined at a level of rigor that is not commonly applied in other areas. Specific 'formal methods,' such as VDM, Z and CSP have been developed to do this. They define collections of mathematical structures with a specific syntax and semantics for use in systems modeling. Several key advances in understanding complex problems in computing have already come about through the development of new mathematical abstractions for the operation of concurrent systems (e.g. Hoare, 1996; Milner, 1989, 1993).

When a computer scientist does produce a formal model of the behavior of a system before it is refined and built, it fulfils a role directly equivalent to that of macrotheory for a psychological architecture. The conceptual parallels run deeper. We can characterize a system of computing interactors in the four component schema used to organize psychological macrotheory. Indeed, the four component framework maps directly onto the schema for applying the mathematics of control theory. A generalized construct of a data array (requirements) on which functions (capability) operate to update the array according to selection restrictions (co-ordination and control). Configurations can then be an emergent property of what functions can operate on what data types within the array at what time.

The psychological theory (ICS) discussed in the previous section is itself a model of a concurrent computing architecture, albeit one grounded in biological wetware rather than silicon. Recognizing the parallel, we should be able to use the mathematics being developed within computer science to model abstract properties of the behavior of the mental architecture and bind it to an equally abstract model of the behavior of a computer system. To develop a coherent Type 1 model of the resulting system, some macrotheory of the interaction needs to be added – in this case, a syndetic model. In the Type 2 transition (Figure 2), macrotheory from psychology and macrotheory from computer science are both re-represented to form microtheory at the level above in the overlapped hierarchy of abstraction.

Our first attempt at this form of theoretical integration is fully reported in a recent issue of the journal HCI (Duke, Barnard, Duce & May, 1998). Although this contribution relies heavily on an understanding of advanced mathematics, the specific form of mathematics, as we discuss later, is less important than the overall methodology. Working within a large scale European project on the integration of theory within design processes (AMODEUS, 1995), we focused on a number of concrete design scenarios for advanced systems. One of these was a Multimodal Air Travel Information System (MATIS) capable of integrating typed, voice and gestural inputs. The user could say '...Flights from here to there..' and click on reference to the relevant cities by pointing at the appropriate referent with a mouse (Nigay & Coutaz, 1995). The second scenario was a specific form of gestural language. This was designed to be used with the sort of data glove interfaces where the hand movements can be used to issue commands to the computer and where an image of the glove is concurrently rendered on the user's display (Bordegoni & Hemmje, 1993). With well defined configurations for information flow (Figure 7), ICS provides some generic rules for how the cognitive mechanism handles multimodal integration (see Barnard & May, 1995). Since the properties of the two computing systems were known and specifiable, the constraints on key components of the two computing systems were modeled using modal action logic. The equivalent properties of the ICS architecture were also modeled using modal action logic, but adding deontic extensions to capture aspects of the indeterminacy of mental processing.

As with simulation methods, the process forced us to be more precise in our specification of human information processing theory than we had been in our earlier attempts to model cognitive activity. The result of this modeling was two sets of axioms. One set of axioms represents assumptions about the constraints governing computer system behavior and the other set represents assumptions about how the user's mental mechanism would function in the same setting. The axioms illustrate two aspects of the Type 2 transition. They discarded detail from more detailed psychological or system models. They were also selective in that they represented only those parts of the lower level models that were necessary to deal with specific issues. These axioms were microtheories of user and system components. The added value of a Type 2 theory, is the addition of a new superordinate level of organization. This was represented by a third set of axioms that specified assumptions about conditions that must apply in the interaction itself.

The combination of micro- and macrotheory represents a coherent Type 1 syndetic model grounded in specifications of cognitive and computing mechanisms. Like any other form of representation, mathematics can be hard to follow unless people are familiar with the notations. Figure 11 reproduces a sample of axioms. In each case the notation is explained in a sentence. These are intended to illustrate the high level abstractions involved. The first two axioms are drawn from the fourteen used to represent the theoretical requirements that we assumed must hold in the mental mechanism for information arriving at a subsystem to be combined within process operation. Two axioms are included from the MATIS specification and just one of the syndetic axioms needed to bind together the user and system models. The syndetic axiom illustrated deals with the co-ordination of user and system capabilities.

Sample ICS axioms

$$\text{coherent}(t_1, t_2) \Leftrightarrow \text{dest}(t_1) = \text{dest}(t_2) \wedge \forall\ p, q : \text{repr} \bullet p \text{ on } t_1 \wedge q \text{ on } t_2 \Rightarrow p \approx q$$

Stream, t_1 and t_2 are coherent if and only if they have the same destination, and for any representation p available on t_1 and q on stream t_2, p and q are coherent.

$$t \in \text{stable} \Leftrightarrow \forall s_1, s_2 : \text{sources}(t) \bullet \text{coherent}(s_1, s_2) \wedge (t = \text{buffered} \vee \text{sources}(t) \subseteq \text{stable})$$

A transformation 't' is stable if and only if every pair of streams on which it operates are coherent, and either the transformation is buffered, or the input streams are stable.

Sample Axioms from the MATIS Model

$$\text{speech} = X \Rightarrow [\text{speak}(nm,d)]\ \text{speech} = X\text{^}\langle(nm, d)\rangle$$

If the speech stream holds X, then speaking a name/data pair results in a speech stream with that pair appended to X.

$$\text{mouse} = M \Rightarrow [\text{select}(d)]\ \text{mouse} = M\text{^}\langle d\rangle$$

If the mouse stream holds M, then selecting a data item d results in a stream in which d is appended to M.

An axiom from the syndetic model

$$\textbf{per}(\text{read}(d)) \Rightarrow d \text{ in MATIS} \wedge \langle *\text{vis-obj}:, :\text{obj-mpl}:, :\text{mpl-prop}:\rangle \in \text{flows}$$

It is possible to read some data item 'd' if d is part of a field of a query in the display and the cognitive configuration enables reading.

Figure 11. Sample axioms from ICS, MATIS and their syndesis (From Duke, Barnard, Duce & May 1998).

The specification can be very economic, since the syndetic model inherits the model of the user and the model of the computer. In the two examples reported by Duke et al. (1998), the syndetic models of the salient aspects of user-system interaction with the MATIS system and with the gesture language required only three and two syndetic axioms respectively. The specifications are really rather short by comparison to those normally encountered in the formal methods literature, and they are certainly far shorter than the kinds of coding required in full AI simulations. Most importantly, the specifications are not committed to any unnecessary detail about either user mentation or system implementation.

Once expressed in mathematical form, the model can do work. Unlike a simulation, this class of specification does not run and make a prediction. Newell (1990) noted, 'theories don't make predictions, *theorists* do.' In this case, conjectures can be postulated and tested. The abstract model of a syndetic system is used to answer questions formulated by designers and to evaluate claims about the behavior of systems of interactors. In the case of the airline system, we might want to know how the user will cope with this design and develop a set of conjectures about it. The model is used *to derive a formal proof* of that conjecture. Duke (1995), Duke, Barnard, Duce & May (1995, 1998), and Duke & Duce (in press) provide a number of examples of this process.

In the case of the multimodal airline information system, the proof indicated that mental activity would not be well coordinated with the system behavior and that deictic reference using speech and mouse with this particular system would be problematic. Although the functionality was available, the analysis suggested that the likelihood that it would be used was low. As with the detours of the unselected window scenario, the analysis provided an insight into the properties of a behavior trajectory for a system composed of user and computer. It also provided an example of where additional functionality could be shown to be insensitive to wider properties of interaction. The analysis addresses exactly those problems discussed in section 4, but the reasoning is now formal and explicitly grounded in lower level models.

As the underlying cognitive and system models provide the microtheory supporting the syndetic macrotheory, the reason why a particular difficulty occurs can be traced to somewhere in the Type 1 theory for that system of interactors. Alternative design solutions can then be driven by theoretical insight rather than by generate-and-test cycle of ad hoc change. Using mathematics, the consequences of theory can be explored in much the same way as they are in other physical sciences. The reasoning depends upon the theoretical assumptions, and these are not intermingled with the kind of detail that can lead to problems in linking underlying theory to consequence in simulation methodologies (Cooper, Fox, Farringdon & Shallice, 1996).

Similarly, it is expensive to generate new models afresh for each new system and application. What Duke et al. (1998) show is that axioms developed in one context can be re-used to model generically related circumstances. Once in place, a body of psychological theory, such as ICS, can be combined with models of different computer systems. In the Duke et al. (1998) case, the reasoning about pose-formation with a data glove has significant fragments in common with the reasoning about the MATIS system. Where similar abstract conditions apply in the models of computer systems, then the relevant bodies of theory are re-used to support the development and exploration of alternative syndetic systems. Although the specifications and proofs presented by Duke et al (1998) cover several pages, they rely on inferences that could, in principle, be carried out using the current generation of theorem provers such as MURAL (Jones et al. 1991) or PVS (Owre et al., 1995). It should not ultimately be necessary to have a tame professor of formal methods in every design team's cupboard, or in every psychology laboratory.

It will take a good deal of work for this approach to mature to a point where we have models that come with appropriate scientific guarantees of their validity. The approach does not depend upon the specific psychological theory we have used – ICS. Other models could be represented in the same way. Likewise, we are fully aware of the limitations of modal action logic – this form of mathematical representation has limited expressive power and can only capture part of what we would wish the mathematics to accomplish within the framework of syndetic modeling. The overall approach is not reliant on a specific mathematical calculus.

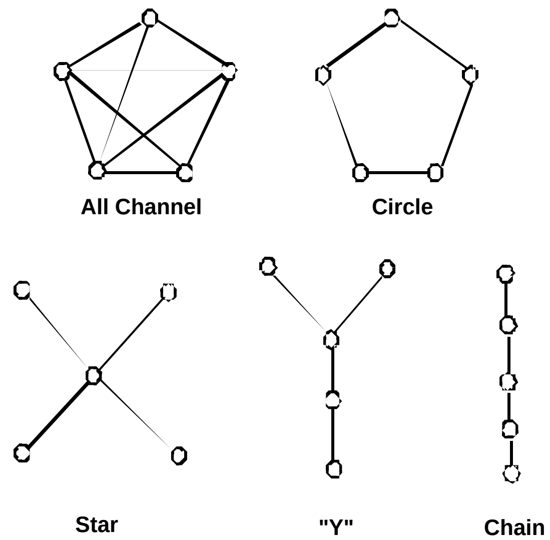

Figure 12. Centralized and de-centralized communication networks.

Indeed, ICS has itself been modeled using an altogether different formal method, LOTOS (Bowman & Faconti, 1999). The body of mathematics that is evolving in computer science is not static, but moving forward to encompass problems that were not previously tractable. To deal adequately with the human interactor, we might need to marry up several different forms of mathematics, including, for example, the mathematics of signal processing, the mathematics of control theory, and yet other forms of logic for computational systems, such as interval temporal logics.

7. Extension to Higher order systems of interactors

An important part of our argument is that four generic classes of constraint need to be modeled at all levels of systems analysis. So far we have illustrated how macrotheory can be explicitly specified for users and for systems. We have also illustrated how it can be combined in a formal mathematical model of a syndetic system. However, to support wider integration across the field of HCI, it needs to be shown that very similar forms of analysis might also hold for a characterization of behavior trajectories for groups and yet larger organizations.

An example of how principles of similar form and content may also apply to the behavior of groups originates in research on communication nets, carried out in the 1950s and 1960s. Starting with Leavitt (1951), social psychologists carried out research on tasks performed by collaborative groups working in configurations where they could only communicate along the kinds of constrained paths shown in Figure 12. These are analogous to our representation of the potential configurations for a cognitive architecture (Figure 7). Pre-dating computer technologies, this paradigm can be regarded as paper-and-pencil supported co-operative work. The only interactors in this setting that have the capability to change representations or physical states are the people. They communicate by passing written messages through slots in re-configurable walls. The paper, writing implements and walls are, of course, subordinate interactors in this setting and they too have definable capabilities (paper affords writing etc., slots afford message passing). However, our concern is not with the micro-properties of how this medium might differ from a more modern email counterpart, but with characterizing how properties of the configurations relate to interactor capabilities, requirements and their dynamic control and co-ordination. In this setting the basic interactor is 'Human agent+paper+pencil' and it has the capability to construct task relevant messages and pass them through one or more slots.

Leavitt (1951) made use of a very simple task. Each person had a set of different symbols and their collective task was to identify the particular symbol that was common to all of them.

Leavitt found that, of the circle, star and 'Y' configurations, the circle configuration gave rise to most errors, took most time and involved the greatest number of messages exchanged. While least efficient according to conventional criteria, people actually liked being in this configuration more than the other ones. It is neither hard to generate explanations of why the more centralized star configuration was more efficient, nor why it was disliked more. An essential requirement for the group level interactor is to bring information together, and the centralized networks have, by their very structure, a locus of control through which all information is channeled. It is easy to see how the interaction trajectories for the centralized and decentralized nets might differ in terms of the segments and phases of interaction needed to describe them. Although any given interactor in the circle network might have the capability to determine the required answer, the messages have to pass through more links, and the co-ordination of exchanges, and of decision making capability, would be more involved.

Most people in the network may have disliked their restricted status and the fact that they could only communicate with those who were adjacent to them. The point to take from this example is not the specification of particular a trajectory or its properties. What is more important are the findings demonstrating that the relationships are not constant. With more complex tasks, Shaw (1954) found that decentralized networks gave rise to better performance. The requirements of bringing information together, and the issues that must be resolved in doing so, exceed the capabilities of the single individual acting as the locus of control. They have to interleave information and control transactions and keep track of everything that is going on. In contrast, activity in decentralized networks supports concurrent activity on different parts of the problem, ultimately resulting in fewer exchanges and less time.

Clearly, in order to model this example, including the alternative outcomes of participant satisfaction, the attributes and principles of a four component analysis would need to be captured with more precision and completeness. Our reasoning about why different trajectories are followed from a particular point of departure to outcome is nonetheless framed by *inter-relationships* that hold between properties of configurations, capabilities, requirements and the dynamic control and co-ordination of the system itself. No amount of task description, knowledge analysis, or considerations of the limits on human cognition would provide such an analysis without referencing some other factors not normally considered to be within their scope.

In the case of yet higher order systems of interactors, the very terminology of the four component framework would, at least in some contexts, have a strong sense of familiarity. Military strategists tend to think in terms of how their forces are configured, what their capabilities are, what requirements must be met for them to use that capability and, of course, how command and control is to be exercised. Preliminary attempts apply the four component framework to such systems have suggested that there might indeed be interesting similarities between the behavior of such systems and those of more traditional topics in HCI. For obvious reasons, our final concrete scenario is historic rather than current.

On August 2nd 216 BC, a Carthaginian army commanded by Hannibal engaged a Roman army commanded by Varro at the Battle of Cannae. By the end of the day Hannibal's forces had slaughtered over 60,000 Romans, whilst losing only 6,000 of their own, very much smaller force. The three main phases representing the behavior trajectory for this highly concurrent system of interactors are shown in Figure 13. The progress of the battle can be succinctly summarized[1]. In the opening phase of the battle, Hannibal advanced a thin salient of his infantry toward the Romans; and the Romans advanced to meet them. At the end of this phase Hannibal executed a pre-conceived tactical withdrawal. In the second phase the Roman army, believing that they were already in sight of victory, was drawn into Hannibal's trap as he maneuvered his infantry around the flanks of the advancing Romans. At the end of the second phase of the battle, Hannibal's cavalry, who had engaged the Roman cavalry to the rear during the first phase, now disengaged and completed the encirclement of the Roman

[1] details derived from Dupuy & Dupuy (1993). Collins Encyclopaedia of Military History

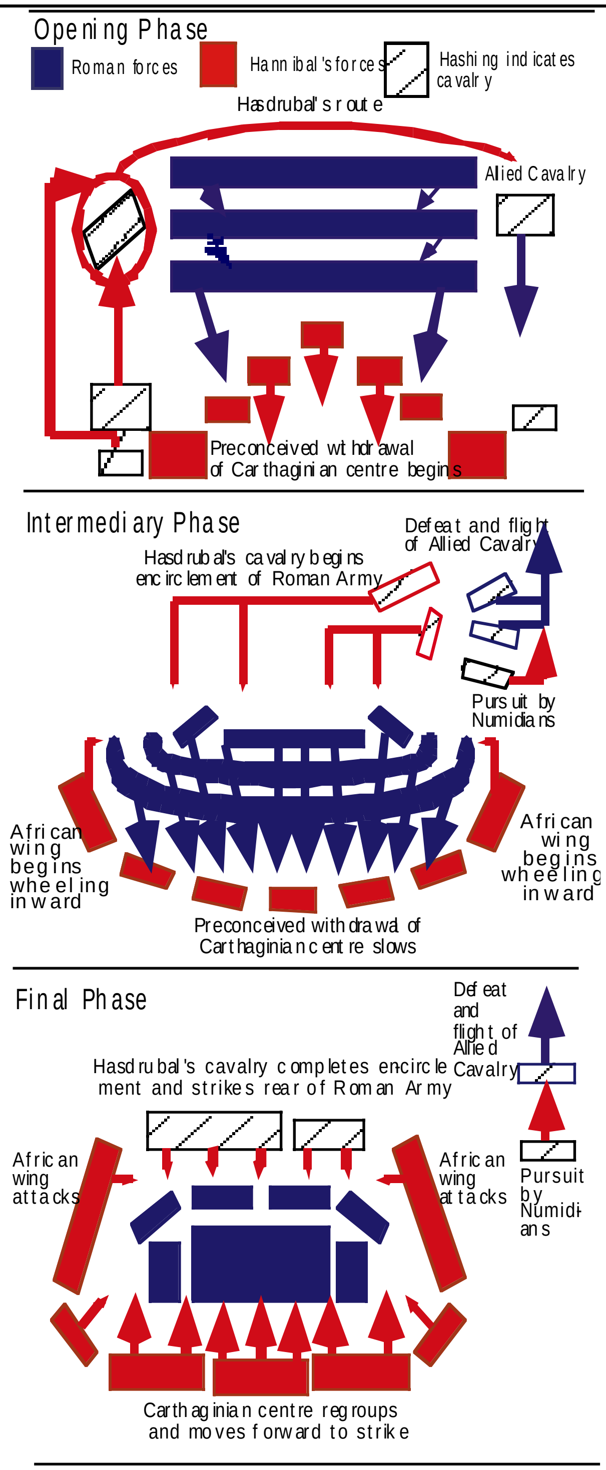

Figure 13. The Battle of Cannae.

army. In the third phase of the battle, the Romans became a herd of panic-stricken individuals, their force structure losing all coherence and unity at the point when they realized that they had been enveloped. Hannibal's numerically inferior force quite literally cut them to pieces.

The block diagrams specify the different configurations for the various interactors within this system, the different units of infantry and cavalry. Given that their infantry outnumbered the Carthaginians by more than two to one, the Romans should have won. Hannibal had negated the Roman's numerical superiority in two ways. First, once encircled, a significant proportion of the Roman units was trapped behind the line of engagement. They could not be configured to interact physically with the Carthaginians. Second, being only human, the Romans panicked and the individual units lost the ability to fight as constituents within a higher level force structure. The loss of co-ordination and control within basic units of organization maps upward to a degradation in the capability of interactors at the higher level of systems analysis.

Some of the connections to our earlier discussions are at least intriguing. For example, the training wheels system developed by Carroll and Carrithers (1984) denied novice users access to software functionality. As at Cannae, the 'potential' of subsequent behavior trajectories was reduced by denying use of capability. In both instances the effect was to facilitate the trajectory towards an intended outcome, albeit faster learning in one instance and more effective slaughter in the other. The informally presented analysis of the unselected windows scenario (section 4), and the more formal modeling of interactions with the multimodal air travel information system (section 6), both involved reasoning about dynamic control and co-ordination of simple user and system exchanges. When mapped upwards to the concerns of designers of the higher order system, the conclusions may need to be captured, not in the detailed model, but in a Type 2 mapping. As in the case of Cannae, if coordination and control is degraded in a lower order system, then it impacts the capability of the higher order system. One sort of window improves capability of the combined system relative to the existing variant. It simplifies behavior trajectories by reducing the likelihood of detours. An envisaged design for multimodal integration of speech and mouse action in airline inquiries is unlikely to be used and hence have an important impact the capability of the higher order system. Only when it could be argued that an envisaged design would facilitate behavior trajectories for the higher order system might further investment in that aspect of software development be justified.

If those operating at different levels of systems analysis and on different topics were able to relate their own contributions to some common vehicle of expression, such as taking about generic constraints on interaction, then communication should be enhanced. Individual researchers might even gain substantial benefit from advances made by those working on quite different systems of interactors. Those working on modeling computational systems, cognitive systems, social systems or military systems might improve their own tools and models on the basis of greater interdisciplinary coordination and reciprocated insight.

9. CONCLUSION

We have argued that the future course of theory in HCI might not be best served by the unconstrained development of more and more local theories and models specifically tailored only to the meet needs of different levels of analysis or different software applications. Because of the different systems of scientific semantics adopted, such an approach makes it hard to realize connections – either within or between levels of systems analysis.

We have offered some arguments for developing macrotheories. These provide connective tissue to bind together different concerns within a level of explanation. We also argued that macrotheories were needed to support a different kind of theory. The second form of theory provides the connective tissue between the systems of scientific semantics adopted at adjacent levels of analysis. A specific framework was proposed for the consistent organization of macrotheory across levels of analysis. This assumed that the behavior of any system was determined by four classes of constraint. At each level of analysis we would expect the content expressed in a theory or model to be quite distinct and testable in its own right. We used a particular cognitive theory and a particular mathematical route to building models of syndetic systems to support the wider argument that such developments are a tractable research proposition. The four classes of constraint are dependent upon neither our psychological model nor a particular form of mathematics. We used these tools because they are the ones that we are most familiar with. The general argument could equally well have been based upon other cognitive models or notations. We also sought to illustrate how the framework might be applied across a broad range of issues. Concrete scenarios were used from both past and current technologies, from window design to the design of multimodal or data glove interfaces. In addition to interactions, at the human computer interface, we also briefly considered the behavior of groups and organizations. We concluded by examining how the framework might help express some simple parallels between interactions at different levels of system analysis.

Hannibal's plan for the Battle of Cannae provided military theorists with an oft-cited model of tactical perfection. Some two thousand years later, General Schwarzkopf remarked that he essentially re-used Hannibal's model for the battle of Cannae when he directed operation Desert Storm. An abstract model that is re-usable over a couple of millennia, and from the technologies of swords and shields to those of tanks and missiles, is a not insignificant achievement. As HCI moves into the next millennium, the development of new forms of theory, or a viable mathematics for modeling systems at different levels of analysis, still poses a whole range of exciting and important challenges. Before reliable and enduring models are achieved, research may well go up numerous blind alleys. Nonetheless, we believe the case for the development of models and theories is just as strong, if not stronger, as it was when Newell (1990) and Card, Moran & Newell (1983) laid out their earlier visions for how theories might be integrated and applied to the practice of HCI.

10. References.

AMODEUS (1995) *Assaying Means of Design Expression for Users and Systems. AMODEUS-2 Project Final Report,* D13, Barnard, P.J., Bernsen, N.O., Coutaz, J., Darzentas, J., Faconti, G., Hammond, N. H., Harrison, M.D., Jørgensen, A.H, Löwgren, J., May, J., Young, R.M. CEC Brussels, pp43. `http://www.mrc-cbu.cam.ac.uk/amodeus/abstracts/d/d13.html`

ANDERSON, J.R. (1993). *Rules of the mind*. Hillsdale, NJ. Erlbaum

BARNARD, P.J. (1985). Interacting cognitive subsystems: A psycholinguistic approach to short term memory. In A. Ellis, (Ed.), *Progress in the psychology of language*. London: Lawrence Erlbaum Associates, 197-258.

BARNARD, P.J. (1987) Cognitive resources and the learning of dialogues. In J. M. Carroll (Ed) *Interfacing Thought* (pp. 112-128) Cambridge, Mass: MIT Press.

BARNARD, P.J. (1991). Bridging between basic theories and the artifacts of human-computer interaction. In J.M. Carroll (Ed.) *Designing Interaction. Psychology at the Human-Computer Interface.* Cambridge: Cambridge University Press, 103-127.

BARNARD, P (1999). Interacting Cognitive Subsystems: modeling working memory phenomena within a multi-processor architecture. In Miyake, A. & Shah, P. (Eds.), *Models of Working Memory: Mechanisms of active maintenance and executive control.* New York: Cambridge University Press, 298-339.

BARNARD, P.J. AND HARRISON, M.D. (1989). Integrating cognitive and system models in human-computer interaction. In A. Sutcliffe and L. Macaulay (Eds.), *People and Computers V*. Cambridge: Cambridge University Press, 87-103.

BARNARD, P.J. AND HARRISON, M.D. (1992). Towards a framework for modeling human-computer interactions. In J. Gornostaev (Ed.), *Proceedings of the East-West International Conference on Human Computer Interaction, EWCHI '92*. Moscow: ICSTI, 189-196.

BARNARD, P. & MARCEL, A. J. (1984). Representation and understanding in the use of symbols and pictograms. In R. Easterby and H. Zwaga (Eds.), *Information Design: The Design and Evaluation of Signs and Printed Material*, John Wiley: Chichester, 37-75.

BARNARD, P. AND MAY, J. (1993). Cognitive modeling for user requirements. In P.F. Byerley, P.J. Barnard & J. May (Eds*.), Computers, Communication and Usability: Design Issues, Research and Methods for Integrated Services* (pp. 101-145). Amsterdam: Elsevier.

BARNARD, P. & MAY, J. (1995). Interactions with advanced graphical interfaces and the deployment of latent human knowledge. In F. Paternó (Ed.), *Eurographics Workshop on Design, Specification and Verification of Interactive Systems* (pp. 15-49). Berlin: Springer-Verlag.

BARNARD, P. AND MAY, J. (1999). Representing cognitive activity in complex tasks. *Human Computer Interaction*. 14,93-158.

BARNARD, P., GRUDIN, J. AND MACLEAN, A. (1989). Developing a science base for the naming of computer commands. In J.B. Long and A. Whitefield (Eds.), *Cognitive Ergonomics and Human Computer Interaction.* (pp. 95-133) Cambridge: Cambridge University Press.

BARNARD, P.J., MACLEAN, A. & HAMMOND, N.V. (1984) User representations of ordered sequences of command operations. *Proceedings of Interact '84: First IFIP Conference on Human-Computer Interaction*,. 434-438, London: IEE,

BARNARD, P.J., WILSON, M. & MACLEAN, A. (1987) Approximate modeling of cognitive activity with an expert system: a strategy for the development of an interactive design tool. *Proceedings of CHI+GI '87*, 21-26, New York: ACM.

BELLOTTI, V., BLANDFORD, A., DUKE, D., MACLEAN, A., MAY, J. & NIGAY, L. (1996). Interpersonal access control in computer-mediated communications: A systematic analysis of the design space. *Human Computer Interaction, 6,* 357-432.

BLANDFORD, A. & DUKE, D. (1997). Integrating user and computer system concerns in the design of interactive systems. *International Journal of Human-Computer Studies, 46,* 653-679.

BLANDFORD, A., HARRISON, M. & BARNARD, P. (1995). Using interaction framework to guide the design of interactive systems. *International Journal of Human-Computer Studies, 43,* 101-130.

BORDEGONI, M. & HEMMJE, M. (1993). A dynamic gesture language and graphical feedback for interaction in a 3d user interface. *Computer Graphics Forum, 12(3),* 1-11.

BOWMAN, H. & FACONTI, G. (1999) Analysing cognitive behaviour using LOTOS and MEXITL. *Formal Aspects of Computing,* 11:132-159..

CARD, S.K., MORAN T.P. & NEWELL, A. (1983) *The psychology of human computer interaction* Hillsdale, NJ: Lawrence Erlbaum

CARROLL, J. M. & CARRITHERS, C. (1984) Blocking learner errors in a training wheels system. Human Factors, 26, 377-389.

CARROLL, J. M.& ROSSON, M. B. (1992) Getting around the task-artefact framework: how to make claims and design by scenario. ACM Transactions on Office Information Systems, 10(2), 181-212.

CHAPANIS, A. (1975). Interactive human communication. Scientific American, 232, 36-42.

CHAPANIS, A. (1996). Human factors in systems engineering. John Wiley.

CLARK, H.H. (1996) Using Language. Cambridge: Cambridge University Press.

COOPER, R., FOX, J., FARRINGDON, J., & SHALLICE, T. (1996): A systematic methodology for cognitive modeling. *Artificial Intelligence*, 85, 3-44.

DIX, A. (1991). *Formal Methods for Interactive Systems.* Academic Press.

DUKE, D. (1995). Reasoning about gestural interaction. *Computer Graphics Forum, 14(3),* 55-66.

DUKE, D.J., BARNARD, P.J., DUCE, D.A. AND MAY, J. (1995). Systematic development of the human interface. In *APSEC'95: Second Asia-Pacific Software Engineering Conference,* IEEE Computer Society Press, 313-321.

DUKE, D.J., BARNARD, P.J., DUCE, D.A. AND MAY, J. (1998). Syndetic modelling. *Human Computer Interaction*, 13, 337-393

DUKE D.J. AND DUCE, D.A . (in press) The Formalisation of a Cognitive Architecture and its Application to Reasoning about Human Computer Interaction. *Formal Aspects of Computing*.

DUKE, D. & HARRISON, M. (1993). Abstract Interaction Objects. *Computer Graphics Forum, 12(3),* C-25 – C-36.

DUKE, D. & HARRISON, M. (1994). A theory of presentations. In M. Naftalin, T. Denvir & M. Bertran (Eds.*), FME'94: Industrial Benefit of Formal Methods,* volume 873 of Lecture Notes in Computer Science (pp. 271-290). Berlin: Springer-Verlag.

DUKE, D. & HARRISON, M. (1995a). From formal models to formal methods. In R.N. Taylor & J. Coutaz (Eds.), *Software Engineering and Human-Computer Interaction: ICSE'94 Workshop on SE-HCI: Joint Research Issues,* volume 896 of Lecture Notes in Computer Science (pp. 159-173). Springer-Verlag.

DUKE, D. & HARRISON, M. (1995b). Interaction and task requirements. In P. Palanque & R. Bastide (Eds.), *DSV-IS'95: Eurographics Workshop on Design, Specification and Verification of Interactive Systems* (pp. 54-75). Wien: Springer-Verlag.

DUPUY, R.E. & DUPUY, T. N. (1993). *The Collins Encyclopedia of Military History*. HarperCollins.

GREEN, A.J.K. AND BARNARD, P.J. (1990). Iconic interfacing: The role of icon distinctiveness and fixed or variable screen location *Human-Computer Interaction – INTERACT '90*. 457-462. Amsterdam: Elsevier Science Publishers, B.V.

HAMMOND, N.V. AND BARNARD, P. (1982). Usability and its multiple determination for the occasional user of interactive systems. In M.B. Williams (Ed.), *Pathways to the Information Society*, North Holland: Oxford, 543-548,

HARRISON, M.D. AND BARNARD, P.J. (1993) On defining requirements for interactions. In A. Finkelstein (Ed.), *Proceedings of IEEE International Workshop on Requirements Engineering* New York: IEEE, 50-54.

HARRISON, M. & TORRES, J. (Eds.). (1997). *Design, Specification and Verification of Interactive Systems'97*. Wien: Springer-Verlag.

HOARE, C. (1996). How did software get so reliable without proof? In M.-C. Gaudel & J. Woodcock (Eds.), *FME'96: Industrial Benefit and Advances in Formal Methods,* volume 1051 of Lecture Notes in Computer Science (pp. 1-17). Springer-Verlag.

JONES, C., JONES, K., LINDSAY, P. & MOORE, R. (1991). *MURAL: A Formal Development Support System*. Springer-Verlag.

LANDAUER, T.K. (1987). Relations between cognitive psychology and computer systems design. In J. M. Carroll (Ed.), *Interfacing thought: Cognitive aspects of human-computer interaction* (pp. 1-25). Cambridge, MA: MIT Press.

LEAVITT , H. J. (1951). Some effects of Certain Communication Patterns on Group Performance. *Journal of Abnormal and Social Psychology*, 46, 38-50.

LEE, W-O. (1992) The Effects of Skill Development and Feedback on Action slips. To appear in: Monk, A., Diaper, D. & Harrison, M.D., *People and Computers VIII.* Cambridge: Cambridge University Press, 73-86

MARR., D. (1982). Vision, New York: W.H Freeman and Company

MACLEAN, A., YOUNG, R., BELLOTTI, V. & MORAN, T. (1991). Questions, options, and criteria: Elements of design space analysis. *Human-Computer Interaction, 6(3&4),* 201-250.

MALONE, T. W. & CROWSTON, K. (1994). The interdisciplinary study of coordination. ACM Computing Surveys, 26 87-119.

MAY, J. & BARNARD, P. (1995a). The case for supportive evaluation during design. *Interacting with Computers, 7*, 115-144.

MAY, J., & BARNARD, P.J (1995b) Cinematography and Interface Design *Human-Computer Interaction: Interact'95.* 26-31. London : Chapman & Hall.

MAY, J. & BARNARD, P.J. (1997) Modelling Multimodal Interaction: A theory-based technique for design analysis and support. In S. Howard, J. Hammond & G. Lindegaard (eds) Human-Computer Interaction INTERACT '97. Chapman & Hall: London. pp 667-668.
`http://www.shef.ac.uk/~pc1jm/guide/guide97.pdf`

MAY, J., BARNARD, P. & BLANDFORD, A. (1993). Using structural descriptions of interfaces to automate the modeling of user cognition. User Modeling and User Adaptive Interfaces, 3, 27-64.

MAY, J., SCOTT, S. & BARNARD, P. (1995). *Structuring Displays: A Psychological Guide.* Eurographics Tutorial Notes PS95 TN4, ISSN 1017-4656. Geneva: European Association for Computer Graphics.

MAY, J., TWEEDIE, L. & BARNARD, P. (1993) Modeling User Performance in Visually Based Interactions. In J.L. Alty, D. Diaper and S. Guest (Eds.), *People and Computers VIII: Proceedings of the HCI '93 Conference.* (pp 95-110) Cambridge: Cambridge University Press.

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

MILNER, R. (1989). *Communication and Concurrency*. Series in Computer Science. London: Prentice Hall International.

MILNER, R. (1993). Elements of interaction – 1993 Turing Award Lecture. *Communications of the ACM, 36(1),* 78-89.

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

NEWELL, S. & CARD. S.K. (1985) The prospects for science in human computer interaction. *Human Computer Interaction, 1*, 209-242.

NEWELL. A, & SIMON. H. A. (1972). Human problem solving. Englewood Cliffs NJ: Prentice Hall.

NIELSEN, J. (1993). *Usability Engineering*. London. Academic Press.

NIGAY, L. & COUTAZ, J. (1995). A generic platform for addressing the multimodal challenge. *Proceedings of the CHI '95 Conference on Human Factors in Computing Systems*, 98-105. Addison-Wesley.

NORMAN, D, (1983) Design principles for human-computer interaction. In *Proceedings of CHI '83: Human Factors in Computing Systems* (pp. 1-10). New York: ACM.

OWRE, S., RUSHBY, J., SHANKAR, N. & VON HENKE, F. (1995). Formal verification for fault-tolerant architectures: Prolegomena to the design of PVS. *IEEE Transactions on Software Engineering, 21(2),* 107-125.

PATERNÓ, F. & PALANQUE, P. (Eds.). (1997). *Formal Methods in Human Computer Interaction.* Springer-Verlag.

RUDISILL, M., LEWIS, C., POLSON, P. & MCKAY, T. (1996). *Human-Computer Interface Design: Success Stories, Emerging Methods and Real World Context.* San Francisco: Morgan Kaufmann.

SHAW, M. E. (1954). Some effects of problem complexity upon problem solving efficiency in different communication nets. *Journal of Experimental Psychology*, 48, 211-217.

SUCHMAN, L.A. (1986). *Plans and Situated Actions*. Cambridge: Cambridge University Press.

TEASDALE, J.D. & BARNARD, P. (1993). *Affect, Cognition and Change: Re-modeling Depressive Thought*. Lawrence Erlbaum Associates.

VAN MULKEN, S., ANDRÉ, E. & MÜLLER, J. (1998) The persona effect: how substantial is it? In H. Johnson, L. Nigay & C. Roast (Eds.) *People and Computers XIII*, Berlin: Springer-Verlag. 53-66.

VEINOTT, E. S., OLSON, J. S., OLSON, G. M. & FU, X. (1997). Video Matters! When Communication Ability is Stressed, Video Helps. *Proceedings of ACM CHI 97 Conference on Human Factors in Computing Systems* v.2 p.315-316.

WALKER, J.H., SPROULL, L. & SUBRAMANI, R. (1994) Using a Human Face in an Interface In Proceedings of CHI'94, pp. 85-91. ACM: New York.

Acknowledgement:

This paper is based on the plenary address given at HCI '98, Sheffield by the first author. Some of the ideas were also developed in collaborations with others on the ESPRIT AMODEUS-2 project and with the UK Defence Evaluation and Research Agency. Both are gratefully acknowledged.