

Oxford University Computing Laboratory  
1 Wellington Building  
118 St. Aldelme Road  
Oxford OX1 3QD

Formal Aspects  
of  
Human-Computer Interaction

Gregory D. Abowd

Trinity College

June 1991

*Thesis submitted for the degree of Doctor of Philosophy  
at the University of Oxford*

# Formal Aspects of Human-Computer Interaction

Gregory D. Abowd  
Trinity College, Oxford

Submitted for the degree of Doctor of Philosophy, Trinity Term, 1991

## Abstract

This thesis provides a constructive application of formal methods to the study of human-computer interaction. Specifically, we are interested in promoting a principled approach to the analysis and design of interactive systems that will accompany existing heuristic techniques. Previous formal approaches have concentrated on general and abstract mathematical models of interactive systems, proving that psychologically valid principles of interaction can be expressed in a language suitable for computation. These abstract models, however, are too far removed from an actual design practice which is strongly influenced by common sense and liable to break down in the face of complexity. Our efforts are focussed toward retaining the mathematical grounding of previous formalisms while providing additional insight and direction for design practice.

We introduce a unifying framework for the informal description of a user, a system and the interface that sits between them. This interaction framework provides the context for our research and motivates the properties of interaction that we wish to formalize. We adopt the view of an interactive system as a collection of agents based on the stimulus-response model. We provide a mathematical model of the agent capable of expressing interactive properties relating the goals of interaction with the visible consequences of that interaction. We also provide a language for agents which allows a natural expression of an agent's internal state-based behaviour and its external event-based behaviour. We contribute further to practical design issues by introducing templates to relate a task analysis to a specification of a system to support the tasks and an interface to adequately portray that functionality to the user. Finally, we initiate the formal investigation of multiagent architectures. This concludes the mapping of properties on abstract models of interactive systems down to properties on more implementation-based models.

## Acknowledgments

I would like to thank my supervisor, Bernard Sufrin, for the initial inspiration on the topic of this thesis and for directing me to the HCI Group at the University of York for a fruitful working environment.

At York, I have had the opportunity to collaborate with Michael Harrison, Alan Dix and Russell Beale. Those collaborative efforts have resulted in several publications, parts of which have been reproduced in modified form in this thesis. The contents of this thesis represent my own work, but the thoughts and style have been greatly influenced and enhanced by contributions from these three individuals.

In addition, I have also greatly benefitted from repeated conversations with fellow students at the PRG in Oxford, students and staff at York, and colleagues on the AMODEUS project. Research in isolation has never appealed to me and I am grateful to those who have provided ample food for my thoughts and friendship for my heart.

I am particularly indebted to Janet Finlay, Roger Took, Paul Andrews, Chris Johnson and Victoria Miles for their helpful comments on drafts of this thesis.

Finally, I would like to acknowledge the support of the Rhodes Foundation which funded the first two years of my doctoral research and the European Commission which funded the final 18 months while I was working on the AMODEUS project.

## Dedication

*To Richard, Sara, John, Anthony, James, David, Elizabeth, Marypat, Rosemary, Michelle, Stephen, Peter, Paula, Janet, Nina, Tom, Katie, Sara, Tom, Michael, Mark, Maryclaire, David, Paul, Danny, Dennis, Michele, Peter, Kristen, Joseph, Sandy, Anne, Elizabeth, John, Philip, Paula*

*and*

*Meghan.*

*God created us as a family. I have felt your love as if we were separated by only an arm's length, not thousands of miles.*

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>A Framework for Discssing Interaction</b>	<b>7</b>
2.1	An informal definition . . . . .	8
2.1.1	The interactive cycle . . . . .	9
2.1.2	The components of the framework . . . . .	10
2.2	A formal definition . . . . .	12
2.2.1	A simple definition of an agent . . . . .	12
2.2.2	An agent description of an interactive system . . . . .	14
2.3	Conclusions . . . . .	15
<b>3</b>	<b>Background and Related Work</b>	<b>17</b>
3.1	Other interaction frameworks . . . . .	19
3.1.1	Evaluation/Execution Cycle . . . . .	19
3.1.2	Interaction Modelling Framework . . . . .	21
3.1.3	Black box models . . . . .	24
3.1.4	Software architectural models . . . . .	29
3.2	A survey of research within the framework . . . . .	34
3.2.1	Research on <i>articulation</i> . . . . .	34
3.2.2	Research on <i>observation</i> . . . . .	41
3.2.3	Research on <i>performance</i> and <i>presentation</i> . . . . .	43
3.3	Conclusions . . . . .	53
<b>4</b>	<b>Properties of interactive systems: Part I</b>	<b>55</b>
4.1	Properties of translations . . . . .	56
4.1.1	Hutchins, Hollan and Norman distances . . . . .	56
4.1.2	Articulation . . . . .	57
4.1.3	Performance . . . . .	58
4.1.4	Presentation . . . . .	59
4.1.5	Observation . . . . .	59
4.1.6	Assessing overall interaction . . . . .	60
4.2	Formal properties of translations . . . . .	61

4.3	Correspondence between agents . . . . .	62
4.4	Predictability . . . . .	65
4.5	Nondeterminism . . . . .	68
4.6	Synthesis . . . . .	70
4.7	Consistency . . . . .	71
4.8	Conclusions . . . . .	73
<b>5</b>	<b>Refining the agent model . . . . .</b>	<b>75</b>
5.1	Requirements for agents . . . . .	76
5.2	Internal specification . . . . .	78
5.3	External specification . . . . .	88
5.4	Communication . . . . .	91
5.5	Overall Combination . . . . .	94
5.6	Interpretations of agents . . . . .	96
5.6.1	Internal interpretation . . . . .	96
5.6.2	External interpretation . . . . .	100
5.7	Conclusions . . . . .	101
<b>6</b>	<b>A language for describing agents . . . . .</b>	<b>103</b>
6.1	Notations for agents . . . . .	104
6.1.1	The standard Z notation . . . . .	104
6.1.2	Object-oriented notations and Z . . . . .	109
6.1.3	Other concurrent notations . . . . .	116
6.2	The agent language . . . . .	116
6.2.1	A language for external specifications . . . . .	121
6.3	Using the agent language . . . . .	124
6.3.1	A toy reactor . . . . .	124
6.3.2	Some input devices . . . . .	127
6.3.3	A window . . . . .	129
6.4	Conclusions . . . . .	135
<b>7</b>	<b>Properties of interactive systems: Part II . . . . .</b>	<b>137</b>
7.1	Relating Display and Result . . . . .	139
7.2	Templates . . . . .	141
7.2.1	Agent restriction . . . . .	142
7.2.2	Result and display templates . . . . .	143
7.2.3	Equivalence and indistinguishability revisited . . . . .	144
7.3	Predictability and Consistency . . . . .	146
7.4	Synthesis . . . . .	152
7.5	Result initiated interaction . . . . .	158
7.6	Conclusions . . . . .	159

<b>8</b>	<b>Interactive system architectures</b>	<b>161</b>
8.1	From abstract principles to concrete architectures . . . . .	162
8.2	Multiagent models . . . . .	163
8.3	Applying formal methods to levels of abstraction . . . . .	165
8.4	Local correspondence . . . . .	168
8.5	Assessing the graphical interface to a text editor . . . . .	171
8.5.1	Single buffer . . . . .	174
8.5.2	The buffer manager . . . . .	179
8.5.3	Deriving a description of the interface . . . . .	180
8.5.4	Conclusions on the case study . . . . .	182
8.6	Conclusions . . . . .	183
<b>9</b>	<b>Conclusions</b>	<b>185</b>
9.1	Summary of the thesis . . . . .	185
9.2	Contributions of this thesis . . . . .	186
9.3	Future work . . . . .	188
	<b>Bibliography</b>	<b>191</b>
	<b>Appendices</b>	<b>205</b>
<b>A</b>	<b>Use of the Z Notation</b>	<b>207</b>
<b>B</b>	<b>Some theorems on the refined agent model</b>	<b>213</b>
<b>C</b>	<b>Detailed semantics for the agent language</b>	<b>219</b>

# List of Figures

2.1	Four phases of interaction between <i>User</i> (U) and <i>System</i> (S) . . . . .	10
2.2	The general interaction framework . . . . .	10
2.3	Translations between components . . . . .	11
3.1	Relationship between Norman's execution/evaluation cycle and the interaction framework . . . . .	21
3.2	Barnard and Harrison's interaction framework. . . . .	22
3.3	The PIE model. . . . .	25
3.4	The red-PIE model. . . . .	26
3.5	The effects space of the PIE model within the interaction framework. . . . .	26
3.6	Heuristic versus formal reasoning within HCI. . . . .	30
3.7	The Seeheim model of a UIMS . . . . .	32
3.8	The language model of UIMS within the interaction framework . . . . .	32
3.9	The Arch/Slinky model . . . . .	33
3.10	Relationships between levels of PIEs . . . . .	46
4.1	Hutchins, Hollan and Norman's distance concepts within the interaction framework . . . . .	57
4.2	Agent correspondence . . . . .	63
4.3	Levels of predictability . . . . .	67
4.4	The user's model of the system . . . . .	67
4.5	Predictability and synthesis . . . . .	70
5.1	Pictorial representation of an agent . . . . .	76
5.2	Synchronous composition of independent agents . . . . .	78
5.3	Interleaving composition of dependent agents . . . . .	79
6.1	The "onion skin" view of a windowing system . . . . .	112
6.2	The object view of a windowing system . . . . .	112
6.3	Window manager/window relationship . . . . .	133
7.1	The red-PIE model. . . . .	140
7.2	Ambiguous object selection in SuperPaint . . . . .	147



7.3	Agent diagram of the paint layer . . . . .	148
7.4	Agent diagram of the object layer . . . . .	149
7.5	Agent diagram of the co-existing paint and object layers of SuperPaint	152
8.1	The Model-View-Controller paradigm of Smalltalk . . . . .	164
8.2	The MVC hierarchy . . . . .	165
8.3	The PAC model . . . . .	166
8.4	The PAC hierarchy within an interactive system . . . . .	167
8.5	The graphical interface of Spy . . . . .	173
8.6	The graphical interface of Ten . . . . .	174
8.7	Agent representation of multibuffer editor . . . . .	175
8.8	The single buffer agent . . . . .	175
8.9	Agent representation of functional core with operations. . . . .	181
8.10	Agent representation of Spy's interface. . . . .	181
8.11	Agent representation of Ten's interface. . . . .	182

# Chapter 1

## Introduction

It is well known that ninety-nine percent of the world's problems are not susceptible to solution by scientific research. It is widely believed that ninety-nine percent of scientific research is not relevant to the problems of the real world. Yet the whole achievement and promise of modern technological society rests on the minute fraction of those scientific discoveries which are both useful and true.

C. A. R. Hoare in foreword to *Systematic Software Development Using VDM* [90].

Indeed the way that people use machines is of key importance. The most significant advances in computer science will be those facilitating this interaction.

From T. Winograd and F. Flores, *Understanding Computers and Cognition* [169, page 137].

The study of Human-Computer Interaction (HCI) is a multidisciplinary one aimed directly at providing scientific answers to the real problems of designing more usable computer systems. We limit ourselves in this thesis to a close examination of the relationship between two of the major contributors to HCI research—computer science and psychology. Each of these disciplines represent a respectable scientific field with concerns that subsume those of HCI. However, the main goal of HCI research is to produce methods by which the collective knowledge of psychology and computer science can be accumulated and applied to the construction of real interactive systems.

In practice, we are hard-pressed to find true interdisciplinary work significant to both fields. That which bears great psychological weight brings with it little of

practical significance to computer science, whereas that which emphasizes issues in the design of computer systems often lacks psychological validity. One reason for such a bias is that computer science and psychology are two camps which speak a language entirely unto themselves. Crucial to the symbiosis of computer science and psychology is the development of a common language between them. We are not so bold as to suggest the development of such a common language in this thesis. Rather, we hope to demonstrate how a language of design in computer science, in our case derived from the formal notations and methods familiar to software engineering, can and should be directly influenced by the vast body of psychological knowledge on usability.

HCI, as a discipline, has not reached sufficient maturity to enable software engineers to predict the usability consequences of their design decisions. One of the reasons for the lack of applicability of HCI research to real design has been noted by Harrison and Thimbleby [75]:

We believe that much of the early work in HCI has been encumbered by a lack of appropriate abstractness or applicability to the design process.

A formal approach enables development of software whose function can be proved correct and reliable. The ease with which a human user can access this functionality is an important non-functional requirement of the software. As with other requirements on software, mathematical formulation of properties which increase usability would allow greater assurance of software usability earlier in the design process. There have been attempts in the past to show that it is possible to formulate psychologically valid properties of an interactive system in a mathematical language, but only at an abstract and general level [160, 48, 75, 161]. We take the lead from these previous examples and show that it is possible to formulate interactive properties at a more constructive and concrete level of detail which more readily reflects how the interactive system is built from a collection of cooperating objects. Hence, we maintain the ability to discuss interactive properties precisely and formally.

Some researchers have already provided a somewhat more constructive approach to design and analysis of interactive systems [14, 15, 158], but we would like to enhance their work by a stronger link to a methodology linked with empirical psychological evidence of how users perceive the tasks that computers are designed to support. To achieve this, we describe a method of design that is motivated by empirical psychological understanding of interaction. The bulk of HCI formalisms which we will describe in Chapter 3 are supported by psychological theories about how humans understand the tasks they perform on computer. The design of future interactive systems must obey the lessons learned from past mistakes, and they must take advantage of increased psychological evidence to make a good product even better. In order to do that, we need a way to incorporate the psychological evidence

---

in the computational language of design. This cannot be achieved with the abstract and general models of interactive systems mentioned above and further discussed in Chapter 3 because those formalisms do not lend themselves well toward a construction of the interactive systems to embody the usability knowledge. We aim to provide a constructive computational formalism and method which more readily captures interactive properties and knowledge of the user at the source. The formalism we present—the agent model—provides the means for directly capturing empirical evidence, such as given by a task analysis, on interaction for some application domain.

The origins of formal methods are in computer science, so it is a fair criticism that we are biased toward the system side in our HCI research. The bias is intended, as this thesis is to be considered first and foremost as a contribution to the field of computer science, and, more specifically, to the subcomponent concerned with formal methods and its application to the design and analysis of interactive systems. Indeed, it has been precisely the success of formal methods of software engineering in general that tempts us to consider it a worthy candidate for the establishment of engineering for the user interface. Ultimate acceptance of the ideas put forth in this thesis, however, depends upon their validation by the community of psychologists active in HCI research and use by software engineers in the development and analysis of interactive systems.

There are two major contributions arising from this thesis, and they are reflected in the title. To begin with, we are interested in the promotion of formal techniques which can be used in real design. One of the major criticisms of formal techniques arise in an industrial setting where economic forces reign supreme and rigorously principled design practices do not provide ample benefit for the costs they incur. There have been isolated cases where formal development has proven an economic advantage, but for the most part, it is fair to say that formal techniques cost too much for the benefits they are *understood* to provide. We aim to show that the benefits of formal techniques have not been fully appreciated in their application toward more non-functional requirements.

This leads us to the second contribution of this thesis. We have chosen to promote the benefits of a formal approach in HCI, more specifically, in the design and analysis of interactive systems. In attempting to justify the use of formal methods for design in general, we have shown its particular advantage in HCI research where there is great need for a bridge between psychological theory and practical design, a bridge which the formal approach can provide. However, it is not only the case that formal methods can promote HCI research, for the development of the agent language in this thesis a clear case of the converse in which HCI knowledge has influenced the development of a formal approach.

When reading this thesis, it is important to keep in mind its dual purpose. We are interested in formality, but not just for formality's sake. That which we present

formally is motivated by HCI considerations. Furthermore, we are interested in HCI and the considerations which arise in the design of interactive systems, but not just for the sake of HCI. Those principles which we promote to increase the usability of a system are ones which we can capture formally within a rigorous software engineering notation. We feel this focus is necessary to ensure that our research is, as Hoare implores, both scientifically grounded and relevant.

### Overview of thesis

In order to assess the multidisciplinary needs of HCI it is necessary to have an overall view of interaction that is separate from, yet sympathetic towards, both psychology and computer science. In this thesis it is hoped that some bridge may be forged between the research in both disciplines. A suitable introduction, therefore, should provide an overall view of interactive system development and analysis. We refer to this overall view as an *interaction framework*, and its description is the subject of Chapter 2.

The interaction framework provides context for assessing previous research in HCI as well as our own. There have been previous attempts at defining such context, and in Chapter 3 we will discuss these other frameworks and show how our interaction framework has been influenced by them and attempts to extend them. The interaction framework also provides a systematic way of reviewing research in HCI. In our review of Chapter 3, we will highlight the major contributions to HCI research, emphasizing how the psychological content has not fully crept into system design.

Another purpose for the interaction framework is to motivate the kinds of formalisms necessary to express properties of interactive systems. To this end, the framework naturally corresponds to an agent-based modularization of components describing the system, user and interface. Properties that affect the overall usability of an interactive system can be described qualitatively as features of the translations that occur between components of the framework. A formal model of an agent provides ammunition for a first attempt at formalizing those qualitative features of the translations. In Chapter 4, we summarize the qualitative properties of translations in the framework and formalize them in terms of a simple model of the agent, given in Chapter 2.

The simple agent model is inadequate for two reasons. First, it is not rich enough to express interactive properties which relate the goals or end results of interaction to the more immediate perceivable information provided at the interface. Psychological formalisms essentially relate tasks that a system should support to the results that those tasks affect. The simple agent model's inability

to capture constructively result information prevents its use in an overall psychological/computational design method. Secondly, it is not a design notation. We cannot use the simple agent model to describe complex systems by the composition of smaller and simpler subcomponents.

In order to address these two inadequacies, we digress from the specific application of HCI to concentrate attention in Chapters 5 and 6 on the development of a refined agent model and associated language. Chapter 5 defines the new agent model in terms of three perspectives—the internal, state-based specification of the agent, its external event-based specification and a communication specification which links internal operations to external events. The refined model is shown to obey essential compositional properties to allow for a modular design approach. In Chapter 6 we justify the need for a new formal language to describe agents, which is a hybrid notation combining a model-oriented notation similar to Z or VDM and a process algebra notation similar to CSP or CCS, capitalizing on the familiarity and expressiveness of each.

Armed with greater detail of the structure of an agent and a language for describing agents, we resume in Chapter 7 with the formal treatment of interactive system design. The concentration in this chapter is toward showing how notions of result and display allow salient description of interactive properties. In addition, the introduction of templates allows for a method of design specifically geared towards the immediate incorporation of psychological evidence in interactive system descriptions. Result templates embody task analytic information which guide the initial description of a system's functional core. Display templates are chosen to correspond to the result templates and satisfy some interactive properties, such as predictability, honesty, consistency and others.

It is easier to formally express interactive properties when we remain at the abstract and general level. We are then able to ignore the clutter of implementation detail which is not entirely relevant to the expression of the properties. However, we cannot ignore the inevitable refinement towards executable systems, and so we must consider interactive system architectures and how the abstract principles and properties can be mapped into a more realistic design platform. In Chapter 8 we describe the relationship between the abstract level of the interaction framework and the more concrete multiagent architectures that have been used to describe the structure of an interactive implementation. The formal agent model allows us to describe more precisely the features of the heuristically-based multiagent architectures and the design methods they imply in order to assess their value for preserving the properties expressed at the abstract level.

We conclude in Chapter 9 with a summary of the results and contributions of this thesis, along with an agenda for future research.

In addition to the main body of this thesis, we provide three appendices. We make extensive use of the Z notation throughout the thesis, and we will assume familiarity with the standard notation as provided by Spivey [152]. However, for stylistic reasons we have deviated from the standard use in some situations. In addition, it is often the case that we will need to introduce some special notation to make the expression of some predicate more concise. In order not to detract from the flow in the main body of the thesis, in Appendix A we have described our stylistic conventions which deviate from standard Z, along with any notation that is not defined in [152]. Chapter 5 contains the most concentrated sections of formalism in the thesis. We have relegated to Appendix B the proofs of some theorems on the refined agent model, leaving only the outline of their proofs in the thesis body. Finally, in Appendix C we provide greater detail on the semantics for the agent language in terms of the agent model than was deemed appropriate for the body of Chapter 6.

## Chapter 2

# A Framework for Discussing Interaction

We present a general interaction framework which will allow the analysis of the interaction between a user and a system to be expressed in one, unified language. Our intent in this thesis is that the system be some computerized application, but this assumption does not affect the interaction framework. It is also a common interpretation that by distinguishing between user and system we are restricted to single-user applications. This restriction is not an underlying assumption in development of the framework, but rather results from one's interpretation of what system and user represent. The emphasis in the framework is in developing a view of interaction from a single user's perspective. In a multiple user application, such as a multi-party conferencing system, from the point of view of any one user the rest of the users form part of the system.

The interaction framework will be used as a bridging device to provide a common ground for both psychological and computational discussion of interaction and interactive systems. It is important that it be understood at some level by both psychologists and computer scientists. Therefore, it must be free from the jargon of both fields and open to accurate interpretation based on common sense. By making common sense principles explicit within the model we open the path to their automatic inclusion in future design.

### Overview of chapter

This chapter proceeds with an informal description of the major components and translations in the interaction framework in Section 2.1. A simple stimulus-response model of an agent will provide a formal model for the framework in Section 2.2. Both the informal and formal agent descriptions will be used to express properties of interactive systems which attempt to qualify and quantify usability throughout the remainder of this thesis.



## 2.1 An informal definition

The purpose of an interactive system is to aid a user in accomplishing *goals* from some application *domain*. A domain defines an area of expertise and knowledge in some real-world activity. Some examples of domains are graphic design, authoring and process control in a factory. A domain consists of concepts which highlight its important aspects. In a graphic design domain, some of the important concepts are geometric shapes, a drawing surface and a drawing utensil. *Tasks* are operations to manipulate the concepts of a domain. A goal is the desired output from a performed task. For example, one task within the graphic design domain is the construction of a specific geometric shape with particular attributes on the drawing surface. A related goal would be to produce a solid red triangle centred on the canvas.

Our definitions of goal, task and domain generally agree with the approach to problem solving presented by Newell and Simon [119]. We will apply the general term *task analysis* for the identification of the problem space for the user of an interactive system in terms of the domain, goals and tasks. In Chapter 7, we commit to an even more rigid definition of task analysis as a mapping from tasks in the user's goal structure to a set of features or attributes that are intended to represent the action of that task in the system.

The identification of goals and tasks in a problem space is crucial to the work of most analytic approaches to HCI, including our own, for they determine the starting point for analysis of a design, as Lewis points out [101]. There are those who object to the use of goals and tasks as fixed starting points for analysis. Whiteside and Wixon [166] argue that their inclusion should only be as dynamic reflections of broader environmental issues in HCI and not as static entities from which all analysis can proceed with sound grounding. Carroll agrees, pointing out that the usefulness of task analysis to designers is minimized without due attention to such contextual information, usually lacking in most theoretical approaches [35]. A more drastic opinion is put forth by Suchman [157], who states that an understanding of social interaction, not found in existing task analysis techniques, should be the driving force for any theory of single user HCI.

Our belief in the context of this thesis is that an awareness of the goals and tasks of a user in a particular domain form a crucial guide to the assessment of the computer system designed to support the interaction. Therefore, we are not as concerned with the overall theory of cognition and social environment which identifies the goals and tasks as we are concerned with how a definition of task can be used to aid the formal development of the computer system.

Typically, the concepts used in the design of the system and the description of the user are separate; they are considered separate components, and so we refer to them separately as the *System* and the *User*, respectively. The *System* and *User* each have a domain-specific language in which the concepts can be expressed.

These languages treat both *System* and *User* as state machines with operations that can transform the underlying state. The *System's* language we will refer to as the *core language* and the *User's* language we will refer to as the *task language*.

The core language contains system attributes, describing concepts in the domain relevant to the *System* state. The task language contains psychological attributes, describing concepts in the domain relevant to the *User* state. At the most abstract level, both system and psychological attributes are not constrained by any implementation details. Within a formal approach, all that is required is that these attributes be represented by some mathematical object. And once the domain has been adequately captured mathematically, it can be manipulated and reasoned about with rigour.

At the formal and abstract level it is possible to attain a close correspondence between the descriptions of the *System* and *User*, especially in a user-centred design practice which uses the psychological attributes to determine the system attributes of interest. This is the purpose for a task analysis method—to produce some description of the user's understanding of the domain so that a tool can be properly designed for work in the domain. Though there may be a close correspondence between the system and psychological attributes, the *User* does not directly interact with the system attributes. Rather, the interaction is with a representation of the system attributes that is constrained to a far less expressive language of the *physical interface*, exemplified for the most part by two-dimensional displays with primitive sound features and limited tactile facilities.

### 2.1.1 The interactive cycle

The communication between the *User* and the *System* follows a cycle of execution and evaluation, as explained by Norman [123, 124, 125]. The *User* formulates a goal and then must decide the task to perform in order to achieve the goal. The task is executed upon the *System* and the result of the operation is evaluated to see if it agrees with the original goal. This gives four main phases to the interaction—formulation, execution, evaluation and assessment, as shown in Figure 2.1. We will further discuss the execution/evaluation cycle of interaction in the context of previous HCI research in Chapter 3.

Since the result of user-centred design as described above is that the task language of the *User* and the core language of the *System* are closely related, interaction between *User* and *System* is fairly straightforward, since the translation between the two languages can and should be trivial. The simplicity of interaction implied by this close correspondence between the abstract *System* and *User*—what the *User* wants to do the *System* can do—is misleading, because it is often the case that there is a mismatch between the *User's* high-level task language and the low-level entities of the physical interface which the *User* must manipulate in order




Figure 2.1: Four phases of interaction between *User* (U) and *System* (S)

to achieve the desired goal. There is also a mismatch between the task language and that which the *User* interprets from observations of the physical interface when determining if the goal has been achieved. These two mismatches are referred to, respectively, as the gaps of execution and evaluation by Hinchins, Hollan and Norman [85].

### 2.1.2 The components of the framework

In order to attain a more realistic description of interaction, therefore, we break down the interaction between user and machine into four main components, as shown in Figure 2.2. The nodes represent the four major components in an inter-




Figure 2.2: The general interaction framework

active system—the *System*, the *User*, the *Input* and the *Output*. Each component has its own language which is used to express its purpose in the interaction. In addition to the *User*'s task language and the *System*'s core language which we have already introduced, there are languages for both the *Input* and *Output* components

to represent those separate, though possibly overlapping components. *Input* and *Output* together form the system interface. Note that we distinguish between the system interface and the physical interface. The physical interface is that part of the system which, as its name suggests, is in direct contact with the user in the physical world. Therefore, the physical interface is viewed as a subset of the interface in our framework. The input and output languages do not in most cases map very directly onto concepts in the domain. Yet, the interface's position between *System* and *User* mandates that it be an effective mediator for the tasks in the domain of the application. Therein lies the major challenge in interactive system design.

As the interface sits between the *User* and the *System*, there are four steps in the interactive cycle, each corresponding to a translation from one component to another, as shown by the labelled arcs in Figure 2.3. The *User* begins the




Figure 2.3: Translations between components

interactive cycle with the formulation of a goal and task to achieve that goal. The only way way the user can manipulate the machine is through the *Input*, and so the task must be articulated within the input language. The input language is translated into the core language as operations to be performed by the *System*. The *System* then transforms itself as described by the operation translated from the *Input*; the execution phase of the cycle is complete and the evaluation phase now begins. The *System* is in a new state, which must now be communicated to the *User*. The current values of system attributes are rendered as concepts or features of the *Output*. It is then up to the *User* to observe the *Output* and assess the results of the interaction relative to the original goal, ending the evaluation phase and, hence, interactive cycle.

It is easiest to think of the interactive cycle as a true alternation between execution and evaluation, but this is not always the case. Every action by the *User* may not be followed by an evaluation, and it is very possible that the *User* will be required to observe changes to the *System* that were not directly prompted by

actions performed by the *User*. Therefore, a strict interpretation of the interaction framework in terms of an alternating execution/evaluation cycle is not intended in our presentation. However, for explanatory purposes, this interpretation is not overly harmful.

## 2.2 A formal definition

In the previous section, we presented an informal introduction to the interaction framework. In this section, we will provide a formal definition of the framework. The formal model of the interaction framework provides the foundation for the work of this thesis, motivating the formulation of abstract properties of interaction in Chapter 4, the agent model of Chapter 5 and the more constructive interactive properties of Chapter 7. We view an interactive system as a collection of communicating interactive *agents*. This section proceeds by giving a simple formal definition of an agent and a description of each of the components in the interaction framework in terms of agents. We then combine the different agents for a description of a complete interactive system as suggested by the interaction framework and translations between components suggested by Figure 2.3.

### 2.2.1 A simple definition of an agent

The description of an agent serves two purposes—to give its current state and to describe how that state is transformed as its interaction with other agents proceeds. We make a distinction, therefore, between the *state* of the agent and the *events* in which it participates in cooperating with other agents. We present two given sets to represent the set of all possible states of an agent, *State*, and the set of all possible events an agent can participate in, *Event*.

$$\{State, Event\}$$

Further details of the state and event sets will be delayed until the refinement of the agent model described in Chapter 5.

The link between the state of an agent and the events in which it participates is given by the agent's *behaviour*. An agent is based on the stimulus-response model, which has been argued to form part of the user's and designer's model of an interactive system [41]. An agent participates in a stimulus event which triggers a change in the internal state. After the transition, the agent responds with events which will affect other agents in the system. From this description, we choose to model the agent's behaviour as a relation between stimulus-state pairs and response-state pairs. In a given state, receipt of a single stimulus will result in a new state and a sequence of observed response events. Two views of this behaviour

concentrate on different aspects of it. One view gives the state transformation triggered by a stimulus event. The other view gives the pairing between stimuli and responses. This stimulus-response information is an external description of an agent, whereas the state transformation mapping gives internal information on the agent. The simple formal agent model is given by the schema *Agent* below

$  \begin{array}{l}  \textit{Agent} \\  \textit{states} : \mathbf{P} \textit{ State} \\  \textit{inits} : \mathbf{P} \textit{ states} \\  \mathcal{B} : (\textit{Event} \times \textit{states}) \leftrightarrow (\textit{seq Event} \times \textit{states}) \\  \textit{transform} : \textit{Event} \rightarrow (\textit{states} \leftrightarrow \textit{states}) \\  \textit{stimresp} : \textit{Event} \leftrightarrow \textit{seq Event} \\  \hline  \forall \textit{stim} : \textit{Event}; \textit{resp} : \textit{seq Event}; s, s' : \textit{states} \\  \bullet ((\textit{stim}, s), (\textit{resp}, s')) \in \mathcal{B} \Leftrightarrow ( (\textit{stim}, (s, s')) \in \textit{transform} \\  \qquad \qquad \qquad \wedge (\textit{stim}, \textit{resp}) \in \textit{stimresp} )  \end{array}  $
--

We can define two interpretation relations on sequences of stimulus events, or *programs*. The first,  $I_{-}^{\text{int}}$ , gives the internal interpretation of the program, i.e., the possible state that the agent can be in after participating in the program. Each stimulus event in the program represents a state transition relation, and so the overall state transition relation for the program is the sequential composition of the individual state transition relations. Restricting the domain of this overall transition relation to the initial states of the agent ( $A.\textit{inits}$ ) gives the possible transitions for the agent.

$$\begin{array}{l}
 \textit{transformExtend}_{-} : \textit{Agent} \rightarrow \textit{seq Event} \rightarrow (\textit{State} \leftrightarrow \textit{State}) \\
 \textit{transformExtend}_{A}(\textit{stims}) = (A.\textit{inits}) \triangleleft (\% / (\textit{stims} \% A.\textit{transform}))
 \end{array}$$

The internal interpretation function relates the program to the possible final states.

$$\begin{array}{l}
 I_{-}^{\text{int}} : \textit{Agent} \rightarrow \textit{seq Event} \leftrightarrow \textit{State} \\
 (\textit{stims}, s) \in I_{A}^{\text{int}} \Leftrightarrow \\
 \exists s_0 : A.\textit{inits} \bullet (\textit{stims}, (s_0, s)) \in \textit{transformExtend}_{A}
 \end{array}$$

After each input event, the agent responds with a sequence of response events, as given by *stimresp*. A program of inputs, *prog*, is therefore related to a sequence of responses, *respobs*, derived from the responses of each input event in *prog*. The second interpretation relation,  $I_{-}^{\text{ext}}$ , provides this derived overall stimulus-response behaviour.

$$\begin{array}{|l}
I_A^{ext} : Agent \rightarrow seq\ Event \leftrightarrow seq\ Event \\
\hline
\text{dom } I_A^{ext} = \text{dom } I_A^{int} \\
(prog, respobs) \in I_A^{ext} \leftrightarrow \\
\exists respseq : seq\ seq\ Event \\
| ( \#respseq = \#prog \\
\wedge \neg /respseq = respobs) \\
\bullet \forall i : 1.. \#prog \\
\bullet (prog(i), respseq(i)) \in A.stimresp)
\end{array}$$

Note that captured in this definition is the assumption that all of the responses for a given stimulus event are observed before any responses owing to subsequent stimuli. This assumption is not mandatory. We could have replaced the predicate  $\neg /respseq = respobs$  with one such as  $(respseq, respobs) \in interleaves$ , with *interleaves* as defined in Appendix A, or even more sophisticated expressions, to reflect the more general possibility that the order of responses does not fully respect the order of the stimuli. We will disregard that possibility in this thesis, for it is believed to bring more tedium than enlightenment.

### 2.2.2 An agent description of an interactive system

Our initial understanding of an interactive system suggests that it is composed of two agents, representing the *User* and the *System*. In addition, there are translations between the stimuli of one and the responses of the other, which represent the externalized aspects of the execution and evaluation phases. These translations are formalized as relations between the stimuli and responses of the agents. The *execution* translation is a relation between the responses of the *User* agent and the stimuli of the *System* agent. The *evaluation* translation is a relation between the responses of the *System* agent and the stimuli of the *User* agent. A simple interactive system is defined below in terms of the *User* and *System* agents and the translations between them.

$$\begin{array}{|l}
SimpleIntSys \\
\hline
User, System : Agent \\
execution, evaluation : seq\ Event \leftrightarrow seq\ Event \\
\hline
execution \in (ran\ I_{User}^{ext} \leftrightarrow dom\ I_{System}^{ext}) \\
evaluation \in (ran\ I_{System}^{ext} \leftrightarrow dom\ I_{User}^{ext})
\end{array}$$

The translations *execution* and *evaluation* are relations between event sequences, similar to the external interpretation relation for an agent. We can therefore regard *execution* and *evaluation* as specifications of the overall stimulus-response behaviour of agents between *System* and *User*. Constraints on these translations will

be constraints that must be satisfied by the agent which manifests that specification. Therefore, the description of a simple interactive system contains information on four agents, two explicit (*User* and *System*) and two implicit (*execution* and *evaluation*).

The complete interaction framework builds from this simple definition by adding the *Input* and *Output* agents along with translation relations labelled as those in Figure 2.3. The *articulation* and *performance* relations are composed to yield the *execution* translation and the *presentation* and *observation* relations are composed to yield the *evaluation* translation. As was the case for *execution* and *evaluation*, these translations provide specifications for additional implicit agents in the system.

<p><i>InteractionFramework</i></p> <p><i>SimpleIntSys</i></p> <p><i>Input, Output : Agent</i></p> <p><i>articulation, performance,</i></p> <p><i>presentation, observation : seq Event ↔ seq Event</i></p> <p><i>execution = articulation ; I<sub>Input</sub><sup>ext</sup> ; performance</i></p> <p><i>evaluation = presentation ; I<sub>Output</sub><sup>ext</sup> ; observation</i></p> <p><i>articulation ∈ (ran I<sub>User</sub><sup>ext</sup> ↔ dom I<sub>Input</sub><sup>ext</sup>)</i></p> <p><i>performance ∈ (ran I<sub>Input</sub><sup>ext</sup> ↔ dom I<sub>System</sub><sup>ext</sup>)</i></p> <p><i>presentation ∈ (ran I<sub>System</sub><sup>ext</sup> ↔ dom I<sub>Output</sub><sup>ext</sup>)</i></p> <p><i>observation ∈ (ran I<sub>Output</sub><sup>ext</sup> ↔ dom I<sub>User</sub><sup>ext</sup>)</i></p>
--

## 2.3 Conclusions

From our above discussion, we can see that the interaction framework contains at least partial information on eight different agents—the four major components of the *User*, *Input*, *System* and *Output*, and external specifications on four agents which represent the translations between the stimuli and responses of the major components. Since agents are intended as a compositional and constructive model for an interactive system, the actual agent description of a complete interactive system will contain many separate agents. Our point in this thesis is that we can view everything in the interactive system as an agent. Therefore, properties of interaction which we will present in the remainder of this chapter and throughout the thesis can be expressed as properties on agents.

The justification for our division of an interactive system into four major components and four translations between them is the subject of the next chapter, in



which we will use the framework to establish the context of previous HCI research and establish the further contributions of this thesis for HCI.

# Chapter 3

## Background and Related Work

The last chapter provides context for the application of formal methods in HCI by describing a general interaction framework. In this chapter, we will review previous research into establishing the general context of HCI research and formalizing different aspects of interaction. Before launching into the review, we will highlight two of the major conclusions which the review supports.

The first conclusion is that a unified framework, such as presented in Chapter 2, helps determine what psychological information is available to feed into design. Though there is a basic divide between the precision with which reasoning is possible on the two sides of interaction—the human (or user) and the computer (or system)—it is valuable and instructive to view both within the same formal (and informal) framework. The system side deals with objects that can be quantified and reasoned about mathematically. This is the premise upon which the application of formal methods in software engineering is based. The system attributes mentioned in Chapter 2 which are used to describe domain concepts are intended to have executable realizations in refined versions of the system. They may initially be presented as abstract concepts, but the whole purpose of refinement work is to realize an abstraction in concrete detail while preserving the properties of the abstraction.

On the user side, the psychological attributes represent attempts to describe phenomena whose very existence is itself a research question. Though several formalisms exist which provide quantitative predictive power for analysis of the user's side of interaction, the conclusions they provide are questionable from both a psychological and design perspective. However, empirical and theoretical psychological evidence is able to support some assumptions that we can make about the human as user which we can then incorporate into our design process. Incorporating these assumptions about the user explicitly in the design process allows their removal or alteration if they are found to be invalid.

From the design perspective, there is a crucial symbiotic relationship between

system and psychological formalisms. The empirical psychological evidence for interactive behaviour provides data which the system formalism can manipulate. The interactive design method requires both formalisms; hence the need for one framework which unifies the two.

A second conclusion supported is that the bridge between psychological and computer science research is not heavily travelled by researchers in formal methods and software engineering. Though the vast majority of software developed is interactive and could therefore be aided by a rigorous theory of interactive system design, very few designers are engaged in such principled design. The majority of literature in software engineering that covers interactive system development relies on heuristic reasoning about good design. Though a major reason for this is the lack of acceptance of formal notations in general in design practice [44], there is little work done under the name of formal methods which gives fair notice to the consequences of including user considerations in design or specification. We see two reasons for this. First, most of the formalisms which have been offered come from researchers who are mainly psychologists and not, therefore, concerned directly with the design implications of their formalisms. Second, many of the formal techniques available and in use in industry do not provide enough descriptive power to naturally express an interactive system in the way the designer (and the user) perceives it. This last topic is a major consideration in our development of the agent model and its associated specification language in Chapters 5 and 6.

As a result of these conclusions, we can see the purpose of this thesis. We aim to provide a theory of interactive system design that both makes its psychological assumptions explicit for means of validation and is within the grasp of the software engineering profession (or at least practitioners of formal methods within software engineering). By building on and extending previous research in both HCI and formal methods, we present a method for interactive system design which unifies previously separated considerations about the user, the system and the interface which separates them.

### Overview of chapter

In Section 3.1, we will relate the interaction framework of Chapter 2 to previous attempts to define the context of HCI research. Our interaction framework is not the first attempt at breaking up the interaction between a user and a computer into stages. We will present some other frameworks for interaction that predate and influence our own, explaining how our interaction framework extends their work. In Section 3.2, we will present a survey of the research applied to various aspects of HCI. In this section, we will describe in more detail some of the formal and informal research which presents a more narrow focus than the frameworks discussed in Section 3.1. We have tried to classify the different approaches according

to how they fit into our interaction framework.

## 3.1 Other interaction frameworks

We identify four major categories of general and informal theories that have inspired our framework, and we will discuss each in this section. The purpose of frameworks, such as ours, is “not to reveal dramatic new truths,” as Norman points out [123]. Rather, the purpose is to provide insight into the implications each stage within the framework has on the design of interactive systems.

We will emphasize how each of the frameworks below compares to our framework. We hope this serves as a suitable justification for the introduction of yet another view of interaction which will guide the formal approach of the remainder of this thesis.

### 3.1.1 Evaluation/Execution Cycle

Probably the most obvious influence on our interaction framework has come from the execution/evaluation cycle of interaction. This view of interaction is made explicit in much of the literature on HCI and it is implicit in nearly everyone’s common sense understanding of the interaction between human and computer. The human user formulates a plan of action which is then executed at the computer interface. Upon completing the execution of some plan, or part thereof, the user observes the computer interface to evaluate the result of the recently executed plan and to determine the further course of action. In Chapter 2, we acknowledged the seminal work of Norman [124]. A similar division of the interaction cycle has been made explicit by Card, Moran and Newell [32], and we will examine their contribution further in Section 3.2 as it is more detailed in the formalism which it presents. The work by Norman which we reference is mainly qualitative, and so most resembles the presentation of our interaction framework. His views of the interaction between user and computer have been criticized as too simplistic, but we view the real value of his views in the direct appeal to common sense. Arguments that it over-simplifies a complex topic fail to see its purpose as a readily understandable overview of human-computer interaction accessible to those lacking a formal psychological education. Consequently, Norman’s model is accessible to the computer scientist interested in designing a more usable system. Unfortunately, Norman’s model does not consider the system’s contribution to the execution and evaluation cycle as much as the user’s contribution. Our interaction framework is intended to address this disparity.

Norman initially outlined four stages of the user’s activities—intention, selection, execution and evaluation [123]. In later work [124, 125, 85], the interactive cycle can be seen as divided into two major phases, execution and evaluation. Each

of these phases is then subdivided further into different stages of the interaction that can be examined for their particular influence on the effectiveness of the overall interaction. The seven stages mentioned are [124, page 41]:

1. Establishing the Goal
2. Forming the Intention
3. Specifying the Action Sequence
4. Executing the Action
5. Perceiving the System State
6. Interpreting the State
7. Evaluating the System State with respect to the Goals and Intentions

These stages are further related by Norman, as he symmetrically divides the previous execution/evaluation cycle. Hence, perceiving the system state is seen to be the evaluative equivalent of executing the action, interpreting the state is the evaluative equivalent of specifying the action sequence and evaluating the system state with respect to the goals and intentions is the evaluative equivalent of forming the intention.

Figure 3.1 portrays the relationship between the execution/evaluation cycle and our interaction framework. The obvious comment, as we have already mentioned, is that the execution/evaluation cycle does not consider the system beyond its interface (hence the shaded left-hand side in Figure 3.1). Norman simply represents the system as the world of physical activity, analysis of which stops at the physical interface. There is much greater detail on the user's side of the interface, as the translation stages from *User* to *Input* (articulation) and from *Output* to *User* (observation) in our framework are each further divided into three substages in Norman's model.

Norman's model serves two purposes. It provides the first step needed to introduce a computer scientist to the purpose of psychological work in assessing the usability of interactive systems. It also provides an outline for previous and future theoretical or empirical research by psychologists trying to describe how the user interacts with the system and how that interaction can be assessed. This interaction is divided into two phases. One is concerned with the formulation of a plan and its performance. The other phase is concerned with the observation of the results of previously performed plans and their assessment with respect to the original plan. As a rule, formal psychological research into human-computer interaction, therefore, is divided roughly into two areas, one to address the cognitive aspects of formulation and execution of the plan and one to address the cognitive aspects of

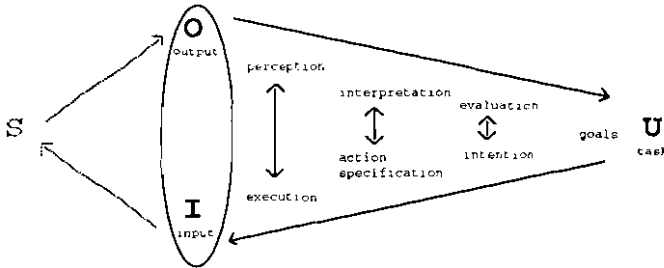


Figure 3.1: Relationship between Norman's execution/evaluation cycle and the interaction framework

perceiving and assessing the consequences of the executed plan [30]. One notable exception to this rule is found in the research by Barnard on Interacting Cognitive Subsystems (ICS) [17, 18], which is a model of human cognition and performance that can address both evaluation and execution in the context of human-computer interaction.

Our interaction framework, therefore, is an attempt to extend Norman's model with a necessary and complementary component which more fully addresses issues of the system. The extension provides a first step needed to introduce a psychologist to the needs of computer science in designing interactive systems. In addition it provides a platform for assessing previous system-based work on interactive systems and provides direction for future formal research on interactive system design.

Norman's lucid account of the execution/evaluation cycle provides a clear insight into what factors affect the translation between languages between the *User* and the interface (*Input* and *Output*). We will discuss in Section 3.2 how Norman's qualitative account can assist in assessment of more formal accounts of these translations. The addition of the *System* and its relationship to the interface in our interaction framework is viewed as a necessary and complementary view that enhances Norman's model. With proper attention focused on the system side, we can also assess previous work on interactive system development and propose how a formal approach can better aid the precise development of more usable systems.

### 3.1.2 Interaction Modelling Framework

Barnard and Harrison [20, 19] have also presented an interaction framework but for a different purpose than either our framework or the execution/evaluation model. The purpose of their framework is to direct research on incorporating distinct models of the system and user, which already exist, by means of a separate and new interaction model. It is for this reason that we distinguish this framework from

ours by referring to the Barnard and Harrison version as an interaction modelling framework. Figure 3.2 is a graphical representation of their framework, taken from [19].

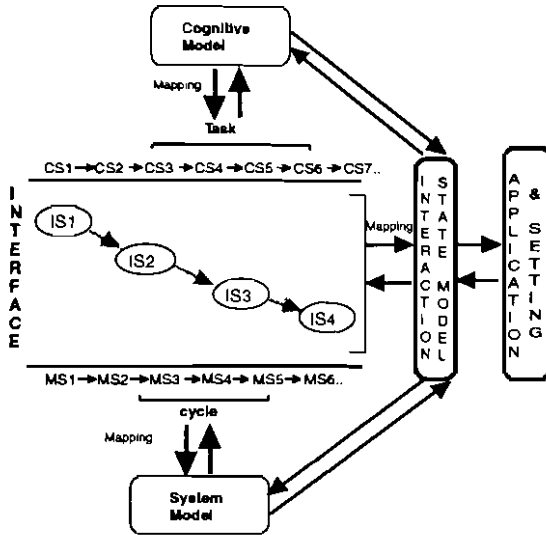


Figure 3.2: Barnard and Harrison's interaction framework.

Barnard and Harrison identify some major problems of system and user models of interactive behaviour. System models make implicit psychological assumptions about the user that can be invalid. For example, the PIE model version of predictability [54, 48], which we will discuss later, makes the assumption that a system designer can *a priori* determine what effects will be perceived and understood by the user in terms of their task language. User models make implicit assumptions about the practice of system design that are impractical. For example, a psychologically valid notion of consistency of an input languages is presented by Payne and Green's TAG notation [127, 128], but their device representation does not consider that system design must take into account more than just input language consistency. It is not so much a problem that the separate modelling domains make assumptions about the other. The problem is when the modelling technique depends on implicit assumptions that it cannot change. In an ideal situation, each modelling domain (via the modellers) must be able to validate the assumptions made by the other modelling domain and also incorporate the results of the other domain within its

own work. How can this be achieved?

In Barnard and Harrison's view, there is a common feature that can be extracted from both models of the system and models of the user. Each can describe its subject in terms of a state-based machine that undergoes transitions. The system states are *machine states*, and the user states are *cognitive states*. Any instance of interaction between user and system results in a sequence of transitions in both the machine and the user. Figure 3.2 depicts these transitions by the sequence of machine states  $MS_1, MS_2, \dots$  and cognitive states  $CS_1, CS_2, \dots$ . There is no clear link made between machine state  $MS_i$  and cognitive state  $CS_i$ . The intention of the mappings from models to state transitions in Figure 3.2 is to show that the system and user models are not wholly determined by what they say about state transitions. Rather, those models each say something about the behaviour of the machine or cognitive state of the user that can be mapped onto state transitions. In turn, particular instances of the machine or user behaviour can be mapped into the respective model for interpretation relative to abstractions within the model, as we will discuss next.

Within a system or user model abstractions can be defined in terms of the transitions of the machine or cognitive state. For example, there is the notion of cycle within the system model and the notion of task in the user model. In trying to establish the relationship between the abstractions made by the system modellers and those of the user modellers, they suggest a third model, called the interaction model. The reasoning is that any relationship between system and user modelling concepts would only be relevant where the two models meet, that is, at the human-computer interface. Therefore, a model of the interface would be the suitable location for the cross-fertilization of the different concepts of system and user modelling without undue emphasis on either side of the interface. This interaction model would also have an interpretation based on states, labelled by 'interaction states', and instances of transitions would be labelled  $IS_1, IS_2, \dots$ , as in Figure 3.2. The interaction model is fed by information about system and user models but in addition can also embody explicitly notions that are usually lacking within the system and user models, such as the domain knowledge.

Finlay *et al.* have attempted to apply the interaction modelling framework to the analysis of exemplar systems, such as an automated teller machine, by making explicit an event structure used to link the system and cognitive models [57]. The significance of the events used by Finlay *et al.* is that they can both be handled by the respective system and cognitive modellers. Events are labels for operations or state transitions. At the interface between user and computer, the same event can be linked to operations on both the user and the system. On the system side, events are similar to those described in Chapter 2 for the simple formal model of an agent. On the user side, Barnard's own ICS model is used, and there is a much less convincing connection between events and the ICS way of modelling



the user's cognitive states and operations performed upon them. This points up not a theoretical failing of ICS but rather its inability in its present form of being expressed in a formal notation similar to that used on the system side. More recent work by Barnard and Harrison [22] has concentrated on providing a uniform description of both user and system model structures.

There is a stronger similarity between the interaction modelling framework and the refinement of our own framework introduced in Section 2.2. In that section, we begin to formalize the concepts of the interaction framework. Our refined interaction framework will address directly the problem discussed above of the linking work done by Fiulay *et al.*, because we will provide a common notation for both the user and system models. What would remain is to express a cognitive model such as ICS within that notation.

### 3.1.3 Black box models

A major difference between how an interactive system is viewed by a designer as opposed to a user is that the designer knows about the intricate detail of the computer system and the user does not. A result of this knowledge deficiency is that the user attempts to formulate a model of how the system works. This model is in part determined by the experience the user has in interacting with the system. In these interactions, all that a user will know about the system will come from all of the perceivable information which the system presents. A principle of user-centred design is that the designer try to match the user's perception of the interaction. This principle justifies to some extent why psychologists focus their analysis on only that part of the system which is observable, since that is the only information that the system provides the user. But in attempting to think as the user, it is often hard for the designer to be divorced from the intricate system detail.

This knowledge discrepancy causes a problem with the principles of user-centred design because the designer knows too much about how a system functions in order to objectively judge such properties as predictability as the user would perceive them. One solution is to analyze the system by forgetting (temporarily) the internal details. Essentially, a user views the computer system as a *black box*, noticing only that which is presented as input to the system and that which is produced as the effect inspired by the input. If the designer also adopted this black box view of the system, then a fairer assessment of the interactive properties of the system could then proceed.

A more important use of a black box model would precede a full specification and implementation. In this case general domain-independent principles can be used to guide the development. Dix *et al.*[52] show how such principles can be used to guide the design of a small programming environment. Monk and Dix [113] use a semi-formal black box model to examine how action-effect rules can provide insight

into the application of general principles of predictability, simplicity, consistency and reversibility within design.

As an aside, these approaches to specification and design suggest a revised software lifecycle, as presented by Harrison [71], in which a model of interaction, of which a black box model is but one example, is used as a link from the requirements phase to the specification phase.

*requirements* → *interaction model* → *specification* → *implementation*

The notion of a black box is quite a common one in software engineering and system development; it appears under the name of data abstraction and the abstract data type [102] or information hiding [126]. It is also quite common in other sciences; for example, the transfer function in control theory. The classic example of a black box approach to interactive system analysis is the *PIE model*, first presented by Dix and Runciman [54], and later greatly expounded upon by Dix [48, 49]. The model considers user input to the system as coming from a set of programs  $P$  and output as being from a set of effects  $E$ . The system is modelled as a black box whose entire functionality is described by an interpretation function  $i$  which takes programs to their effects. Figure 3.3 shows the original PIE model.




Figure 3.3: The PIE model.

The simple PIE model can be refined to distinguish between two separate features of the effect space—the *result* and the *display*. The result deals with the final end-product of the interaction between human and computer, whereas the display deals with the intermediate and ephemeral aspects of the effects. Important properties of the interaction can be characterized as relationships between the display—which the user usually can perceive directly and continually—and the result—which is of most importance to the user but is not always directly nor continually perceivable. In fact, one of the most common heuristic guidelines for interactive system design is termed *WYSIWYG* (“What you see is what you get”), which can be deciphered as a comparison between the instantaneous display and the end result of the interaction. Figure 3.4 depicts this refinement, aptly named the *red-PIE model*. We will discuss the red-PIE model in more detail in Chapter 7.

An important feature of the PIE models is the flexibility of the effect space as given by this definition from Dix [48, p. 40]:




Figure 3.4: The red-PIE model.

**E** - the effects space. The set of all possible effects the system can have on the user. This may be thought of in different ways, and at different levels. For example, it may be regarded as the actual display seen by the user, or as the entire text of a document being edited, perhaps even the entire store of information available to the user.

The program set is also flexible. It may be regarded as the keystrokes and mouse movements, or as more abstract domain-specific operations for the system or the user. Within our interaction framework, this would mean that particular effect and program spaces lie on a continuum between the *System* and the *User*, as depicted in Figure 3.5. With such a flexible definition of the effects and programs spaces,




Figure 3.5: The effects space of the PIE model within the interaction framework.

it is possible to write many interactive properties and pitch them at varying levels of generality. The statement of properties within the PIE model have increasing psychological significance as both effects and programs move towards the *User*. Unfortunately, the increased psychological significance is accompanied by a decreased design significance, as the constructs bear less of a resemblance to constructs which can be represented in a real design.

There is a close connection between the PIE model and the simple agent model defined in Section 2.2. We have defined the internal interpretation relation,  $I_{int}^{-1}$ , which links sequences of stimulus events to achievable states. The sequences of stimulus events are the basic constructs of the program set. The state description of the agent is the definition of the effect set. Therefore,  $I_{int}^{-1}$  is very similar to the interpretation mapping of the PIE model. The importance of this connection between the PIE model and agents is that we are able to incorporate all of the work on PIE models which express interactive properties into our work with agents. In Chapter 4, we will give examples of how some classes of abstract interactive properties first expressed in terms of the PIE model can be expressed within the agent model.

The PIE model demonstrates how a black box model can be used to formalize interactive properties. It is important to highlight some of the problems as well as the advantages of the black box model. The abstractness of a black box model has several major advantages for use in interactive system design and analysis. It is a very simple model that is uncluttered by the details of system implementation. This makes it possible to express properties of interaction that are implementation and domain independent. It can be represented mathematically, so the statements of properties that it allows are precise and they provide a possibility for proof. As shown by Monk and Dix [113] it can also be used as a semi-formal aid to design. An initial encroachment of valid psychological assumptions into the system modelling practice can be achieved with the black box model. This marks the beginning of a bridge between psychological theory and software engineering.

Most of the drawbacks of the black box model arise because it is so abstract. Formal techniques have been criticized for being too far removed from the practice of software engineers designing real systems. The techniques that the black box model present are not constructive, that is, they do not provide enough support for practising designers in going from an initial specification to an implementation. As far as this criticism is concerned, the black box is a move in the wrong direction. Nobody would want to specify an entire system in terms of a PIE! And though a black box model provides the ability to precisely express properties of interaction at various levels, depending on the meaning attached to the effects space  $E$ , proof of these properties in a real specification would be a nightmare. The black box model has no way of managing the proof obligations of large specifications because it makes no attempt to modularize the description. It is in order to address these criticisms that we provide the more constructive theory of agents. We admit that having provided a more constructive theory we are still far away from providing a compositional proof system as we would like. However, the refined agent model is a move in the right direction.

But there are more serious psychological criticisms of the approach that underlies models such as the PIE. A predictability property, which mathematically

states that the result of future input to the system is determined by the current effect presented to the user, claims to capture a psychologically significant feature of an interactive system at a very high level. But predictability is not all that is important for determining the usability of an interactive system. In Chapters 4 and 7, we will provide other examples of general interactive properties relevant at varying levels of detail as well, but the question remains as to whether a 'complete' set of principles exists such that satisfaction of those principles guarantees a usable system. Our formalisms do not address this issue of completeness.

There is skepticism that the mathematically formulated properties capture the intent of their psychological cousins. This is a general criticism concerning requirements capture. Dix and Harrison [51] speak about the 'formality gap' that exists between requirement, usually couched in natural language terms, and formal specification, expressed using mathematics. The purpose of the interaction models, such as a black box model like the PIE, is to partially bridge this gap by allowing a formalism specifically geared towards psychological or HCI issues that may arise in the requirements. But the important point is that the formality gap can never be fully bridged because there will always be a translation from natural language to the formalism. This then demands that the translation be made as readily understandable as possible to the one who must verify its correctness. Since someone with psychological insight would best be suited to verify psychologically grounded requirements for usability, the interaction models must be geared to their understanding.

The black box model assumes that the relevant properties for interactive design can be expressed from first principles for any system. This assumption is reinforced by the placement of the interaction modelling component in the revised lifecycle mentioned above. There are two criticisms of this assumption. First, some believe as Carroll and Rosson [36] that design is essentially empirical 'not because we don't know enough yet, but because in a design domain we can never know enough.' Another somewhat related issue is that expressing general and abstract properties at a high level of design assumes there are features of an interactive system which can be factored out in determining the system's usability. Carroll criticizes this view by noting, along with others, that there is potentially 'infinite detail' at many levels of description in an interactive system, all of which plays a critical role in determining the usability of the system [35]. These views should not be considered a refutation of the value of a formal 'get it right the first time' approach. Rather, they should be considered as a warning that a formal approach alone is not sufficient for the design of usable interactive systems and it should be complemented by other techniques, such as rapid prototyping and empirical evaluation [7].

There are those who believe that an iterative approach to design with fast and easy prototyping is the only way to create usable systems. Lewis classifies these people in the "process is paramount" category of his partitioning of HCI researchers

[101]. However, initial design decisions (step 0 in the iterative approach) greatly affect the final result and the number of iterations necessary to get there. So it is advocated that some semi-formal reasoning be available to the designer in order to make good first guesses at the design based on certain general principles.<sup>1</sup> A black box model of interaction can allow greater confidence in initial design decisions because the properties it best supports are very general and most easily dealt with at an abstract level, which is where most design begins. And placed in a more formal design process, the assurances of step 0 design decisions can be made with greater confidence.

In summary, the importance of the PIE model is in its bridging of the 'psychological gap' in HCI. The PIE model formed the initial inspiration for the work by Sufrin and He [158] in describing interactive processes in a formal notation for subsequent analysis. This work in turn inspired the agent theory of this thesis. Barnard and Harrison's arguments motivating an interaction modelling framework suggests that there is no existing model of the interface and that all modelling techniques should fit squarely into either category of system or user model. We question that assumption, since the black box model directly discusses the model of the interface, though it does not go so far as to propose one interaction state, preferring instead to allow levels of abstraction between the system and the user, all of which represent some kind of interface. The very notion of an interaction state runs counter to the intention of the black box model because it is precisely the abstraction away from underlying state that produces the black box.

### 3.1.4 Software architectural models

The emphasis in this thesis is on the incorporation of psychologically valid claims about human-computer interaction into the design of interactive systems, *by means of a formal approach*. There is, however, another way to incorporate psychological principles into design, and it is by far a more common approach. That would be by a *heuristic* approach in which reasoning based on experience and empirical evidence drives the common sense applied in the design process. Figure 3.6 demonstrates the different reasoning directions one can take in applying psychological knowledge of HCI issues to the design of interactive systems.

We assume that both formal and informal advocates have access to the same body of psychological knowledge about human-computer interaction. This knowledge embodies information which describes what is meant by effective interaction between user and system. The advocates of formal methods try to capture this information in the same mathematical language which is used to describe the system. The aim of this activity is to produce a mathematically constructed usable

---

<sup>1</sup>Lewis also makes this point, but adds that "process is paramount" people doubt the contributions that formal models can provide in attaining best first guesses.

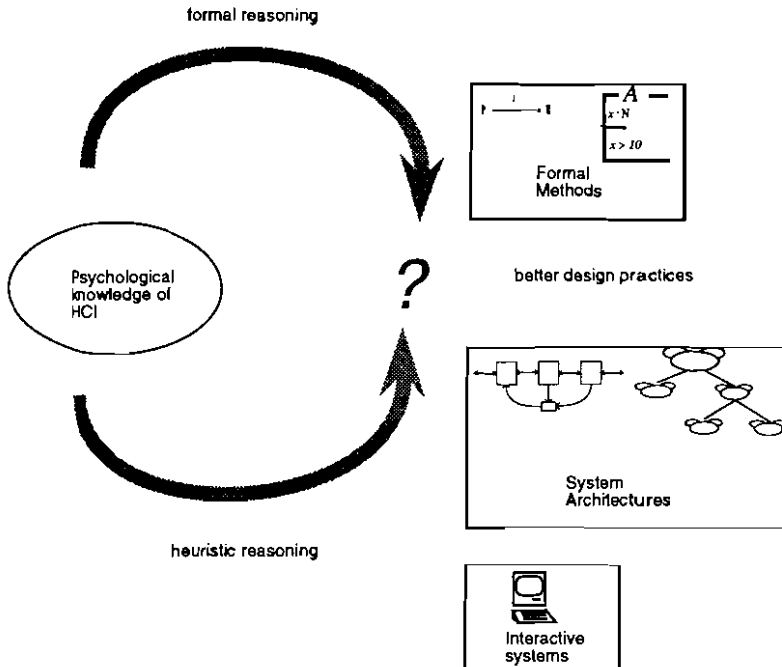


Figure 3.6: Heuristic versus formal reasoning within HCI

system which can then be refined into an executable version which preserves the usability properties. The greater majority of researchers and practitioners, take a far more pragmatic approach. They use common sense in reasoning backwards from the myriad of examples of interactive systems to arrive at abstractions from design that capture why it is that some systems are usable and others are not. In Lewis' taxonomy, these are the "systems are paramount" researchers in HCI.

Some of the abstractions provided by pragmatists we classify as interactive architectures which are themselves expressions of frameworks for interaction. We can further subdivide this class into two (somewhat overlapping) subclasses. One subclass deals with the significant contributions of the User Interface Management System, or UIMS. The other subclass deals with the more recent emergence of multiagent architectures.

An important feature of both subclasses is that they are mainly concerned with

design of the user interface. Restriction to the details of the user interface assumes the existence of the underlying application. There is an enormous economic argument in favour of this approach because it recognizes that there are many competent computerized systems in existence whose major failing is not in performing functions relevant to its domain but in providing an interface through which users can easily access that functionality. Interactive architectures have been developed with the understanding that the user interface can and should be a separate module of the interactive system. In practice, however, such a clean separation is not easy. Nor is it clear that such a separation is wholly desirable.

We contrast this with the theme of this thesis, which is interactive system development. We do not assume the existence of the underlying application in our formal analysis. On the contrary, we depend upon a system core that has been properly and formally specified in accordance with a task analysis on the domain of the application. This process ensures the possibility that user interaction, which will have to occur through an interface mediator, will be effective since the system is designed to perform the tasks that the user will require. Despite the contrast between our approach and that of the restricted user interface design, there is still good reason why we can consider it. User interface design can be viewed as a proper subset of interactive system development. If we can assume the existence of a system core which performs the tasks desired for the interactive system, then we can proceed to the description of the user interface which should preserve the task information it presents from the system to the user.

Unfortunately, the majority of existing systems do not adhere to our requirement that they respect the output of a task analysis for their domain of application. Recent work on the ESPRIT project FOCUS [58] has suggested the use of a back-end management system whose purpose it is to impose a relevant task structure to existing applications for which a new user interface is desired. In theory, the back-end manager will allow a principled design of the user interface, but there is no clear indication that their development will always be more cost effective than complete system development from the ground up.

The earliest work on user interface management systems dates back to work by Newman on the Reaction Handler [120], but it was not until many years later that the term UIMS came into use [150]. A result of the Seeheim Workshop on User Interface Management Systems was one of the first standardized architectures for UIMS, called the Seeheim model [129], shown graphically in Figure 3.7.

The main parts of the UIMS are a presentation component, a dialogue controller and an application interface. These roughly correspond to the lexical, syntactic and semantic levels of Foley and van Dam's model [59], which we call the language model. The fourth unnamed box in the diagram recognizes that for proper semantic feedback and efficiency reasons, it is sometimes necessary or desirable to circumvent the dialogue component and provide direct application interface to presentation






Figure 3.7: The Seeheim model of a UIMS

communication. Figure 3.8 shows how the language model fits within the interaction framework. Thus we can see the the language model gives a more complete account of the interface in our framework, but does not treat the user and system on an equal footing with the interface.




Figure 3.8: The language model of UIMS within the interaction framework

The Seeheim model has recently been revised [23], as shown in Figure 3.9, to include domain specific information and recognize the large number of toolkits available for the design of interactive objects. This new model—referred to as either the Arch or Slinky model—is intended to reflect the ability to capture domain information in interactive objects that lie close to the user.

One of the drawbacks to a model such as the Seeheim model is that although it does separate tasks within the user interface to enhance portability, it does not provide for a modular approach to the development of each part. The designer is given no aid to the structured development of a complex user interface by means of the composition of smaller and simpler interface modules. An additional model must be superimposed on the Seeheim model to facilitate a modular development scheme,

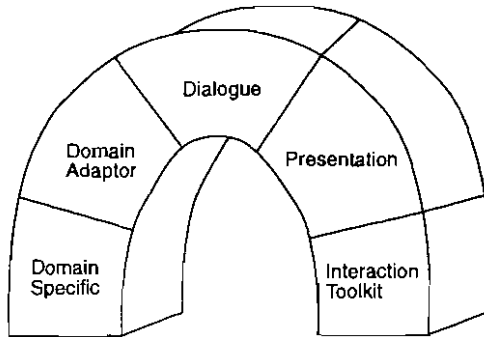


Figure 3.9: The Arch/Slinky model

and many UIMS incorporate such a model for the development of each component in the UIMS. For example, Green discusses the use of recursive or augmented transition networks, BNF notation, and event handler notations as alternative notations for development of dialogue controllers [62]. Various presentation device notations exist, such as PHIGS and GKS, with which toolkits can be constructed for handling of logical input and output mechanisms. A more general notation based on the input and output device spaces defined by Card *et al.* could replace these [31]. Also, the application interface can have its own notation, though its job as a translator from interface to application semantics strongly urges a similar model to that used to develop the application (if such a model exists!).

The apparent freedom of choice for models and notations within each component in the UIMS model suggests another problem for the designer. One of the major obligations of a user interface designer is to ensure that the interface adequately reflects the tasks of the application so that a user is free to interact with the task and not with the interface. Whereas this chore is simplified by breaking the proof into three stages representing semantic, syntactic and lexical correspondence between application and physical interface, it is not clear that the job is at all simplified by imposing very different models at each layer. It is also argued that the strict separation between layers in the Seeheim or language model does not support the handling of semantic feedback [149]. It is for these reasons that some more recent UIMS have utilized a single model in which the various components of the Seeheim model or language model can be expressed. One such example is GWUIMS [150], which is viewed as a UIMS based on the object-oriented model rather than an object-oriented implementation of another UIMS model.

Models such as GWUIMS mark the move in our discussion toward the second subclass of interactive architectures, namely multiagent architectures. Multiagent

architectures recognize that the key to building complex interactive systems is by use of a compositional model in which smaller, simpler modules can be created and composed to form larger, more complex modules. We will only mention some of those in this thesis, namely GWUIMS, the MVC model for Smalltalk [61, 29, 96] and Coutaz's own PAC model [39, 38, 40]. A somewhat related architecture has been presented by Took, in which the interaction medium—the surface—is treated as an independent agent with which both user and application interact [163, 162]. Like the language models, these multiagent architectures recognize a separation between presentation and application, but there is no uniform argument in the multiagent architecture literature as to how this separation should be done.

What separates these compositional multiagent models from the more general compositional process models of software engineering is their emphasis on the special considerations of interaction in the intended design. One of the purposes of the general agent model in this thesis is to demonstrate how a general and more formal compositional model can be used to describe interactive architectures. Our conclusion is that heuristic and formal approaches to design are complementary and a connection between these formerly distinct approaches is the key to an effective design practice. This theme is continued in Chapter 8, in which we investigate the connection between the formal agent model and multiagent architectures.

## 3.2 A survey of research within the framework

We will now use the interaction framework to provide a structured review of the previous formal and informal research relevant to interactive system design. Section 3.2.1 concentrates on work that has addressed the *articulation* translation from *User to Input*. In Section 3.2.2 we will outline the work done on the *Output to User* translation (*observation*). Section 3.2.3 discusses work pertaining to the translation from *Input to System (performance)* and from *System to Output (presentation)*. We combine the work on these two translations because most work on one of them has been also applied to the other.

Much of the material in this section has been culled from state-of-the-art reviews by Abowd *et al.* [6] and Abowd, Dix and Harrison [1] and in a tutorial on formal methods in HCI by Harrison and Abowd [69]. Other good reviews on analytical models have been produced by Reisner [134] and Butler *et al.* [30].

### 3.2.1 Research on *articulation*

The formalisms which address the translation from *User to Input* have been developed by either psychologists or computer scientists whose main interest is in understanding user behaviour. Therefore, the emphasis in this work has been on formulating some model of the user—a *user model*. By user model in this context

we mean the designer's model of the user (see [122]). The user model is then analyzed to describe features of the interaction. This may involve either an additional description of the input (as in Kieras and Polson's CCT) or simply be done by using the input description as the user model (as is done in Reisner's use of the BNF notation). A useful summary of these formalisms may be found in the paper by Green, Schiele and Payne [63] who classify them in respect to how well they describe the *competence* and *performance* of the user. A *task* or *goal* is basic to both approaches. In practice all the notations that deal with competence and performance incorporate aspects of both to a greater or lesser degree. Quoting from Simon [151]:

Competence models tend to be ones that can predict legal behaviour sequences but generally do this without reference to whether they could actually be executed by users. In contrast, performance models not only describe what the necessary behaviour sequences are but usually describe both what the user needs to know and how this is employed in actual task execution.

Simon goes on to classify these notations (and cognitive models in general) in a 3-dimensional space, representing various trade-offs made by their designers. We choose a simpler classification here, partitioning the formalism into two main categories:

- hierarchical representations of the user's task and goal structure; and
- descriptions of the dialogue as a language, or formal grammar.

### Goals, Operations, Methods and Selection (GOMS)

The best example of the hierarchical representations is the Goals, Operations, Methods, and Selection notation (*GOMS*), based on the Model Human Processor and Keystroke-Level model presented by Card, Moran and Newell [32]. This assumes the user has a hierarchical decomposition of goals into subgoals. The goal decomposition may be deterministic or may involve choice among different strategies for achieving the goal. At the leaves of the resulting goal tree are the basic operations that the user carries out to achieve the lowest level subgoals. Granularity of the resulting analysis is given by these basic operations. So for example at a coarse level `edit-document` may be regarded as basic, whereas for finer grained analysis `press-cursor-up-key` may be terminal. Analysis of the goal structure provides measures for determining user performance. For example, the stacking depth of subgoals can indicate short term memory requirements. The models of the users' mental processes implied by GOMS are very idealized and do not apply to error conditions that will arise in the attempt to satisfy goals.

### Formal grammars

Representative of the linguistic approach is Reisner's use of *BNF* notation to describe the dialogue between user and system as a formal grammar [133, 135, 134]. This views the dialogue at a purely syntactic level, ignoring the semantics of the language. Typically, grammar rules ignore computer output and the emphasis for analysis is on the complexity of the input. There are well known techniques for analysing the complexity of grammars, and these can be used to give a crude measure of the difficulty of a dialogue, however the interpretation of such measures is severely complicated by such things as familiarity with (perhaps complex) grammatical forms, clear mode changes etc. Others have used state transitions diagrams, most notably Jacob [89, 88], and added actions to grammar rules, which include output among the grammar's terminals.

### Task Action Grammar (TAG)

Payne and Green have developed a notation called Task Action Grammar (*TAG*) [127, 128, 142] to describe the *consistency* of the input language by describing the user's task structure and the action sequences that accomplish those tasks. The task structure description conforms with the hierarchical goal structure of GOMS. The central role of the task in this formalism is at the expense of attention to the user. The view of the user is just as the goal structure; no consideration is given to how that structure may be modified by the results of previous interactions as observed in the output.

TAG uses parameterised grammar rules to emphasise consistency and world knowledge (e.g., up is the opposite of down). For example, assume the following is an adequate description of the user's knowledge of how to draw a graphic object, such as a rectangle, in Apple MacDraw, as shown below.

```

select rectangle tool           then
place mouse at one corner of the desired rectangle then
depress button                 then
drag to opposite corner        then
release button.
```

This is represented in the TAG notation by means of task production rules. The use of *semantic features* such as *objecttype* allows a more general description of the task by use of user-relevant attributes, similar to the psychological attributes discussed in Chapter 2. The following is a TAG schema rule description of the task for drawing a rectangle or square (a square is a constrained rectangle).

```

task[effect = add, objecttype = rectangle,
constraint = any, selecttool = any] :=
selecttool[objecttype = rectangle] + draw[constraint]
```

```

draw[constraint = yes] :=
    press SHIFT + place mouse ...

draw[constraint = no] :=
    place mouse ...

```

Within this claimed mental representation of the system it then becomes possible to analyse notions of consistency. Since a square is just a constrained form of a rectangle, its creation should be consistent, or similar to, the creation of a rectangle. This is captured above in the first task production rule by means of the *any* value for the semantic feature *constraint*. Hence, a desired consistency is captured by the generic task description. Here consistency is related to the user's understanding. Consequently, there are clear design implications. What TAG does not attempt is to provide any linkage between appropriate task action grammars and possible designs.

Recent developments of the notation (by Howes and Payne [83]) include attempts to make good some of the limitations of TAG, in particular to include display information and flow information (no state is implied by TAG). A pessimistic view of these developments is that they make an already cumbersome notation worse. The possibility of developing informative notions of consistency becomes even more remote. Although this notion is purported to be a competence formalism [63], it is clear that the breakdown of the task into action also has performance implications. Attempts to scale TAG to substantial applications have not been satisfactory (see for example Schiele and Green [142]).

Despite these criticisms, the methodology of TAG is extremely important when considering the link between psychological and computational HCI research. The semantic features that are the basis for TAG schemas are based on empirical evidence of how users perceive the tasks they perform and the interfaces with which they interact. We will take this identification of task information from the user's perspective and use it as the foundation for a user-centred design and analysis method described in Chapters 7 and 8.

### Command Language Grammar (CLG)

Taking the lead from linguistic theory and parsing, it is often suggested that several levels of grammars ought to be used. Moran's Command Language Grammar (CLG) [114] is probably the most well known example of this. CLG uses four—lexical, syntactic, semantic and task. It is more design oriented than most other similar approaches, the task level being described first, obtained presumably from a task analysis, then the semantic level, formalising the entities, before moving on to the more concrete levels. Various rules are given for checking consistency within and between the levels, although these are rather loose and incomplete. This approach comes closer to viewing the entire system as involved in the interaction

rather than just the surface dialogue. Unfortunately, it has been found unwieldy to use in practice [148, 147], and the notation used is particularly arcane, looking very much like LISP.

### Cognitive Complexity Theory (CCT)

*Cognitive Complexity Theory (CCT)*, first presented by Kieras and Polson [93]) combines the goal hierarchy and dialogue grammar approaches. It has two parallel descriptions. User goals are expressed as *production rules* à la GOMS. The system grammar is given by a generalized transition network (*GTN*). The production rules are a sequence of rules of the form

*if condition then action*

where *condition* is a statement about the contents of working memory. If the condition is true then the production is said to have fired. An *action* may consist of one or more elementary actions. The *user program* is written in a LISP-like language and generates actions at the keystroke level that have associated performance characteristics. This user program can be executed and assessed empirically and analytically. In addition, mismatches between it and the system grammar can be found and a dissonance measure produced. The GTNs which describe this system grammar are in the form of diagrams representing the dialogue states with arcs representing the possible transitions on user actions. The difference from simple state transition diagrams is that the nodes may be hierarchically decomposed. This system part of CCT could be executed in the same way as a grammar to give a crude dialogue prototype.

### Problems with goal structures and grammars

The formation of the goal hierarchy is largely a post-hoc technique and runs a very real risk of being defined by the dialogue rather than the user. Knowles [94] attempts to rectify this by producing a goal structure based on a pre-existing manual procedure. She thus hopes to obtain a natural hierarchy. In addition, she criticizes the mechanical measures of complexity because they do not take into account issues such as user knowledge.

In addition, as more display-oriented systems encourage less structured methods for goal achievement. Instead of well defined plans, the user is seen as performing a more exploratory task, recognizing fruitful directions and backing out of others. Typically, even when this exploratory style is used at one level we can see within it and around it more goal-oriented methods. So for example, we might consider the high level goal structure below.

```
WRITE_LETTER ==> FIND_SIMILAR_LETTER + COPY_IT + EDIT_COPY
```

However, the task of finding a similar letter would be exploratory, searching through folders and recognizing possible places may not be well represented as a goal structure at all. Similarly, the actual editing would depend very much on non-planned activities. If we drop to a lower level again, goal hierarchies become more applicable. For instance, during the editing stage we might have the following sub-dialogue for deleting a word.

*delete.word* ⇒ *select.word* + *click.on.delete*  
*select.word* ⇒ *move.mouse.to.word.start* + *depress.mouse.button*  
                   + *move.mouse.to.word.end* + *release.mouse.button*  
*click.on.delete* ⇒ *move.mouse.to.delete.icon* + *click.mouse.button*

Thus goal hierarchies can partially cope with display oriented systems by appropriate choice of level, but the problems do emphasize the prescriptive nature of the cognitive models underlying them.

Grammar techniques were initially developed to examine command-based and keystroke dialogues. One problem in applying them to mouse driven window systems is determining the lowest level lexical structure. Pressing a cursor key is a reasonable lexeme, but moving a mouse one pixel is less sensible. In addition, pointer dialogues are more dependent on the display. Clicking a cursor at a particular point on the screen has a meaning dependent on the current screen contents. This problem can be partially resolved by regarding operations such as *select-region-of-text* or *click-on-quit-button* as the terminals of the grammar. If this approach is taken, the detailed mouse movements and parsing of mouse events in the context of display information, such as menubars, are abstracted away. This means that any prototyping of the dialogue will be at a similarly abstract level or require "Wizard of Oz" techniques to mock up the full interface.

### Programmable User Model (PUM)

Goal structures such as those provided by GOMS, TAG, CCT and CLG form rudimentary user models, none of which are very good at handling user error. More recent research by Young *et al.* [173, 175] has investigated the possibilities of programmable user models (*PUMs*) which can more directly address the question of error in order to further aid design of the interface. This research involves executing programs in the SOAR cognitive architecture [98] to perform *scenarios*—typical examples of user interaction with the machine—to determine usability consequences of a given system design for accomplishing given tasks. An advantage of this approach over the others is that a detailed description of user procedures is not necessary. Rather, a *knowledge analysis*, embodying the user's understanding of the functionality of the system from its intended interface, characterizes the possible behaviours of the user. The executable cognitive architecture then uses minimal problem solv-



ing techniques, familiar to the AI community, to highlight usability consequences and possible behavioural errors which would result from a means-end analysis.

The PUM methodology relies on levels of description, much like CCT and CLG. Two levels of interest are the task level and the device level. The task level description is device independent, and problem solving in that space is usually straightforward. Solutions determined at the task level are then mapped into the device level, which is given by the knowledge analysis. The PUM provides a trace of user actions at the device level and the goals that those actions satisfy. This trace of user behaviour can be compared to the designer's intended behaviour. A discrepancy would indicate to the designer a possible problem that may need correction. A simple and effective example in of this procedure in text editing is given in [173].

### Graphical or diagrammatic approaches

A major criticism of formal techniques is that they are not accessible to the average designer. On the principle that many people find graphical notations easier to use, there have been many different notations proposed. Most of the hierarchical and grammar notations can be given a graphical form. In addition, there are data-flow diagrams, state transition diagrams (of many flavours), Jackson System Development (*JSD*) diagrams and simple flow diagrams. Diagrammatic notations are often used in conjunction with other notations and may have automatic support. For instance, Marshall's diagrammatic notation [105] (see below) links Harel statecharts to VDM.

Sutcliffe [159] has used JSD process structure diagrams to describe tasks. He then analyzes these in order to highlight possible problems such as memory limitations (rather like GOMs). Similarly Walsh *et al.*[165] have integrated task analysis techniques with JSD. They point out that these notations are already heavily used for the software development side, and therefore they form a common language. JSD diagrams can be used quite simply as a model of the dialogue, being a particular form of grammar.

### Conclusions on articulation

This section has discussed the different psychological and soft computer science models of the *User to Input* translations. All these approaches are still at the research stage. However, the general idea that producing a description of how the user is to accomplish expected tasks in parallel to the actual system development seems useful. It is generally agreed that the form of the modelling is not nearly as important as the discipline it enforces on the designer.

### 3.2.2 Research on *observation*

In this section, we overview some of the analytical methods used to assess the translation from *Output* to *User*, the *observation* translation in our framework. This is perhaps the most important translation as far as determining overall usability of an interactive system, and yet it is the most elusive. This is not surprising, since at the very core of this analysis, we are trying to determine how individuals understand that which they perceive. It is one thing to empirically test the perceptual capabilities of an average human user [112], and from that provide some model of the user as an information gatherer, but it is far more to explain how that perceived information is transformed into knowledge about the surrounding world. This topic is certainly beyond the scope of this thesis, but it is interesting to note that research which has occurred in this area. Research in this area is now generally viewed as the next major challenge for psychological research in HCI.

Most of the information on analytical models of perception has been taken from the review done at the CHI'88 Workshop on Analytical Models [30], and therefore our comments are very brief. A model not usually grouped with the analytical models is Barnard's Interacting Cognitive Subsystems (ICS) which we discuss here because of its contribution to visual processing and its apparent amenity with the agent model.

#### Display Analysis Program (DAP)

Tullis has produced a computer program, the *Display Analysis Program (DAP)*, which takes as input the actual displays for a system and produces a listing of improvements that can be made in the design of the screen layout to improve the time requirement for location of specific text units on the display. His work ignores the semantics and task structure of the display. He provides support for the variance in locating textual units in terms of certain selected characteristics of the layout and perceptual attributes that can be objectively assessed by a computer.

#### ANets

Just as CCT and other notations mentioned above use a two layer description of goal structure and device structure to discuss the translation between the two, so does Chechile's approach to modelling comprehension of displays rely on a two level description, one of real world knowledge and one of display knowledge. Each of these descriptions are given in terms of *ANets*, augmented forms of semantic networks. The world knowledge network represents the users general knowledge about the display format and the domain of the interactive application. This would comprise knowledge about how concepts relevant in the domain (the psychological attributes in our terminology) would be represented in the display (the display attributes).

The display knowledge represents information about actual displays as snapshots in the dynamic interaction.

### Cognitive Environment Simulation (CES)

Roth *et al.* have developed a symbolic processing model of the inferencing and evaluation procedures of nuclear power plant operators. The aim of the model is to provide predictions of situations and properties of the environment and the information provided by an interactive system that will lead to errors in assessment and intention formation. CES is an example of knowledge-based simulation models of human performance whose objective is to explicitly present domain goals and the knowledge necessary to support those goals. CES is not based on a cognitive architecture, as ANets are.

### Interacting Cognitive Subsystems (ICS)

Barnard attempts to incorporate two separate psychological traditions in describing his cognitive architecture of *Interacting Cognitive Subsystems (ICS)* [17, 18, 21]. One is the architectural and general-purpose information-processing approach of *Short Term Memory (STM)* research. The other is the computational and representational approach characteristic of psycholinguistic research and AI problem-solving literature.

ICS provides an architecture for perception, cognition and action built up by the coordinated activity of nine smaller subsystems, five comprising a peripheral subsystem in contact with the physical world and four comprising a central subsystem. Each of the nine subsystems is specialized for handling some aspect of external or internal processing. For example, one peripheral subsystem is the visual system for describing what is seen in the world. An example of a central subsystem is one for the processing of propositional information, capturing the attributes and identities of entities and their relationships with each other (a particular example is that propositional information represents “‘knowing’ that a particular word has four syllables, begins with ‘P’ and refers to an area in central London.”)

A subsystem is described in terms of its typed inputs and outputs along with a memory store for holding typed information and transformation functions for processing the input and producing the output. So ICS can be seen to have a natural description in terms of agents, but just how close the correspondence is between the ICS subsystems and the agents presented in this thesis is an interesting question which, unfortunately, remains open.

Though ICS is purported to provide a model of perception, cognition and action, the type of information it provides on the action side is not the same as provided by the models for action described in our discussion of models for *articulation* above. ICS is not intended to produce a description of the user in terms of sequences of

actions that are performed. ICS provides a more holistic view of the user as an information processing machine. The emphasis is in determining how easy particular procedures of action sequences become as they are made more automatic within the user. The lack of quantitative output from ICS makes it less practical than the other *articulation* models. However, ICS does provide competent and understandable analyses of how visual information is perceived and transformed by the user, and this information provides substantial advice to designers wishing to enhance the proceduralization of user behaviour, since proceduralized behaviour is assumed to be less prone to error and, hence more effective.

The main purpose of ICS is not so much as a model for generating a description of a user's behaviour, as one views the purpose of a model such as PUMS. Rather, its purpose is as a classification method for the wealth of empirical psychological evidence on user behaviour. Thus, we would expect ICS to perform better as a rationalization of user behaviour based on empirical evidence, whereas PUMS we would expect to produce better predictive information.

### Conclusions on *observation*

Research on the translation from *Output to User* is not as progressed as research on the translation from *User to Input*. This is perhaps due to a lack of empirically validated psychological theory concerning perception and understanding, or at least a failure to apply those theories to questions of human-computer interaction. Contributions from this research, however sparse, will bear a great significance in directing design of interactive systems.

### 3.2.3 Research on *performance and presentation*

Research in this section focuses on features of the interaction which are directly affected by the *System* and the interface (*Input and Output*). We have bundled together the *performance* and *presentation* translations because most models devised to deal with one also deal with the other. The work in this section fits along a continuum from very abstract models (such as the PIE) to very constructive notations (such as Alexander's SPI). We will try to present this review from the more abstract towards the constructive.

### Extensions to the PIE model

We have already presented the *PIE* and *red-PIE* models in Section 3.1. In order to study more specific areas of interactive behaviour, further refinements to the basic model have been introduced.

### Handle spaces

A client-server relationship is a common way to view the independent execution of several programs (the clients) under one controlling program (the server). This view can be readily extended to situations in which a user plays the role of client or server.

An example of the user as client is seen in a multi-windowing system, where different windows represent different tasks and so the user treats them as independent interactive systems. Each window can be treated as a red-PIE, and so the multi-windowing system is represented as a collection of PIEs. The user is able to direct input to particular windows, i.e., the user can select a window for input. This selection can be modelled by giving each window as a red-PIE a *handle* for distinguishing its input. Alternatively, there is an injective function from handles to PIEs. The overall display and result functions will depend on the set of active handles, or active windows, yielding a red-PIE description of the multi-windowing system. Since we assume the user considers the separate windows as independent, interference between windows is undesirable. We can therefore formulate a constraint on the collection of window PIEs, called *result independence* by Dix, which would ensure that the contribution from interaction with one window to the overall result be separable from that of any other window. Handle spaces and result independence are dealt with in more detail by Dix and Harrison. [50, 48].

Situations in which the user is the client (and hence, there are many users) fall under the research theme of computer supported cooperative work (CSCW), a topic which is outside the scope of this thesis. A formal methods approach to CSCW has not yet been attempted, though we suspect that a treatment analogous to handle spaces would be enlightening.

### Nondeterminism

There are several notions of nondeterminism relevant to interactive systems. Nondeterminism arises from the loss of knowledge, either deliberate or not, of some aspects of the whole system. The user, who for reasons of inexperience or a conscious decision to remain ignorant, does not know exactly what aspects of the system determined its behaviour. It can arise that different situations which visually appear the same react differently to the same commands. So, for example, in a graphical drawing package the selection algorithm will take into account a structural or temporal hierarchy (into layers, perhaps) of the objects which may not be fully manifested in the visual presentation. If two objects appear to overlap and the mouse cursor is placed in a location which is 'covered' by both objects, a subsequent mouse click to select will choose one or the other object depending on which is considered on top, or it may choose both, if the two objects have been grouped as one. The point is that there is usually no visual indication to aid the

user in predicting the outcome of select; the algorithm is nondeterministic as far as the user is concerned even though it may have a perfectly deterministic description at some level of detail.

Dix has discussed this nondeterminism [47] in interactive systems and how it can be modelled in a PIE. In Chapter 4 we extend Dix's ideas in terms of agents when we investigate this user nondeterminism with respect to predictability.

### Temporal models

One feature of an interactive system which greatly affects its perceived usability is the avoidance of display lag wherein the current display does not adequately reflect the state of the system resulting from all prior user input. Dix [48, 49] has discussed display lag and whether it can be avoided in any system. The conclusion is that rather than chase the "myth of the infinitely fast" machine, HCI research can concentrate on what it means to make a system usable that will admit the inevitable delays. A simple extension to the PIE model in which a single null input event is introduced in the program language  $P$  leads to a definition of steady-state behaviour and the consequences on predictability of allowing buffered user input. One of the most important results of this work has been a formal representation of the requirements that can be placed on a system which experiences display lag in order to expose information that must be available to the user.

### Levels of system description

As we explained in Section 3.1, a PIE analysis can be given at varying levels of abstraction. As Carroll has pointed out [35] important interactive features become apparent depending on the level of detail in the system description. Though we do not adhere to Carroll's further and fatalistic belief that there are an infinite number of such levels—each of grave importance to the overall usability of a system—and, hence all attempts to model such properties will be hopelessly incomplete, it is important to consider how the varying levels of detail can be related.

We can view the description of an interactive system as a collection of PIEs, each representing the system at different level of abstraction and each capturing relevant interactive properties along the way. The system description is completed by providing mappings between the program spaces (parsers, as Dix calls them) and effect spaces (projection or embedding mappings) of the various levels. The result space will be associated to a more abstract level PIE as the concrete level PIEs usually deal with more immediate and temporary features of the interaction (such as editing a command line or displaying a pop-up menu). Each PIE will contribute to the overall display. Figure 3.10 depicts these relationships in a two level description.




Figure 3.10: Relationships between levels of PIEs

### State display model

Another obvious way to separate into levels is by considering the interface separately from the underlying system, as is suggested in the interaction framework. Harrison and Dix have considered such a state display model [73] in an attempt to formalize notions behind direct manipulation interfaces. This model considers the underlying system and the display as separate machines, similar to agents. The link between these machines is referred to as *state display conformance* in which operations performed on the display by the user are adequately mirrored by operations performed on the state. State display conformance can be used to assess the quality of a graphical interface, as Ahowd, Dix and Harrison have shown [2]. The theory of data refinement is close to that of state display conformance and we will discuss this in terms of agent correspondence in Chapter 4.

### Templates

One of the problems with system models is that although they describe interactive behaviour they have no conception of how the user sees the system. We demonstrate what we mean by an example from [20] in the context of the predictability principle. One way [53, 158, 9] of making a principle of predictability precise is to require that if the effect of any two input programs are the same, then no future experimentation will betray any difference in effect between the systems. Viewing programs as sequences of inputs, we can formalize this as below.

$$\forall p, q \in P \bullet i(p) = i(q) \Rightarrow \forall r \in P \bullet i(p \frown r) = i(q \frown r)$$

This notion stresses that the *effect* is sufficient to determine the equivalence of distinct system states. From a user's point of view the fact that the two effects are

identical may not be sufficient. A stronger requirement of predictability (which as it happens may be too strong for any realistically complicated interactive system) is that the equivalence of the effects will depend on the *user's perception* of whether the effects are equivalent. If we regard the display as what the user perceives then it may be more appropriate to define predictability as below.

$$\forall p, q \in P \bullet \text{display}(i(p)) = \text{display}(i(q)) \Rightarrow \forall r \in P \bullet i(p \wedge r) = i(q \wedge r)$$

This says that if the viewable effects are the same, then the obtainable effects are indistinguishable (though not necessarily visibly indistinguishable). By indistinguishable we mean that no experimentation by means of input  $r$  would betray a difference.

We can go further than this and add structure to the model to incorporate claims about, say, user attention. Certain components of the display are more likely to be noticed in making decisions about the next action than others. Some parts of what is seen of the system will be different in a way that is irrelevant to the future of the application (e.g., more general system status information and the time and date).

One way to add structure to the model is through the use of *templates* as suggested by Harrison, Roast and Wright [74]. Instead of maintaining a single display and result function, we can allow for many display and result mappings on the effect space, each of which we will refer to as either a *display template* or a *result template*. It is beneficial to formulate properties of a design in terms of the relationship between the ephemeral and immediate display and the end result. In essence, that is precisely what the predictability property above is doing. Since we have access to a result function in a red-PIE, it would probably be beneficial to rewrite the above predictability property including the result.

$$\begin{aligned} \forall p, q \in P \bullet \text{display}(i(p)) = \text{display}(i(q)) \Rightarrow \\ \forall r \in P \bullet \text{result}(i(p \wedge r)) = \text{result}(i(q \wedge r)) \end{aligned}$$

This says that if the viewable effects are the same, then the obtainable end results are indistinguishable (though still not necessarily visibly indistinguishable).

The main use of templates by Harrison, Roast and Wright was in an investigation of cyclic properties of a bibliographic database. A structure that is particularly common in menu-based systems is a *cycle*. A main menu presents an initial set of options, each leading to a dialogue sequence. The sequence affects some change in the result only at its conclusion, when the main menu is redisplayed and the cycle is complete. Hence in the database, a *select* entry in the main menu will lead to a cycle that results in the selection of a reference given a particular name or date or source. Recognition of the cycle is especially useful to the designer if the user understands and relies on it for effective interaction. It then becomes important to ensure that some display template at the start and end of the cycle is clearly



significant to the user and that whatever effect the system has on the associated result template takes place at the end. The cycle is then, in some circumstances, a mirror of the task as it would apply in the user model. Claims about the interaction from a task point of view, perhaps the identification of possible error situations, can then be formulated in terms of the cycle and experimentally observed through evaluation of actual use.

Templates are determined from empirical evidence, as were the semantic features of TAG. The use of templates marks an initial encroachment of solid psychological information into the system modeling component of HCI. A more constructive view of templates in terms of attributes is one of the main reasons for a refinement of the agent model in Chapter 5. Templates, therefore, form the basis for an interactive design and analysis method described in Chapters 7 and 8.

### Limitations and application of abstract models

The biggest danger of any formal approach is that the designer may attach more credence to it than it deserves. This is particularly true of abstract models of the sort described here. An unwary user of the techniques may believe that formal consistency with some of the principles and models described above was *sufficient* for a system to be usable. The fallacy of this is obvious to anyone with an understanding of HCI, but formalisms can be both seductive and blinding. It is therefore essential to understand the limitations of the abstract techniques above.

Rather than being sufficient conditions for usability, the various formal statements of principles tend to be *necessary* conditions. A formalization of the familiar slogan "what you see is what you get" (*WYSIWYG*) as done by Dix [48] and Thimbleby [161] and Sufrin and He [158] says that it is possible to work out the end result of a system by interactively examining it. This is essential if the system is to be usable at all, but does not tell us how easy it is to work out the result, or how visually and spatially faithful a representation of the result we see on the screen. The latter of these problems is perhaps easy for the designer to verify, and mistakes will be obvious, but the former requires a deep understanding of human cognition that is unlikely to be formalizable to the same extent.

Some of the psychological issues are just beyond the scope of the models, and one can simply note that even when the system has passed the formal tests, more human centred analysis must be applied. Indeed, the introduction of templates was specifically geared to marrying together the abstract principled approach with an evaluative approach. In order to satisfy the formal principle the designer would have to give the relevant template functions. These can then be validated either empirically or by a human factors specialist. Similarly, the *strategies* by which the user can investigate the system state [48] are a form of user program, and can thus be validated against executable cognitive models, or again by direct evaluation. Often

the simple fact that the operations and deductions that the user must perform have been explicitly stated as part of the formal proof will be sufficient to see whether they are reasonable or not.

The other major non-formalisable part of the use of these models is deciding which principles are applicable and desirable to a particular application, and also at what degree of abstraction to apply them. A system may be understood at various levels of abstraction, such as concrete keystrokes and mouse actions, syntactic units, or semantic commands. The models can often be applied to the system at each level. Some of the properties will be universally applicable to all systems at all levels, but in general the designer will be more selective. So for instance, in a command based operating system, one expects to have a total view of the current command being edited, that is, it obeys a very strong visibility principle. However, when the command is submitted (entering the carriage return key) the results of it on the file system will be far less visible, usually requiring explicit commands to view files, directories etc. Arguably, in this example, the semantic level could do with being more visible too, but it is a design decision as to what degree of visibility is required at which level.

In short, as with any method or model the domain of applicability of abstract models must be born in mind when they are used or evaluated.

A further and related issue involves the refinement of one abstract level in order to more closely approximate an executable system while still preserving the interactive properties. Having produced specifications that satisfy desirable properties, there is no guarantee that the structures used in the specification designed to match the user model will be well-suited to implementation.

The major aim of the above models has been to *define* useful properties. There is then the issue of actually building systems that satisfy them.

Having produced specifications that satisfy desirable properties, it is probable that the specification structures designed to match the user model are very ill-suited to implementation. This conflict will arise with *any* specification of interactive systems. If a designer leans towards efficient implementation structures, then it is likely that user requirements are badly defined, whereas if she leans towards the user then inefficient structures result. Hence the desire in this thesis to provide a more constructive approach via agents that is at once amenable to the kind of analysis of the abstract models but also sympathetic to software engineering.

But our agents are certainly not the first constructive approach to the application of formal methods in HCI, and we will review those now.

### Dialogue specification

A problem with the use of general specification techniques is that they are too general. The dialogue of the interaction will not be separate in the notation, instead

it must be modelled. The dialogue component may be packaged and used within the standard notation, allowing the free mixture of interface and standard forms. This gives maximum expressiveness, but at the cost of losing the dialogue/application separation which is frequently seen as desirable. In particular it makes it hard to analyse the dialogue structure as a separate part of the system.

There have been various attempts to add dialogue specification components to standard notations. These may be simply sugarings that are then translated into the underlying notation to give them semantics, or have a separate level of semantics given them. In either case, the actual concrete notation makes a clear separation between the two styles of specification.

### **EPROL**

Hekmatpour and Ince [79], for instance, have a separate user interface design component in their *wide-spectrum language EPROL*. This interface component seems rather disappointing however, being simply a teletype forms and menu description such as may be included with many data-base languages or fourth generation languages. The dialogue is apparently described entirely within the main specification language and may thus be easily obscured.

### **Statecharts and VDM**

Marshall [105] has merged a graphical dialogue specification technique based on Harel's statecharts [67, 68] with *VDM*. This includes standard constructs such as sequence, choice and iteration in the dialogue, each terminal dialogue "box" is related to a piece of VDM specification. She also suggests that the user's actions in this can be represented by a parallel diagrammatic/VDM description, but in her examples this diagram consists of a single box, so the claim is not supported. This exposes the fact that the diagrammatic notation does not support parallel activities (such as multi-window dialogues). It would be quite easy to add such a construct at the diagrammatic level, but the meaning when translated into VDM semantics would not be clear. The actual acceptance of input is handled by "shared" global variables with the user "process" and is hardly clean. Another problem, is that each piece of VDM works on global variables, making it difficult to trace the semantic impact of particular user actions without analysing the pieces of VDM in detail, which runs counter to the desires of dialogue separation.

A graphical notation has the considerable advantage that it is easier to comprehend initially. A disadvantage of a graphical notation is that it is difficult to formulate properties, or theorems, within a strictly graphical notation. Since our work emphasizes the use of a formal notation for the precise expression of interactive properties that can then be used to analyze and develop an interactive system, this latter disadvantage far outweighs the advantage of comprehensibility. It is a

reasonable objective in the future to extend our agent model to include a graphical representation.

## SPI

Alexander [10, 11, 12] has designed an executable specification/prototyping language around *CSP* and *me-too* called *SPI*, (specifying and prototyping interaction). *Me-too* is an executable specification language based on VDM and implemented under several dialects of LISP. The *CSP* forms the dialogue specification part, whereas the *me-too* supplies the semantics. This is rather similar to the way *CCS* and *ACT-ONE* are combined in *LOTOS*[87, 164].

*SPI*'s dialogue component is called *eventCSP*, it includes most *CSP* constructs, sequential action, choice, iteration, and most importantly parallel composition. The parallel composition makes it possible to express concepts such as the choice between mouse and keyboard input. The expression of choice is based on the occurrence of events and is thus more clear. It inherits drawbacks from *CSP* however, such as the lack of direction in events, it is not evident in the syntax whether an event is due to external input, produces external output or is an internal synchronisation between parallel processes. This can be confusing in dialogue design when there is an obvious direction of control flow. However, the problem is largely mitigated for user I/O by the judicious choice of event names. It is thus only internal events that remain confusing. The structure of possible events is static too. This would make it hard to deal with the dynamic creation of windows for instance. This lack of dynamic configuration (and related lack of parameterisation) is common to many dialogue languages, it would be easy to add to most, but would typically reduce the possibility of analysis of the dialogue structure.

The semantic part of *SPI* is called *eventISL*. Although it is based upon *me-too*, it is intended to operate with various languages, in particular a *C* version is available. The host language independent part consists of several elements: a clause giving the global values needed for the event, a pre-condition expressing when the event can occur, output and input parts. The host language part, simply describes what updates to global values are possible. The globals used and updated are made explicit and thus tracing the effects of events is easier. It would be possible to use other specification notations such as an algebraic notation or *Z* as the host language, with a subsequent sacrifice in executability, a choice we consciously make (at least initially) in designing the agent language.

*SPI* has a prototyping tool for use when only the *eventCSP* dialogue description has been produced. This allows the designer to examine possible event traces. Later full prototypes using the *me-too* version of *eventISL* or the *C* version can be executed. One drawback with the implementations is that they do not offer the full parallelism of the *CSP*. This is because the underlying languages they

were built upon did not allow full multiplexed, non-blocking I/O. They fake the nondeterminism as long as they can for internal events, but when one of several choices of user input device are possible, the system makes an arbitrary choice. Most versions of C on UNIX or PCs have system calls for non-blocking I/O, so it should be possible to rectify this, at the cost of some loss of portability.

### Anderson's work

Anderson has built on the claims of the PIE model with a more constructive refinement of it [14, 15]. He provides explicitly an input language generated by a context-free grammar, whereas the input language of the PIE model imposes no such structure on the dialogue. A state description is separated from the output, with a meaning function mapping input constructs to state transitions and a display function mapping the resulting state to its output presentation. As such, Anderson's model contains three of the components found in our own interaction framework, though we represent both *Input* and *Output* as state machines and he does not. The important contribution of Anderson's work is that he provides concrete examples of a system, a file browser, and translates the abstract properties expressed as theorems on the specification and then proves their satisfaction.

### Interactive processes in Z

The work of this thesis is based on another constructive refinement of the PIE model given by Sufrin and He's description of the design, analysis and refinement of interactive processes [158]. He [76, 77] and Josephs [92] independently gave descriptions of a state-based version of CSP in which the events of the language were coupled to transitions of an underlying state machine. Sufrin and He provided the description of the state machine in the Z language and encapsulated a CSP-like syntax within Z for the event specification. The importance of this work can be seen in several ways.

The PIE model served as the motivation for the interactive processes. One of the requirements for interactive processes was that they would be able to express the same types of properties expressible with a PIE. Sufrin and He convincingly satisfied this requirement, ultimately providing a very systematic and intriguing categorization of WYSIWYG-like properties in terms of the display and result mappings. Our own derivations of interactive properties in Chapters 4 and 7 are based on this technique. The link between the CSP and Z is similar to one suggested by SPI and mentioned earlier, but this particular work showed more of the advantage of such an approach by showing its use in formulating properties on the specification that could then be the subject of proof and refinement.

Z is becoming increasingly familiar in both academic and industrial circles, so the presentation of interactive processes in Z showed convincingly that HCI

concerns could be expressed in the language of software engineering, not so much as a separate and disparate module in development but as an integral component with the software development lifecycle.

### Conclusions on performance and presentation

An immediate conclusion arising from this review of system based work on formal methods in HCI is that except for the Sufrin and He model, there has been little effort in the last four years on incorporating the abstract principles motivated by the PIE model into a more concrete design practice. One of the principal aims of this thesis, therefore, is to investigate a further constructive formal method and its links to interactive system design and analysis.

He and Josephs' work on state-based CSP strongly suggests that the monolithic presentation of interactive processes can and should be made as compositional as a process algebra allows. Outside of interactive system design, there has been an increasing interest in the development of model-oriented specification techniques which more adequately address the modularization necessary for the description of large systems. In Chapter 6 we will discuss several of the attempts to extend the Z notation to capture object-oriented techniques. The agent language developed in this thesis is an attempt to address both the need for better modular specification notations and also provide a flexible and consistent notation to express internal state-based behaviour as well as external event-based behaviour. The resulting notation will allow us to formally express interactive properties at the task level. This implies that the agent model can be used in an interactive design method which is linked to task analysis information.

## 3.3 Conclusions

We have seen how the interaction framework of Chapter 2 was influenced by pre-existing general accounts of the structure of interactive systems. The contribution of the framework is two-fold. It first provides a uniform language for describing the three major components of an interactive system, the *User*, the *System* and the mediating interface (in terms of *Input* and *Output*). The unified treatment of these three components opens up the possibility for a cross-fertilization of research on both the system side and user side of interactive system analysis and design.

The second contribution of the framework is that it provides context for research in HCI. Having completed this contextualization we can see where the work in this thesis fits in the general scheme of interactive system analysis and design. We intend to provide a constructive model of the *System*, *Input* and *Output* in terms of the agent model and its associated language. This model will provide the means for a compositional description of realistic interactive systems. We will concentrate in

our descriptions on revealing those features of the design which highlight properties of the interaction between the user and the system. In doing so, we will provide formal accounts of existing interactive design heuristics which can lead to a more principled design practice.

In the next chapter we will describe both formally and informally interactive properties that are derived from the interaction framework. In order to relate these properties to a design and analysis methodology, we will need to provide greater detail on the structure of agents in the refined model of Chapter 5. Having refined the agent model, we will provide a design language for agents in Chapter 6 which matches more closely than the standard Z notation how a designer thinks of describing an agent. In Chapters 7 and 8 we will link the refined agent model and notation to a design methodology for generating descriptions of interactive systems and analyzing existing system based on empirical evidence of how user's understand the tasks the system supports.

## Chapter 4

# Properties of interactive systems: Part I

The purpose of this thesis is to provide a means for designing interactive systems which can be analyzed with respect to desirable properties to enhance usability. We have established the context of this research in the previous two chapters. In this chapter we will investigate the kinds of properties which influence usability.

We present a catalogue of interactive properties, discussing each entry in terms of the interaction framework and/or the simple agent model. This catalogue is not intended to be complete in the sense that it lists every property of an interactive system that could possibly affect usability. The artillery with which we are equipped at this point is not powerful enough and it is even doubtful that such a complete catalogue could be compiled. We first inspect the interaction framework to uncover properties of the translations which affect overall usability, providing examples within real systems. The properties of interest for translations in the framework concern the ease associated to the translation and the coverage of the translation, and we can attempt to formalize those notions. Other interactive properties have been discussed in the PIE model literature and the Sufrin and He model which pertain to the relationship between the *System* and the *Input* and *Output* components of the framework. Assuming a compositional model of the agent, which we will define explicitly in Chapter 5, these properties can be defined over a single agent.

### Overview of chapter

In Section 4.1 we will give an informal account of properties of translations within the interaction framework. The intent of this informal section is to motivate how the coverage and ease of a translation affects usability. In Section 4.2, we begin the formal treatment with the definition of a definedness ordering on translations. Though this ordering directly addresses coverage of a translation, we see its use



mainly as the basis for a refinement ordering on the external behaviour of agents. We do not pursue refinement of agents any further in this thesis. In Section 4.3, we address the ease of a translation as a correspondence relation between agents similar to treatments of data refinement. The remaining sections of this chapter apply the Sufrin and He method for classifying interactive properties as relationships between input history and state or response history. We conclude the chapter by indicating how the agent model must be altered to accommodate more salient descriptions of interactive properties relating results and displays.

## 4.1 Properties of translations

Figure 2.3 depicts four separate translations between the components of the interaction framework. Properties of those translations have consequences toward the overall usability of the interactive system. In this section we will discuss qualitatively how those translations affect usability before attempting to formalize properties of the translations in the next section.

### 4.1.1 Hutchins, Hollan and Norman distances

Hutchins, Hollan and Norman have used the execution/evaluation cycle of interaction as the basis for an informal understanding of direct manipulation [85]. Crucial to their qualitative account of the ease of use that results from direct manipulation interfaces are two notions of distance—*semantic* and *articulatory*—both of which are minimized in effective interactive systems. They classify input and output languages in the more general category of interface languages. Any expression within an interface language has both a meaning and a form. Semantic distance concerns the translation between the user's intentions and the meaning of the interface language. It is measured by the expressiveness (Is it possible to say what one wants to say in this language?) and conciseness (Are simple task expressions translated into complicated input expressions?) of the input language. The articulatory distance is a measure of the correspondence between the meaning and the form. This can be measured in terms of the link between the structure of the input or output terms and their intended meaning. For example, input to a graphical drawing package will be assisted by gesture input with a mouse, or touch-sensitive screens will be used to provide more natural indicative input (e.g., move *that* item to *there*). Monitoring a variable in a process control system may be aided by output in the form of a meter reading or a continuous graph, instead of a semantically equivalent table of time-stamped values.

The translation from user intention to user input is influenced by both the semantic and articulatory distances associated with the input interface language, as the user translates from the goal in the task language to the meaning of that goal

and its subsequent form within the input interface language. Similarly, semantic and articulatory distances affect the translation from user-perceived system output to user assessment. Semantic and articulatory distance measures are applied to the execution as well as the evaluation phase of interaction. They attempt to answer questions about the possibilities enabled by the translation and the ease of the translation the user and the interface. Figure 4.1 portrays the Hutchins, Hollan and Norman sense of articulatory and semantic distances within our interaction framework.

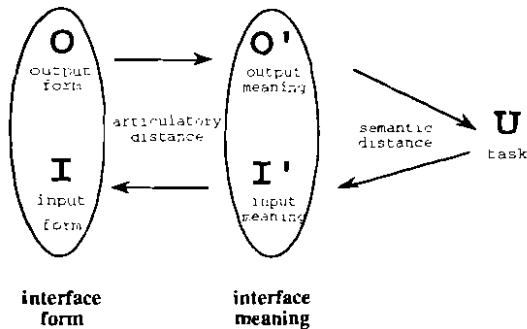


Figure 4.1: Hutchins, Hollan and Norman's distance concepts within the interaction framework

Had Norman's initial execution/evaluation cycle been more detailed on the system side, it may have prompted Hutchins, Hollan and Norman to have extended the definition of articulatory and semantic distance to the correspondence between the system and the input and output languages of the interface.

Since Hutchins, Hollan and Norman have not provided a quantifiable metric to these notions of articulatory and semantic distance, we feel uncomfortable referring to them as distances. It is for this reason that we no longer refer to the descriptions of properties on translations as distances.<sup>1</sup>

#### 4.1.2 Articulation

The *User's* formulation of the desired task to achieve some goal needs to be articulated in the input language. The tasks are responses of the *User* and they need to be translated to stimuli for the *Input*. As pointed out above, this articulation

<sup>1</sup>In an earlier paper [5], qualities attached to each translation in the framework were referred to as distances

is judged in terms of the coverage from tasks to input and the relative ease with which the translation can be accomplished. The task is phrased in terms of certain psychological attributes that highlight the important features of the domain for the *User*. If these psychological attributes map clearly onto the input language, then articulation of the task will be made much simpler. An example of a poor mapping that is common in our everyday lives occurs in a large room with overhead lighting controlled by a bank of switches<sup>2</sup>. Many times it is desirable to control the lighting such that only a certain section of the room is lit. We are then faced with the puzzle of determining which switches control which lights. The consternation resulting from repeated experimentations with the switches to achieve the desired lighting effect can be traced to the difficulty of articulating a goal such as "Turn on the lights in the front of the room" in the input language, which consists of a linear row of switches which may or may not be oriented in a manner suggestive of their operation.

Examples in which articulation affects the ease or possibility of interaction with a computer systems are apparent as well. Much of the allure of virtual reality systems is achieved by novel input devices, such as a data glove, which are specifically geared toward easing articulation. Speech recognizers make it possible for disabled users to input text when typing via a standard keyboard is impossible. Systems which track the eyes of the user can make pointing tasks possible when input through a mouse or touch are not possible.

### 4.1.3 Performance

At the next stage, the responses of the *Input* are translated to stimuli for the *System*. Of interest in assessing this translation is whether the translated input language can reach as many states of the *System* as is possible using the *System* stimuli directly. This is similar to the notions of reachability defined by Dix and Runciman [54, 48] in the PIE model. For example, the remote control units for some compact disc players do not allow the user to turn the power off on the player unit, hence the off state of the player cannot be reached using the remote control's input language. On the panel of the compact disc player, however, there is usually a button which controls the power. The ease with which this translation from *Input* to *System* takes place is of less importance because the effort is not expended by the user. However, there can be a real effort expended by the designer and programmer. In this case, the ease of the translation is viewed in terms of the cost of implementation.

---

<sup>2</sup>This example comes courtesy of Donald Norman [125]

#### 4.1.4 Presentation

Once a state transition has occurred within the *System*, the execution phase of the interaction is complete and the evaluation phase begins. The new state of the *System* must be communicated to the *User*, and this begins by translating the *System* responses to the transition into stimuli for the *Output* component. This presentation translation must preserve the relevant system attributes from the domain in the limited expressiveness of the output devices. The ability to capture the domain concepts of the *System* within the *Output* is a question of possibility for this translation.

For example, while writing a paper with some word processing package, it is necessary at times to see both the immediate surrounding text where one is currently composing, say the current paragraph, and a wider context within the whole paper that cannot be easily displayed on one screen (say the current chapter). When moving files between the directories in a hierarchically arranged file system, such as Unix<sup>3</sup>, it is necessary for the user to know the current position in the directory structure in order to use relative path names effectively.

The Information Visualizer developed at Xerox PARC uses three dimensional presentation techniques to provide more context in displaying hierarchical information that would not otherwise be possible in two dimensions [33, 138]. They give the example of the display of the hierarchy of job positions in a large company. A standard two dimensional display of such an organizational chart would be far too large to display on a display screen. A three dimensional representation in the form of a cone-tree allows the whole chart to appear on the display with a readable portion on top. The advantage of the cone-tree is that it more easily displays the context within which a portion of the hierarchy is viewed. At all times it is possible to trace the path along the hierarchy from a viewed portion to the root.

The importance of this translation is that differences between two *System* states be present in the *Output*. Hidden differences result in an increased burden on the user who is trying to assess the result of previous input relative to a specific task goal while also trying to predict the outcome of future interactions based on the current output. Lack of ambiguity between distinct *System* states and the *Output* is at least a necessary condition for overall predictability and proper goal assessment.

#### 4.1.5 Observation

Ultimately, the user must interpret the output to evaluate what has happened. The response from the *Output* is translated to stimuli for the *User* which trigger assessment. The observation translation will address the ease and coverage of this final translation. For example, it is difficult to tell the time accurately on an

<sup>3</sup>Unix is a registered trademark of AT&T Laboratories.

unmarked analogue clock, especially if it is not oriented properly. It is difficult in a tty interface to determine the result of copying and moving files in a hierarchical file system. Typesetting a report using one of the popular typesetting programs available today is made virtually impossible without some previewing facility which allows rapid (and tree-saving) feedback to assess progress.

#### 4.1.6 Assessing overall interaction

The interaction framework is presented as a means to judge the overall usability of an entire interactive system. In reality, all of the analysis that is suggested by the framework is dependent on the current task (or set of tasks) in which the *User* is engaged. This is not surprising since it is only in attempting to perform a particular task within some domain that we are able to determine if the tools we use are adequate. In my own experience, I have found that different text editors are better at different things. For a particular editing task, I choose the text editor that I believe is best suited for interaction relative to the task. The best editor, if I am forced to choose only one, is the one that best suits the tasks that I most frequently need to do. Therefore, it is not too disappointing that we cannot extend the interaction analysis beyond the scope of a particular task.

A simple example of programming a VCR from a remote control shows that all four translations in the interaction cycle can affect the overall interaction. Ineffective interaction is indicated by the user not being sure the VCR is set to record properly. This could be because the user has pressed the keys on the remote control unit in the wrong order; this can be classified as an articulatory problem. Or maybe the VCR is able to record on any channel but the remote control lacks the ability to select channels, indicating a coverage problem for the performance translation. It may be the case that the VCR display panel does not indicate that the program has been set, a presentation problem. Or maybe the user does not interpret the indication properly, an observational error. Any one or more of these deficiencies would give rise to ineffective interaction.

Throughout this section, we have been referring to the two features of translations, the ease with which they happen and the possibilities or coverage they provide. With overall interaction, we can make similar distinctions between the ease of assessing the result of the previous interaction with the intended goal and the ability to achieve that goal. It is desirable to make assessment as easy as possible. However, if attainment is difficult at the same time, then the user is forced to traverse the interactive cycle many times, and this is not desirable. The best system would maximize ease of assessment and goal coverage.

More examples of how the interaction framework can be used to assess the overall effectiveness of interaction have been provided by Abowd and Beale [4, 5]. We will move on now to show how coverage and ease of translations can be formalized.

## 4.2 Formal properties of translations

A translation takes sequences of one language of events—the source language—to sequences in another language of events—the target language. In the interaction framework, the translations of *articulation*, *performance*, *presentation*, and *observation* each have their source and target languages fixed by some component. We introduce the generic set of translations with fixed source and target languages.

$$\text{Translation}[S, T] = S \leftrightarrow T$$

We refer to elements in the domain of a translation as source elements and elements in its range as target elements. We can define properties of the translation relation in terms of its *possibility*, *coverage* and *ambiguity*.

Possibility indicates how much of the source language is translatable. A translation  $t$  is said to be less partial than  $t'$  if every source element of  $t'$  is also a source element of  $t$ . The source elements of a translation are exactly the domain of the translation, so possibility can be phrased as the following predicate on domains:

$$\text{dom } t' \subseteq \text{dom } t.$$

Coverage indicates how much of the target language is expressible as a translation of source elements. A translation  $t$  is said to be more expressive than  $t'$  if every target element of  $t'$  is also a target element of  $t$ . The target elements of a translation are exactly the range of the translation. Coverage is therefore the same as the level of surjectivity of the translation. This property can be phrased as the following predicate on ranges:

$$\text{ran } t' \subseteq \text{ran } t.$$

For a given translation,  $t$ , a source element can be translated into many different target elements. In that sense, the translation is ambiguous. If the translation is unique for all source elements, then it is unambiguous. Two translations over the same source and target languages can be compared to determine if one is less ambiguous than the other. We say that  $t$  is less ambiguous than  $t'$  if every source element of  $t'$  has fewer possible translations in  $t$ . We represent this as the following predicate on image sets of source elements:

$$\begin{aligned} \forall es : (\text{dom } t' \cap \text{dom } t) \\ \bullet \quad t(\{ es \}) \subseteq t'(\{ es \}). \end{aligned}$$

We can combine the properties of possibility, coverage and ambiguity to create a partial ordering on translations which will give a comparative measure of definedness for translations on fixed source and target languages. A translation  $t$  is said to



*Correspondence* concerns the relationship between the state and stimuli of two agents. We assume there is a relationship between the underlying state spaces of the two agents and we investigate whether the operations that can be performed on one agent are sufficiently mirrored as operations on the other agent. With respect to interactive systems, a similar feature was introduced as *state display conformance* by Harrison and Dix [73] and has been further discussed by Abowd, Dix and Harrison [2] with respect to graphical interfaces. The notion of state display conformance is very similar to the notion of data refinement (or data reification), which has been treated by He *et al.* [78] and Coenen *et al.* [37] in a general relational format. We will use the ideas of data refinement in forming our definition of correspondence.

The purpose of data refinement is to characterize how the function of an abstract description of a state machine can be adequately captured in a more concrete representation. A concrete representation is supposed to be indistinguishable from its abstract counterpart, in the sense that an outside observer interacting with the concrete representation cannot distinguish it from the abstract one. The concrete representation is then a suitable *simulation* of the abstract [76, 170]. We are not concerned here with the abstract-concrete refinement. Rather, we are concerned with the relationship that exists between the *System* and the interface (*Input* along with *Output*) of an interactive system. The observer in this case is the *User*, and we want to stipulate that operations performed on the interface which the *User* observes accurately reflect operations being performed on the *System*. In this section, we will develop a formal definition of correspondence and in Chapter 7 we will investigate further its relevance to interactive systems.

We will describe correspondence as a constraint between two agents, say *A* and *B*. Figure 4.2 gives a picture of the correspondence we intend to describe. We want




Figure 4.2: Agent correspondence

agent *B* to adequately mirror some functionality of agent *A*, i.e., we want operations performed on *A* to be matched by operations performed on *B*. The matching is done in terms of a given relationship which holds between the state descriptions of



$A$  and  $B$ , and we will call this relation *retrieve* (of type  $B.states \leftrightarrow A.states$ ). This relation determines from the state of  $B$  what the state of  $A$  should be, and it is understood as the interpretation  $A$  can make of the state of  $B$ .

To define this correspondence, we begin by asserting that the initial states of  $A$  and  $B$  are related by the *retrieve* relation. Then for any state that  $A$  can attain there must have been a way for  $B$  to attain a related state. And for any future state that  $A$  can attain,  $B$  must be able to attain a related state. Note that we do not demand that at all times  $A$  and  $B$  are in related states. The transitions of  $A$  are what determine the states that  $B$  must eventually attain, and not the other way around.

The schema *AgentCorr* gives the formal definition motivated by the above description of correspondence.

$AgentCorr$
$Agent^I$ $Agent^J$ $retrieve : states^J \leftrightarrow states^I$
$inits^J \times inits^I \subseteq retrieve$
$\forall (tr_I, s_I), (tr_J \cap tr'_J, s'_J) : I_{\theta Agent^I}^{int}$ $\quad   \quad (s_I, s'_J) \in tr'_J \circ transform^I$ <ul style="list-style-type: none"> <li>• <math>\exists (tr_J, s_J), (tr_J \cap tr'_J, s'_J) : I_{\theta Agent^J}^{int}</math>  <math>\quad   \quad (s_J, s'_J) \in tr'_J \circ transform^J</math>  <ul style="list-style-type: none"> <li>• <math>\{(s_I, s_J), (s'_J, s'_J)\} \subseteq retrieve</math></li> </ul> </li> </ul>

We define *correspond\_* as a relation between agents indexed by a retrieving relation between their state sets. The pair of agents  $(D, R)$  is in *correspond<sub>ret</sub>* if agent  $R$  corresponds to agent  $D$  as defined by the schema *AgentCorr* with retrieving relation *ret*.

$correspond\_ : (State \leftrightarrow State) \rightarrow (Agent \leftrightarrow Agent)$
$correspond_{ret} = \{$ $\quad AgentCorr$ $\quad   \quad retrieve = ret$ $\quad \bullet \quad \theta Agent^I \mapsto \theta Agent^J \}$

We have given in this section a very coarse treatment of correspondence. In Chapter 8 we will see how correspondence can be treated more constructively in the refined agent model using templates. The purpose of the constructive treatment is to formally link the empirical psychological evidence which identifies useful template information with the heuristic motivation behind multiagent architectures.

## 4.4 Predictability

Users interact with any system in order to satisfy certain goals. Fundamental to goal satisfaction is the ability to plan interaction which will achieve the goal. In a graphic drawing package, the user may want to select one or more of a collection of overlapping figures in order to perform some similar operations to them. In this case it is necessary that the user be able to predict how to select the figures of interest. *Predictability* is a property of interaction concerning the degree of confidence with which a user can determine the effect subsequent task execution will have on the achievement of the goal. Therefore, predictability is necessary for effective interaction and is a crucial measuring stick by which to gauge the usability of an interactive system. In the following, we will formalize how predictability can be expressed on agents and characterize levels of predictability within the interaction framework.

A formal treatment of predictability was first given with the PIE model by Dix and Runciman [54]. Predictability in the PIE model is defined by an unambiguous interpretation function between programs and effects. When two programs lead to the same effect, then future experimentation will betray no difference between them. Letting  $P$  represent the set of programs,  $E$  the effects set and  $I$  the interpretation function between programs and effects ( $I : P \rightarrow E$ ), a predictable PIE satisfies the following predicate.

$$\forall p, q, r : P \mid I(p) = I(q) \bullet I(p \wedge r) = I(q \wedge r)$$

This says that the current effect is enough to predict future behaviour.

In terms of the agent model, predictability can be seen as an observable relationship between the stimuli and responses. Based on the Sufrin and He approach to classifying interactive properties as relationships between input histories and state or response histories, we can model this predictability (and other interactive properties) within the simple agent model.

Viewing an agent as another representation for a PIE, we can see that the sequences of stimuli are the programs and the history of responses are the effects. From the definition of an agent we derived an external interpretation relation,  $I_A^{ext}$ , which provides a relationship between a program (sequence of stimuli) and the resulting responses.  $I_A^{ext}$  is a relation, not a function, and so our formulation of predictability is slightly more cumbersome than for the PIE model. One program can produce more than one possible sequence of responses. Two different programs,  $p$  and  $q$ , which yield the same possible responses are said to be *externally equivalent*, and we write this as  $p \overset{ext}{\approx}_A q$ .

$$\left| \begin{array}{l} \_ \overset{ext}{\approx} \_ : Agent \rightarrow (seq\ Event \leftrightarrow seq\ Event) \\ \_ \overset{ext}{\approx}_A \_ \subseteq (\text{dom } I_A^{ext}) \times (\text{dom } I_A^{ext}) \\ p \overset{ext}{\approx}_A q \Leftrightarrow I_A^{ext}(\{\{p\}\}) = I_A^{ext}(\{\{q\}\}) \end{array} \right.$$

Programs that are initially externally equivalent may cease to be so after the next stimulus event is received and the agent transforms itself and responds. If any such experimentation on equivalent programs betrays no such difference, then the programs are said to be *externally indistinguishable*, and we write  $p \overset{ext}{\approx}_A q$ .

$$\left| \begin{array}{l} \_ \overset{ext}{\equiv} \_ : Agent \rightarrow (seq\ Event \leftrightarrow seq\ Event) \\ \_ \overset{ext}{\equiv}_A \_ \subseteq (\text{dom } I_A^{ext}) \times (\text{dom } I_A^{ext}) \\ p \overset{ext}{\equiv}_A q \Leftrightarrow \forall r : seq\ Event \bullet (p \frown r) \overset{ext}{\approx}_A (q \frown r) \end{array} \right.$$

A predictable agent would be one in which all externally equivalent programs are externally indistinguishable. Since it is trivially true that all externally indistinguishable programs are externally equivalent (let  $r = ()$  in the definition above), predictable agents can be characterized by the equality of the relations  $\overset{ext}{\approx}_A$  and  $\overset{ext}{\equiv}_A$ .

$$\left| \begin{array}{l} \text{Predictable} : \mathbf{P}\ Agent \\ \hline A \in \text{Predictable} \Leftrightarrow (\_ \overset{ext}{\approx}_A \_) = (\_ \overset{ext}{\equiv}_A \_) \end{array} \right.$$

In the same way that the PIE model can address varying levels of abstraction by suitable definition of the program and effects sets, we also can apply predictability at different levels within the framework. Three levels naturally arise from the framework, and we can investigate what predictability means at each of these levels. Figure 4.3 depicts the various levels we will discuss.

At the lowest level, we consider stimuli and responses of the *System* alone. Predictability at this level we call *algorithmic predictability*. This is the simplest form of predictability possible and it is a necessary condition for the remaining levels. We assume that algorithmic predictability is satisfied so that all programs issued to the *System* which are externally equivalent are also externally indistinguishable.

At the next level, we consider the physical interface. Predictability at this level is *perceivable predictability* because it is here that contact with the *User* is possible. Perceivable predictability says that the information provided by the history of *Output* responses is enough to determine how future programs received by the *Input* will affect the *Output*. It is therefore possible that the *User* could predict the outcome of future input.




Figure 4.3: Levels of predictability

At the highest level, we move beyond the world of keystrokes and displays to include how the *User* observes the *Output* and understands the *Input*. *Interactive predictability* means that that which the *User* understands of the history of *Output* responses is enough to determine the effect of future tasks. Many systems can satisfy perceivable predictability and yet remain unusable because of the way the *Output* responses are observed by the *User*.

One feature of the agent model which we rely on is its compositionality. This will be further developed in Chapter 5. For now, we will assume that components of the framework, which are defined as agents, and the translations, which are external descriptions of agents, can be composed to form more complex agents. If we combine all of the components and translations of the interaction framework except for the *User* we have a user's model of the system, or, more precisely, the designer's model of the user's model of the system, depicted in Figure 4.4. The external interpretation

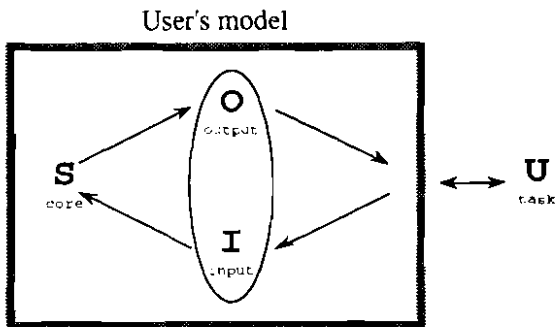


Figure 4.4: The user's model of the system

relation for this agent, which we refer to as  $I_{User's\ model}^{ext}$ , relates information at the level of how *Input* stimuli and *Output* response are understood by the *User*. Hutchins, Hollan and Norman refer to this level of information as the meaning of the input and output interface languages[85].

A colloquial expression for this interactive predictability would be

*What you have done up to now determines what you will see from now on.*

There are two separate aspects of this statement to consider, what is seen and what is remembered.  $I_{User's\ model}^{ext}$  makes explicit the assumptions about what is seen and how the *User* understands the responses from the *Output*. For example, when using a calculator, it may be a safe assumption that numbers displayed are interpreted in base 10, so imagine the surprise and consternation when the user enters 9+4 and gets back the response D! What is remembered is important because sometimes the order of stimuli and responses contains information crucial for predictability. For example, in an interactive drawing package, the order in which overlapping objects are created can determine the effect of clicking the mouse button to select objects when the selection cursor is located in the overlapping region.

Predictability as defined above assumes both perfect memory and equivalence over the entire history of the interaction. This is not satisfactory since users forget and similarity of future responses does not depend on exact similarity of past responses. Predictability is a feature of an interactive system which enhances the interaction. Formalizing what it means to be predictable highlights the cost associated in order to realize the benefits of predictability. Our earlier definition has too high a cost, and so we would like to investigate ways of reducing that cost. One way to reduce the cost is by revising the definition of equality so that it is not over entire response histories but over some subset and only on particular attributes associated with the responses as well.

## 4.5 Nondeterminism

Predictability only considered the external behaviour of an agent. The responses from an agent occur after the state transition fired by the stimulus. There is more information available at the state level than at the response level. We can apply the method used for predictability on responses to formalize an internal version of predictability. The internal interpretation relation,  $I_A^{int}$  relates programs to the possible states the agent can be in after participating in the program. Internally equivalent programs result in the same possible states. Internally indistinguishable programs will betray no state difference.

$$\begin{array}{|l}
 \hline
 - \approx_{-}^{\text{int}} \rightarrow \\
 - \equiv_{-}^{\text{int}} : \text{Agent} \rightarrow (\text{seq Event} \leftrightarrow \text{seq Event}) \\
 \hline
 - \approx_A^{\text{int}} \subseteq (\text{dom } I_A^{\text{int}}) \times (\text{dom } I_A^{\text{int}}) \\
 - \equiv_A^{\text{int}} \subseteq (\text{dom } I_A^{\text{int}}) \times (\text{dom } I_A^{\text{int}}) \\
 p \approx_A^{\text{int}} q \Leftrightarrow I_A^{\text{int}}(\{p\}) = I_A^{\text{int}}(\{q\}) \\
 p \equiv_A^{\text{int}} q \Leftrightarrow \forall r : \text{seq Event} \bullet (p \hat{\ } r) \approx_A^{\text{int}} (q \hat{\ } r) \\
 \hline
 \end{array}$$

A *deterministic* system is internally predictable, so that internally equivalent programs are also internally indistinguishable.

$$\begin{array}{|l}
 \hline
 \text{Deterministic} : \text{P Agent} \\
 \hline
 A \in \text{Deterministic} \Leftrightarrow (- \approx_A^{\text{int}} -) = (- \equiv_A^{\text{int}} -) \\
 \hline
 \end{array}$$

The state holds enough information to determine future internal behaviour.

We have modelled the behaviour of an agent as a relation which reflects two types of nondeterminism—true randomness or ‘don’t care’ nondeterminism. True randomness is very rare in computers. ‘Don’t care’ nondeterminism comes from abstracting away details that will be present in any implementation but do not matter for the purpose of the current description. Don’t care nondeterminism is very common and encouraged in abstract specification. It is the main reason we have modelled behaviour as a relation rather than as a function. The definition of deterministic above essentially rules out internal randomness. It does not rule out nondeterminism that arises from abstraction.

Another form of nondeterminism arises from how the user perceives the system. To a naive user even the simplest of interactive systems is nondeterministic. An example used by Dix [48, 49] and Thimbleby [161] is of a prime number generator. Even though the algorithm for deriving the sequence of increasing primes is deterministic, if the user does not know that the system is producing primes, the sequence 2, 3, 5, 7, 11, ... may appear somewhat random. It is only through experience interacting with the system that the user can begin to understand its deterministic functionality. And this user-perceived nondeterminism arises naturally even for an expert, as uncertainty about the past interaction history leads to uncertainty of the current state. Therefore, even though an agent’s behaviour may satisfy the definition of deterministic above, it can still seem nondeterministic to its observer.

Equivalence was based on all state information. Only part of the state information may be important for a given task, and so we may want to analyze whether the agent was deterministic relative to that part of the state. To do this, we will need a way to highlight specific information about the state. The simple agent model does

not allow us to constructively define such a restricted view, and so we define the state of the refined agent model in Chapter 5 in terms of attribute-value mappings. Then in Chapter 7, we define restrictions on agents by means of templates, which are simply formed as subsets of the agent's state attributes.

## 4.6 Synthesis

Predictability assumes that the user has some knowledge of how the system works, i.e., that the user has internalized a model of how the system works. Expert users will have a more complete internalized model than novices, which reflects their greater experience. Most users, expert or novice, will continually be presented with situations which are new to them and for which their model is not able to provide enough information to predict future outcomes of subsequent input. In these cases, the user will experiment with the system by offering some input and trying to detect the effects which result, hence building up more information in their own model. We call this process of internalizing input-output relationships *synthesis*. Synthesis is complementary to predictability. A predictable system is one in which it would be possible for the user to internalize a model that would be of benefit to future interaction. Predictability is done by reasoning forward in time based on information available at present. Synthesis involves reasoning back in time in order to determine how the present information was achieved from the past. Figure 4.5 shows how predictability and synthesis relate to the interaction framework.




Figure 4.5: Predictability and synthesis

The observer of an agent only has access to the stimulus-response behaviour of the agent. Since the stimuli cause state transitions which in turn determine the responses, effective synthesis depends on changes to the underlying state being

reflected in the responses. If a program changes the underlying state, then the subsequent responses should indicate this. In order for the observer to have the chance of correctly synthesizing the effect of the previous input, the response should occur before any other changes are made to the internal state. An *honest agent* will satisfy the property that externally equivalent programs are also internally equivalent. A weaker property would only require that externally indistinguishable programs be internally equivalent; such agents are deemed *trustworthy*<sup>4</sup>.

$$\begin{array}{|l}
 \hline
 \text{Honest, Trustworthy : P Agent} \\
 \hline
 A \in \text{Honest} \Leftrightarrow \_ \overset{ext}{\approx} \_ A \_ \subseteq \_ \overset{int}{\approx} \_ A \_ \\
 A \in \text{Trustworthy} \Leftrightarrow \_ \overset{ext}{\equiv} \_ A \_ \subseteq \_ \overset{int}{\approx} \_ A \_ \\
 \hline
 \end{array}$$

A good example of the difference between honesty and trustworthiness can be seen with search and replace facilities in most text editors. A single search for a string of text followed by the replacement with another string of text is honest, for the replaced text is usually displayed on screen right where the old text was located. Most global search and replace facilities do not display all of the changes made, leaving the user to confirm the changes made by subsequent browsing of the text. Global facilities such as this can only ever be as good as trustworthy.

## 4.7 Consistency

*Consistency* is prominent in the HCI and ergonomic literature [113, 142]. It is put forth as a cause for increased learnability of a system [142, 93, 32]. In reading different accounts of consistency, there is clearly a debate about what exactly is meant by consistent at the psychological level (witness, for example the debate on consistency between Grudin [64] and Reisner [136]). Such informal arguments highlight even more the importance of a formal framework in which different formulations of consistency can be compared and contrasted.

As an example of how we can formalize consistency, we will take a simple definition. Our informal definition of consistency is that the same input in similar situations has similar effects. This is a generalization of predictability in which the same input in the same situation has the same effect. The criteria for deciding similarity are crucial, just as the criteria for equivalence were for predictability. And as was done for predictability, we can discuss consistency internally and externally. Consistency of an agent is defined with respect to the similarity criterion,

<sup>4</sup>Sufrin and He originally used the names honest and trustworthy to apply to individual programs which satisfy conditions somewhat similar to those described here. Their definition, however, demanded that programs which were not indistinguishable internally were either not equivalent (for honesty) or not indistinguishable (for trustworthy) externally.



which is modelled as an equivalence relation. Internal consistency depends on a similarity relation, say  $\sim$ , which is an equivalence relation on the states of the agent. The state transition relation associated to any program must then respect the equivalence classes of  $\sim$ .

$$\begin{array}{|l}
 \text{IntConsistent} : \mathcal{P}(\text{Agent} \times (\text{State} \leftrightarrow \text{State})) \\
 (A, \sim) \in \text{IntConsistent} \Leftrightarrow \\
 \quad (\text{equivalence}(\sim) \\
 \quad \wedge \sim \subset A.\text{states} \times A.\text{states} \\
 \quad \wedge \forall p : \text{seq Event}; E1, E2 : \text{classes}(\sim) \\
 \quad \bullet E1 \triangleleft (\text{transformExtend } p) \subseteq E2
 \end{array}$$

A similar definition applies for external consistency, with the only difference being that the similarity relation is defined on the histories of responses for the agent. The interesting question to ask now is whether for an externally consistent agent  $A$  with respect to response similarity criterion  $\sim$ , is there a corresponding state similarity criterion  $\sim'$  such that  $A$  is internally consistent with respect to  $\sim'$ ? Within the interaction framework, it would be possible for the *User* to detect the external consistency of the interface and it would be useful if that consistency mapped down to some consistency internal to the *System* which related to the domain of application. Likewise, the *User* may rely on some internal consistency of the *System* which would then be expected to be manifested externally at the interface.

Whereas this definition of consistency is simple to express, we must exercise caution in understanding what it actually is saying. First of all, we had to assume that there was a way to partition the state or responses by some nontrivial similarity criterion. And we assume that this similarity criterion is relevant to the observer for external consistency. If such similarity criteria exist, then an advantage of this definition of consistency is that it is cumulative. By this we mean that if we have two similarity criteria then their conjunction is also a similarity criterion (though there again is no guarantee that it will be relevant to the observer). Therefore, consistency as we have defined it also has many levels of application. A word of caution, however, since compounded equivalence relations will tend to have many more and smaller equivalence classes, until the compounded similarity criterion gives no more information than the trivial one.

Another possible disadvantage of this definition is its application to all possible programs. It seems more realistic to assume that consistency properties will be relevant for subsets of programs. Finally, consistency is defined with respect to the similarity criteria. Therefore, an agent consistent by one similarity criterion may be inconsistent by another, suggesting a trade-off to be made during design.

We have only considered one formalization of consistency here. An obvious other consistency notion would have similar input in the same situation resulting

in similar behaviour, but we have not investigated this possibility yet.

## 4.8 Conclusions

In this chapter we have provided both qualitative and quantitative evidence for properties of interactive system which enhance the interaction between user and system. The qualitative accounts have appealed mainly to a common sense understanding of how the various translations in the framework impact on the overall effectiveness of interaction. Two important features of any translation were identified as its coverage or possibility of converting one language to another and the ease with which this conversion can occur. The interaction framework provided a mechanism for applying heuristics for effective interaction uniformly.

Though we strongly urge the use of heuristic reasoning in the design and analysis of interactive systems, our main emphasis in this thesis is on transferring that heuristic knowledge onto a more rigorous software engineering platform based on formal methods. The attempts to formalize various interactive properties in terms of the simple agent model have shown that the formalism can capture the intended meaning of those properties. But more importantly, by making our definitions of the properties such as predictability and consistency formal, we can more easily see what it is about those properties that we have *not* captured in the formalism.

For example, we can see that our definition of predictability assumes we can judge equality over complete response histories for an agent, and this seems too stringent a requirement on the observer of the agent, especially when that observer is a user. We require a way to express predictability so that it embodies more realistic constraints on the observer.

Though the agent model breaks open the black box of the PIE model somewhat, allowing expression of properties such as correspondence and honesty, it is clear that we have not opened the box wide enough. We still have not provided an easier way to describe the behaviour of an agent in order to accommodate proof of satisfaction of various interactive properties. We have only slightly motivated the need for a compositional agent model, and this needs further investigation.

One of the advantages of the red-PIE model and of Sufrin and He's model of interactive processes was that they both distinguished that which is discernible by the user (the display) from that which the user wants to achieve from the interaction (the result). We have made light of this distinction in the simple agent model. Displays are immediate and ephemeral indications from a system and we have equated those with responses of an agent. Results are more permanent and we have equated them with the underlying state of the agent. This is not satisfactory because displays are more like state machines than a sequence of responses would indicate and the state of an agent may contain information that is not relevant to

State-Response relationship	Interactive Property
$(-\overset{ext}{\approx}_A -) = (-\overset{ext}{\equiv}_A -)$	predictability
$(-\overset{int}{\approx}_A -) = (-\overset{int}{\equiv}_A -)$	deterministic
$(-\overset{ext}{\approx}_A -) \subseteq (-\overset{int}{\approx}_A -)$	honesty
$(-\overset{ext}{\equiv}_A -) \subseteq (-\overset{int}{\approx}_A -)$	trustworthiness

Table 4.1: Interactive properties expressed as state-response relationships

the end result from the user's point of view. We will need a way to more faithfully represent these crucial notions of display and result.

We have adopted Sufrin and Ile's approach to classification of interactive properties in terms of equivalence and indistinguishability because the structure of the classification suggests the possible properties that can be expressed. Table 4.1 lists the properties we have discussed in terms of equivalence and indistinguishability of state and response (internal and external). A similar classification approach which takes into account result and display information will be the subject of Chapter 7.

Thus we have set the agenda for the remainder of this thesis. We begin with a refinement of the agent model in Chapter 5 followed by a definition of a design language for describing agents in Chapter 6. The state of the agent will be described in terms of attributes and so we will be able to extract partial information, or templates on the state. We will then show how the refined agent model allows the expression of further interactive properties in Chapter 7. The introduction of attributes and templates will lead to a more direct treatment of the result-display relationships that will take a similar form to the state-response relationships. The description of agents in the refined model is more closely linked to empirical psychological evidence of how users understand the tasks the agents are designed to support. The interactive properties will also be based on that task-specific evidence, providing for a user-centred methodology for the design and analysis of an interactive system. In Chapters 7 and 8 we will show how the methodology is implied by the description of further interactive properties and the investigation of heuristic multiagent architectures.

# Chapter 5

## Refining the agent model

In preceding chapters we have discussed an overall framework for the description of an interactive systems and how it provides a basis for the description of general properties of interaction expressed within a formal model. The purpose of this chapter is to refine the model of the agent used in the previous chapters. The next chapter then provides a design language for agents.


The agent model we present in this thesis is directly motivated by the formal presentation of the interaction framework. We desire a language that will allow the discussion of interactive properties and with which it will be possible to describe a complete system. The formal interaction framework provided a means for a precise discussion of interactive properties, but it did not allow a realistic description of a complex system. What is required is a compositional model, and this is what the refined agent model is intended to provide.

The refinement of the agent model in this chapter is a necessary digression. We have seen the need for more detail of the structure of an agent in order to describe interactive properties more closely related to the way a user perceives interaction with respect to the goals of tasks. The refinement is also necessary from a software engineering point of view. To bring HCI considerations into earlier stages of design it is not sufficient to merely provide an abstract mathematical model for describing interactive properties. The model must be expressible by some design language and accompanied by some method for directing the description. In this chapter, we address the design issue by providing a compositional agent model.

### Overview of chapter

In Section 5.1 we will formulate the requirements placed on our agent model for its use in the design and analysis of interactive systems. One of the major requirements is that the refined model supports a modular description of complex agents by the composition of smaller and simpler agents.

The description of the refined agent model is separated into three stages representing the internal state-based specification, the external event-based specification and the communication specification, which links events to internal operations. Hence, the outline for the refined agent model is as below.



In Sections 5.2, 5.3 and 5.4 we will address separately the internal, external and communication specifications of the refined agent model. In each section, the model for the particular specification will be given and then it will be shown how separate specifications can be composed associatively. There are two ways we will compose agents, by interleaving and by synchronization. In Section 5.5, we will unite the separate specifications and show how interleaving and synchronization are defined between agents.

In Section 5.6, we will show how the refined agent model relates to the simpler model of Chapter 2 by defining the internal and external interpretation functions.

## 5.1 Requirements for agents

An agent has state and a means for communicating stimulus and response events. A pictorial view of an agent is given in Figure 5.1. The stimulus and response events




Figure 5.1: Pictorial representation of an agent

are communicated along input and output channels. There are several requirements

on the refined model which motivate its development in this chapter and the next. Here we will discuss requirements relating to compatibility with the simple agent model and compositionality.

We have already given a simple definition of an agent in Chapter 2 which is adequate for expressing a variety of interactive properties at the abstract level. The new model must be consistent with the older model. The two most important features of the simpler model were the derived internal and external interpretation relations,  $I_A^{int}$  and  $I_A^{ext}$ , so we must be able to derive similar relations with the new model. It turns out that though we can express the external interpretation function,  $I_A^{ext}$ , we will not use it in formulating properties relating result and display in Chapter 7.

In Chapter 4, we assumed that agents could be composed to form more complex agents. We will need to demonstrate to what degree the refined agent model allows the construction of complex agents through the combination of existing agents. Given two agents,  $A$  and  $B$ , there will be two ways of combining them to achieve a third agent.

We may want to treat  $A$  and  $B$  as independent agents that can only affect each other's private state by means of synchronous message passing along common communication channels. The agent derived by this synchronous composition we will denote by  $compose_{sync}(A, B)$ . There are various conditions that are imposed on  $A$  and  $B$  so that they may be synchronously composed. They must have separate state spaces with no common attributes. Any common channels must be an input for one agent and an output for the other. The common channels become synchronized, meaning that communications that occur along them must be participated in by both component agents simultaneously. The synchronized channel is no longer visible outside the composed agent. Figure 5.2 is a graphic depiction of synchronous composition.

We may want to treat  $A$  and  $B$  as dependent agents. In this case, some part of their internal states is shared between them. Such dependent composition we describe by interleaving their descriptions to obtain another agent, denoted by  $compose_{int}(A, B)$ . There are conditions we impose on interleaving composition as well. The separate agent's states may be defined over common attributes and they may have common input or output channels. The effect of operations defined for both agents is constrained such that they only alter the values of common attributes. Figure 5.3 graphically depicts this composition by interleaving. The overlapping circles are meant to indicate shared state information. Note that interleaved agents cannot communicate via synchronous channels, so an input for one agent cannot be an output for the other agent. In this sense, the agents are unlinked.

In Section 5.5, we present two main theorems concerning the associativity of synchronous and interleaving composition. The proof of these theorems is dealt with in terms of smaller theorems on the associativity of composition for the in-




Figure 5.2: Synchronous composition of independent agents

ternal, external and communication specifications of the previous sections. The associativity of composition will allow any number of agents to be composed either all synchronously or all by interleaving without worrying about the order of the composition.

## 5.2 Internal specification

An agent has its own private and persistent state. The possible set of states for an agent is described by the values associated to some finite set of attributes. We assume a set of attribute names and a universal set of values. A state is a finite function from attributes to values.

$$[\mathcal{A}, \mathcal{V}]$$

$$\text{State} == \mathcal{A} \mapsto \mathcal{V}$$

A *state set* is a set of states, all of whose elements are defined over the same set of attributes. In addition, the attributes of a state set have *types* associated with them. A *type* is a nonempty subset of  $\mathcal{V}$ . The values of an attribute are restricted to being in that attribute's type.




Figure 5.3: Interleaving composition of dependent agents

<p><i>Stateset</i></p> <p><i>attrs</i> : <math>F \mathcal{A}</math></p> <p><i>type</i> : <math>\mathcal{A} \leftrightarrow \mathbf{P}_1 \mathcal{V}</math></p> <p><i>states</i> : <math>\mathbf{P}_1 \text{State}</math></p> <hr/> <p><i>attrs</i> = dom <i>type</i></p> <p><math>\forall s : \text{states}, a : \text{attrs} \bullet s(a) \in \text{type}(a)</math></p> <p><i>states</i> <math>\subseteq (\text{attrs} \rightarrow \mathcal{V})</math></p>
---

The internal specification of an agent gives the state set along with operations defined on that state set. Each operation is labelled by a *message*, obtained from a set of all possible messages. The internal specification is completed by giving the set of initial states.

[*Message*]

<p><i>InternalSpec</i></p> <p><i>Stateset</i></p> <p><i>inits</i> : <math>\mathbf{P}_1 \text{State}</math></p> <p><i>operations</i> : <i>Message</i> <math>\leftrightarrow</math> {<i>states</i> <math>\leftrightarrow</math> <i>states</i>}</p> <p><i>messages</i> : <math>\mathbf{P} \text{Message}</math></p> <hr/> <p><i>inits</i> <math>\subseteq</math> <i>states</i></p> <p><i>messages</i> = dom <i>operations</i></p>
--



## Combining internal specifications

Two separate internal specifications can be combined to create a new internal specification. One constraint on the combination concerns the separate state sets—they must agree over all common attributes. If  $s_1$  is a state of one internal specification and  $s_2$  is a state of the other internal specification, and the values of  $s_1$  and  $s_2$  agree over their common attributes, i.e.,  $s_1$  and  $s_2$  are *compatible*, then they define a state in the combined internal specification. A simple way to express this compatibility constraint takes advantage of the functional representation of a state and says that two states are compatible if overriding one by the other is equal to the functional union of the two.

$$\begin{array}{l} \text{compatible} : \text{State} \leftrightarrow \text{State} \\ \hline (s_1, s_2) \in \text{compatible} \Leftrightarrow s_1 \oplus s_2 = s_1 \cup s_2 \end{array}$$

The schema *JoinStatesets* characterizes the combination of two state sets to form a third state set whose elements are the union of such compatible pairs of states. We have stipulated that the state sets are nonempty, so state sets which are to be joined must have at least one pair of compatible states. We choose the name *join* because it is suggestive of the similar operation for relational databases. Joining is also similar to schema signature combination in Z, where components with the same name from two different schemas must be of the same type and take on the same values [153]. Note that the *I* and *J* are used in the following schema to decorate the argument state sets and ' is used to decorate the state set resulting from the join.

$$\begin{array}{l} \text{JoinStatesets} \\ \hline \text{Stateset}^I \\ \text{Stateset}^J \\ \text{Stateset}' \\ \hline \forall a : \text{attrs}^I \cap \text{attrs}^J \bullet \text{type}^I(a) = \text{type}^J(a) \\ \exists s_I : \text{states}^I; s_J : \text{states}^J \bullet (s_I, s_J) \in \text{compatible} \\ \text{attrs}' = \text{attrs}^I \cup \text{attrs}^J \\ \text{type}' = \text{type}^I \cup \text{type}^J \\ \text{states}' = \{ s_I : \text{states}^I; s_J : \text{states}^J \\ \quad | (s_I, s_J) \in \text{compatible} \\ \quad \bullet s_I \cup s_J \} \end{array}$$

For convenience, we can represent the above schema definition as a binary operation *join* on *Stateset*. We will frequently make use of this mechanical technique

for characterizing a combination via a schema and then converting it to a binary operation.

$$\left| \begin{array}{l} \text{join} : (\text{Stateset} \times \text{Stateset}) \leftrightarrow \text{Stateset} \\ \text{join} = \{ \text{JoinStatesets} \\ \quad \bullet (\theta \text{Stateset}^I, \theta \text{Stateset}^J) \mapsto \theta \text{Stateset}^I \} \end{array} \right.$$

As we described in Section 5.1, we will want to compose several agent descriptions and we want that composition to be associative. The proof of the associativity of the overall agent composition functions is distributed throughout this and the next two sections. The first theorem below shows that joining of states sets satisfies associativity.

**Theorem 5.1**

$$\begin{array}{l} \forall S1, S2, S3 : \text{Stateset} \\ | ( (S1, S2) \in \text{dom join} \\ \quad \wedge (\text{join}(S1, S2), S3) \in \text{dom join} ) \\ \bullet \text{join}(\text{join}(S1, S2), S3) = \text{join}(S1, \text{join}(S2, S3)) \end{array}$$

**PROOF OF THEOREM 5.1:**

The proof proceeds in two stages. First, under the conditions of the hypothesis, namely that

$$\begin{array}{l} (S1, S2) \in \text{dom join} \\ \wedge (\text{join}(S1, S2), S3) \in \text{dom join} \end{array}$$

we must show that

$$\begin{array}{l} (S2, S3) \in \text{dom join} \\ \wedge (S1, \text{join}(S2, S3)) \in \text{dom join} \end{array}$$

After proving that  $(S1, \text{join}(S2, S3)) \in \text{dom join}$ , we must show the equality of the different joins. Both stages of this proof rely on a lemma on state compatibility which says that if the union of two compatible states is compatible with a third state, then each of those states is compatible separately with the third state. Full details of the proof of this lemma and the stages of this theorem are provided in Appendix B.

◇ END OF PROOF OF THEOREM 5.1

To complete the definition of composition of internal specifications, we must show how the operations of two separate internal specifications are coalesced to derive the operations of the combined internal specification. The explanation behind this derivation is relatively straightforward. The messages of the new internal

specification fall into one of three sets—messages unique to one specification, or unique to the other, or common to both. If a message is unique to one of the internal specifications, then it is lifted to an operation which is consistent with the original specification and behaves like the identity transition elsewhere. If a message is common to both internal specifications, then it is lifted to an operation which is consistent with both of the original specifications. The derivation of the operations set reflects the three mutually exclusive conditions on messages of the combined internal specification, and hence is a bit large. We could have been more concise with the description of *operations*, but this large definition mirrors the way we regard the operations set in the associativity proofs which follow.

The schema *CombineInt* defines how two internal specifications can be combined to form a third internal specification.

<p><i>CombineInt</i></p> <p><i>InternalSpec</i><sup>I</sup></p> <p><i>InternalSpec</i><sup>J</sup></p> <p><i>InternalSpec</i><sup>I</sup></p> <p><i>JoinStatesets</i></p> <p><math>\exists s_0^I : \text{inits}^I; s_0^J : \text{inits}^J \bullet (s_0^I, s_0^J) \in \text{compatible}</math></p> <p><math>\text{inits}' = \{ I : \text{inits}^I; s_0^J : \text{inits}^J \mid (s_0^I, s_0^J) \in \text{compatible} \bullet s_0^I \cup s_0^J \}</math></p> <p><math>\text{messages}' = \text{messages}^I \cup \text{messages}^J</math></p> <p><math>\forall m : \text{messages}' - \text{messages}^J</math></p> <ul style="list-style-type: none"> <li>• <math>\text{operations}'(m) =</math> <math display="block">\{ s, s' : \text{states}'</math> <math display="block">\mid ( (\text{attrs}^I \triangleleft s, \text{attrs}^I \triangleleft s') \in \text{operations}^I(m)</math> <math display="block">\wedge (\text{attrs}^I \triangleleft s, \text{attrs}^I \triangleleft s') \in \text{id}(\text{KnotIattrs} \rightsquigarrow \mathcal{V}) \}</math> <ul style="list-style-type: none"> <li>• <math>s \mapsto s'</math></li> </ul> </li> </ul> <p><math>\forall m : \text{messages}^J - \text{messages}^I</math></p> <ul style="list-style-type: none"> <li>• <math>\text{operations}'(m) =</math> <math display="block">\{ s, s' : \text{states}'</math> <math display="block">\mid ( (\text{attrs}^J \triangleleft s, \text{attrs}^J \triangleleft s') \in \text{operations}^J(m)</math> <math display="block">\wedge (\text{attrs}^J \triangleleft s, \text{attrs}^J \triangleleft s') \in \text{id}(\text{KnotJattrs} \rightsquigarrow \mathcal{V}) \}</math> <ul style="list-style-type: none"> <li>• <math>s \mapsto s'</math></li> </ul> </li> </ul> <p><math>\forall m : \text{messages}^I \cap \text{messages}^J</math></p> <ul style="list-style-type: none"> <li>• <math>\text{operations}'(m) =</math> <math display="block">\{ s, s' : \text{states}'</math> <math display="block">\mid ( (\text{attrs}^I \triangleleft s, \text{attrs}^I \triangleleft s') \in \text{operations}^I(m)</math> <math display="block">\wedge (\text{attrs}^J \triangleleft s, \text{attrs}^J \triangleleft s') \in \text{operations}^J(m) \}</math> <ul style="list-style-type: none"> <li>• <math>s \mapsto s'</math></li> </ul> </li> </ul> <p><b>where</b> <math>\text{KnotIattrs} = \text{attrs}' - \text{attrs}^I</math></p> <p><math>\text{KnotJattrs} = \text{attrs}' - \text{attrs}^J</math></p>
--

The first predicate stipulates that one pair of compatible states must be from the initial states of each, to ensure that there is an initial state of the composed specification. The last three predicates cover the three cases for messages of the new internal specification.

The schema *CombineInt* describes the most general combination of internal specifications we will need. This composition can be represented as a function derived from the schema above. Given two internal specifications, which satisfy the constraints of *CombineInt* above, *Icompose* yields the unique internal specification which is their combination.

$$\begin{array}{|l}
\hline
Icompose : (InternalSpec \times InternalSpec) \leftrightarrow InternalSpec \\
\hline
Icompose = \{ CombinCInt \\
\bullet (\theta InternalSpec^I, \theta InternalSpec^J) \mapsto \theta InternalSpec^C \} \\
\hline
\end{array}$$

When interleaving agents, associativity of *Icompose* is dependent upon how the separate agents are defined to behave over any common messages. We stipulate that common messages can only change the values of shared attributes. In practice, this constraint will be satisfied because the common messages between interleaved agents will only be output messages defined to be the identity transition on all attributes. This general condition of message compatibility is characterized in the following schema. We express it mathematically by saying that the operations associated to common messages do not affect attributes not in common between the agents.

$$\begin{array}{|l}
\hline
MessageCompatible \\
\hline
InternalSpec^I \\
InternalSpec^J \\
notIJ : \mathbf{P} \mathcal{A} \\
\hline
notIJ = (attribs^I \cup attribs^J) - (attribs^I \cap attribs^J) \\
\forall m : messages^I \cap messages^J; \\
(s_I, s'_I) : operations^I(m); \\
(s_J, s'_J) : operations^J(m) \\
\bullet ( (notIJ \triangleleft s_I, notIJ \triangleleft s'_I) \in id(notIJ \leftrightarrow \mathcal{V}) \\
\wedge (notIJ \triangleleft s_J, notIJ \triangleleft s'_J) \in id(notIJ \leftrightarrow \mathcal{V})) \\
\hline
\end{array}$$

$$\begin{array}{|l}
\hline
messagecompatible : InternalSpec \leftrightarrow InternalSpec \\
\hline
messagecompatible = \{ MessageCompatible \\
\bullet \theta InternalSpec^I \mapsto InternalSpec^J \} \\
\hline
\end{array}$$

In the case of synchronous composition, associativity is satisfied because the separate state sets are forced to be disjoint. The associativity conditions for the composition of internal specifications are summarized in the following two theorems. Theorem 5.2 treats the case for interleaving composition and Theorem 5.3 covers the case for synchronous composition. The proofs are only outlined; full proofs such as those provided in Appendix B could be provided.

**Theorem 5.2** *Message compatible internal specifications can be composed associatively, i.e.,*

$$\begin{aligned}
& \forall I1, I2, I3 : \text{InternalSpec} \\
& | ( (I1, I2) \in \text{dom } I\text{compose} \\
& \quad \wedge (I\text{compose}(I1, I2), I3) \in \text{dom } I\text{compose} \\
& \quad \wedge \{(I1, I2), (I2, I3), (I1, I3)\} \subseteq \text{messagecompatible} ) \\
& \bullet I\text{compose}(I\text{compose}(I1, I2), I3) = I\text{compose}(I1, I\text{compose}(I2, I3))
\end{aligned}$$

PROOF OF THEOREM 5.2:

Theorem 5.1 allows us to prove that if the first two predicates in the hypothesis are satisfied, then the composition can be formed in either order and the state sets and initial states are equivalent. We examine the two cases for messages, ones that are not shared and ones that are.

**Case 1:** The message is not shared by any two of  $I1, I2, I3$ .

There are three subcases to investigate, but the argument is the same for each. Suppose, for example, the message belongs to  $I1$  and not  $I2$  or  $I3$ . We have the operations defined as below. We adopt a shorthand notation, so, for example,

$$\text{operations}^{I(12)3}$$

will stand for

$$I\text{compose}(I\text{compose}(I1, I2), I3).\text{operations}$$

The proof proceeds. For any message  $m : \text{messages}^{I1}$ , we have

$$\begin{aligned}
& \text{operations}^{I(12)3}(m) = \\
& \{ s, s' : \text{states}^{I(12)3} \\
& \quad | ( (\text{attrs}^{I12} \triangleleft s, \text{attrs}^{I12} \triangleleft s') \in \text{operations}^{I12}(m) \\
& \quad \quad \wedge (\text{attrs}^{I12} \triangleleft s, \text{attrs}^{I12} \triangleleft s') \\
& \quad \quad \quad \in \text{id}(\text{attrs}^{I3} - \text{attrs}^{I12} \leftrightarrow \mathcal{V})) \\
& \bullet s \mapsto s' \}
\end{aligned}$$

By definition of  $\text{operations}^{I12}$ , and since  $m \in \text{message}^{I1}$  and Theorem 5.1 tells us that the state sets are equivalent, the above set comprehension is equivalent to the following.

$$\begin{aligned}
& \{ s, s' : \text{states}^{I1(23)} \\
& \quad | ( (\text{attrs}^{I1} \triangleleft s, \text{attrs}^{I1} \triangleleft s') \in \text{operations}^{I1}(m) \\
& \quad \quad \wedge (\text{attrs}^{I2} \triangleleft \text{attrs}^{I1} \triangleleft s, \text{attrs}^{I2} \triangleleft \text{attrs}^{I1} \triangleleft s') \\
& \quad \quad \quad \in \text{id}(\text{attrs}^{I2} - \text{attrs}^{I1} \leftrightarrow \mathcal{V})) \\
& \quad \quad \wedge (\text{attrs}^{I12} \triangleleft s, \text{attrs}^{I12} \triangleleft s') \\
& \quad \quad \quad \in \text{id}(\text{attrs}^{I3} - \text{attrs}^{I12} \leftrightarrow \mathcal{V})) \\
& \bullet s \mapsto s' \}
\end{aligned}$$

And the last two predicates can be collapsed to obtain the following.

$$\begin{aligned} & \{ s, s' : \text{states}^{I1(23)} \\ & \quad | \quad ( \text{attrs}^{I1} \triangleleft s, \text{attrs}^{I1} \triangleleft s' ) \in \text{operations}^{I1}(m) \\ & \quad \quad \wedge ( \text{attrs}^{I1} \triangleleft s, \text{attrs}^{I1} \triangleleft s' ) \in \text{id}(\text{attrs}^{I23} - \text{attrs}^{I1} \rightsquigarrow \mathcal{V}) \\ & \quad \bullet \quad s \mapsto s' \} \end{aligned}$$

This final set comprehension is  $\text{operations}^{I1(23)}(m)$ , as desired. A similar proof holds in the other cases, and so we have shown that the operation sets defined are equivalent.

**Case 2:** The message is shared by 2 or more of I1, I2, I3.

Again, there are several subcases, but each is argued similarly. Assume, for example, that  $m$  is a message for I1 and I3, but not I2. The operations are defined as

$$\begin{aligned} & \text{operations}^{I(12)3}(m) = \\ & \{ s, s' : \text{states}^{I(12)3} \\ & \quad | \quad ( \text{attrs}^{I12} \triangleleft s, \text{attrs}^{I12} \triangleleft s' ) \in \text{operations}^{I12}(m) \\ & \quad \quad \wedge ( \text{attrs}^{I3} \triangleleft s, \text{attrs}^{I3} \triangleleft s' ) \in \text{operations}^{I3}(m) \\ & \quad \bullet \quad s \mapsto s' \} \end{aligned}$$

By definition of  $\text{operations}^{I12}$ , and since  $m \in \text{message}^{I1}$  and Theorem 5.1 tells us that the state sets are equivalent, the above set comprehension is equivalent to the following.

$$\begin{aligned} & \{ s, s' : \text{states}^{I1(23)} \\ & \quad | \quad ( \text{attrs}^{I1} \triangleleft s, \text{attrs}^{I1} \triangleleft s' ) \in \text{operations}^{I1}(m) \\ & \quad \quad \wedge ( \text{attrs}^{I2} \triangleleft \text{attrs}^{I1} \triangleleft s, \text{attrs}^{I2} \triangleleft \text{attrs}^{I1} \triangleleft s' ) \\ & \quad \quad \quad \in \text{id}(\text{attrs}^{I2} - \text{attrs}^{I1} \rightsquigarrow \mathcal{V}) \\ & \quad \quad \wedge ( \text{attrs}^{I3} \triangleleft s, \text{attrs}^{I3} \triangleleft s' ) \in \text{operations}^{I3}(m) \\ & \quad \bullet \quad s \mapsto s' \} \end{aligned}$$

By assumption,  $\text{operations}^{I1}(m)$  and  $\text{operations}^{I3}(m)$  have no effect on attributes not in common, so the middle clause can be rewritten, substituting I3 for I1, to obtain

$$\begin{aligned} & \{ s, s' : \text{states}^{I1(23)} \\ & \quad | \quad ( \text{attrs}^{I1} \triangleleft s, \text{attrs}^{I1} \triangleleft s' ) \in \text{operations}^{I1}(m) \\ & \quad \quad \wedge ( \text{attrs}^{I2} \triangleleft \text{attrs}^{I3} \triangleleft s, \text{attrs}^{I2} \triangleleft \text{attrs}^{I3} \triangleleft s' ) \\ & \quad \quad \quad \in \text{id}(\text{attrs}^{I2} - \text{attrs}^{I3} \rightsquigarrow \mathcal{V}) \\ & \quad \quad \wedge ( \text{attrs}^{I3} \triangleleft s, \text{attrs}^{I3} \triangleleft s' ) \in \text{operations}^{I3}(m) \\ & \quad \bullet \quad s \mapsto s' \} \end{aligned}$$

This, in turn can be collapsed, by the definition of *operations*<sup>I23</sup>.

$$\{ s, s' : \text{states}^{I1(23)} \mid \begin{aligned} & ( \text{attrs}^{I1} \triangleleft s, \text{attrs}^{I1} \triangleleft s' ) \in \text{operations}^{I1}(m) \\ & \wedge ( \text{attrs}^{I23} \triangleleft s, \text{attrs}^{I23} \triangleleft s' ) \in \text{operations}^{I233}(m) \end{aligned} \\ \bullet s \mapsto s' \}$$

This last set comprehension is equivalent to *operations*<sup>I1(23)</sup>, as desired.

◇ END OF PROOF OF THEOREM 5.2

**Theorem 5.3** *Internal specifications with no common attributes can be composed associatively, i.e.,*

$$\begin{aligned} & \forall I1, I2, I3 : \text{InternalSpec} \\ & \{ \text{disjoint} \{ I1.\text{attrs}, I2.\text{attrs}, I3.\text{attrs} \} \\ & \quad \wedge ( I1, I2 ) \in \text{dom } \text{Icompose} \\ & \quad \wedge ( \text{Icompose}(I1, I2), I3 ) \in \text{dom } \text{Icompose} \\ & \bullet \text{Icompose}(\text{Icompose}(I1, I2), I3) = \text{Icompose}(I1, \text{Icompose}(I2, I3)) \end{aligned}$$

PROOF OF THEOREM 5.3:

Equivalence of state sets follows from an identical argument as above. There are many cases to investigate to determine the equivalence of the operations set. The arguments are similar for all of these cases, so we will provide one example. Assume that the message satisfies

$$m \in ( I1.\text{messages} \cap I3.\text{messages} ) - I2.\text{messages}.$$

The operation defined for that  $m$  is given by the following set comprehension.

$$\begin{aligned} & \text{operations}^{I(12)3}(m) = \\ & \{ s, s' : \text{states}^{I(12)3} \mid \begin{aligned} & ( \text{attrs}^{I12} \triangleleft s, \text{attrs}^{I12} \triangleleft s' ) \in \text{operations}^{I12}(m) \\ & \wedge ( \text{attrs}^{I3} \triangleleft s, \text{attrs}^{I3} \triangleleft s' ) \in \text{operations}^{I3}(m) \end{aligned} \\ & \bullet s \mapsto s' \} \end{aligned}$$

By the definition of *operations*<sup>I12</sup>( $m$ ), this expands to

$$\begin{aligned} & \{ s, s' : \text{states}^{I123} \mid \begin{aligned} & ( \text{attrs}^{I1} \triangleleft s, \text{attrs}^{I1} \triangleleft s' ) \in \text{operations}^{I1}(m) \\ & \wedge ( \text{attrs}^{I2} \triangleleft \text{attrs}^{I1} \triangleleft s, \text{attrs}^{I2} \triangleleft \text{attrs}^{I1} \triangleleft s' ) \\ & \quad \in \text{id}(\text{attrs}^{I2} - \text{attrs}^{I1} \mapsto \mathcal{V}) \\ & \wedge ( \text{attrs}^{I3} \triangleleft s, \text{attrs}^{I3} \triangleleft s' ) \in \text{operations}^{I3}(m) \end{aligned} \\ & \bullet s \mapsto s' \} \end{aligned}$$



We can rewrite the middle clause because the attribute sets are disjoint.

$$\begin{aligned} & \{ s, s' : \text{states}^{I123} \\ & \quad | ( (\text{attrs}^{I1} \triangleleft s, \text{attrs}^{I1} \triangleleft s') \in \text{operations}^{I1}(m) \\ & \quad \quad \wedge (\text{attrs}^{I2} \triangleleft \text{attrs}^{I3} \triangleleft s, \text{attrs}^{I2} \triangleleft \text{attrs}^{I3} \triangleleft s') \\ & \quad \quad \in \text{id}(\text{attrs}^{I2} - \text{attrs}^{I3} \leftrightarrow \mathcal{V}) \\ & \quad \quad \wedge (\text{attrs}^{I3} \triangleleft s, \text{attrs}^{I3} \triangleleft s') \in \text{operations}^{I3}(m)) \\ & \bullet s \mapsto s' \} \end{aligned}$$

By the definition of  $\text{operations}^{I23}$ , this can be rewritten as follows.

$$\begin{aligned} & \{ s, s' : \text{states}^{I12(3)} \\ & \quad | ( (\text{attrs}^{I1} \triangleleft s, \text{attrs}^{I1} \triangleleft s') \in \text{operations}^{I1}(m) \\ & \quad \quad \wedge (\text{attrs}^{I23} \triangleleft s, \text{attrs}^{I23} \triangleleft s') \in \text{operations}^{I23}(m)) \\ & \bullet s \mapsto s' \} \end{aligned}$$

This last set comprehension is equal to  $\text{operations}^{I1(23)}(m)$ , as desired.

◇ END OF PROOF OF THEOREM 5.3

### 5.3 External specification

The internal specification implies an ordering on the sequence of state transitions in which an agent can participate. If a designer wants to specify the order of operations to satisfy some constraint, it is not always easy to do by means of the internal specification. In other words, some constraints are not naturally expressible in terms of state transitions. This is a similar conclusion to that reached by Lamport [99], Sufrin and He [158] and Took [162].

Another way to express constraints on the state transitions is to view the constraints as external, meaning that an agent has both an internal and external specification. The agent participates in events and the external constraints are expressible as predicates on event participation. *Process algebras* are specifically designed to accommodate such event descriptions. The standard examples of process algebras are Hoare's Communicating Sequential Processes (CSP) notation [81, 82], developed at Oxford, and Milner's Calculus for Communicating Systems (CCS) developed at Edinburgh.

We adopt the traces model of a CSP process as the basis for the external specification as a nonempty prefix-closed set of sequences of events.

<i>ExternalSpec</i> $alphabet : \mathcal{P} \text{Event}$ $traces : \mathcal{P} \text{seq alphabet}$
$\{ \} \in traces$ $prefix\_closed (traces)$

The definition of the predicate `prefix_closed` is defined in Appendix A.

It would be possible to incorporate more complex process algebra models. We do not ignore the significant variety of process algebra models, most notably the family of algebras derived from Milner's CCS. For the purposes of this thesis, however, a demonstration of the incorporation of a process algebraic technique with the model-oriented axiomatic technique is all that is important, and in doing so we limit ourselves to the simplest case of the CSP traces model. More sophisticated models would allow expression of more sophisticated external constraints. Within the traces model, only *safety* constraints may be expressed, which would allow explicit expressions to rule out undesirable behaviour. *Liveness* constraints which make explicit that desirable behaviour will happen, require a model such as the failures-divergence model of CSP [28]. Some timing constraints may require even more sophisticated process algebras that incorporate a model of time and concurrency beyond the interleaving semantics of the traces or failures-divergence models. Examples in the CSP family are the various timed models of Reed [132] and Davies and Schneider [46, 144, 45]. A timed version of CCS has been provided by Moller and Tofts [111].

## Combining external specifications

One of the main features of a process algebra is its compositionality. We will discuss in a later section how that compositionality is achieved in isolated descriptions of external specifications. Then, we will provide both a constructive trace notation, which is described in detail in Appendix C and a more implicit descriptive technique of a predicate calculus on traces. In this section, we will concentrate on how to combine external specifications at the agent level, i.e., when the external specification is linked with an internal and communication specification. In this case we have two desired combinations—*independent combination by interleaving* and *dependent combination by synchronization*.

*Interleaved combination* is used to allow the behaviours of the separate external specifications to occur in the combined specification. The traces resulting are all interleavings of pairs of traces from the two separate specifications. The description of this interleaving is exactly the same as the interleaving of processes in CSP [82], which we also further define in Appendix C.

$ \begin{array}{l} \textit{InterleaveExt} \\ \textit{ExternalSpec}^I \\ \textit{ExternalSpec}^J \\ \textit{ExternalSpec}' \\ \hline \textit{alphabet}' = \textit{alphabet}^I \cup \textit{alphabet}^J \\ \textit{traces}' = \{ s : \textit{seq alphabet}' \\ \quad   \exists t : \textit{traces}^I; t' : \textit{traces}^J \bullet s \textit{ interleaves } (t, t') \\ \quad \bullet s \} \end{array} $
---

We can represent this combination as a binary operation,  $Ecompose_{int}$ .

$ \begin{array}{l} \textit{Ecompose}_{int} : (\textit{ExternalSpec} \times \textit{ExternalSpec}) \rightarrow \textit{ExternalSpec} \\ \hline \textit{Ecompose}_{int} = \{ \textit{InterleaveExt} \\ \quad \bullet (\theta \textit{ExternalSpec}^I, \theta \textit{ExternalSpec}^J) \mapsto \textit{ExternalSpec}' \} \end{array} $
---

Synchronous combination, as its name suggests, allows for the synchronizing of events between two external specifications. Its definition is similar to that for synchronous parallel combination in CSP, which we have also further defined in Appendix C.

$ \begin{array}{l} \textit{SynchronizeExt} \\ \textit{ExternalSpec}^I \\ \textit{ExternalSpec}^J \\ \textit{ExternalSpec}' \\ \hline \textit{alphabet}' = \textit{alphabet}^I \cup \textit{alphabet}^J \\ \textit{traces}' = \{ t : \textit{seq alphabet}' \\ \quad   ( t \overset{\textit{seq}}{\triangleright} \textit{alphabet}^I \in \textit{traces}^I \\ \quad \quad \wedge t \overset{\textit{seq}}{\triangleright} \textit{alphabet}^J \in \textit{traces}^J ) \\ \quad \bullet t \} \end{array} $
---

Sequence range restriction,  $\overset{\textit{seq}}{\triangleright}$ , is defined in Appendix A. It operates much like a filter in a functional programming language. We can represent this combination as a binary operation,  $Ecompose_{sync}$ .

$ \begin{array}{l} \textit{Ecompose}_{sync} : (\textit{ExternalSpec} \times \textit{ExternalSpec}) \rightarrow \textit{ExternalSpec} \\ \hline \textit{Ecompose}_{sync} = \{ \textit{SynchronizeExt} \\ \quad \bullet (\theta \textit{ExternalSpec}^I, \theta \textit{ExternalSpec}^J) \mapsto \textit{ExternalSpec}' \} \end{array} $
--

It is well known that both of these operators are associative [82, pp. 70,120].

## 5.4 Communication

The internal specification induces an ordering on the state transitions of the agent and this ordering can be represented by the message labels associated to the transitions via the function *operation*. The external specification gives directly the ordering on events that can occur during *traces* of the behaviour. In the communication specification for an agent, we link the messages of the internal specification to the events of the external specification. Events are communications of messages along channels.

To accommodate this within the agent model, we will introduce explicit channels that are meant as point-to-point, one-way, synchronous communication paths between agents, as in the programming language *occam* [86, 91]. Though this is a fairly restrictive means of communication, it has been noted that this restriction makes possible tractable algebraic techniques for reasoning rigorously about an *occam* program [25, 139, 140], which we can then incorporate into the agent development method. We assume a set of all possible channel identifiers.

$[ChannelID]$

An event is a message communicated along some channel. We can represent the refined *Event* type as a schema which indicates the channel name and the message.

$Event == ChannelID \times Message$

We provide some shorthand notation for accessing the channel and message content of an event.

$chan : Event \rightarrow ChannelID$ $mess : Event \rightarrow Message$	$chan = fst$ $mess = snd$
--	------------------------------

A channel is associated to a set of message which can be communicated across it. There are four kinds of channels possible. A channel can be used for the input of messages or the output of messages. The combination of an agent which uses a channel for input and an agent which uses the same channel for output creates a synchronized channel. Finally, there is a completely internal channel which handles internal communication within an agent.

We define a function *events* which yields the set of all possible events that can occur along a channel.

$events : (ChannelID \times P Message) \rightarrow P Event$	$events(c, MS) = \{ m : MS \bullet (c, m) \}$
---	---

The communication specification gives all input, output and synchronized channels. These channels are all distinct. Each agent also has a special channel for internal events, which is called  $\tau$ . From the channel information, we can derive the alphabet of events in which the agent can participate and the set of messages for which operations must be defined. These two derived sets will match the corresponding sets in the external and internal specifications.

<p><i>Communication</i></p> <p><math>inputs, outputs, syncs : ChannelID \rightarrow \mathbf{P} Message</math></p> <p><math>internal : \mathbf{P} Message</math></p> <p><math>alphabet : \mathbf{P} Event</math></p> <p><math>messages : \mathbf{P} Message</math></p> <p>disjoint (dom <math>inputs</math>, dom <math>outputs</math>, dom <math>syncs</math>, <math>\{\tau\}</math>)</p> <p><math>alphabet = events(\{inputs\}) \cup events(\{outputs\}) \cup</math>  <math>events(\{syncs\}) \cup events(\tau, internal)</math></p> <p><math>messages = mess(alphabet)</math></p>
--

### Combining communication specifications

Two communication specifications are *channel compatible* if no synchronized channel of one is either an input, output or synchronized channel of the other.

<p><i>ChannelCompatible</i></p> <p><math>Communication^I</math></p> <p><math>Communication^J</math></p> <p><math>dom syncs^J \cap (dom inputs^I \cup dom outputs^I \cup dom syncs^I) = \emptyset</math></p> <p><math>dom syncs^I \cap (dom inputs^J \cup dom outputs^J \cup dom syncs^J) = \emptyset</math></p>
---

Two communication specifications are *unlinked* if they are channel compatible and no input channel of one is an output channel of the other.

<p><i>Unlinked</i></p> <p><math>ChannelCompatible</math></p> <p><math>(dom inputs^I \cap dom outputs^J) = \emptyset</math></p> <p><math>(dom inputs^J \cap dom outputs^I) = \emptyset</math></p>
--

Two communication specifications which are unlinked can be combined to yield a third communication specification. Unlinked communication specifications can

share the same input channels or the same output channels. The combined specification forms the union of the message sets associated to each channel identifier. This combination of communication specifications will be used to form the interleaved composition of agents.

<p><i>InterleaveComm</i></p> <hr/> <p><i>Unlinked Communication'</i></p> <hr/> <p><math>internal' = internal^I \cup internal^J</math></p> <p><math>syncs' = syncs^I \cup syncs^J</math></p> <p><math>inputs' = \{ c : \text{dom } inputs^I \cup \text{dom } inputs^J</math>  <ul style="list-style-type: none"> <li>• <math>c \mapsto (inputs^I(c) \cup inputs^J(c)) \}</math></li> </ul> </p> <p><math>outputs' = \{ c : \text{dom } outputs^I \cup \text{dom } outputs^J</math>  <ul style="list-style-type: none"> <li>• <math>c \mapsto (outputs^I(c) \cup outputs^J(c)) \}</math></li> </ul> </p>
--

We can represent this combination as the binary operation,  $Ccompose_{int}$ .

<p><math>Ccompose_{int} : (Communication \times Communication) \rightarrow Communication</math></p> <hr/> <p><math>Ccompose_{int} =</math>  <math>\{ InterleaveComm</math>  <ul style="list-style-type: none"> <li>• <math>(\theta Communication^I, \theta Communication^J) \mapsto Communication'</math></li> </ul> </p>
---

Two communication specifications which are channel compatible and which have no common input channels nor common output channels are said to be *linkable* and they can be synchronized. The common input/output channel pairs are made synchronous channels. The messages which are allowed on the synchronized channel are ones which are allowed as input by one agent or output by the other.

<p><i>Linkable ChannelCompatible</i></p> <hr/> <p><math>(\text{dom } inputs^I \cap \text{dom } inputs^J) = \emptyset</math></p> <p><math>(\text{dom } outputs^I \cap \text{dom } outputs^J) = \emptyset</math></p>
--

<p><i>SynchronizeComm</i></p> <p><i>Linkable</i></p> <p><i>Communication'</i></p> <hr/> <p><math>internal' = internal^I \cup internal^J</math></p> <p><math>syncs' = syncs^I \cup syncs^J</math></p> <p style="padding-left: 2em;"><math>\cup \{ c : (\text{dom } inputs^I \cap \text{dom } outputs^J)</math></p> <p style="padding-left: 4em;"><math>\bullet c \mapsto (inputs^I(c) \cup outputs^J(c)) \}</math></p> <p style="padding-left: 2em;"><math>\cup \{ c : (\text{dom } inputs^J \cap \text{dom } outputs^I)</math></p> <p style="padding-left: 4em;"><math>\bullet c \mapsto (inputs^J(c) \cup outputs^I(c)) \}</math></p> <p><math>inputs' = \text{dom } syncs' \triangleleft (inputs^I \cup inputs^J)</math></p> <p><math>outputs' = \text{dom } syncs' \triangleleft (outputs^I \cup outputs^J)</math></p>
---

We can represent the synchronized composition of communication specifications as the binary operation,  $Ccompose_{sync}$ .

<p><math>Ccompose_{sync} : (Communication \times Communication) \mapsto Communication</math></p> <hr/> <p><math>Ccompose_{sync} =</math></p> <p><math>\{ SynchronizeComm</math></p> <p style="padding-left: 2em;"><math>\bullet (\theta Communication^I, \theta Communication^J) \mapsto \theta Communication' \}</math></p>
--

By arguments similar to those given to show that internal composition is associative, we can also show that communication composition operators are also associative.

## 5.5 Overall Combination

Having specified the components of the refined agent model and how they are individually composable, we can now define two operations to compose agents. The two composing functions will represent an interleaving composition, in which the two agents act as one agent by interleaving their individual behaviour, and a synchronizing composition, in which communication will be synchronized between common channels.

Interleaving composition is specified in terms of the general internal specification composition and the interleaved composition of the external and communication specifications. We include the constraint on message compatibility defined earlier so that the conditions of Theorem 5.2 are upheld and the resulting composition operator,  $compose_{int}$ , is associative.

<i>InterleaveCompose</i> <i>Agent<sup>I</sup></i> <i>Agent<sup>J</sup></i> <i>Agent'</i> <i>MessageCompatible</i> <i>CombineInt</i> <i>InterleaveExt</i> <i>InterleaveComm</i>
---

The binary operation for interleave composition is  $compose_{int}$

$compose_{int} : (Agent \times Agent) \rightarrow Agent$ $compose_{int} =$ $\{ InterleaveCompose$ $\bullet (\theta Agent^I, \theta Agent^J) \mapsto \theta Agent' \}$
--

Arguments in the preceding three sections lead to the proof of the associativity of  $compose_{int}$ .

**Theorem 5.4** *Interleaving composition, when defined, is associative, i.e.,*

$$compose_{int}(compose_{int}(A1, A2), A3) = compose_{int}(A1, compose_{int}(A2, A3)).$$

Synchronizing composition is specified in terms of the general internal specification composition and the synchronized composition of the external and communication specifications. We add the constraint that the attribute sets of the agents to be composed must be disjoint. This is to ensure that the conditions of Theorem 5.3 are upheld in order that  $compose_{sync}$  be associative.

<i>SynchronizeCompose</i> <i>Agent<sup>I</sup></i> <i>Agent<sup>J</sup></i> <i>Agent'</i> <i>CombineInt</i> <i>SynchronizeExt</i> <i>SynchronizeComm</i> <hr/> $attrs^I \cap attrs^J = \emptyset$
--

The binary operation for interleave composition is  $compose_{sync}$



$$\begin{array}{|l}
\hline
compose_{sync} : (Agent \times Agent) \leftrightarrow Agent \\
\hline
compose_{sync} = \\
\{ SynchronizeCompose \\
\bullet (\theta Agent^I, \theta Agent^J) \mapsto \theta Agent' \} \\
\hline
\end{array}$$

Arguments in the preceding three sections lead to the proof of the associativity of  $compose_{sync}$ .

**Theorem 5.5** *Synchronous composition, when defined, is associative, i.e.,*

$$\begin{array}{l}
compose_{sync}(compose_{sync}(A1, A2), A3) = \\
compose_{sync}(A1, compose_{sync}(A2, A3)).
\end{array}$$

Since these composition operators are associative, we can define operations which compose arbitrary nonempty sequences of agents by means of the folding operations of standard functional programming [24]. Refer to Appendix A for definitions of the folding operations in standard Z.

$$\begin{array}{|l}
\hline
composeall_{int}, \\
composeall_{sync} : seq_1 Agent \leftrightarrow Agent \\
\hline
composeall_{int} = foldl1 compose_{int} \\
composeall_{sync} = foldl1 compose_{sync} \\
\hline
\end{array}$$

The composition operators are not mutually associative, but in practice this will not matter. Interleaving is used to build up a specification of an agent that will have a private state. Synchronization is used to connect the separate agents, i.e., ones which have a private state.

## 5.6 Interpretations of agents

The communication specification now provides the link necessary for combining an internal specification and an external specification. What remains is to describe the refinements of the internal and external interpretation relations.

### 5.6.1 Internal interpretation

We want to derive the internal interpretation relation  $I_A^{int}$  for an agent which maps traces of events that an agent can participate in to the possible states that the agent can be in after participating in the trace. Following the approach of Chapter 2, we first extend the *operations* mapping from the internal specification. However, since

that mapping uses messages to index the state transitions and not events, we must also extract the message content of the event.

$$\left| \begin{array}{l} \text{opExtend}_- : \text{Agent} \rightarrow \text{seq Event} \rightarrow (\text{State} \leftrightarrow \text{State}) \\ \text{opExtend}_A(\text{trace}) = A.\text{inits} \triangleleft (\text{g}/(\text{trace} \text{ ; mess ; } A.\text{operations})) \end{array} \right|$$

A difference to note from the development of Chapter 2 is that every event has an associated state transition, instead of just input events. We refer to sequences of events in this definition as traces instead of programs to alert the reader of the difference. It is the role of an event which distinguishes stimuli (input) from response (output). The domain of *opExtend* manifests the constraints on event sequencing which arise from the internal specification. The external specification is already expressed in terms of event sequences.  $I_A^{\text{int}}$  must be consistent with both.

$$\left| \begin{array}{l} \underline{I_A^{\text{int}}} : \text{Agent} \rightarrow (\text{seq Event} \leftrightarrow \text{State}) \\ (\text{trace}, s) \in I_A^{\text{int}} \Leftrightarrow (\exists s_0 : A.\text{inits} \bullet (\text{trace}, (s_0, s)) \in \text{opExtend}_A \\ \wedge \text{trace} \in A.\text{traces}) \end{array} \right|$$

It is important to note in this definition, that the overall interpretation is given as a conjunction of internal and external constraints. This was the main motivation for the hybrid notation of the Sufrin and He model. In that formal model, however, the external description was defined to constrain the transitions possible from the internal description, but not vice versa. However, it was clearly the *intent* of those authors to allow the internal specification to constrain the traces possible from the external description. We mention this point as another example of how a formal approach can uncover inconsistencies between informal requirements (the intent mentioned above) and the deliverable (the actual formalism).

The *historics* of an agent, written  $\mathcal{H}[A]$ , give a recording of the events in which it can participate and are derived from the domain of  $I_A^{\text{int}}$ .

$$\left| \begin{array}{l} \mathcal{H}[_] : \text{Agent} \rightarrow \mathcal{P}(\text{seq Event}) \\ \mathcal{H}[A] = \text{dom } I_A^{\text{int}} \end{array} \right|$$

There are three properties which must be satisfied by the internal interpretation relation and its derivatives:

- the set of histories of an agent is prefix closed, so that for an agent to participate in a sequence of events it must have participated in each prefix;
- the internal interpretation relation respects the history of interaction, so that any state the agent can attain must have been reached from a previously attainable state; and

- the internal interpretation relation is nonempty, so that every agent has some behaviour.

We will prove each of these properties of the internal interpretation relation in turn as theorems.

**Theorem 5.6** *The set of traces of an agent is prefix closed, i.e.,*

$$\forall A : \text{Agent} \bullet \text{prefix\_closed } \mathcal{H}[[A]].$$

**PROOF OF THEOREM 5.6:**

To prove this theorem, we need Lemma 5.1, which states that *opExtendA* obeys a nice compositional property.

**Lemma 5.1**

$$\begin{aligned} \forall A : \text{Agent}; s : A.\text{states}; s_0 : A.\text{inits}; hd, tl : \text{seq Event} \\ \bullet (hd \frown tl, (s_0, s)) \in \text{opExtend}_A \Leftrightarrow \\ \quad \exists s' : A.\text{states} \\ \quad \bullet ( (hd, (s_0, s')) \in \text{opExtend}_A \\ \quad \wedge (tl, (s', s)) \in (tl \text{ ; mess ; } A.\text{operations})) \end{aligned}$$

**PROOF OF LEMMA 5.1:**

Because of layout considerations, in the following argument we have substituted the expression *Aops* for the correct expression *A.operations*.

$$\begin{aligned} (hd \frown tl, (s_0, s)) \in \text{opExtend}_A & \quad \text{[assumption]} \\ \Leftrightarrow (s_0, s) \in \text{;}/(hd \frown tl \text{ ; mess ; } Aops) & \quad \text{[defn. of } \text{opExtend}_A, s_0 \in A.\text{inits}] \\ \Leftrightarrow (s_0, s) \in \text{;}/((hd \text{ ; mess ; } Aops) \frown (tl \text{ ; mess ; } Aops)) & \quad \text{[; dist. over } \frown] \\ \Leftrightarrow (s_0, s) \in (\text{;}/(hd \text{ ; mess ; } Aops)) \text{ ; } \text{;}/(tl \text{ ; mess ; } Aops) & \quad \text{[defn. of ;/]} \\ \Leftrightarrow \exists s' : A.\text{states} & \quad \text{[defn. of ;]} \\ \bullet ( (s_0, s') \in (\text{;}/(hd \text{ ; mess ; } Aops)) \\ \quad \wedge (s', s) \in (\text{;}/(tl \text{ ; mess ; } Aops))) & \\ \Leftrightarrow \exists s' : A.\text{states} & \quad \text{[defn. of } \text{opExtend}_A] \\ \bullet ( (hd, (s_0, s')) \in \text{opExtend}_A \\ \quad \wedge (tl, (s', s)) \in (tl \text{ ; mess ; } Aops)) & \end{aligned}$$

◇ END OF PROOF OF LEMMA 5.1

Continuing with the proof of Theorem 5.6, let  $tr \in \mathcal{H}[A]$  and  $hd \dashv tr$ , i.e.,  $hd \frown tl = tr$ . We need to show that  $hd \in \mathcal{H}[A]$ .

$$\begin{aligned}
& hd \frown tl \in \mathcal{H}[A] && \text{[assumption]} \\
& \Leftrightarrow hd \frown tl \in \text{dom } I_A^{int} && \text{[defn. of } \mathcal{H}[A] \text{]} \\
& \Leftrightarrow \exists s : A.\text{states}; s_0 : A.\text{inits} && \text{[defn. of } I_A^{int} \text{]} \\
& \quad \bullet (hd \frown tl, (s_0, s)) \in \text{opExtend}_A \\
& \quad \quad \wedge hd \frown tl \in A.\text{traces} \\
& \Leftrightarrow \exists s' : A.\text{states}; s_0 : A.\text{inits} && \text{[ Lemma 5.1} \\
& \quad \bullet (hd, (s_0, s')) \in \text{opExtend}_A && \text{prefix\_closed}(A.\text{traces}) \text{]} \\
& \quad \quad \wedge hd \in A.\text{traces} \\
& \Leftrightarrow hd \in \text{dom } I_A^{int} && \text{[defn. of } I_A^{int} \text{]} \\
& \Leftrightarrow hd \in \mathcal{H}[A] && \text{[defn. of } \mathcal{H}[A] \text{]}
\end{aligned}$$

◇ END OF PROOF OF THEOREM 5.6

**Theorem 5.7** *The internal behaviour of an agent is history-closed, i.e.,*

$$\begin{aligned}
& \forall A : \text{Agent}; (tr, s) : I_A^{int} \\
& \quad \bullet \exists s' : A.\text{states}; tr', tr'' : \text{seq Message} \\
& \quad \quad | \quad tr' \frown tr'' = tr \\
& \quad \quad \bullet (tr', s') \in I_A^{int} \\
& \quad \quad \quad \wedge (s', s) \in (tr'' \# \text{mess}; A.\text{operations})
\end{aligned}$$

PROOF OF THEOREM 5.7:

Assume we have an agent  $A$  with  $(tr, s) \in I_A^{int}$  and  $tr' \frown tr'' = tr$ . If  $tr' = tr$ , then we can satisfy the conditions of the theorem by letting  $s' = s$  and  $tr'' = ()$ . So assume that  $tr' \neq tr$ . By Lemma 5.1, we know that

$$\begin{aligned}
& \exists s' : A.\text{states}; s_0 : A.\text{inits} \\
& \quad \bullet (tr', (s_0, s')) \in \text{opExtend}_A \\
& \quad \quad \wedge (tr'', (s', s)) \in \text{opExtend}_A
\end{aligned}$$

Since  $tr \in A.\text{traces}$  and  $A.\text{traces}$  is prefix closed by definition, we conclude that  $(tr', s') \in I_A^{int}$ , as desired.

◇ END OF PROOF OF THEOREM 5.7

**Theorem 5.8** *For any agent  $A$ , the internal interpretation relation is nonempty.*

## PROOF OF THEOREM 5.8:

Since  $A.inits$  is nonempty, this theorem is satisfied by demonstrating that, for all agents  $A$  and initial state  $s_0 : A.inits$ , the element  $(\langle \rangle, s_0)$  is in  $I_A^{int}$ . First, we show that  $opExtend_A \langle \rangle = id\ State$ .

$$\begin{aligned} opExtend_A \langle \rangle &= \mathcal{S}/(\langle \rangle_{seq\ Message} \mathcal{S}\ mess_{\mathcal{S}} A.operations) && [\text{defn. of } I] \\ &= \mathcal{S}/(\langle \rangle_{State-State}) && [\text{defn. of } \mathcal{S}] \\ &= id\ State && [\text{defn. of } \mathcal{S}/] \end{aligned}$$

Since  $(s_0, s_0) \in id\ State$ , we have satisfied the first predicate in the definition of  $I_A^{int}$  because

$$(\langle \rangle, (s_0, s_0)) \in opExtend_A.$$

Finally,  $\langle \rangle \in A.traces$  by the definition of  $ExternalSpec$ . Therefore, the internal interpretation relation of an agent is never empty, since it contains  $(\langle \rangle, s_0)$ .

◇ END OF PROOF OF THEOREM 5.8

### 5.6.2 External interpretation

In Chapter 2, we also defined an external interpretation relation,  $I_-^{ext}$  to reflect the overall stimulus-response behaviour of an agent. The importance of this interpretation relation was twofold. It provided an external specification of an agent, so that the translations of the interaction framework could be viewed as specifications of agents themselves. It also was used to define some interactive properties, such as predictability, honesty and trustworthiness.

External specification is now explicit in the agent model, but it includes more information than the overall stimulus-response and it also includes event behaviour that could be excluded by the internal specification.  $I_A^{ext}$  is defined on the traces of  $A$ , applying only to legal behaviours of the agent, and it filters out the input and output events. The stimuli of an agent are the events that can occur along the input channels. The responses are the events that can occur along the output channels.

$$\left| \begin{array}{l} stimuli, responses : Agent \rightarrow \mathbf{P}\ Event \\ \underline{I_-^{ext}} : Agent \rightarrow (seq\ Event \leftrightarrow seq\ Event) \\ stimuli(A) = events(|A.inputs|) \\ responses(A) = events(|A.outputs|) \\ I_A^{ext} = \{ t : \mathcal{H}[A] \\ \quad \bullet t \triangleright^{seq} stimuli(A) \mapsto t \triangleright^{seq} responses(A) \} \end{array} \right.$$

Because  $\mathcal{H}[[A]]$  is prefix closed, we can show that both the domain and range of  $J_A^{ref}$  are prefix closed as well.

## 5.7 Conclusions

In this chapter, we have refined the model of the agent. The specification of an agent has been split into three parts to define the internal state-based behaviour, the external event-based behaviour and the communication specification which links the internal and external specifications. We have described two composition operators on agents. The first corresponds to the interleaved composition of dependent agents which will aid in the gradual development of complex agents sharing attributes and messages. The second composition is a synchronization of independent agents which share no attributes and communicate via synchronized message passing.

The agent model presented is good for theoretical use. We can express and prove general interactive properties of agents using this model. In Chapter 7 we will use the refined model as a basis for the reformulation of the interactive properties expressed in Chapter 4 and introduce some new ones. The agent model, however, is not a very good design notation. We will justify this claim in the next chapter and present a language for the description of agents which can map back into the agent model.

## Chapter 6

# A language for describing agents

Up to this point in the thesis, the only formal notation we have used has been Z. We believe that strict adherence to Z as a design notation for agents is not desirable. The way a designer conceptualizes an agent's behaviour must be more directly captured in the agent language than standard Z allows. We will, therefore, provide a new language for the description of agents. We want the new language to be flexible. The properties we express on an agent relate the events and the state, so it is important that we can describe the event-state behaviour. As described in the last chapter, sometimes it is easier to describe such a desired behaviour via the internal state description linked with a communication description and sometimes it is easier to do directly via the external event description.

The agent language is a formal notation, and there are many existing formal notations which are increasingly being used in both academic and industrial circles. We want to take advantage of the familiarity with those existing notations. However, most of the notations are either more suited to the state-based description or the event-based notation, and so they do not alone satisfy the expressiveness requirement described above. Following on from Sufrin and He's model of interactive processes [158], we propose a hybrid notation which marries a model-oriented descriptive technique for the internal description and a process algebraic technique for the external description. Elsewhere, we have described a version of the agent language which uses an algebraic notation for the internal description [69], but the semantics for such a language we have not defined.

### Overview of chapter

In Section 6.1, we justify the need for an agent notation different from standard Z and we give an overview of some other agent-like notations. In Section 6.2, we provide a language for describing agents and we outline the mapping from that syntactic domain to the semantic model, further details of which can be found in Appendix C. The best way to explain how the agent language is used is by example,

and so we offer several examples in Section 6.3. We have placed a heavy emphasis on familiarity with the new notation, which has resulted in a hybrid language resembling established model-oriented languages, such as Z or VDM, and process algebra notations, such as CSP or CCS. Any language exerts an influence on its users—some things will be easy to express within the language and some things will be difficult. Such an influence is not a bad thing, as long as we are aware of it and we recognize the limitations it implies. At the end of this chapter we will summarize the limitations of the agent language.

## 6.1 Notations for agents

Although the model presented in Chapter 5 will be adequate for deriving the interactive properties to be discussed in Chapter 7, it does not satisfy the requirement of natural expressiveness. Before we present the agent language, we explore in this section some other possible languages for our agents and other object-oriented formalisms.

### 6.1.1 The standard Z notation

We have used the Z notation to present all of the formalisms so far in this thesis, and we will continue to use Z as the means for mathematical expression and reasoning within the agent model. However, we do not favour Z as the design notation for agents for two main reasons.

Our first criticism of Z is the lack of modularity it provides in standard use. Whereas in most tutorials on the Z language a modular development approach is advocated, any modularity in resulting specifications is left to the reader to extract. The principal features of an internal specification of an agent—the state space, the initial states, and the operations on the state space—are separately describable in Z. When we want to bundle these components together, it is possible in Z, but not natural to the standard Z development. The possibility was demonstrated by Sufrin and He [158], so to justify our argument, we will provide an example of the specification of a window in a multi-windowing environment using their model of a process, given below.

[E]



$\frac{\text{SandHProcess}[S]}{\alpha : \mathbf{P} E}$ $\hat{\beta} : E \leftrightarrow (S \leftrightarrow S)$ $\text{Trace} : \mathbf{P}(\text{seq } E)$ $t : \mathbf{P} S$
$\alpha = \text{dom } \hat{\beta}$
$\forall s, t : \text{seq } E \bullet s \cap t \in \text{Trace} \Rightarrow s \in \text{Trace}$

For our purposes, we can equate the type  $E$  with the type *Message*, and  $S$  with *State*.

We will begin by giving a  $\mathcal{Z}$  description of the functionality of the individual window. A window has two representations, depending on whether it is open or closed. When closed it takes the form of an icon, which is one from a set of all possible icons, denoted *ICON*.

[*ICON*]

When open, the window covers a rectangular region on a finite coordinate plane representing the visual display. This plane is commonly viewed as a pixel plane with boundaries in the horizontal ( $x$ -axis) and vertical ( $y$ -axis) directions. The type *PIXEL* will represent points in this finite coordinate plane.

$x_{max}, y_{max} : \mathbf{N}$

$\text{PIXEL} == 0 \dots x_{max} \times 0 \dots y_{max}$

For our purposes, a window is completely defined when we have the following information on it:

- the icon to represent it when closed and its position in *PIXEL* space
- the position and extent of its rectangular region in *PIXEL* space when open
- an indication of whether the window is open or closed

The schema type *WindowState* describes the state of such a window. Each aspect mentioned above is represented directly by a component of the schema type *WindowState*.

$\text{WindowState}$ $\text{icon} : \text{ICON}$ $\text{iconpos}, \text{winpos}, \text{winsize} : \text{PIXEL}$ $\text{status} : \text{open} \mid \text{closed}$
--

We could add some constraints on the window state. For example, we could constrain the size of the window such that the whole window is contained in *PIXEL* space. For simplicity, we ignore such constraints for the moment.

When a window is created, we can stipulate that it satisfies certain constraints beyond any given as the state invariant of *WindowState* (had we given any). Initialization constraints are by convention detailed by the schema *WindowInit*, and the purpose of this schema is to give the subset of all possible states in which a window can be initially. This schema only contains a copy of the state after the initialization (the window is assumed not to exist before initialization). In our example, we will stipulate that a window begins with status *open*.

$\begin{array}{l} \textit{WindowInit} \\ \hline \textit{WindowState}' \\ \hline \textit{status}' = \textit{open} \end{array}$
---

Some of the normal operations performed upon a window would be to open or close it, to move it (when open or closed) or to resize it when open. We indicate a window operation by a schema description containing the declaration  $\Delta \textit{WindowState}$  which contains two copies of window schema binding to represent the window before (using undashed component names) and after (using dashed component names) the operation. We can specify the operation to open a window by requiring that the status of the window indicate that it is closed before the operation and open afterwards. No other component of the window is changed.<sup>1</sup>

$\begin{array}{l} \textit{OpenWindowOp} \\ \hline \Delta \textit{WindowState} \\ \Xi(\textit{WindowState} \setminus \{\textit{status}\}) \\ \hline \textit{status} = \textit{closed} \\ \hline \textit{status}' = \textit{open} \end{array}$
--

This operation definition is equivalent to the following expanded schema.

<sup>1</sup>Note the use of the  $\Xi$  convention along with schema hiding ( $\setminus$ ) to give such framing conditions. The hidden components are precisely those we wish the operation to be able to change; the other remain the same. Though not widely publicized, this technique of explicitly naming the framing conditions of operations is good practice [106, 27].

$OpenWindowOp$ $WindowState$ $WindowState'$
$status = closed$ $status' = open$ $icon' = icon$ $iconpos' = iconpos$ $winpos' = winpos$ $winsize' = winsize$

Similar descriptions can describe other operations on a window, such as closing, moving or resizing. As an example of an operation which takes an argument, we specify how a window is moved. Either the open window or iconified window position can be altered, so we define the general move operation in two parts to cover the two different cases. The overall window repositioning operation is then the disjunction of those two separate operations. When the window is closed, the argument to the repositioning operation indicates the new position for the icon.

$MoveWindowClosed$ $\Delta WindowState$ $\Xi(WindowState \setminus \{iconpos\})$ $newpos? : PIXEL$
$status = closed$ $iconpos' = newPos?$

When the window is open, the argument indicates the new position for the rectangular region.

$MoveWindowOpen$ $\Delta WindowState$ $\Xi(WindowState \setminus \{winpos\})$ $newpos? : PIXEL$
$status = open$ $winpos' = newPos?$

$$MoveWindowOp \hat{=} MoveWindowClosed \vee MoveWindowOpen$$

The window is represented as a process in the following manner. The state set is given by *WindowState* and the initial states ( $\iota$ ) are characterized by *Windowinit*. The message set contains elements that can be linked to the schema operations realized as relations on *WindowState*.

$$\begin{array}{l}
 E ::= \text{open} \\
 \quad | \text{close} \\
 \quad | \text{move}\langle\langle \text{PIXEL} \rangle\rangle \\
 \quad | \text{resize}\langle\langle \text{PIXEL} \rangle\rangle \\
 \quad \vdots \\
 \\
 \left[ \begin{array}{l}
 \text{Window} : \text{SandHProcess}\{\text{WindowState}\} \\
 \text{Window}.\alpha = E \\
 \text{Window}.\iota = \{ \text{Windowinit} \bullet \theta \text{Window}' \} \\
 \text{Window}.\hat{\beta}(\text{open}) = \\
 \quad \{ \text{Open WindowOp} \bullet \theta \text{WindowState} \mapsto \theta \text{WindowState}' \} \\
 \forall p : \text{PIXEL} \\
 \bullet \text{Window}.\hat{\beta}(\text{move}(p)) = \\
 \quad \{ \text{Move WindowOp} \\
 \quad \quad | \text{newpos?} = p \\
 \quad \bullet \theta \text{WindowState} \mapsto \theta \text{WindowState}' \} \\
 \vdots
 \end{array} \right.
 \end{array}$$

Sufm and He [158], showed that it is even possible to give the external specification within  $Z$ , and we offer further proof of that in Appendix C with a full semantics for a constructive trace language based on CSP. So there is no feature of a process (or agent) which cannot in principle be represented in  $Z$ . But this representation is achieved through rather roundabout measures, none of which is difficult, but all of which seem unnecessary to the description of an agent.

Our second criticism of  $Z$  as the agent notation again arises because  $Z$  is not specifically geared to describe functioning entities (like an agent or process) in isolation within a system of other entities.  $Z$  has a limited ability to express communication between schemas, sequential composition and piping operators being defined in the schema calculus having only the limited possibility of communicating with one other schema. The stimulus-response model dictates that the agents be able to communicate to an arbitrary number of other agents in response to any stimulus received. Though extensions to  $Z$  which we discuss below remedy the earlier problem about bundling internal information into one object, no previous formalism adequately addresses this failing.

### 6.1.2 Object-oriented notations and Z

Our notion of an agent is somewhat similar to that of an object in object-oriented programming parlance. With the advent of object oriented programming languages, there has been a change in the way many designers view the systems that they build [43]. A system can now be naturally viewed as a collection of objects which pass messages that cause changes in their neighbours, that is, designers have adopted a stimulus-response view of their systems. In response to the increasing acceptance of object-oriented programming notations in industry, researchers have attempted to provide a formal notation to represent the largely informal and intuitively appealing concepts of objects.

The roots of object-oriented programming can be traced to data abstraction [102], in which only the means for transforming an underlying data structure, not the procedure, are made apparent to its user. Though data abstraction within an algebraic framework was initiated at least as far back as 1978 by Guttag and Horning [65] and Goguen *et al.*[60], it has only been investigated more recently within a model-oriented axiomatic approach. In this section we will discuss a few of the techniques distinct from our agent model that have been offered, in a roughly chronological order. These techniques have all influenced the development of the agent model, though none seems to address the issues around communication as well as the agent model.

#### Promotion

The description of a window and the operations that can be performed on it is easiest when done in isolation, that is, without consideration of any other windows which may coexist. The schema definitions describe the kinds of transformations on windows that we want to be possible. The operation definitions describe transitions on all possible windows; given *any* element of type *WindowState*, the operation definitions provide a description of when and how that element can be transformed to another element in *WindowState*. This is subtly different from an interpretation of what operations on windows as objects in an interactive windowing environment represent. In that case, we create a window and that marks the beginning of its existence as an entity (or object or process or agent). This entity has characteristics (or components or attributes) which fully describe it at any point during its existence and there are operations which can be performed on the entity to change those characteristics *but without changing the identity of the window within the system*.

This last point is very important point because it embodies another criticism of Z, namely that it does not allow for identification of the objects in a system so that changes to one object can be isolated. We do not feel, however, that this criticism is valid, as the standard Z development method provides *promotion* as a

means of objectifying isolated specifications in order to incorporate them directly in a larger global system description. The technique of promotion is common amongst Z practitioners, being first explained by Morgan and Sufrin in their description of the Unix filing system [117], and given a more rigorous mathematical treatment by Woodcock [171].

A fundamental characteristic of an object is its identity. To promote the local window description into the more global window manager level, we need a way of identifying the different instances of windows, all of which are of the same type *WindowState*. We introduce a set of window identifiers, *WINID*, so that the identification of individual windows is modelled as a partial function from identifiers to the schema type *WindowState*. The windows that the manager “knows” about are precisely those in the domain of this function. The window manager also keeps track of the current selected window or windows, to which all future input is directed. The window manager may contain other information as well, but for our present purpose we need not bother with any further detail.

[*WINID*]

$\begin{array}{l} \textit{WindowManager} \\ \textit{windows} : \textit{WINID} \rightarrow \textit{WindowState} \\ \textit{known} : \mathbf{P} \textit{WINID} \\ \textit{selected} : \mathbf{P} \textit{WINID} \\ \hline \textit{known} = \text{dom } \textit{windows} \\ \textit{selected} \subseteq \textit{known} \end{array}$
--

We have already defined operations such as *MoveWindowOp* on a single, isolated window because they were more conveniently described in that context, rather than in the context of the window manager. We can promote the local operations via the promotion schema *WindowPromote*.

$\begin{array}{l} \textit{WindowPromote} \\ \Delta \textit{WindowManager} \\ \textit{win}^? : \textit{WINID} \\ \Delta \textit{Window} \\ \hline \textit{win}^? \in \textit{known} \\ \theta \textit{Window} = \textit{windows } \textit{win}^? \\ \textit{windows}' = \textit{windows} \oplus \{ \textit{win}^? \mapsto \theta \textit{Window}' \} \\ \textit{known}' = \textit{known} \\ \textit{selected}' = \textit{selected} \end{array}$
--

This promotion schema will allow us to embed the local window operations as operations of the window manager with minimal changes to their definition. The advantage to this method is that in the definition of the window operations we did not have to worry about properties of the window manager, which makes their definitions not only simpler but more natural to express. For example, when defining the move window operation it should not concern the specifier what other windows are known to the window manager. Operations on windows can be defined directly as operations at the level of the window manager via this promotion schema.

$$WMOpenWindowOp \hat{=} (OpenWindowOp \wedge WindowPromote) \setminus \Delta WindowState$$

$$WMCloseWindowOp \hat{=} (CloseWindowOp \wedge WindowPromote) \setminus \Delta WindowState$$

$$WMMoveWindowOp \hat{=} (MoveWindowOp \wedge WindowPromote) \setminus \Delta WindowState$$

$$WMResizeWindowOp \hat{=} (ResizeWindowOp \wedge WindowPromote) \setminus \Delta WindowState$$

In our example, we did not want the global state to affect the local operations. Though the promotion technique can cope with such interference by the global state by suitable parameterization of the global operations, it is more in line with the spirit of data abstraction to avoid such interference. A limitation on promoted operations is that they only affect that part of the global state which was promoted. It is for this reason that in the operation promotion schema *WindowPromote* that we have explicitly stated that the *known* and *selected* components remain unchanged because these are aspects of the global window manager state that are separate from the individual windows. Some operations on windows are only relevant at the level of the window manager. Examples of such operations are creation and destruction of windows. Whereas the conditions for an initialized window or a terminable window were described in isolation from the window manager, the operation of creating and destroying a window can only be defined at the more global level. Whereas the lone window did not exist before creation and after termination, the window manager existence subsumes that of any window it manages.

The main advantage of promotion is that the effects of operations can be isolated to the smallest part of a complex system. This is of great advantage for describing the functionality of a complex system since the specification takes on a compositional look. The power of promotion for layering the description of a complex system is the major contribution first provided by Morgan and Sufrin. A further advantage of promotion is given by its more formal treatment by Woodcock in which we can see that the modularization allows for a proof management system

to prove the properties of a complex specification. Isolation of proof obligations is also important in our agent language.

Promotion captures the identity of objects, but it does not address very adequately the second point about object-oriented notations, that of conceptualization. As Woodcock points out [171], promotion advocates an “onion skin” approach to system development. The windowing system is viewed in layers of functionality, as shown in Figure 6.1. The innermost layer represents the single window, followed




Figure 6.1: The “onion skin” view of a windowing system

by the multi-windowing environment, the abstract interface (in terms of logical input and output devices) and the physical interface. By contrast, an object view of the windowing system may look like Figure 6.2. The onion skin view shows the




Figure 6.2: The object view of a windowing system

inheritance relationship, whereas the object view shows the instance relationship and communication.

In the next few sections, we will summarize some of the formalisms more in line with the object view.



## Schuman and Pitt

Schuman and Pitt [146], first suggested in 1987 a variant on the Z notation which was specifically designed to meet the needs of object-oriented system design and conform to its “pragmatic appeal”. Schuman, Pitt and Byers [145] followed three years later with an interpretation of classes in their earlier model as concurrent processes. Though they recognized the need to relate state declaration with initialization and operations more strongly than is done with Z (or VDM), they deliberately did not enforce that relationship syntactically, as we desire to do.

They note that one of the prime features of a concurrent specification techniques is the ability to reason about the behaviour of the individual processes (or objects) and about combinations of the several processes. They take minimalist approach to process behaviour, characterizing a process in terms of trace/implication pairs. Traces are exactly the same as for our agents above, that is, a record of events in which the process has participated. The events are operation names which are defined by a pre- and postcondition semantics on the components of the state. Implications are that information on the state of a process which can be inferred as a result of the postconditions of the operations in the trace. Communication is in terms of data flow, so synchronization results naturally from conjunction of pre and postconditions. They claim this to represent the most natural form of concurrency, a point which is very debatable. An advantage of our agents is in what we believe to be a more natural facility for expressing the communication between the separate components.

Schuman, Pitt and Byers provide two means for composing, inheritance and instantiation. Inheritance occurs when a process  $C$  can incorporate and extend the process (or class) information of  $A$  and  $B$  to define a new process which is a subclass of both  $A$  and  $B$ . Instantiation occurs when  $C$  incorporates named versions of  $A$  and  $B$  in its definition. The mathematical distinction between inheritance and instantiation is that between normal union and disjoint union, and this distinction is similar to the conditions for independent synchronization and interleaving in our agent model, though we do not need only use normal union on attribute/value mappings to obtain the composition in both cases.

## Hall's object-oriented conventions in standard Z

Hall has used standard Z with special conventions designed specifically to deal with object-oriented considerations in software development. The main convention is the use of special components in the state definition as a means of object identification. This identifier is referred to as *self* in the state definition. We did the same in promotion for the window example by introducing a set of window identifiers, *WINID*. Hall's convention suggests that this be performed from the start. Operations on the state of an object are not intended to change its identity, and so this can be

made explicit by the addition of the predicate  $self' = self$  for all operations.

Hall is also quick to point out the lack of state/operation bundling in Z, but rather than suggest an alternative, he says that the convention is to assume that all operations defined are the only ones possible for the object. This seems reasonable, but it does not address the problem of bundling.

The semantics of objects are again in terms of the traces of operations in which the object can participate. He gives examples of how a mechanical procedure can convert a schema definition of an operation to its relational equivalent. This procedure is identical to the one we showed earlier in converting the window specification to an element in *SandHProcess*, except that Hall uses only the input parameters as the index to the operation. In the agent model this would be equivalent to modelling the operations function as a relation between operations and state transitions instead of a function, and this decision appears very dubious. Hall admits that this mechanical definition cannot be defined in Z, as we have noted, and so he calls upon a convention to perform the mechanical conversion.

The set of conventions advocated by Hall has been used with apparent success in the specification of "a substantial piece" of software written in Objective-C. A very interesting case study was performed using the Model-View paradigm of Smalltalk. We will discuss the formalization of such interactive architectures in Chapter 8.

## Object-Z

Research at Queensland has resulted in the development of Object-Z, a true extension to Spivey's semantics for standard Z [34, 55, 56]. Object-Z introduces a class structure whose explicit purpose is to bind together a state description with its invariant, initialization and operations. The semantics of a class is given by the events in which the objects (variable instances) of a class participate. An event is defined as one of the class operations along with the before and after state bindings associated to that instance of the operation application. A trace of such events is called a history.

A relatively minor extension to Spivey's semantics [153] allows for the description of classes. One of the main features of the extension is the addition of roles which are used to distinguish the use of various identifiers as either input or output arguments to an operation definition or before or after values of the state. The *pre* and *post* roles are used to define legal histories of a class instance, since *post* values of the state after one operation must be the *pre* values of the state before the next operation.

Beyond the constraints on state transitions that are defined by the operations, there is also the ability to express constraints on the histories of a class instance explicitly using a temporal logic predicate language. However, this relationship between external specification over histories and the normal internal specification

is not completely formalized in [55]. The authors claim that this will allow the specification of liveness properties, in which assertions can be made about what good properties will occur in the system. Without a clear relationship between the external and internal specifications, such a statement is hard to back up. The external specification may express some liveness criterion, but that behaviour could very well be disallowed by the internal description.

Object-Z provides a simple and very useful extension to Z by the addition of classes. What remains to be done with Object-Z is clear. Just as there is a schema calculus in standard Z which allows for the manipulation and composition of schemas, there needs to be a class calculus in Object-Z which addresses the composition of, and communication between, classes. We feel that our agent model addresses some of those questions.

## Whysall and McDermid

Whysall and McDermid have also proposed a means of adding more structure to Z specifications to make them more amenable to object-oriented description via data abstraction [167, 168]. The description of objects is split into two parts, the *export* and the *body* specifications. They make explicit, as we do, that one of the main advantages of such structure is to allow composition of modules, both for the purpose of description of large systems and for the delegation of proof obligations for properties of the large systems. These two specifications roughly correspond to our external and internal interpretation relations on agents.

The export specification gives the input/output trace behaviour of the object. The notation provided allows for this specification in an algebraic language which is based on observational equivalence of method traces. Observational equivalence in this case is similar to external indistinguishability ( $\overset{ext}{\equiv}$ ) we defined in Chapter 4. The body specification dictates how the methods behave on the underlying state of the object. This is similar to our internal specification. Whysall and McDermid define the relationship which must hold between an object's export and body specifications, namely that they cannot be inconsistent with each other. Traces of method invocations allowed by the export specification must be allowed by the body specification.

## Z++

Work at Oxford has recently resulted in an object-oriented extension to Z called Z++ [100]. A separate notation has been added which performs precisely the bundling of state specification and operations. It has been shown how this notation maps down to more conventional Z. This notation does not consider communication at all between separate objects because it has been developed based on a three-

level design of systems in which each level uses operations of the level below, like the “onion skin” model associated to promotion discussed earlier. Z++ is a *wide spectrum* language because it contains both nonexecutable specification statements and potentially executable procedural statements (in UNIFORM [154]).

### 6.1.3 Other concurrent notations

We have explicitly incorporated a process algebra model of the external specification of agents because it is a simple compositional model which addresses concurrency. Process algebras gain their simplicity and power from their deliberate disregard of state information. But as Josephs points out [92], this is not always a positive feature in the design of complex systems, as the success of model-oriented specification techniques such as Z and VDM has shown. Both Josephs and He [76, 77] have developed state-based versions of communicating processes in which the transitions on the explicit state space define the traces of behaviour. It was He’s work which influenced the hybrid approach of the Sufrin and He model of interactive processes and our own extension of that model in terms of agents.

Notations for concurrency abound; an adequate review of these notations is not within the aims and scope of this thesis. We take time here to mention the Josephs and He models as signposts for the trend in recent years to develop decent models of concurrency with explicit state. Our agent model is such an attempt which has arisen out of the special needs to express interactive properties of a system design. One other attempt of note is the work by Morgan and Woodcock [170] in which a weakest precondition semantics is used to define CSP-like combinators for a concurrent extension to Dijkstra’s guarded command language.

## 6.2 The agent language

In this section, we will describe the language for agents and outline how the notation maps to the model of agents. The agent language is as far as possible a mixture of Z and CSP, so there is no need to give a detailed semantics for the language in the body of this thesis. The interested reader is directed to Appendix C for further details of the semantics for the agent language.

As we have stated, there are two ways to combine agents—by the interleaving of agents which may share common events or attributes, or by the synchronization of agents with disjoint attribute sets. Existing agent definitions can be combined, therefore, by interleaving them or synchronizing them. In addition, we can constructively define an agent by giving the internal, external and communications specifications directly. The examples in Section 6.3 show how and when each method of agent description is used.

An interactive system is a mapping from agent identifiers to the set of agents in *Agent*. We introduce a set of possible agent identifiers.

[*AgentID*]

*IntSys* == *AgentID*  $\leftrightarrow$  *Agent*

The system semantic function,  $\mathcal{S}[\_]$ , takes an existing interactive system and an agent language expression and produces a new interactive system. The agent language description represents either the synchronization or interleaving of existing agents, interleaved with an additional 3-part description of a new agent (internal, external, communication specification), or a completely new 3-part description of an agent. The following is a BNF-like description of the agent language syntax. Square brackets are used to indicate an item which is optional.

<i>AgExp</i> ::=	<b>agent</b> <i>AgentID</i>	- synchronization
	<b>synchronizes</b> <i>AgentIDList</i>	
	[ <b>with</b> <i>3PartSpec</i> ]	
	<b>endagent</b> <i>AgentID</i>	
	<b>agent</b> <i>AgentID</i>	- interleaving
	<b>interleaves</b> <i>AgentIDList</i>	
	[ <b>with</b> <i>3PartSpec</i> ]	
	<b>endagent</b> <i>AgentID</i>	
	<b>agent</b> <i>AgentID</i>	- 3-part specification
	<i>3PartSpec</i>	
	<b>endagent</b> <i>AgentID</i>	

The system semantic function  $\mathcal{S}[\_]$  is defined structurally over the elements in *AgExp*. For synchronized combination, the expression

```

agent A1
synchronizes AS
with Spec
endagent A1

```

maps the (fresh) agent identifier *A1* to the synchronous composition of the agents indicated by the sequence of (distinct) agent identifiers *AS*, if such a composition is allowed by *composeall<sub>sync</sub>*. This may then be interleaved with the agent defined by the 3-part specification *Spec*, according to the semantic operator  $\mathcal{A}g[\_]$  discussed later.

$$\begin{array}{|l}
\mathcal{S}[-] : (IntSys \times AgExp) \mapsto IntSys \\
\forall A1 : AgentID; AS : seq_1 AgentID; \rho : IntSys \\
| ( A1 \notin \text{dom } \rho \\
\wedge (AS \ ; \ \rho) \in \text{dom } compose_{all, sync} \\
\bullet \mathcal{S} \left[ \begin{array}{l} \text{agent } A1 \\ \text{synchronizes } AS \\ \text{endagent } A1 \end{array} \right] = \\
\rho \oplus \{A1 \mapsto compose_{all, sync}(AS \ ; \ \rho)\} \\
\forall A1 : AgentID; AS : seq_1 AgentID; Spec : 3PartSpec \ \rho : IntSys \\
| ( A1 \notin \text{dom } \rho \\
\wedge AS \ ; \ \rho \in \text{dom } compose_{all, sync} \\
\wedge (compose_{all, sync}(AS \ ; \ \rho), Ag[Spec]) \in \text{dom } compose_{int} \\
\bullet \mathcal{S} \left[ \begin{array}{l} \text{agent } A1 \\ \text{synchronizes } AS \\ \text{with } Spec \\ \text{endagent } A1 \end{array} \right] = \\
\rho \oplus \{A1 \mapsto compose_{int}(compose_{all, sync}(AS \ ; \ \rho), Ag[Spec])\}
\end{array}$$

Note that because of the associativity of  $compose_{sync}$  (Theorem 5.5), the order of the agent identifiers in  $AS$  does not matter.

For interleaved combination, the expression

```

agent A1
interleaves AS
endagent A1

```

maps the fresh identifier  $A1$  to the interleaved product of the known agent definitions in  $AS$  and the 3-part specification  $Spec$ , if given.

$$\left. \begin{array}{l}
\forall A1 : AgentID; AS : seq_1 AgentID; \rho : IntSys \\
| ( A1 \notin \text{dom } \rho \\
\quad \wedge AS \ddagger \rho \in \text{dom } \text{compose}_{int}) \\
\bullet S \left[ \begin{array}{l}
\text{agent } A1 \\
\text{interleaves } AS \\
\text{with } Spec \\
\text{endagent } A1
\end{array} \right] = \rho \oplus \{ A1 \mapsto \text{compose}_{int}(AS \ddagger \rho) \} \\
\forall A1 : AgentID; AS : seq_1 AgentID; Spec : 3PartSpec; \rho : IntSys \\
| ( A1 \notin \text{dom } \rho \\
\quad \wedge (AS \ddagger \rho) \sim Ag[Spec] \in \text{dom } \text{compose}_{int}) \\
\bullet S \left[ \begin{array}{l}
\text{agent } A1 \\
\text{interleaves } AS \\
\text{with } Spec \\
\text{endagent } A1
\end{array} \right] = \\
\rho \oplus \{ A1 \mapsto \text{compose}_{int}((AS \ddagger \rho) \sim Ag[Spec]) \}
\end{array} \right\}$$

Note that because of the associativity of  $\text{compose}_{int}$  (Theorem 5.4), the order of the agent identifiers in  $AS$  does not matter.

A stand alone specification of an agent,

```

agent A1
Spec
endagent A1

```

maps the fresh identifier  $A1$  to the agent  $Ag[Spec]$ .

$$\left. \begin{array}{l}
\forall A1 : AgentID; Spec : 3PartSpec; \rho : IntSys \\
| ( A1 \notin \text{dom } \rho \\
\quad \wedge Spec \in \text{dom } Ag[-]) \\
\bullet S \left[ \begin{array}{l}
\text{agent } A1 \\
Spec \\
\text{endagent } A1
\end{array} \right] = \rho \oplus \{ A1 \mapsto Ag[Spec] \}
\end{array} \right\}$$

The semantic function  $Ag[-]$  maps a 3-part description from the agent language to its representative element in *Agent*. This three part description of an agent is given by separate internal, external and communication languages. The three part specification of an agent is given by an internal, external and communication language.

```

3PartSpec ::= internal IExp
             communication CExp
             external EExp

```

The agent semantic operator,  $Ag[-]$ , is defined in terms of semantic operators for each of the sublanguages. The domain of  $Ag[-]$  is the set of combinations of internal, communication and external expressions which yield a valid agent description when they are mapped to their respective specifications in the model.

$$\begin{array}{l}
 Ag[-] : 3PartSpec \rightarrow Agent \\
 \mathcal{I}[-] : IExp \rightarrow InternalSpec \\
 \mathcal{C}[-] : CExp \rightarrow Communication \\
 \mathcal{E}[-] : EExp \rightarrow ExternalSpec \\
 \hline
 \text{dom } Ag[-] = \{ IE : IExp; CE : CExp; EE : EExp; Agent \\
 \quad \mid ( \theta InternalSpec = \mathcal{I}[ IE ] \\
 \quad \quad \wedge \theta Communication = \mathcal{C}[ CE ] \\
 \quad \quad \wedge \theta ExternalSpec = \mathcal{E}[ EE ] ) \\
 \quad \bullet \left( \begin{array}{l} \text{internal } IE \\ \text{communication } CE \\ \text{external } EE \end{array} \right) \} \\
 Ag \left[ \left( \begin{array}{l} \text{internal } IE \\ \text{communication } CE \\ \text{external } EE \end{array} \right) \right] = \mu Agent \\
 \quad \bullet ( \theta InternalSpec = \mathcal{I}[ IE ] \\
 \quad \quad \wedge \theta Communication = \mathcal{C}[ CE ] \\
 \quad \quad \wedge \theta ExternalSpec = \mathcal{E}[ EE ] )
 \end{array}$$

The template for this part of the agent language is shown below.

```

agent AgentID
internal
  types
    typedeclarations
  attributes
    attribute value bindings
  invariant
    predicate on state bindings
  initially
    predicate on state bindings
  operations
    op1(typed argument list)
      changes (explicit framing condition)
      pre precondition on state and arguments
      post postcondition on before/after state and arguments
    op2 ...
    :
communication

```



```

inputs typed input channels
outputs typed output channels
 $\tau$  internal events
external
  constructive trace description
satisfying
  trace predicate
endagent AgentID

```

For the body of this thesis, we feel there is a strong enough intuitive link between these languages and their models to warrant passing over the detail of the denotational semantics. We will, however, go into more depth on the external language.

### 6.2.1 A language for external specifications

There is a distinction about how an external specification can be produced—*explicitly* via a constructive language of traces, or *implicitly* via a predicate language on traces. The constructive language of traces is usually what is provided in the notation of a process algebra, such as CSP. In this language, the constructors provide a way to build up large process specifications in terms of smaller and simpler processes. For example, the construction

$$a \rightarrow P$$

is supposed to represent the external specification which first engages in event  $a$  and then behaves like  $P$ . The complete constructive language is very similar to a subset of CSP as defined by Hoare [82].

$ConEEExp ::= stop\langle\langle P \text{ Event} \rangle\rangle$	– deadlock
$run\langle\langle P \text{ Event} \rangle\rangle$	– the total behaviour
$skip\langle\langle P \text{ Event} \rangle\rangle$	– successful termination
$\langle\langle Event \rangle\rangle \rightarrow ConEEExp$	– prefix composition
$ConEEExp \square ConEEExp$	– choice composition
$ConEEExp ; ConEEExp$	– sequential composition
$ConEEExp \parallel ConEEExp$	– synchronous parallel composition
$ConEEExp \parallel\!\!\! \parallel ConEEExp$	– asynchronous parallel composition
$f(ConEEExp)$	– process relabelling
$\mu X : \langle\langle P \text{ Event} \rangle\rangle \bullet F(X)$	– guarded recursion

These syntactic constructs are mapped to elements in the external specification space via the semantic operator  $\mathcal{E}_{con}[\ ]$ . The function  $\mathcal{E}_{con}[\ ]$  is defined in terms of two functions, one which yields the alphabet of the construct,  $\alpha[\ ]$  and one which yields the trace set,  $\mathcal{T}[\ ]$ .

$$\left[ \begin{array}{l} \mathcal{E}_{con}[-] : ConEEExp \rightarrow ExternalSpec \\ \alpha[-] : ConEEExp \rightarrow \mathbf{P} Event \\ \mathcal{T}[-] : ConEEExp \rightarrow \mathbf{P} seq Event \\ \hline (\mathcal{E}_{con}[ES]).alphabet = \alpha[ES] \\ (\mathcal{E}_{con}[ES]).traces = \mathcal{T}[ES] \end{array} \right.$$

We define the prefix composition operator as follows.

$$\left[ \begin{array}{l} \_ \rightarrow \_ : (Event \times ConEEExp) \rightarrow ConEEExp \\ \hline \text{dom}(\_ \rightarrow \_) = \{ e : Event; P : ConEEExp \\ \quad | e \in \alpha[P] \\ \quad \bullet (e, P) \} \\ \alpha[e \rightarrow P] = \alpha[P] \\ \mathcal{T}[e \rightarrow P] = () \cup \{ t : \mathcal{T}[P] \bullet \langle e \rangle \wedge t \} \end{array} \right.$$

As done by Hoare [82, p. 134], we introduce some syntactic conventions to handle input and output. The expression

$$c?m \rightarrow P_m$$

is equivalent to the choice of every possible message that can occur along channel  $c$  followed by the behaviour of some predefined description given by  $P_m$ . For example, if there were only two messages that could occur on channel  $c$ , say  $m_1$  and  $m_2$ , then we would have the following syntactic equivalence.

$$c?m \rightarrow P_m \stackrel{def}{=} (c.m_1 \rightarrow P_{m_1} \\ \quad \square \\ \quad c.m_2 \rightarrow P_{m_2})$$

For output, we will use  $c!m \rightarrow P$  as a syntactic equivalent for  $(c, m) \rightarrow P$ . For both input and output, the arguments for messages will not be given with type information, as it is assumed the order and types of the arguments for any given message is as defined in the internal and communication specifications.

In Appendix C we give the full denotational semantics for this constructive trace specification language, along with an operational semantics to aid in the intuition behind each constructor.

The advantage of a constructive language for the external specification is that it can lead naturally to an implementation. Whereas this is a definite advantage if the desire is to produce a rapid prototyping tool from the agent language, or even a full-blown programming environment, it is not the only purpose of the agent language, nor is it the primary purpose. The primary purpose is to be able to

provide a description of the components of an interactive system in such a way as to satisfy the constraints imposed by interactive properties such as those described in Chapter 4. With that purpose in mind, we admit that it is not always so simple, or possible, to describe constraints on the external behaviour using the constructive trace language.

An example of an implicit description of an external specification would be by formulation of predicates on the trace set. In CSP, these are referred to as behavioural specifications on the traces[46]. They are used as criteria for judging whether a given CSP expression satisfies some requirement. In the external specification of agents, we will allow trace predicates to describe additional constraints on the communication of events along with the constructive language.

Providing a very powerful predicate language, such as the modal mu-calculus introduced by Pratt [131] and Kozen [95], is beyond the scope of this thesis. Just as we have given a simple example of incorporating a constructive language within the agent model, so motivating further work on incorporating more complex languages, we also provide a simple example of a predicate language and leave it open to incorporate more complex languages. A trace predicate, then, is simply a set of functions with a prefix closed domain.

$$\left| \begin{array}{l} \text{ImpEEExp} : \mathbf{P}(\mathbf{P} \text{ Event} \times (\text{seq Event} \rightarrow \text{Bool})) \\ (A, TP) \in \text{ImpEEExp} \Leftrightarrow \text{prefix\_closed } A(TP(\text{true})) \end{array} \right|$$

Just as there was a semantic function,  $\mathcal{E}_{\text{con}}[\_]$ , which mapped the constructive external language to the *ExternalSpec*, we will also have a function,  $\mathcal{E}_{\text{imp}}[\_]$  to map the predicate language to an *ExternalSpec* element.

$$\left| \mathcal{E}_{\text{imp}}[\_] : \text{ImpEEExp} \rightarrow \text{ExternalSpec} \right|$$

The overall external specification is obtained by intersecting the trace sets derived from the explicit constructive language and the implicit trace predicate language.

$$\text{ESpecLanguage} ::= \text{ConEEExp} \times \text{ImpEEExp}$$

$$\left| \begin{array}{l} \mathcal{E}[\_] : \text{ESpecLanguage} \leftrightarrow \text{ExternalSpec} \\ \text{dom } \mathcal{E}[\_] = \{ A : \mathbf{P} \text{ Event}; \text{con} : \text{ConEEExp}; \text{imp} : \text{ImpEEExp} \\ \quad \mid (\alpha[\text{con}]) = A \\ \quad \quad \wedge \text{fst}(\text{imp}) = A \\ \quad \quad \bullet (\text{con}, \text{imp}) \} \\ \mathcal{E}[(\text{imp}, \text{con})].\text{alphabet} = \alpha[\text{con}] \\ \mathcal{E}[(\text{imp}, \text{con})].\text{traces} = \mathcal{T}[\text{con}] \cap \mathcal{E}_{\text{imp}}[\text{imp}].\text{traces} \end{array} \right|$$

### 6.3 Using the agent language

Our presentation of the agent model will be aided by some examples. We would like to demonstrate in these examples how the definition of an agent in the new language more closely matches the way a designer understands it. The purpose of these examples is both to make clearer the meaning of the language and to show how various development techniques can be used to build up agent descriptions incrementally. The first example is adapted from Took's thesis and involves the description of a very crude nuclear reactor. The second example involves a description of parts of a simple windowing system.

#### 6.3.1 A toy reactor

Took [163] describes a primitive nuclear reactor, in which there are damping rods to control the reaction and coolant to take away the heat generated by the reaction. There is only one relevant attribute for the rods which indicates whether they are up, enabling the reaction, or down, disabling the reaction. Initially, the rods are down. Two operations can be performed on the rods, either lifting them or dropping them, after which the rods are up or down, respectively. The rods respond to being lifted or dropped by informing some other agent of the rod position after the change. The agent definition of the rods is given below.

```

agent rods
internal
  types
    Rodpos ::= up | down
  attributes
    position : Rodpos
  invariant
    true
  initially
    position = down
  operations
    lift()
      changes (position)
      pre true
      post position' = up
    drop()
      changes (position)
      pre true
      post position' = down
  inform(pos : Rodpos)
    pre pos = position
communication

```

```

inputs rodin : lift(), drop()
outputs rodout : inform(pos : Rodpos)
 $\tau$  :
external
 $\mu X \bullet ((rodin, lift()) \rightarrow rodout!inform(pos) \rightarrow$ 
 $(rodin, drop()) \rightarrow rodout!inform(pos) \rightarrow X)$ 
satisfying
 $true$ 
endagent rods

```

This simple example can be used to explain the notation and several conventions we will adopt in its use. There is no state invariant (beyond that implied by the typing information of the state attributes) and we have indicated this by the predicate *true* in the invariants section of the definition. The pre- and postconditions for some operations are also listed as *true*, meaning they are as weak as possible. A *true* predicate in the satisfaction predicate for the external description means that there are no further constraints on the external definition beyond those imposed by the constructive definition and implied by the internal specification. In the future, we will usually omit these sections with predicates that are *true*.

The communication description lists the input and output channels together with the messages that can be sent along those channels. Any internal messages would have been listed after the internal channel identifier  $\tau$ . There are no internal messages in this description. In the future, if there are no messages associated to a channel, it need not appear in the description. If there are no input or output events for the agent, then that section can be omitted as well.

The external specification of the rods shows the hybrid nature of the agent language. We intend to prohibit two consecutive lifts or drops of the rods. We could have easily defined the preconditions for the operations for the *lift()* and *drop()* messages to accommodate this. However, we chose in this example to use the constructive trace language of the external specification to express this constraint. Though it does not matter in such a simple example, there may be reasons to avoid expressing such ordering constraints by definition of preconditions. We also enforce the constraint that an output message informing of the new rod position must be performed between each input message.

The coolant for the reactor is very similar to the rods (when viewed as an agent, that is). The coolant has only one attribute of interest which indicates the level of the coolant as either in or out. Initially, the coolant is in. Operations are defined to add or remove the coolant, after which the coolant is in or out, respectively. The agent definition of the coolant is given below.

```

agent coolant
internal
types

```

```

    CoolantLevel ::= in | out
  attributes
    level : CoolantLevel
  initially
    level = in
  operations
  add()
    changes (level)
    pre level = out
    post level' = in
  remove()
    changes (level)
    pre level = in
    post level' = out
  inform(l : CoolantLevel)
    pre l = level
  communication
    inputs rodin : add(), remove()
    outputs rodout : inform(l : CoolantLevel)
  external
  satisfying
     $\forall t : \text{traces}$ 
    •  $\#(t \mid \{\text{rodout}\}) \leq \#(t \mid \{\text{rodin}\}) \leq \#(t \mid \{\text{rodout}\}) + 1$ 
  endagent coolant

```

Again, no two *add()* or *remove()* messages can be received consecutively. In this example we have used the precondition technique to specify this constraint. The external description need only ensure that each input event is followed by an *inform* event on the *rodout* channel, which is described by a satisfaction predicate on the trace set. The lack of a constructive external description means that its constraint is as weak as possible, equivalent to  $\text{run}_A$ , where  $A$  is the alphabet of the agent.

The reactor agent interleaves the rod and coolant agent with the additional constraint that at no time can the level of the coolant be *out* whilst the damping rods are *up*. The interleaved combination of *rods* and *coolant* does not prohibit this, and so we must add this in as a separate invariant of the interleaved combination.

```

  agent reactor
  interleaves (rods, coolant)
  with internal
  invariant
     $\neg (\text{rods} = \text{up} \wedge \text{level} = \text{out})$ 
  endagent reactor

```

The behaviour of the agent *reactor* is any interleaving of the separate behaviours of the rod and coolant agents except for those which would lead to the forbidden

state.

### 6.3.2 Some input devices

#### The keyboard

The keyboard is composed of a collection of buttons, one for each labelled key. A button is a simple transducer, converting inputs of presses and releases by the user to outputs of ups and downs to the window manager. Such a button is defined below. It has one attribute of interest representing the status of the button (*up* or *down*). Two input operations correspond to the press action and the release action. After each, the new button status is sent as a response to the window manager.

```

agent button
internal
  types
    BStatus ::= up | down
  attributes
    bstatus : BStatus
  initially
    bstatus = up
  operations
    press()
      changes (bstatus)
      pre bstatus = up
      post bstatus' = down
    release()
      changes (bstatus)
      pre bstatus = down
      post bstatus' = up
    inform(bs : BStatus)
      changes ()
      pre bs = bstatus
  communication
    inputs buttin : press(), release()
    outputs buttout : inform(bs : BStatus)
external
   $\mu X \bullet \text{buttin}?x \rightarrow \text{buttout}!\text{inform}(b) \rightarrow X$ 
endagent button

```

We can define a relabelling operation on agents, so that for any agent identifier  $A1$ , the expression

$$A1[\text{old}_1, \text{old}_2, \dots / \text{new}_1, \text{new}_2, \dots]$$

is mapped to the agent identified by  $A1$  with the name  $old_1$  replaced by  $new_1$ ,  $old_2$  replaced by  $new_2$ , etc. We assume that the mapping from old names to new names is injective. Furthermore, for any agent identifier  $A2$  we will write  $new::A2$  as shorthand for renaming all attributes and channels with the prefix  $new$  added. The keyboard would be represented as the synchronized combination of keys, each of which behaves as defined by *button* with the attribute and input and output channels appropriately renamed. We could have chosen to either synchronize or interleave in this example since they would both result in the same agent definition. We choose to synchronize in this case because it corresponds with our notion of the keys as separate entities on the keyboard. In the design space notation of Card, Mackinlay and Robertson [31], this composition is classified as “layout composition”, referring to the collocation of two or more devices (the keys in our example) on different places of a common panel (the keyboard in our example).

```

agent keyboard
synchronizes { a::button, b::button, ... }
endagent keyboard

```

### The mouse

The mouse also makes use of the *button* agent for each of its buttons. In addition, the mouse can transform movements in a 2-dimensional plane to values in the *Pixel* plane. This transformation is an unchangeable attribute of the mouse agent, *res*, and it corresponds to what Card, Mackinlay and Robertson refer to as a “resolution function”. Below, we describe the movement portion of a mouse as the agent *MouseMove*.

```

agent MouseMove
internal
types
  mousezmax, mouseymax : N
  MousePlane == 0..mousezmax × 0..mouseymax
attributes
  currentmove : Pixel
  res : MousePlane → Pixel
operations
  sweep(delta : MousePlane)
  changes (currentmove)
  post currentmove' = res(delta)
  mousemove(delta : Pixel)
  changes ()
  pre delta = currentmove
communication
  inputs gesture : sweep(delta : MousePlane)
  outputs mouseout : mousemove(delta : Pixel)

```



```

external
   $\mu X \bullet gesture?x \rightarrow mouseout!mousemove(pd) \rightarrow X$ 
endagent MouseMove

```

The mouse agent is the synchronized combination of the movement agent and the buttons agents, of which we assume there are three.

```

agent Mouse
synchronizes  $\langle 1_m::button, 2_m::button, 3_m::button, MouseMove \rangle$ 
endagent Mouse

```

Card, Mackinlay and Robertson have produced a generative design space for describing a large class of input devices. The example of the keyboard and mouse above show how we can realize some points in that space and it would be an interesting exercise to generate more of that space in the agent language. Our limited experience suggests that the slightly unnatural way of expressing stimulus-response behaviour may be a limiting factor in the success of such an exercise.

### 6.3.3 A window

We can represent an individual window in a windowing system based on the description given in Section 6.1. First we will allow some global type definitions which will be available to all agent definitions in this section. These will allow us to talk about points in the window space in terms of pixels, and images in terms of pixel maps.

```

xmax, ymax :  $\mathbb{N}$ 
Pixel ==  $0..xmax \times 0..ymax$ 
Bit ::= 0 | 1
PixelMap == PIXEL  $\leftrightarrow$  Bit

```

We will assume appropriately defined operations on elements in *Pixel*, such as addition and a natural ordering  $\leq$ . We give an incremental description of the window which mirrors the development method used with Z. The window state is first, indicating the status of the window as open or closed, the icon associated to the window when closed and its position, and the size of the window when open and its position. Initially, the window is open, and the other attributes are set to some default value.

```

agent WindowState
internal
  types
    winsize_def : Pixel
    icon_def, blank : PixelMap
  attributes

```

```

    icon, contents : PixelMap
    iconpos, winpos, winsize : Pixel
    status : open | closed
invariant
    winpos + winsize ≤ (xmax, ymax)
initially
    status = open
    winpos = iconpos = (0,0)
    winsize = winsizedef
    icon = icondef
    contents = blank
operations
    showwin(c : PixelMap)
        changes ()
        pre status = open
           c = contents
    showicon(i : PixelMap)
        changes ()
        pre status = closed
           i = icon
communication
outputs
    winout : showwin(c : PixelMap), showicon(i : PixelMap)
endagent WindowState

```

The agent *WindowState* will be able to continuously send output messages of its blank window contents. Because the messages *showwin(c)* and *showicon(i)* will be used by many other agents to be defined below which interleave with *WindowState*, we must show in each case that the effect on attributes not in common is the identity transition. In practice, the messages that are shared between agents represent output responses and these are usually defined as identity transitions, so we will satisfy the constraint trivially.

Opening and closing the window is described by interleaving *WindowState* with an agent description of those operations.

```

agent OpenClose
interleaves (WindowState)
with
internal
    operations
    open()
        changes (status)
        pre status = closed
        post status' = open
    close()

```

```

    changes (status)
    pre status = open
    post status' = closed
communication
  inputs
    winin : open() , close()
endagent OpenClose

```

Moving the window is described as a single operation. The postcondition covers the two cases in which the window is open or closed.

```

agent Move
interleaves { WindowState }
with
internal
  operations
    move (pos : Pixel)
    changes (winpos, iconpos)
    post status = closed  $\Rightarrow$  ( winpos' = pos
                                 $\wedge$  iconpos' = iconpos)
    status = open  $\Rightarrow$  ( iconpos' = pos
                         $\wedge$  winpos' = winpos)
  communication
  inputs
    winin : move(pos : Pixel)
endagent Move

```

Finally, resizing a window can only occur when the window is open.

```

agent Resize
interleaves { WindowState }
with
internal
  operations
    resize (s : Pixel)
    changes (winsize)
    pre status = open
    post winsize' = s
  communication
  inputs winin : resize(s : Pixel)
endagent Resize

```

All of the above agents can be combined to give an agent description of a window. We add the constraint overall that requires each event on the *winin* channel to be followed by an event on the *winout* channel. The event on the output channel either indicates the iconic form of the window or displays the contents, depending

on whether the window is closed or open. In order to guarantee this behaviour we introduce internal events  $(\tau, isopen())$  and  $(\tau, isclosed())$  to decide which view of the window is output as a response to the last input.

```

agent Window
interleaves {OpenClose, Move, Resize}
with
internal
  operations
  isopen()
    pre status = open
  isclosed()
    pre status = closed
  communication
   $\tau : isopen(), isclosed()$ 
external
  winout!showwin(c)  $\rightarrow$ 
   $\mu X \bullet (winin?x \rightarrow$ 
     $((\tau, isopen()) \rightarrow winout!showwin(c) \rightarrow X$ 
     $\square$ 
     $(\tau, isclosed()) \rightarrow winout!showicon(c) \rightarrow X))$ 
endagent Window

```

### A window manager

Having defined the individual window, we want to investigate how to define the window manager as an agent which synchronizes with and coordinates the activity of the individual windows. Figure 6.3 gives a graphical view of the agent relationship we want to capture.

Interaction with any window is coordinated via the window manager. In the manager we describe here, we will assume that windows are nonoverlapping in *Pixel* space and that commands to the window manager are directed to the currently selected window (if there is one). The currently selected window is determined by the position of the mouse cursor. These assumptions have been made to make the description simpler. We could easily relax these restrictions, but the corresponding description of the window manager would be more complex. We will not give a complete description of the window manager either. Instead, we will give an example of how it coordinates activity by describing how the window opening command is defined as an operation on the window manager.

The attributes of concern for the window manager include:

- the position of the mouse cursor;
- the set of *PixelMaps* associated to active windows in the system;

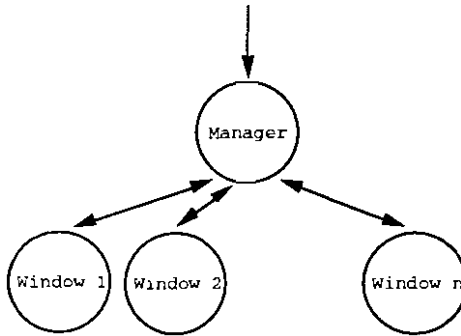


Figure 6.3: Window manager/window relationship

- a function for determining to which window a point in Pixel space belongs; and
- the currently selected window, which is at most one window.

Identification of windows will be by a set of window identifiers, as was proposed by the promotion solution discussed earlier.

{WinID}

We will not concern ourselves with initialization considerations in this description, as it is not important to the example of opening windows. Below is the agent description of the attributes and invariants for the window manager. The constraint of nonoverlapping windows is embodied in the definition of *pick* as a function.

**agent** *WindowManager*

**internal**

**attributes**

*mousepos* : Pixel  
*windows* : WinID  $\leftrightarrow$  PixelMap  
*known, selected* : P WinID  
*pick* : Pixel  $\rightarrow$  WinID

**invariant**

*known* = dom *windows*  
 $\#selected \leq 1$   
 $selected \subseteq known$   
 $pick = \{ w : known; p : Pixel$   
     |  $p \in \text{dom}(windows(w))$   
     •  $p \mapsto w \}$   
 $selected = pick(\{mousepos\})$

The window manager will receive as input messages from a potential user which can change the mouse position or open the current window. From individual windows, the manager receives as input the *showwin(c)* message that was mentioned earlier. The response message of importance in this example is the message sent to the selected window to open it. We also describe an internal message *isselected(w : WinID)* which decides if the window labelled by the identifier *w* is the currently selected window. This internal event will allow us to give the desired external description of the manager's behaviour. The following is the definition of the operations for the window manager.

```

operations
movemouse(dp : Pixel)
  changes (mousepos)
  post mousepos' = mousepos + dp
showwin(c : PixelMap)
  changes (windows, selected, pick)
  pre selected ≠ ∅
  post windows' = windows ⋈ { s : selected • s ↦ c }
openwindow()
  changes ()
  pre selected ≠ ∅
open()
  changes ()
isselected(w : WinID)
  changes ()
  pre w ∈ selected

```

The communication specification for the window manager has one channel *wmin* for input received from the user and one input channel from every window. Our description of the window manager is statically determined, that is, all communications channels that will ever be used must be declared at once since there is no means of adding channels.<sup>2</sup> We therefore define an input channel from every window, even if that channel is never used. Likewise, we define an output channel from the window manager to each window. The communication description for the window manager follows.

#### communication

##### inputs

```

wmin : movemouse(dp : Pixel), openwindow()
∀ w : WinID • w::winout : showwin(c : PixelMap)

```

##### outputs

```

∀ w : WinID • w::wmin : open()
τ : isselected(w : WinID)

```

<sup>2</sup>It should be possible to use the agent model dynamically, allowing channels to be communicated as well, but we have not investigated this as yet

The external specification of the window manager is more complicated than we have seen previously. This is because the operations defined do not imply much ordering information. When an *openwindow()* message is sent along the *wmin* channel, it is meant to open the currently selected window, determined by the mouse position. The window manager determines which window is selected by the internal event  $(\tau, isselected(w : WinID))$  which binds the value of *w* so that the manager can send the *open()* message along the channel *w::winin*, and then receive the new pixel image for the newly opened window along the *w::winout* channel, as desired. The selected window is changed by movement of the mouse. The external specification for the the window manager is given below.

```

external       $\mu X \bullet ((\Box_{w, WinID}((\tau, isselected(w)) \rightarrow OPEN_w))$ 
                 $\square$ 
                 $(wmin, movemouse(dp)) \rightarrow X)$ 
where
   $OPEN_w = ((wmin, openwindow()) \rightarrow (w::winin, open()) \rightarrow$ 
             $(w::winout, showwin(c)) \rightarrow X)$ 
endagent WindowManager

```

## 6.4 Conclusions

In this chapter we have presented a language for the description of agents which is more suited to design than the model notation of Chapter 5. The language we have chosen borrows heavily from standard Z and CSP. This has two advantages. First, the semantics of the notation is similar to the semantics for those two notations. Second, familiarity with those notations should increase familiarity with the agent notation.

There are disadvantages, however, to our approach. The agent model is based on a stimulus-response model, but there is no natural means for expressing the connection between stimulus and response, which should occur at the operation description level. Rather, we have had to describe all inputs (stimuli) and outputs (responses) as independent messages and leave it to the pre- and postconditions of their operation definitions or the external specification to describe any relationship between them. Since output messages usually have no effect on the internal state of an agent (at least not in the examples we have given), the relationship between stimulus and response is left entirely up to the external specification. A more serious criticism is that this is not necessarily a flaw in the notation we have chosen but an undesired bias in the agent model itself.

This lack of stimulus-response connectedness is probably the most important feature of the Z language that we have lost in the agent language. Z provides a very natural way of grouping stimulus and response behaviour with a state transition. Unfortunately, Z does not provide the general communication mechanism we would

like, which is why we had to abandon it. A possibility we would like to investigate would be the enhancement of Object-Z by provision of a more flexible facility for expressing object to object communication.

Despite these criticisms, the agent language is useful for descriptions of realistic interactive systems, and its hybrid approach and the composition operators do allow for a systematic description of complex behaviour, as desired.



## Chapter 7

# Properties of interactive systems: Part II

In this chapter we will use the refined agent model to formally express interactive properties at a more concrete level than was done in Chapter 4. In that chapter, we showed how a simple stimulus-response model could be used to formulate properties relevant to interaction with an agent. The class of properties which we could express were very abstract, much like the properties expressible in the PIE model. Properties such as predictability, consistency, honesty and trustworthiness, as expressed in the simple agent model, are not very closely linked with the interactive properties suggested by the interaction framework. The framework suggests that the *Interface* (*Input* composed with *Output*), should effectively mediate between the tasks of the *User* and the functions of the *System*.

Therefore, the task is seen as providing scope for properties which attempt to describe how the *Interface* can serve as a good mediator. We define the role of task analysis in the reformulation of interactive properties. We define the output of task analysis as an identification of *System* and *Interface* attributes which are understood by the *User* to portray the effect of a given task. The *System* attributes form a *result template* of information relevant to the achieved goals of the interaction. The *Interface* attributes form a *display template* of visible information that the *User* understands in terms of how it reflects the values of attributes in the result template. Templates are used to restrict access to information of an agent by any of its observers. The interactive properties we will formulate in this chapter will be relationships between the input histories and a result or display template value history.

Since the properties are formulated with respect to the result and display templates, they are task-specific. Hence, we will be able to describe, for example, what it means for a graphics package to be display predictable with respect to selection of objects or show how a direct manipulation operating system interface lacks hon-

esty in supporting the task of directory creation. Though we will not present any of the examples in this chapter in complete detail, we will present enough of the detail within the agent language so that an informal understanding of the properties can be gained with confidence that a completely formal description could also be provided.

The reformulation of task-specific interactive properties and the examples justify our claim that the refined agent model and language take us one step closer towards realizing a design practice guided by principled reasoning. In a design situation, the task analysis provides the result template information for each task that an agent should support. The *System* agent is then specified and can be verified to check that it satisfies certain result only properties, such as result predictability. The duty of the designer is then to determine for each task which display templates can be chosen and specified to satisfy the constraints imposed by the result-display properties, such as honesty. In analysing an existing design, the task analysis provides the display template information as well as the result template information and the design can be tested to see if it satisfies any of the interactive properties, and if not why it fails.

The main reason we will be able to speak more concretely is that the we know more about how the agent functions. The internal structure of the agent is no longer just a set of states about which we know nothing further. We can now look inside that state and describe it in terms of its attributes, and these attributes can be determined based on the level of the description necessary to address a given task. Just as the agent view of an interactive system allows us to separate its description into more manageable computational units, so too does the attribute and template view of an agent allow us to separate its description into more manageable perceivable units.

### Overview of chapter

We begin in Section 7.1 with a discussion of the relationship between display and result and how previous formalisms have treated this relationship. In Section 7.2 we define a template as a means of isolating aspects of the display and result and relating them between composed agents. The remaining sections of this chapter reformulate the properties discussed in Chapter 4 using the display and result template information to derive more task-relevant definitions. Several examples are discussed along the way to make clear the use of the reformulated interactive properties, especially in the analysis of existing designs.

## 7.1 Relating Display and Result

As we have said, the purpose of interaction between human and computer is for the user to attain certain goals within some application domain. The *results* of interaction, then, are the achieved goals. The product of a task analysis is a description of the user's assumed goals, or an identification of the desired results of interaction. A principled, user-centred design methodology begins with this information and describes the system initially as an abstract machine whose state description is well suited to the task structure of the user. Hence, there should be a close match between the results as defined by the task analysis and the state of the system.

Even for systems which have not been designed initially in this way, it is possible to reverse engineer an abstract description of the system which concentrates on task analytic information. This is the intent behind the examples of this chapter and the next, in which properties are exemplified by examining how they are missing in existing interactive systems. The motivation for this kind of analysis comes from the scenario methodology [174, 175], which provides examples of user behaviour with real systems as a means of cross-fertilization of different modelling domains in HCI research.

It is outside the scope of this thesis to consider any further the description of the *User*. Since we assume it is possible to produce a description of the *System* in line with the results desired by the *User*, we will use the term result from now on to refer to the end products of user interaction in terms of the state of the *System*. The agent model and its language deal directly with building up a description of the *System* state.

The user does not see the state of the system directly; rather, parts of it are rendered in the display space. A *display* is the immediate and perceivable feature of the system from which the user must interpret the relevant features of the system state. The user constructs a relationship between the perceivable information of the display and the desired but hidden information on the result of the system. Having established this crucial link between the display and the result, we can investigate means for formalizing the relationship.

As mentioned in Chapter 3, the red-PIE model is an extension to the simple PIE model which allows for an abstract discussion of the relationship between result and display. In the red-PIE model, we introduce separate functions on the effects space which separately extract the display and result information, as shown again in Figure 7.1.

In the same vein, the Sufrin and He model of processes was enhanced to discuss the relationship between results and views. In that model, an interactive process is built on top of the original process (represented by the schema type *SandHProcess[S]* in Chapter 6) by including mappings from the state space to the display and result spaces and events which trigger the calculation of the display




Figure 7.1: The red-PIE model.

and result from the value of the state at the time of the event.

In the agent model, we adopt a more constructive approach in order to relate better with actual design. Consequently, we must be more explicit about the display and the result than is necessary in either of the other two more abstract models. Looking to the interaction framework, we see that the results as we have described above are derived from the *System* and the displays from the *Interface*, or more specifically, the *Output*. Relationships between display and result, then, are relationships between agents. We derived one such relationship in Section 4.3, called correspondence. The key to correspondence earlier was the definition of a relationship between the state spaces of two agents, called the *retrieve* relation. The simple agent model provided no way to derive the *retrieve* relation. Attributes and templates will provide such a mechanism that can be derived from empirical evidence. The attributes used in the description of an agent are then viewed as design decisions by the specifier and the justification for those decisions comes from the templates obtained from task analysis. We delay a fuller treatment of correspondence until the next chapter.

Properties like predictability, consistency and honesty were also expressible in the simple agent model. These properties relied on a distinction between internal state behaviour and external response behaviour and classified systems in terms of the equivalence and indistinguishability of stimuli with respect to the state changes or responses they determined. Using a similar approach in this chapter, we will be able to define these properties and others in terms of display and result templates on agents composed as described by the interaction framework. The result templates will highlight relevant parts of the *System* state, whereas display templates will highlight information to be presented potentially as channels of communication to the user. The advantage of this approach will be that it will be made explicit what assumptions an agent design makes about the user's understanding of the interaction, and these assumptions can then be tested to see if they lead to more usable systems.

In sum, the properties we are interested in expressing in this chapter deal with the relationship between the display and the result, as indicated by the additional

unnamed arrow in Figure 7.1. They will be given as properties dictated by the design of an agent-based system and they are intended as gauges for usability in situations where the user depends on the display to determine the underlying result of the interaction or when the result knowledge of the user is needed to determine the meaning of the display.

## 7.2 Templates

The reason for introducing attributes is to enable restricted views of an agent. The restricted views are called *templates* and were first discussed by Harrison, Roast and Wright [74] as a means of focusing attention on parts of the result or display in order to formalize realistic relationships between them. For example, when using a text editor within a single window on a multi-windowing workstation, the user does not usually care about the contents of other windows. They may not even pay attention to some of the information in the text editor window. If a user relies upon certain properties of the interface, such as honesty, then it is important to be able to narrow the scope of information in the system over which the honesty property holds. The importance of templates in design and analysis is that they are not generated by our formalism; rather, they are regarded as input to the formalism which leads to a truly formal *method* of design and analysis.<sup>1</sup>

We mentioned in Chapter 4 how we can view the properties such as predictability in terms of the demands they make on the user. For example, in a graphic drawing package, the user must remember the order in which objects were drawn for predicting the consequence of clicking to select in a region of the 'canvas' occupied by overlapping objects. A Macintosh user must remember the value of the 'items' indicator in the top left-hand corner of a folder window to determine whether clicking on 'New Folder' produced a new folder, since the new icon may not appear in the window [137]. In text editing, I don't usually pay attention to the position of the mouse pointer when I am typing, so the mouse pointer is not in my display template when I predict the result of inserting a character. This does cause problems—the infamous 'unselected windows' scenario—when the mouse is accidentally nudged and my keyboard input is suddenly directed to the wrong window. Notice how the description of each of these examples is given within the scope of some task the user wishes to perform.

---

<sup>1</sup>Most of what is commonly referred to as *formal methods* does not actually represent any method, but just a notation with a formal semantics

### 7.2.1 Agent restriction

A template defines a restricted view of an agent by limiting the information known about the underlying state. Since all that is known about individual points in the state space is contained in the attribute-value mappings, we can define a template as a nonempty, finite set of attributes.

$$\text{Template} ::= F_1 \mathcal{A}$$

A template defines that part of the internal specification which is of interest; it therefore induces a new internal specification for an agent. The attributes of the template must all be valid attributes of the original internal specification. The types of the template attributes remain the same. We only keep state information on the attributes in the template, so states of the original internal specification which agree on all attributes of the template are equated in the new internal specification. For each message of the original internal specification, and for each state transition in the operation associated to that message, restricting the view of those states to the attributes of the template yields a state transition for that message in the new internal specification. A characterization of template restriction on an internal specification is given by the schema *InternalRestrict* below, in which the original internal specification is decorated with  $^I$  and the induced internal specification is decorated with  $^J$ .

$\text{InternalRestrict}$ $\frac{\text{InternalSpec}^I \quad \text{InternalSpec}^J \quad t : \text{Template}}{t \subseteq \text{attrs}^I}$ $\text{attrs}^J = t$ $\text{type}^J = t \triangleleft \text{type}^I$ $\text{states}^J = \{ s : \text{states}^I \bullet t \triangleleft s \}$ $\text{messages}^J = \text{messages}^I$ $\forall m : \text{messages}^J$ <ul style="list-style-type: none"> <li>• <math>\text{operations}^J(m) = \{ (s, s') : \text{operations}^I(m) \bullet (t \triangleleft s) \mapsto (t \triangleleft s') \}</math></li> </ul>
---

A restriction of an agent with respect to a template of attributes is characterized by the schema *Restrict*, in which the original agent is decorated with  $^I$  and the template-restricted agent is decorated with  $^J$ . The only change to the agent is to its internal specification, as given by *InternalRestrict*. The communication and external specifications remain the same.

$Restrict$ $Agent^I$ $Agent^J$ $InternalRestrict$
$\theta Communication^I = \theta Communication^J$ $\theta ExternalSpec^I = \theta ExternalSpec^J$

We write  $A \Downarrow t$  for the agent induced by restriction of  $A$  to the template  $t$ .

$\_ \Downarrow \_ : (Agent \times Template) \mapsto Agent$
$\_ \Downarrow \_ = \{ Restrict \bullet (\theta Agent^I, t) \mapsto \theta Agent^J \}$

## 7.2.2 Result and display templates

We can restrict any agent to a template of a subset of its attributes. However, within the interaction framework, different templates take on a different purpose, and so we distinguish them by name. Templates applied to the *System* agent are called *result templates*. Templates applied to the *Interface* agent we call *display templates*.

Attributes in the description of the *User* are similar to *semantic features* as used in TAG [128] and in the knowledge analysis work of Young and Whittington [173]. A task description for the user highlights the psychological attributes of importance. These psychological attributes are roughly equivalent to *System* attributes. Therefore, given a task, we can isolate the *System* attributes of interest for that task. So the output of some task or knowledge analysis, would yield a mapping from a set of identified tasks to result templates. A designer must then choose attributes of the *Interface*, the display template, which will relate to the result template for the task. We model the output of task or knowledge analysis by functions from some set of tasks to the result and display templates the tasks require.

[TaskID]

$Rtemplate : TaskID \mapsto Template$ $Dtemplate : TaskID \mapsto Template$
$dom Rtemplate = dom Dtemplate$ $Rtemplate(t) \subseteq System.attrs$ $Dtemplate(t) \subseteq Interface.attrs$

From a methodological point of view, the result template information is always provided by the task analysis. Display template information is provided for analysis of an existing system, but it is not provided for the design of a new system, since that is the responsibility of the designer. However, since real design is an iterative procedure and the systems users have in the past experienced influences their understanding of the tasks they will perform with new systems, we will stick with the assumption above that task analysis provides both the display and result template information.

We can now define properties of the interactive system in terms of these task-dependent result and display templates. This provides a principled means of iterative design, because if the result and display templates defined do not satisfy the requirements of the properties we will discuss then they will have to be altered. Typically, the display templates will be altered, as they represent attempts by the designer to effectively portray the result templates at the interface. However, as the structure of tasks is seen by many HCI researchers as dynamic [157, 101, 35], we can allow for the result template information to change as well.

### 7.2.3 Equivalence and indistinguishability revisited

In the latter half of Chapter 4, we introduced the notions of equivalence and indistinguishability as a way of relating different program inputs for an agent. Then, we differentiated between internal and external equivalence and indistinguishability to show how certain classes of properties could be expressed as relationships between internal behaviour and external behaviour. From a software engineering point of view, the distinction between internal and external behaviour, as personified by the interpretation relations  $I_A^{int}$  and  $I_A^{ext}$ , is satisfactory because it can be used as the basis for a refinement calculus on agents, similar to the way that Whysall and McDermid use export and body specifications of objects in their object-oriented use of Z [167, 168]. From an HCI point of view, this distinction is less useful because it does not directly address the relevance of the external or internal information to the goal-directed interaction.

The result/display distinction is directly intended to address this earlier deficiency. Rather than suggest that the following formalizations of interactive properties replace the ones developed in Chapter 4, we suggest that they are complementary, moving more toward analysis of HCI considerations. We point out that the earlier properties need further investigation to make clear their contribution to refinement within the agent model, but it was not the intent of this thesis to pursue that point. Sufrin and He [158] showed how to define refinement on processes and proved the soundness of downward simulation along with a method for stepwise refinement. A similar procedure can be carried out on the agents.

Given a task  $t : TaskID$  and an interactive system  $IS$  defined as a collection



of composed agents, we want to formalize the relationship between the results and displays of  $IS$  with respect to the task  $t$ . If we restrict  $IS$  to the attributes in  $Rtemplate(t)$ , we will capture the result behaviour of the computer. If we restrict  $IS$  to the attributes in  $Dtemplate(t)$ , we get the display behaviour. Histories of either restricted agent are equivalent if they lead to the same possible results or displays and they are indistinguishable if further identical extensions to the histories does not betray the result or display equivalence. Below we give the formal definitions of result and display equivalence and indistinguishability.

$$\begin{array}{|l}
 \hline
 \underline{- \approx -} : (Agent \times Template) \rightarrow \{\text{seq Event} \leftrightarrow \text{seq Event}\} \\
 \underline{- \equiv -} : (Agent \times Template) \rightarrow \{\text{seq Event} \leftrightarrow \text{seq Event}\} \\
 \hline
 \text{dom}(\underline{- \approx -}) = \{ A : Agent, t : P \mathcal{A} \mid t \subseteq A.attribs \} \\
 \text{dom}(\underline{- \equiv -}) = \text{dom}(\underline{- \approx -}) \\
 p \approx_{(A, temp)} q \Leftrightarrow I_{A \downarrow temp}^{nl}(\{\{p\}\}) = I_{A \downarrow temp}^{nl}(\{\{q\}\}) \\
 p \equiv_{(A, temp)} q \Leftrightarrow \forall r : \text{seq Event} \bullet (p \hat{\ } r) \approx_{(A, temp)} (q \hat{\ } r) \\
 \hline
 \end{array}$$

In situations where the agent is clear from context, we will usually abbreviate  $p \approx_{(A, temp)} q$  to  $p \approx_{temp} q$ .

We can now use these more task-oriented definitions of equivalence and indistinguishability to discuss interactive properties similar to Chapter 4. The introduction of task information and the result and display templates allows us to define a scope to the properties discussed in Chapter 4 that was not then possible. In addition, the meanings associated to the properties will bear more significance toward the interaction between *User* and *System*, since they will be couched in terms of task and will be focussed on features relevant to the task.

In addition to defining a scope for properties, attributes allow the definition of some proof obligations. The properties we will discuss below are expressed over histories of interaction, which includes all events in which an agent participates. When an agent is restricted to a given template, some of the events in which the agent participates have no effect on the restricted state; their effect is confined to attributes outside the template and, therefore, they are independent of the template. Such independent events can be neglected in proofs over histories for a given template.

$$\begin{array}{|l}
 \hline
 \underline{independent} : (Agent \times Template) \rightarrow P \text{Event} \\
 \hline
 \text{dom } \underline{independent} = \{ A : Agent, t : Template \mid t \subseteq A.attribs \} \\
 e \in \underline{independent}(A, t) \Leftrightarrow \\
 \#((A \downarrow \text{not\_}t).operations(\text{mess}(e))) = \#(A.operations(\text{mess}(e))) \\
 \text{where } \text{not\_}t = A.attribs - t \\
 \hline
 \end{array}$$

Sometimes, the user believes some events are independent of some attributes, when in fact they are not. For example, in the ‘unselected windows’ scenario mentioned earlier, the problem arises because keyboard events are not independent of mouse position, though the user tends to forget this.

### 7.3 Predictability and Consistency

We fix on a particular instance of the task knowledge for an interactive system given by  $Rtemplate$  and  $Dtemplate$  defined earlier. A task is *result predictable* if result equivalent histories with respect to that task and its associated result template are also result indistinguishable.

$$\begin{array}{|l}
 \text{ResultPredictable} : \mathbf{P}(\text{TaskID}) \\
 \hline
 \text{task} \in \text{ResultPredictable} \Leftrightarrow ( \text{task} \in \text{dom } Rtemplate \\
 \quad \wedge ( - \approx_{rt} - ) \subseteq ( - \equiv_{rt} - ) ) \\
 \text{where } rt = Rtemplate(\text{task})
 \end{array}$$

Technically, the above equivalence and indistinguishability should be indexed by the particular agent, say  $IS$  which represents the composition of a *System* and *Interface* along with the translations between them. We take the liberty of omitting explicit mention of  $IS$  throughout this chapter.

It is sensible that every task of interest be result predictable. Some possible reasons for a task which is not result predictable would be the task’s dependence on pure randomness (rare), or an incomplete task analysis (undesirable).

A task is *display predictable* if display equivalent histories are also display indistinguishable.

$$\begin{array}{|l}
 \text{DisplayPredictable} : \mathbf{P}(\text{TaskID}) \\
 \hline
 \text{task} \in \text{DisplayPredictable} \Leftrightarrow ( \text{task} \in \text{dom } Dtemplate \\
 \quad \wedge ( - \approx_{dt} - ) \subseteq ( - \equiv_{dt} - ) ) \\
 \text{where } dt = Dtemplate(\text{task})
 \end{array}$$

Display predictability is related to the original predictability properties formulated in the PIE except that the addition of task information has confined its scope. We can give a simple and semi-formal example of the violation of display predictability, based on the SuperPaint graphic art package for the Macintosh [143]. SuperPaint supports *layers* for the construction and manipulation of pictures. These layers are independent canvasses upon which pictures can be constructed. The two layers in SuperPaint are the paint layer, in which freehand pictures are constructed and manipulated at the lowest level of screen detail (the pixel), and the object layer, in which text, boxes, circles, etc. can be created and edited as whole entities.

The particular scenario of interest for display predictability involves the task of selection. Figure 7.2 depicts the situation presented to the user. A mouse click on

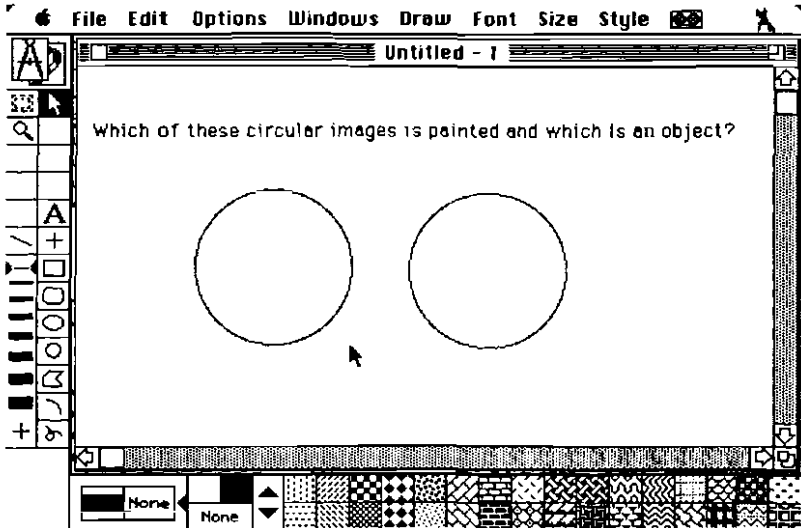


Figure 7.2: Ambiguous object selection in SuperPaint

an object will select it, so if the user knew which visual images were objects, then it would be possible to predict from a display template consisting of the mouse pointer position and the positions of the objects the effect of subsequent mouse clicks. SuperPaint does not contain in its description of the *Interface* attributes which distinguish objects from painted pictures, so the only attribute information that is perceived via a display template by the user is the set of all visual images. Since similar looking visual images can be created in either layer, selection is not display predictable.

We can base the task analysis information for selection on the description given in the user's manual [143]. We will first treat the different layers separately, and then examine their combination. For the paint layer, we have an agent description

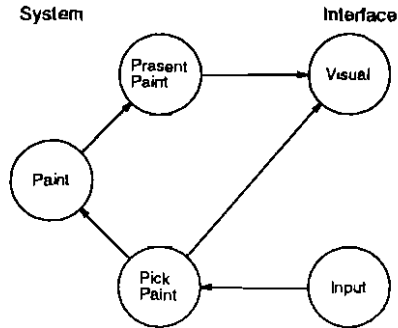


Figure 7.3: Agent diagram of the paint layer

presented graphically by Figure 7.3. The system description is contained in the agent *Paint* and it contains the result information for selection. There is only one image of interest and that is the whole painted image on some finite coordinate plane, called the canvas. Selection is done by dragging out some connected region in the canvas plane, which then selects that region for future operations. The definition of the agent *Paint* for the selection task only is given below.

**agent *Paint***

**types**

$canvas_{max}, canvasy_{max} : \mathbb{N}$   
 $Canvas == 0..canvas_{max} \times 0..canvasy_{max}$   
 $Bit ::= black \mid white$   
 $CanvasMap == Canvas \rightarrow Bit$

**attributes**

$painting : CanvasMap$   
 $pselected : \mathcal{P} Canvas$

**invariant**

$pselected \subseteq \text{dom } painting$

**operations**

$selectregion(region : \mathcal{P} Canvas)$   
**changes** ( $pselected$ )  
**post**  
 $pselected' = region$   
 $boundingbox(region : \mathcal{P} Canvas)$   
**changes** ()  
**pre**  
 $region = pselected$

**communication**

```

inputs paintin : selectregion(region : P Canvas)
outputs paintout : boundingbox(region : P Canvas)
external
   $\mu X \bullet ((\text{paintin}, \text{selectregion}(x)) \rightarrow \text{paintout}! \text{boundingbox}(y) \rightarrow X)$ 
endagent Paint

```

For the object layer, we have a very similar agent description as that for the paint layer, and this is given graphically in Figure 7.4. The system description is

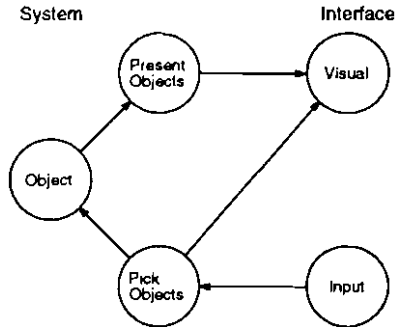


Figure 7.4: Agent diagram of the object layer

contained in the agent *Object*, and it is more complicated than the description of the agent *Paint*. It is necessary now to identify individual graphical objects and the structure implied by grouping of those objects into more complicated objects (we will ignore the further complication of levels associated to overlapping objects). We represent this structure by a straightforward hierarchy in which any object can be linked to at most one parent object. The selectable objects are only those which are not linked to a parent object. The currently selected objects are some subset of the selectable objects. We define only one selection operation, which given some subset of the selectable objects makes that set the currently selected objects. The description of *Object* corresponding to this result template of attributes is given below.

```

agent Object
types
  [ObjID]
attributes
  drawn : P ObjID
  grouping : ObjID  $\leftrightarrow$  ObjID
  objselectable : P ObjID
  objselected : P ObjID

```

```

invariant
  drawn = dom grouping  $\cup$  ran grouping
  objselectable = ran grouping - dom grouping
  objselected  $\subseteq$  objselectable
operations
  selectobjs(ids :  $\mathbf{P}$  ObjID)
    changes (objselected)
    post
      objselected' = region
  showselected(ids :  $\mathbf{P}$  ObjID)
    changes ()
    pre
      ids = objselected
communication
  inputs objectin : selectobjs(ids :  $\mathbf{P}$  ObjID)
  outputs paintout : showselected(ids :  $\mathbf{P}$  ObjID)
external
   $\mu X \bullet ((objectin, selectobjs(x)) \rightarrow paintout!showselected(y) \rightarrow X)$ 
endagent Object

```

The *Visual* agent contains a mapping from some screen coordinate space, *Pixel*, to bit values, which we will restrict to simply the value *black* or *white*. The pixel map can be further divided into regions which the user may identify as an independent image, so that the whole pixel map is considered as a collection of smaller, possibly overlapping, pixel maps. There is an indication of the current layer (object or paint) and the current tool in that layer, both of which give the mode in which input is interpreted. The mouse cursor is also located somewhere in the screen coordinate space. So, for the task of selection, we highlight information described above as a set of attributes describing the state space of the *Visual* agent.

```

agent Visual
internal
  types
    screenxmax, screenyumar :  $\mathbf{N}$ 
    Pixel == 0..screenxmar  $\times$  0..screenyumar
    Bit ::= black | white
    PixelMap == Pixel  $\rightarrow$  Bit
    PaintLayerIcon, ObjectLayerIcon : PixelMap
    PaintTool, ObjectTool :  $\mathbf{P}$  PixelMap
    selectpainttool : PaintTool
  attributes
    visible : PixelMap
    images :  $\mathbf{P}$  PixelMap
    vslected :  $\mathbf{P}$  PixelMap
    layer : PixelMap

```

```

    tool : PaintTool U ObjectTool
    mousecursor : PixelMap
    :
endagent Visual

```

Selection in both the paint and object layer are result predictable given the result templates we have chosen for *Paint* and *Object*. Restricting for the moment to the relationship between the agents *Paint*, *Visual*, and *Input* we can see that the attributes *visible* and *vselected*, which form a display template for the selection task, provide enough information for selection in the paint layer to be display predictable. They adequately portray the same information as the attributes *painting* and *pselected*, which form the result template. This assumes that we can adequately portray the mouse position information which identifies the region of points which will be sent as the argument for the *selectregion* message, and that is the function of the agent *PickPaint* in Figure 7.3.

For the object layer, it is obviously more difficult to ensure display predictability since the display template has no means of portraying the hierarchical structure in the result template (it is slightly less difficult to portray the overlapping ambiguity, but still not trouble free). This problem could be overcome, and display predictability satisfied, if there was a way to present the structural hierarchy in the agent *Visual*. Translating the region of points determined by the mouse into a set of object identifiers is performed by the agent *PickObjects* in Figure 7.4. For a user to be able to predict the outcome of selection, she must have access to as much information as *PickObjects* needs to perform that translation.

We can see that the paint layer is display predictable and the object layer is not. Since the ease with which complicated drawings can be manipulated is enhanced by the facilities given in the object layer, this lack of display predictability is not enough to abandon its use. However, the layering of paint and objects onto the same visual space makes matters even worse. In Figure 7.5 we present a graphical representation of the two layered *SnperPaint* agent description. There is no visible distinction between a painted image from the paint layer and a drawn image from the object layer. So not only does the user have to remember the grouping structure of the objects, she must also remember which images are objects and which ones are just paint. It may be possible to suggest a display strategy to circumvent this display unpredictability, but it seems unnecessary to make the distinction between the layers that has been forced by the system description. Therefore, our semi-formal analysis of this system with respect to display predictability of the selection operator has uncovered a bad design decision.

In Chapter 4, we distinguished between predictability and consistency, since the latter was a generalization of the former in the simple agent model. In the refined model, the informal definition of consistency that we gave—the same input

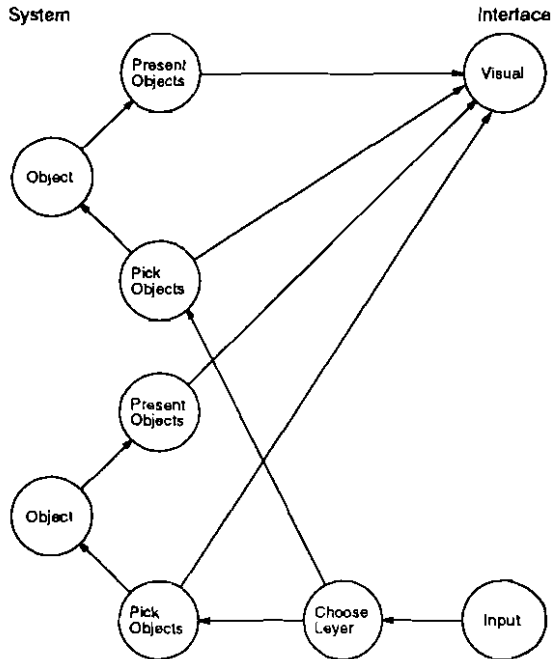


Figure 7.5: Agent diagram of the co-existing paint and object layers of SuperPaint

in similar situations has similar effects—is now very similar to predictability. Earlier, consistency was formalized in terms of a similarity criterion on the state set of an agent. Templates are a concrete way of formulating these similarity criteria. We noted earlier that definition of consistency did not capture task information. Similarity based on result and display templates does.

## 7.4 Synthesis

With the simple model of the agent, we stated that the observer of an agent only has access to the stimulus-response behaviour of the agent. Within the agent-based description of the *System* and *Interface*, this still applies. An agent can only participate synchronously in input or output events with another agent. Between the *User* and the computer, it is a bit artificial to name channels of communication



which serve as inputs to the user. Consequently, we have introduced display templates to represent the information from the computer which can be perceived by the user.

Synthesis describes the process by which the user determines how the effect of previous input on the result template for a particular task is observable via the display template for that task. The computer is honest with respect to a task if changes to the result template are immediately made apparent in the display template. If two histories are display equivalent, then they are also result equivalent. The computer is trustworthy if changes to the result template are eventually made apparent in the display template. Two histories must be display indistinguishable before they are considered result equivalent. Colloquially, honesty is paraphrased as the property:

*If what you see now (display template) is the same, then what you have now (result template) is the same.*

whereas a trustworthiness conforms to the weaker property:

*If all you can possibly see from now on is the same, then what you have now is the same.*

Formally, we would represent these task-centred properties as below.

$$\begin{array}{|l}
 \text{Honest, Trustworthy} : \mathcal{P}(\text{TaskID}) \\
 \text{task} \in \text{Honest} \Leftrightarrow ( \text{task} \in \text{dom Dtemplate} \\
 \quad \wedge (- \approx_{dt} -) \subseteq (- \approx_{rt} -) ) \\
 \text{task} \in \text{Trustworthy} \Leftrightarrow ( \text{task} \in \text{dom Dtemplate} \\
 \quad \wedge (- \equiv_{dt} -) \subseteq (- \approx_{rt} -) ) \\
 \text{where } dt = \text{Dtemplate}(\text{task}) \\
 \quad \wedge rt = \text{Rtemplate}(\text{task})
 \end{array}$$

We can give an example in the agent language which portrays the difference between honesty and trustworthiness, to show how honest interfaces can be more valuable. Our analysis will be conducted on a simplified version of a popular visual filesystem interface—the Macintosh desktop interface. We will expound on the ‘New Folder’ scenario briefly mentioned above and introduced by Roast and Wright [137]. Our description consists of a simple agent definition of a file system and a visual folder. The file system represents the *System* and the folder represents the *Interface*. Each description will only provide enough detail to discuss the scenario. We also will not make explicit the communication between the *System* and *Interface*, partly because that will be covered in the next chapter and partly because we need a definition of stable traces (also covered in the next chapter) to fully formalize the discussion.

The filesystem is a hierarchical arrangement of files and directories. Morgan and Sufrin [117] have provided a specification in  $Z$  of the Unix filing system, and our description is a simplified version of that. Files and directories are identified by elements in the set of all possible file identifiers, *FileID*. File identifiers are mapped to files in the system; we give no further description of files. The hierarchical arrangement is represented by two relations on file identifiers, one giving the unique *parent* of a file identifier and the other giving the set of *children* for a given file identifier. The distinction between files and directories is that directories can have children and files cannot. The hierarchy is acyclic and has a single *root* element, of which all other known files are descendants. At any time, one directory is *current*. We describe the filesystem as an agent below with only one operation, the creation of a new subdirectory under the current directory.

```

agent Filesys
internal
  types
    [FileID, File]
    root : FileID
    emptydir : File
  attributes
    known :  $\mathbf{P}$  FileID
    files : FileID  $\rightarrow$  File
    parent : FileID  $\leftrightarrow$  FileID
    child : FileID  $\leftrightarrow$  FileID
    currentdirectory : FileID
    currentfiles :  $\mathbf{P}$  FileID
  invariant
    known = dom files
    root  $\in$  known
    dom parent = known - {root}
    ran parent  $\subseteq$  known
    child = parent-1
    child*({root}) = known
    currentdirectory  $\in$  known
    currentfiles = child(currentdirectory)
  initially
    known = {root}
  operations
  :
  makesubdirectory(id : FileID)
  changes (known, files, parent, child, currentfiles)
  pre id  $\notin$  known
  post files' = files  $\oplus$  id  $\mapsto$  emptydir
     currentfiles' = currentfiles  $\cup$  {id}

```

```

:
communication
  inputs filesysin : makesubdirectory(id : FileID),...
  outputs ...
  r : ...
external
  ...
endagent Filesys

```

The description of the *Interface* is given (partially) by the Desktop agent. The desktop contains folders, one of which is current. Each folder contains a set of icons, a subset of which are visible at any time (we only describe the visible icons for the current folder in the description below for simplicity). The current folder also displays the number of icons it contains. The only operation we define on the desktop agent is one to create a new folder icon in the current folder.

```

agent VisualInterface
internal
  types
    [FolderID, Icon]
    desktop : FolderID
    emptyfolder : Icon
  attributes
    folders :  $\mathbf{P}$  FolderID
    contents : FolderID  $\leftrightarrow$  Icon
    current : FolderID
    visible :  $\mathbf{P}$  Icon
  invariant
    folders = dom contents
    current  $\in$  dom contents
    desktop  $\in$  folders
    visible  $\subseteq$  contents(current)
  initially
    folders = { desktop }
  operations
  :
  newfolder(fid : FolderID)
  changes (folders, contents, visible)
  pre fid  $\notin$  folders
  post folders' = folders  $\cup$  { fid }
    contents' = contents  $\oplus$ 
      current  $\mapsto$  (contents(current)  $\cup$  { emptyfolder })
    visible  $\subseteq$  visible'
  :

```

```

communication
  inputs folderin : newfolder(fid : FolderID),...
  outputs ...
   $\tau$  : ...
external
  ...
endagent VisualInterface

```

For the task of creating a new folder, the user wants to know if the current directory contains a new subdirectory, so the result template for this task is only the attribute *currentfiles*. When a new folder is created, this attribute value changes, increasing by one as the result of the inclusion of the (fresh) file identifier *id*. The display template corresponding to that result template is given by the attribute *visible*, indicating a subset of the icons which are contained in the current folder. However, the change due to the *newfolder* message does not require that the new icon be in the visible set. Macintosh users will be familiar with situations in which the new folder's icon does not appear in the immediately visible set of icons for the current folder. The system is not honest in this situation, requiring the user to browse through the current folder to observe that there is a new icon for the newly created folder. Hence, trustworthiness is all that can be claimed of this visual interface.

And we might add that we would expect all tasks to be trustworthy, so the visual interface is not gaining us much with this task over any other interface to a hierarchical file system. It might even be said to be worse than a command-based interface because the user of the visual interface is led to believe that if they don't see something change, then it hasn't changed. The visual interface is sly—it leads us to believe it is honest when it is not.

In defence of the Macintosh, there is a way to salvage honesty for this task. The visual interface provides another possible display template in the form of an *items* attribute which gives the count of the number of icons in the current folder. This display template honestly reflects the changes in the *currentfiles* result template.<sup>2</sup> We give the revised agent description of the visual interface below.

```

agent VisualInterface
internal
  types
    [FolderID, Icon]
    desktop : FolderID
    emptyfolder : Icon
  attributes
    folders :  $\mathbb{P}$  FolderID

```

<sup>2</sup>Well, at least it tries to do so. It is possible to manipulate the window of the current folder in such a way that even this honest display template is obscured!

```

    contents : FolderID  $\rightarrow$  Icon
    current : FolderID
    visible : P Icon
    items : N
invariant
    folders = dom contents
    current  $\in$  dom contents
    desktop  $\in$  folders
    visible  $\subseteq$  contents(current)
    items = #(contents(current))
initially
    folders = {desktop}
operations
    :
    newfolder(fid : FolderID)
    changes (folders, contents, visible, items)
    pre fid  $\notin$  folders
    post folders' = folders  $\cup$  {fid}
        contents' = contents  $\oplus$ 
            current  $\mapsto$  (contents(current)  $\cup$  {emptyfolder})
        visible'  $\subseteq$  visible
    :
communication
    inputs folderin : newfolder(fid : FolderID), ...
    outputs ...
     $\tau$  : ...
external
    ...
endagent VisualInterface

```

As Roast and Wright point out, even though this attribute maintains the honesty property, there is no guarantee that it is observed by the user, i.e., the *items* may not be used by the Macintosh user as a display template when creating new folder.

## WYSIWYG

When the task is result predictable and honest, we satisfy the property:

*What you can see now determines all you will be able to get.*

This sounds very much like the popular slogan "What you see is what you get" (WYSIWYG). A weaker, and consequently more realistic, property results from result predictability and trustworthiness and satisfies:

*All you can possibly see determines all you will be able to get.*

These two versions of WYSIWYG are summarized below.

$$\begin{array}{l}
 \text{WYSIWYG}_{strong}, \text{WYSIWYG}_{weak} : \mathbf{P}(\text{TaskID}) \\
 \text{task} \in \text{WYSIWYG}_{strong} \Leftrightarrow ( \text{task} \in \text{dom } D\text{template} \\
 \quad \wedge (- \approx_{dt} -) \subseteq (- \equiv_{rt} -)) \\
 \text{task} \in \text{WYSIWYG}_{weak} \Leftrightarrow ( \text{task} \in \text{dom } D\text{template} \\
 \quad \wedge (- \equiv_{dt} -) \subseteq (- \equiv_{rt} -)) \\
 \text{where } dt = D\text{template}(\text{task}) \\
 \quad \wedge rt = R\text{template}(\text{task})
 \end{array}$$

That these properties formalize the popular WYSIWYG slogan is slightly misleading because they only state that some desirable relationship holds between the result and display templates. Though this is a necessary condition for the interaction taking advantage of WYSIWYG, it is not sufficient since we do not guarantee that the relationship is understood by the user. But we are even more pessimistic than that. The formulation of this property does not even guarantee that the designer is aware of the relationship, because it does not have to be explicit in the agent description. In the next chapter, we will discuss localized correspondence as a way to make explicit the relationship between result and display templates.

## 7.5 Result initiated interaction

The properties of the last section were all display initiated, that is, equivalence or indistinguishability of display templates had implications toward equivalence or indistinguishability of result templates. Properties that are result initiated reverse the situation, so result information has implications toward display information. We do not have names for these properties. The first property satisfies:

*What you can have (or know) now determines what you can see now.*

and is characterized by the following implication, assuming result and display templates  $rt$  and  $dt$  associated to the same task.

$$(- \approx_{rt} -) \subseteq (- \approx_{dt} -)$$

If the task is display predictable, we have the stronger implication satisfying:

*What you can have (or know) now determines all you will be able to see.*

and characterized by

$$(- \approx_{rt} -) \subseteq (- \equiv_{dt} -)$$

The weaker versions of the first property satisfies:

*All you can possibly have (or know) from now on determines what you can see now.*

and is characterized by the following implication.

$$(- \equiv_{rt} -) \subseteq (- \approx_{dt} -)$$

When coupled with display predictability, we have the property satisfying:

*All you can possibly have (or know) from now on determines all you will be able to see from now on.*

and is characterized by the following implication, assuming result and display templates  $rt$  and  $dt$  associated to the same task.<sup>3</sup>

$$(- \equiv_{rt} -) \subseteq (- \equiv_{dt} -)$$

Since we have assumed that result templates bear significance to the task or knowledge structure of the user, it is possible that these last four properties would be considered more relevant within the user modelling domain.

## 7.6 Conclusions

In this chapter, we have investigated how the increased structure of the refined agent model allows the formulation of properties over collections of agents which are relevant to the internal goals of interaction—the results—and the tangible evidence of interaction—the displays. We have introduced templates as a task-centred means of restricting the knowledge of an agent's internal state. This provides us with the means of more clearly defining the scope of properties, confining them to the interactional unit of the task, instead of over the computational unit of the agent, as was done in Chapter 4. Table 7.1 summarizes the task-dependent properties defined in this chapter.

Templates extend the bridge between psychology and computer science, by forcing an iterative relationship between the psychologist and the interactive system designer. The HCI specialist provides the task analytic information, highlighting the tasks and the psychological attributes (or semantic features) related to each

<sup>3</sup>Sufirin and He classify a version of this last property as "goal determines view"

Result-Display relationship	Interactive Property
$(- \approx_{dt} -) \subseteq (- \equiv_{dt} -)$	display predictability
$(- \approx_{rt} -) \subseteq (- \equiv_{rt} -)$	result predictability
$(- \approx_{dt} -) \subseteq (- \approx_{rt} -)$	honesty
$(- \equiv_{dt} -) \subseteq (- \approx_{rt} -)$	trustworthiness
$(- \approx_{dt} -) \subseteq (- \equiv_{rt} -)$	WYSIWYG (weak)
$(- \equiv_{dt} -) \subseteq (- \equiv_{rt} -)$	WYSIWYG (strong)
$(- \approx_{rt} -) \subseteq (- \approx_{dt} -)$	result initiated interaction
$(- \approx_{rt} -) \subseteq (- \equiv_{dt} -)$	result initiated interaction
$(- \equiv_{rt} -) \subseteq (- \approx_{dt} -)$	result initiated interaction
$(- \equiv_{rt} -) \subseteq (- \equiv_{dt} -)$	result initiated interaction

Table 7.1: Task dependent interactive properties expressed as result-display relationships

task. The designer uses this information to formulate the result templates for the system. An agent description of the *System* can begin from that point. In designing the *Interface*, appropriate display templates must be determined for each result template in such a way as to maximize satisfaction of properties described above. If no such display templates can be found to satisfy even the weakest properties above, it may be the case that the task analysis was faulty, or incomplete, in which case some of the process can be iterated. Once some set of satisfactory display templates is settled, the agent description of the *Interface* the process is not complete as the display/result relationships must be subjected to tests to see if they make realistic demands on the user.

We realize that the relationship between display and result implied by the properties of this chapter must be made explicit in a design situation. In the next chapter, we will investigate how agents can be used to capture more explicitly the relationship between display and result.



# Chapter 8

## Interactive system architectures

Our purpose in applying formal methods to the analysis and design of interactive systems is to provide a means of expressing user recognizable structures in a language more closely related to the design practice. Abstract formalisms, such as the PIE model, and more concrete formalisms, such as the Sufrin and He model and the agent model of this thesis, have opened up the possibility of incorporating valid psychological assumptions of usability as design principles in a more rigorous software engineering environment. The intention is that a principled design process will result in the ability to engineer more usable systems. Currently, the majority of psychological knowledge concerning the usability of interactive systems has crept into design by means of common sense heuristics. We aim to show in this chapter how the agent formalism can capture the meaning behind some heuristics of interactive system architectures.

### Overview of chapter

In Section 8.1, we discuss the progression from abstract to concrete in interactive system design and how principles or properties that apply at the abstract level are mapped down to a concrete architecture. In Section 8.2, we describe two multiagent architectures for interactive system design, showing how they can be viewed in terms of the interaction framework. In Section 8.3, we discuss how the agent model allows a formal explanation and comparison of all levels from abstract to concrete, but especially at the concrete architectural level where few formalisms have been brought to bear.

In Section 8.4, we provide an example of how the agent model can formalize heuristic properties of a multiagent architecture, specifically one role of the control component of a PAC agent. To do this, we concentrate on the abstract property of correspondence introduced in Section 4.3. We give a more constructive definition of correspondence with the aid of template information. Correspondence provides an architectural constraint on the relationship between collections of agents in the

### *Interface and System.*

In Section 8.5, we investigate another architectural consideration in designing an interactive system. We use a semi-formal approach to justify how predictability and synthesis can be supported by the analysis of user recognizable structures in a graphical interface. The argument which we support in this extended example is that the agent architecture of the *System* specification should lead to a natural description of an *Interface* which will support the user's understanding of the system functionality. The particular case we examine concerns the possible arrangement of commands in the graphical interface of a multiple file text editor. It does not matter that our example is yet another specification of a text editor, because the scenario we describe concerns more the support of user recognizable structures in an interface than the text editing domain.

## 8.1 From abstract principles to concrete architectures

Thimbleby [160] introduced generative user-engineering principles, called *gueps*, as a means of capturing colloquially understood features of usability within the design process. A guep must satisfy the following four criteria.

- it can be expressed formally;
- it has a colloquial form that is accessible to users;
- it embodies valid psychological guidelines; and
- it is constructive, so that it indicates how the principle can be attained in design.

It is easiest to satisfy the first three criteria at an abstract level in which irrelevant clutter can be removed and concentration can be focussed on the essential features of the interaction. To address the fourth criterion, concrete architectural detail is needed. We identify three levels in the progression from abstract to concrete.

### The conceptual level

The first and most abstract level identifies the context of an interactive system with a gross separation between the human and the computer. In this thesis, the interaction framework provides this conceptual and contextual description, and it has proved beneficial for satisfying the second and third criteria above. Measures applied to the translations within the framework gauge whether a component (*System*, *User*, *Input* or *Output*) can access the behaviour of another component and how easy it is to access that behaviour.

### The layers level

At the second level, a finer grain of separation is introduced for both the human and the computer side of the interaction. At this level we introduce models such as Norman's seven stage model of the user activity and the Seeheim model of user interface management systems, both discussed in Chapter 3. Each of these models presents separation in terms of layers or stages of action. Norman's model certainly embodies valid psychological guidelines, since it was formed out of a psychological model of the user. The Seeheim model advocates the separation of presentation from application, which goes part of the way toward addressing the constructive criterion by isolating those features relevant to usability from those features relevant to functionality.

### The architectural level

At the third and most concrete level, architectural models are introduced to directly address the construction of a system. In this thesis, we are interested in multiagent architectures, broadly defined as the class of architectures which advocates the identification of simple, independent computational units—agents—coordinated in order to produce complex behaviour. Multiagent architectures have several useful constructive features, as summarized by Coutaz [41]. They support iterative design since an agent defines the unit of modularity and can be altered without greatly affecting the behaviour of the rest of the agents in the system. They support distributed or parallel implementations. Considering the user as an agent which can communicate with system agents means that multithreaded behaviour, in which the user is participating in more than one task concurrently, can be more naturally expressed by assigning different agents to the different tasks of the user. In the next section, we describe two different multiagent models.

## 8.2 Multiagent models

### Model-View-Controller

The Model-View-Controller (*MVC*) paradigm is the model used for the design of most Smalltalk interactive programs [107, 29, 96, 66]. On the surface, the MVC model fits very nicely with the interaction framework, as shown in Figure 8.1. An MVC triad represents three of the four components of the framework. The model is the object on which work in the application domain is to be performed, similar to the *System* in our framework. The controller provides the input interface to the user, similar to the *Input*. The view provides the output interface, similar to the *Output*. The user, though not explicit in MVC, is assumed to communicate directly with the controller and makes observations of the view.




Figure 8.1: The Model-View-Controller paradigm of Smalltalk

At the abstract level of the interaction framework, we were not concerned with any further details on the decomposition of components. MVC, on the other hand, is implemented in an object-oriented language, thus providing a natural mechanism for a hierarchical arrangement of MVC triads. Figure 8.2 shows a typical hierarchical arrangement of a collection of windowed interactive applications. Another distinction between the interaction framework and MVC can be seen by the explicit communication link between controller and view; there is no such explicit connection between *Input* and *Output* in the framework. This is another difference that arises from the different levels of abstraction they are intended to address.

### Presentation Abstraction Control

Coutaz has suggested an alternative to the MVC paradigm, called the Presentation Abstraction Control—or *PAC*—model, shown in Figure 8.3 [38, 39, 40]. The abstraction component corresponds to the model of MVC and the *System* of the framework. The presentation component combines both view and controller of MVC, similar to the *Interface* of the framework. In MVC, the view and controller communicate directly with the model, and vice versa. In PAC, this communication is coordinated by the control component to ensure that the state of the presentation component faithfully portrays the state of the abstraction component. The control component also coordinates communication between separate PAC agents hierarchically arranged.

PAC is intended as a high-level design notation for the description of an interface. It does not have an underlying implementation as MVC does. However, PAC is still intended to express a hierarchical relationship between a collection of interface objects. Figure 8.4 gives the relationship between a PAC hierarchy and the other components of an interactive system. The PAC hierarchy can be seen as the dialogue and presentation control.

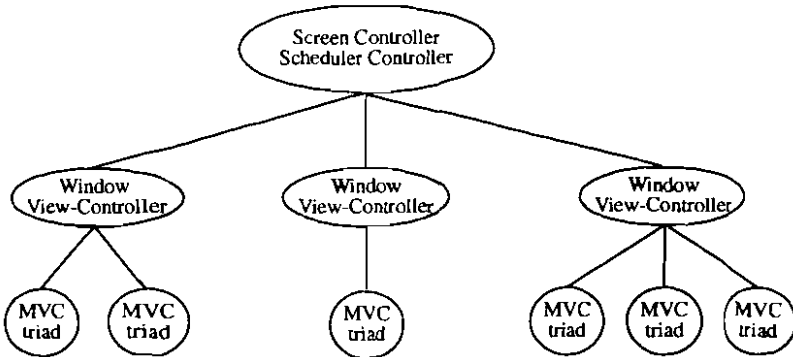


Figure 8.2: The MVC hierarchy

### 8.3 Applying formal methods to levels of abstraction

Having described the three levels progressing from abstract to concrete descriptions of an interactive system, we need to show how properties which arise at the most abstract level are traced down to the concrete, architectural level. We have adopted the agent as the formal means for propagating interactive properties. At each level described above, the inspiration for the properties to be formalized comes from common sense heuristics. Formality, therefore, does not replace intuition; the two are complementary.

At the most abstract level, there are computational formal models, such as the PIE model in which to express many gueps of interest. The agent model can express the same properties as the PIE model by looking at its trace-state behaviour. An example of an abstract psychological model is Norman's execution and evaluation model, which was the primary motivation for the interaction framework and the simple agent model.

At the next level, where there is more separation between components, the formal models need to express relationships between the layers or stages. The state-

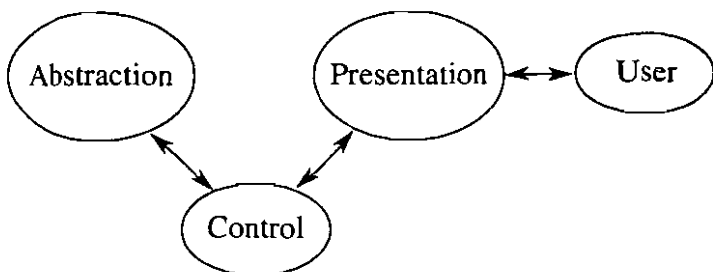


Figure 8.3: The PAC model

display model for the system and CCT for the user are examples of such layered formal models. There is also a need to relate properties which are expressed on the layered models to the abstract properties given by the PIE model, the agent model and the interaction framework. Agents can represent the layers or stages—for example, the state-display model becomes the *System-Interface* model—and so properties expressed as constraints between agents apply to this level.

At the architectural level, there are no existing examples of a formal model which captures interactive properties. The Forest project has produced an agent-based description of the MVC model within a modal action logic (MAL) [66, 141], but the purpose of that formal model was not to investigate the interactive properties of MVC. This absence of formalism at the constructive stage of interactive system design was one of the prime motivations for the refined agent model. One of the problems with the multiagent models described above is that while they do address very strong heuristic notions of usability—for example, catering for multithreaded dialogue between user and system so that the user is freed to direct the interaction—the lack of a formal model does not allow the description of desirable properties such as predictability, consistency, honesty, etc. that were valuable at the more abstract levels. It is also difficult to compare different multiagent architectures when they are expressed informally. The refined agent model allows us to formalize features of a multiagent architecture, for a more honest comparative appraisal.

The first step formalizing a multiagent architecture is to reject the fundamental units of both MVC and PAC. MVC suggests that the model-view-controller triad is the basis for the fundamental unit of interactive system design. PAC suggests the same for the PAC agent. In practice, however, this informal rule is continually broken, to the extent that the separations implied by each model are too artificial to be useful. View-controller pairs without an associated model are common in Smalltalk interactive systems in order to affect flow of control (for example, to express the role of a parent window which administers the change of control between




Figure 8.4: The PAC hierarchy within an interactive system

subwindows [66]). Similarly, in examples of PAC diagrams, it is not uncommon to see PAC agents without a presentation or an abstraction component, or with multiple presentation components. More recent models of PAC by Nigay [121] allow communication between two separate application components to bypass the control components. While we may be surprised by such glaring inconsistencies, it is not surprising that they arise. The PAC and MVC architectures are intended as a guides for designers; they are not intended to be strictly enforced as would be demanded by a formal model.

The agent, as defined in this thesis, is a well-defined building block for interactive system design, unlike the MVC triad or the PAC compound agent.<sup>1</sup> We can now view the MVC and PAC models for what they are—heuristic guidelines for the construction of agents into interactive objects. The MVC triad and the PAC compound agent can be realized as composed agents which can address one or more of the user's tasks. The model or abstraction represents result information. The view-controller pair or the presentation represents display information.

What was not discussed in Chapter 7 was how the relationship between the

<sup>1</sup>We use the term compound agent now to distinguish the PAC unit from our formal agent model.

model and its view or the abstraction and its presentation is to be maintained. Herein lies a major difference between the MVC and PAC models. Correspondence between model (respectively, abstraction) and view (respectively, presentation) must be maintained in MVC (respectively, PAC) in order to support the user's understanding of the result of interaction based on observations of the display. In PAC, the control component makes the correspondence explicit and separate from either presentation or abstraction. In MVC, the correspondence must be maintained either within the model or the view, or both, meaning that model and view are not as separable as may be desired.

But how does the control component of a PAC agent maintain the correspondence? In the following section, we will develop the formal notion of correspondence introduced in Chapter 4 in an attempt to formalize the role of the control.

## 8.4 Local correspondence

Recall in the definition of correspondence from Section 4.3 that two agents are said to correspond if the operations associated with input for one is sufficiently mirrored in the other according to some *retrieve* relation between their state sets. We need to slightly revise the definition of correspondence to relate to the refined agent model. In order for agent  $B$  to correspond to agent  $A$ , we only require that retrieve relation hold for *stable* traces of the two. A stable trace is one which can only be extended by a stimulus event, that is, the agent can only proceed by participating in some input event.

$$\left| \begin{array}{l}
 \text{stable} : \text{Agent} \rightarrow \mathbf{P}(\text{seq Event}) \\
 \hline
 t \in \text{stable}(A) \Leftrightarrow ( t \in \mathcal{H}[A] \\
 \quad \wedge \forall t' : \mathcal{H}[A] \\
 \quad \quad | \quad t - t' \wedge t \neq t' \\
 \quad \quad \bullet \quad t'(\#t + 1) \in \text{stimuli}(A)
 \end{array} \right.$$

The definition for correspondence for the refined agent model is very similar to the one given in Section 4.3 with the addition of stability conditions.



$  \begin{array}{l}  \text{AgentCorr} \\  \text{Agent}^I \\  \text{Agent}^J \\  \text{retrieve} : \text{states}^J \leftrightarrow \text{states}^I \\  \hline  \text{inits}^J \times \text{inits}^I \subseteq \text{retrieve} \\  \forall (tr_I, s_I), (tr_I \frown tr_I', s_I') : I_{\theta \text{Agent}^I}^{\text{int}} \\  \quad   \quad (s_I, s_I') \in (tr_I' \text{ ; } \text{mess} \text{ ; } \text{operations}^I) \\  \quad \quad \wedge \{tr_I, tr_I \frown tr_I'\} \subseteq \text{stable}(\theta \text{Agent}^I) \\  \bullet \exists (tr_J, s_J), (tr_J \frown tr_J', s_J') : I_{\theta \text{Agent}^J}^{\text{int}} \\  \quad   \quad (s_J, s_J') \in (tr_J' \text{ ; } \text{mess} \text{ ; } \text{operations}^J) \\  \quad \quad \wedge \{tr_J, tr_J \frown tr_J'\} \subseteq \text{stable}(\theta \text{Agent}^J) \\  \bullet \{(s_I, s_J), (s_I', s_J')\} \subseteq \text{retrieve}  \end{array}  $
---

We define  $\text{correspond}_-$  as a relation between agents indexed by a retrieving relation between their state sets. The pair of agents  $(D, R)$  is in  $\text{correspond}_{ret}$  if agent  $R$  corresponds to agent  $D$  as defined by the schema  $\text{AgentCorr}$  with retrieving relation  $ret$ .

$  \begin{array}{l}  \text{correspond}_- : (\text{State} \leftrightarrow \text{State}) \rightarrow (\text{Agent} \leftrightarrow \text{Agent}) \\  \hline  \text{correspond}_{ret} = \{ \text{AgentCorr} \\  \quad   \quad \text{retrievc} = ret \\  \quad \bullet \theta \text{Agent}^I \mapsto \theta \text{Agent}^J \}  \end{array}  $
---

Within the interaction framework, we want the *User's* interaction with the *Interface* to be so mirrored by the *System*. Overall, we would want to satisfy the following constraint.

$$\begin{array}{l}
 \exists \text{AgentCorr} \\
 \bullet ( \theta \text{Agent}^I = \theta \text{Interface} \\
 \quad \wedge \theta \text{Agent}^J = \theta \text{System} )
 \end{array}$$

This means that we have to find the *retrieve* relation to interpret information of the *System* within the *Interface*. Motivated by the multiagent architectures, it is desirable to localize the correspondence between *Interface* and *System*, both to make it easier to satisfy the demands of correspondence and to reflect the relationship between interactive objects as seen by the user and their counterparts within the *System*. The display and result templates introduced in Chapter 7 are useful for such localization.

[TaskID]

$$\begin{array}{l}
 Rtemplate : TaskID \rightarrow Template \\
 Dtemplate : TaskID \rightarrow Template \\
 \hline
 \text{dom } Rtemplate = \text{dom } Dtemplate \\
 Rtemplate(t) \subseteq System.attrs \\
 Dtemplate(t) \subseteq Interface.attrs
 \end{array}$$

This task information forms the basis for the local correspondence between *System* and *Interface*. For every task, we require a retrieve relation between the values of the result attributes and the display attributes for that task.

$$\begin{array}{l}
 \underline{taskretrieve} : TaskID \rightarrow (State \leftrightarrow State) \\
 \hline
 \text{dom } taskretrieve = \text{dom } Rtemplate \\
 \text{dom}(taskretrieve(t)) \subseteq \{ s : System.states \bullet Rtemplate(t) \triangleleft s \} \\
 \text{ran}(taskretrieve(t)) \subseteq \{ s : Interface.states \bullet Dtemplate(t) \triangleleft s \}
 \end{array}$$

Given the result and display templates for the interactive system, we can then require that for any task identifier  $t$ , the *System* restricted by the result template for  $t$  corresponds to the *Interface* restricted to the display template for  $t$ . Below we give a formal representation of this local correspondence.

$$\begin{array}{l}
 \underline{LocalCorrespondence} \\
 \hline
 System, Input, Output, Interface : Agent \\
 tasks : P TaskID \\
 Rtemplate : tasks \rightarrow Template \\
 Dtemplate : tasks \rightarrow Template \\
 taskretrieve : tasks \rightarrow (State \leftrightarrow State) \\
 \hline
 Interface = compose_{sync}(Input, Output) \\
 Rtemplate(t) \subseteq System.attrs \\
 Dtemplate(t) \subseteq Interface.attrs \\
 \text{dom}(taskretrieve(t)) \subseteq \{ s : System.states \bullet Rtemplate(t) \triangleleft s \} \\
 \text{ran}(taskretrieve(t)) \subseteq \{ s : Interface.states \bullet Dtemplate(t) \triangleleft s \} \\
 \forall t : tasks; rt : Rtemplate(t); dt : Dtemplate(t) \\
 \bullet (System \Downarrow rt, Interface \Downarrow dt) \in correspond_{taskretrieve(t)}
 \end{array}$$

The PAC compound agent is related to a set of tasks for the user. Given a task, the description of the abstraction agent supports that task in terms of the attributes in the result template. Related attributes for the presentation agent are

then chosen and their values are constrained by a *retrieve* relation. The behaviour of the control agent is in part specified by the local correspondence relationship that it must maintain for each task the PAC compound agent is meant to support. We say in part because the control component also coordinates communication between PAC multiagents and not just between its presentation and abstraction agents.

The *System* is derived from a PAC description as the synchronized composition of all of the abstraction agents. The *Interface* is derived as the synchronized composition of all of the presentation agents. *Perform* and *Present*, the agent manifestations of the *performance* and *presentation* translations are derived from the synchronized composition of all of the control agents. We can further divide the control agent into a performance subagent controlling dialogue flow from presentation to abstraction, a presentation subagent, controlling dialogue flow from abstraction to presentation, and a hierarchical subagent, controlling dialogue flow between other control agents.

In the next section, we will concentrate on an example showing how local correspondence can be used to suggest a natural graphical interface to a *System* description.

## 8.5 Assessing the graphical interface to a text editor

In this section, we use the agent language to motivate a semi-formal support for predictability and synthesis. The detail of strict formality sometimes clouds the insight which it can support, and so seek in this example to demonstrate how the formal properties of predictability and synthesis can be used to formulate heuristic guidelines for architectural constraints for agents.

From the description of the interaction framework, we know that the presentation of the interface has great impact on the effectiveness of the interaction. Given an agent description of the *System*, we can provide a realistic example of how to increase the effectiveness of the graphical interface. The example we use extends from work done on the ESPRIT Basic Research Action project 3066 (AMODEUS) [13, 118] involving scenarios of interesting user behaviour within an interactive system. Within the project, scenarios are used to compare and contrast the effectiveness of system and user modellers in explaining interactive phenomena. The scenario we will describe concerns the grouping of commands in two versions of the graphical interface to a multibuffer text editor. The versions of the editor are Spy (version 9) and its direct descendant Ten (presumably short for Spy version 10). Both were developed at the Rutherford Appleton Laboratories in the United Kingdom.

In this example, the snippet of action involves the invocation of operations

performed on selected text. The advantage of a multiple file, or multibuffer, text editor is that the user can rapidly switch attention between different files. Changes to the text of one file which involve text from another file—for example, the copying of some text from one file to a specified place in the other—are readily provided for in a multibuffer editor. In this example, the user may have many buffers active at once, each containing text from some file in the filesystem, but only one of those buffers is the current buffer—the one to which active editing commands such as insertion of a character are directed. Also, the user may select a contiguous region of text within the current buffer. This selected text may be deleted from the buffer, or it may be copied or moved to any of the open buffers in the editor.

The multibuffer editors on which this scenario is based have many other features, most of which are deserving of further analysis, but for the purposes of this scenario, the functionality of the operations to be performed on the selected text are all that concern us. The challenge posed to HCI modellers is to assess two options for the arrangement in the graphical interface of the operations for deletion, copying and moving of selected text.

The first option is represented in the editor *Spy* (version 9), and is shown in Figure 8.5. In this graphical interface, the operations to copy and move are grouped together on a pop-up menu and the delete operation is separated and appears on the menu bar associated to each buffer. The second option is represented in the editor *Ten*, and is shown in Figure 8.6. In this graphical interface, all of the operations—delete, copy and move—are grouped together on the same pop-up menu.

We begin with an agent description of the *System* component of the multibuffer editor. The specification of the functionality of the multibuffer editors represented by *Spy* and *Ten* is a reverse-engineered description. A separate specification of *Spy* has been presented by Martins exclusively in the *Z* notation [106]. We hope to demonstrate that the agent language provides an equally powerful specification technique with the advantage that it more closely reflects the object-oriented view crucial to the relationship between the *System* and *Interface*.

After presenting a specification of the *System* component of the multibuffer editor, we will move on to discuss the graphical interface. Our assumption again is that both *Spy* and *Ten* share a common *System* description—at least as far as this scenario is concerned—but differ in the arrangement of the graphical interface. The *System* description leads us to the suggestion of a natural graphical presentation of the *System* which we could also be described in terms of agents. This natural graphical interface can then be compared to the actual interfaces for both *Spy* and *Ten*.

We use the agent language to give a limited description of the functionality of the multibuffer editor. Our agent description is limited because we will only define that part of the editor that concerns the scenario described above, i.e., we will only describe enough to facilitate specification of the operations of deletion, copying and

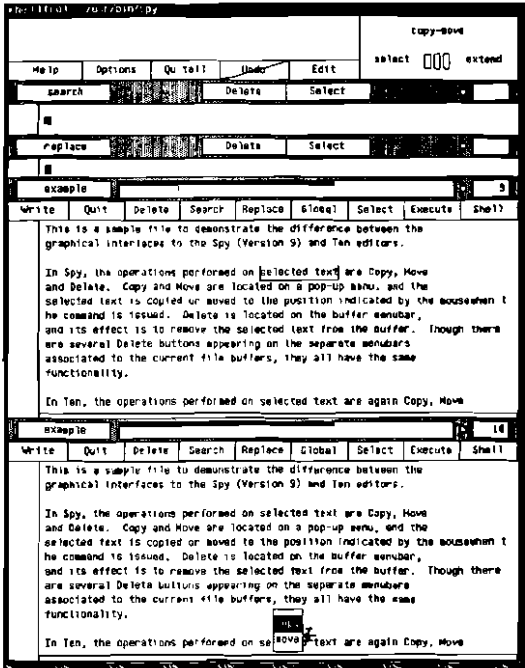


Figure 8.5: The graphical interface of Spy

moving of selected text.

Figure 8.7 is a simple graphical representation of an agent architecture for the multibuffer editor functional core. In Figure 8.7, we represent the *Interface* as a single agent, and the relationship between interface and functional core is represented by the *Perform* and *Present* agents.

A single buffer will contain text and we will describe operations that manipulate text both within and between buffers. Within a buffer, the usual insertion and deletion of single characters is allowed at an insertion point which can be set to any point within the text of the buffer. A section of text within a buffer can be selected for subsequent deletion or it can be copied or moved to any point within the text of any buffer. There is only one selected section of text in the whole multibuffer editor at any one time.

We will now give descriptions for the *System* agents in Figure 8.7.

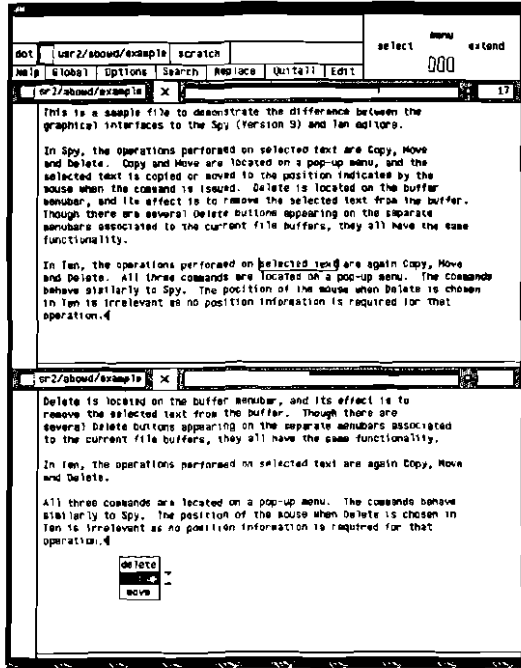


Figure 8.6: The graphical interface of Ten

### 8.5.1 Single buffer

Figure 8.8 is a graphical representation of the single buffer agent. We begin the description of the single buffer by describing its state and how it is initialized. For the purposes of this description, the only attributes of interest involve the main text and the selected text within the main text. We are making the explicit assumptions about the the result template for these editing tasks, and are defusing the *System* as restricted by this result template. Text is a sequence of characters, chosen from a set of all possible characters ( $CH$ ). The selected text, if it exists, is delimited by two natural number indices within the main text. If there is no selected text, i.e., it is of value *null*, then these indices also have the value *null*. A non-*null* insertion point indicates where subsequent characters are inserted into the text. The insertion point is an index within the text that splits the text into two subsequences—the text before the insertion point and the text after the insertion point. Initially, the contents of the buffer is empty. The agent description for this core of the buffer is




Figure 8.7: Agent representation of multibuffer editor




Figure 8.8: The single buffer agent

given below.

```

agent bufferstate
internal
types
  [CH]
  ■ ∉ CH
  null ∉ CH
  Text == seq(CH ∪ ■)
  ∀ t : Text • ( last(t) = ■
                ∧ ■ ∉ ran(front(t)) )
attributes
  text : Text
  before, selected : seq CH ∪ {null}
  insertion, begsel, endsel : ℕ ∪ {null}
invariant
  
```

$$\begin{aligned}
 \text{insertion} = \text{null} &\Rightarrow ( \text{begsel} = \text{null} \\
 &\quad \wedge \text{endsel} = \text{null} \\
 &\quad \wedge \text{selected} = \text{null} \\
 &\quad \wedge \text{before} = \text{null} ) \\
 \text{insertion} \neq \text{null} &\Rightarrow ( \text{insertion} \in \text{dom text} \\
 &\quad \wedge \text{before} = 1..(\text{insertion} - 1) \triangleleft \text{text} \\
 &\quad \wedge \text{begsel} = \text{insertion} \\
 &\quad \wedge \text{endsel} \geq \text{begsel} \\
 &\quad \wedge \text{selected} = \text{begsel}.. \text{endsel} - 1 \triangleleft^{\text{seq}} \text{text} )
 \end{aligned}$$

**initially**  
 $\text{text} = \{ \}$   
 $\text{insertion} = \text{null}$

**operations**  
**communication**  
**external**  
**endagent** *bufferstate*

Text in the buffer is terminated by the special end of text symbol ■, which is not in the set *CH*. The state invariant embodies the link between the insertion point and the selected text. If there is no insertion point for a buffer, then there is no selected text. In addition, we constrain the selected text to occur as a contiguous sequence of text positioned directly after the insertion point.

We can now specify the operations which deal with selection of text. These operations will be collectively specified in the agent *bufferSelect*. The first two operations are marking the beginning of the selection of text and extending the selection of text. These are intended to be initiated by the *User*, via the *Interface*. They are specified in the agent *bufferMarkSelect*. After either a mark or extend event, the buffer manager is informed of the new selected text.

**agent** *bufferMarkSelect*  
**interleaves** (*bufferstate*)  
**with**  
**internal**  
**operations**  
 $\text{beginmark}(n : \mathbb{N})$   
**changes** (*before*, *selected*, *insertion*, *begsel*, *endsel*)  
**pre**  
 $n \in \text{dom text}$   
**post**  
 $\text{insertion}' = n$   
 $\text{begsel}' = \text{endsel}'$   
 $\text{extendmark}(n : \mathbb{N})$   
**changes** (*before*, *selected*, *insertion*, *begsel*, *endsel*)  
**pre**



```

     $n \in \text{dom } \textit{text}$ 
     $\textit{insertion} \neq \textit{null}$ 
post
     $n \leq \textit{begsel} \Rightarrow ( \textit{begsel}' = n$ 
                         $\wedge \textit{endsel}' = \textit{endsel} )$ 
     $n > \textit{begsel} \Rightarrow ( \textit{begsel}' = \textit{begsel}$ 
                         $\wedge \textit{endsel}' = n )$ 
selecttext( $t : \textit{seq } CH$ )
changes ()
pre
     $t = \textit{selected}$ 
communication
inputs
     $\textit{bufferin} : \textit{beginmark}(n : \mathbf{N}); \textit{extendmark}(n : \mathbf{N})$ 
outputs
     $\textit{managerin} : \textit{selecttext}(t : \textit{seq } CH)$ 
external
     $\mu X \bullet ( \textit{bufferin}?x \text{ -- } \textit{managerin}!\textit{selecttext}(t) \rightarrow X )$ 
endagent bufferMarkSelect

```

Notice how the external description of *bufferMarkSelect* is so structured to make the stimulus-response of this agent clear. We also note that this external description would allow extending the selection before the selection has begun. You can do a  $(\textit{bufferin}, \textit{extendmark}(n))$  before a  $(\textit{bufferin}, \textit{mark}(n))$  has been done, but this is disallowed by the internal description of the agent operations (specifically, the precondition of the *extendmark* operation prevents this).

The next operation, unselecting text, is a control operation which will be requested by the buffer manager to ensure that only one of the buffers contains selected text. Its isolated specification is given by the agent *bufferUnselect*.

```

agent bufferUnselect
interleaves (bufferstatc)
with
internal
operations
    unselect()
changes (before, selected, insertion, begsel, endsel)
pre
     $\textit{insertion} \neq \textit{null}$ 
post
     $\textit{insertion}' = \textit{null}$ 
communication
inputs
     $\textit{managerout} : \textit{unselect}()$ 
external

```

```

run
endagent bufferUnselect

```

The final operations at the buffer level describe the insertion of text at an arbitrary position (used for copying and moving) and the deletion of the selected text.

```

agent bufferInsDel
interleaves (bufferstate)
with
internal
  operations
    insert(t : seq CH; n : N)
      changes (before, selected, insertion, begsel, endsel, text)
      pre
        insertion ≠ null
      post
        insertion' = n
        before' = 1..(n - 1) < text
        endsel'..#text'seq < text' = n..#text'seq < text
        selected' = t
    remove()
      changes (selected, endsel, text)
      pre
        insertion ≠ null
      post
        selected' = { }
        endsel'..#text'seq < text' = endsel..#text'seq < text
  communication
    inputs
      managerout : insert(t : seq CH; n : N); remove()
  external
    run
endagent bufferInsDel

```

We can finally combine the above three agents to arrive at a specification of a single buffer agent with selection facilities, which we call *bufferSelect*. We add the display channel at this point for communication to the *Interface*. The implicit external constraint states that every input event to the buffer is followed by an output event to the display.

```

agent bufferSelect
interleaves (bufferMarkSlect, bufferUnselect, bufferInsDel)
with
internal

```

```

operations
  showtext( $t : \text{seq } CH; m, n : \mathbf{N} \cup \text{null}$ )
    pre
       $t = \text{text}$ 
       $m = \text{begsel}$ 
       $n = \text{endsel}$ 
communication
  outputs
    display : showtext( $t : \text{seq } CH; m, n : \mathbf{N} \cup \text{null}$ )
external
  satisfying
     $\forall t : \text{traces}$ 
    •  $\#(t \mid \{\text{display}\}) \leq \#(t \mid \{\text{bufferin}, \text{managerout}\}) \leq \#(t \mid \{\text{display}\}) + 1$ 
endagent bufferSelect

```

We could continue to define further behaviour of a single buffer, but for the example of this scenario, the behaviour of *bufferSelect* is adequate, and so we can refer to it as the *buffer* agent.

```

agent buffer
interleaves (bufferSelect)
endagent buffer

```

### 8.5.2 The buffer manager

The concurrent activity of a collection of buffers will be managed by the agent *manager*. The purpose of this agent is twofold. It will ensure that only one buffer is active, that is, only one buffer contains a non-null insertion point. It will also provide a means by which the selected text is deleted, copied, or moved. We will need a set of buffer labels, *BUFFID*, which will serve as unique identifiers for all of the possible buffers that could be known to the manager. The manager is initialized internally.

```

agent managerState
internal
  types
     $BUFFID == \mathbf{P} \mathcal{V}$ 
     $\text{null} \notin BUFFID$ 
  attributes
     $\text{known} : \mathbf{P } BUFFID$ 
     $\text{current}, \text{last} : BUFFID \cup \{\text{null}\}$ 
     $\text{selected} : \text{seq } CH$ 
  invariant
     $\text{current} \neq \text{null} \Rightarrow \text{current} \in \text{known}$ 
     $\text{last} \neq \text{null} \Rightarrow \text{last} \in \text{known}$ 

```

initially  
      $known = \emptyset$   
 communication  
 external  
 endagent *managerState*

We will not give the detailed description of the buffer manager at this point. The importance of this description is that it would contain events corresponding to the copying and moving of selected text from one buffer to another. Input from the *User* via the interface that is intended to perform these two operations on selected text must be controlled via the manager; interaction directly with the buffer would not be able to control the communication of selected text between buffers. In fact, it is precisely because individual buffer agents cannot control this communication that the buffer manager was introduced. The manager ensures that only one buffer is active at a time and it keeps track of the selected text in order to communicate it to any buffer necessary.

It is not necessary for the buffer manager to control deletion of the selected text, since that operation can be defined at the buffer level. However, a design decision for these text editors has placed the delete command at the control of the buffer manager.

### 8.5.3 Deriving a description of the interface

As we suggested earlier, the agent description of the functional core suggests a natural agent description of its graphical interface. Figure 8.9 is a more detailed agent representation of the functional core. This figure shows how operations relating to selected text are distributed in the agent hierarchy. The copy, move and delete commands are located within the buffer manager, and the other operations on the selected text—marking the beginning and extending the selected region—are located at the buffer level.

Predictability, as described earlier in this thesis, centres around the future results of the interaction being determinable based on knowledge of what the results are currently, and likewise for the displays. Synthesis concerned whether the user was able to determine changes to the result based on observed changes to the display. If the result and display behaviour is predictable, then to satisfy properties associated to synthesis (honesty and trustworthiness are the ones we have discussed explicitly in this thesis), we need to ensure a correspondence between the structures the user recognizes from the graphical interface and the result structures the designer intended the display to portray.

With this in mind, we suggest that the agent description depicted in Figure 8.9, and given formally earlier, leads to a natural description of the graphical interface which corresponds to the agent description of the *System*. We can contrast this with




Figure 8.9: Agent representation of functional core with operations.

an agent description of the interfaces to Spy and Ten based on the visual appearance of the graphical interfaces shown in Figures 8.5 and 8.6. In Figure 8.10, we give the agent description of the Spy interface. Here it is seen that the delete operation




Figure 8.10: Agent representation of Spy's interface.

is distributed to the presentation of each buffer agent. User of this system will falsely connect the functionality of the delete selected text operation to the buffer containing the selected text. But experimental interaction with the interface shows that pressing *any* delete button will always delete the selected text.

In Ten (see Figure 8.11), since the delete command has been included in the pop-up menu with copy and move, it is no longer falsely associated to the single buffer. We say this because the pop-up menu is not connected graphically to any individual buffer, but rather roams independently with the mouse, whose input to the system is constrained to the region in which the editor is located.




Figure 8.11: Agent representation of Ten's interface.

### 8.5.4 Conclusions on the case study

Though the graphical interface of Ten corresponds more closely to the functional description of the text editor when restricted to the tasks of deleting, copying and moving selected text, there still remain questions about its implementation. By placing the delete operation on the pop-up menu, deletion was removed from its ill-advised connection to the individual buffer. However, pop-up menus attached to the mouse are linked with the positional information associated to the mouse as well. Whereas this information is necessary for the copy and move operations, because their operations depend on positional information (the destination of the moved or copied text), it is not necessary for deletion, which is an operation independent of the mouse position. This argument leads us to suggest that the best location for delete would have been in some overall text editor menu, or even attached to some key on the keyboard. Arguments for placing deletion on the pop-up menu centre around minimizing the motor movement and distraction of the user. Since deletion will often take place shortly after the mouse has been used to mark and extend the selected text, it would be easiest and least distracting to the user to allow them to use the mouse again to issue the deletion command. We cannot sufficiently support these kinds of keystroke level or attention arguments within our formal interaction models. As a result, we admit that our formal techniques can only truly be effective in concert with other psychological modelling techniques which will provide decision support for design where our models cannot.

There are several reasons why this particular multibuffer text editor provides a very interesting case study. For the purposes of the scenario, the two versions are functionally equivalent. Their only difference lies in the arrangement of commands in their graphical interfaces. There is a significant difference in the two graphical interfaces, and one must ask the reason for the massive changes. Ultimately, the changes were made in order to increase the amount of screen real estate dedicated

to the display of the buffer contents.<sup>2</sup> Our analysis has tried to trace the usability consequences of this design decision.

It turns out for this particular of command organization, the later version (Ten) was an improvement on usability. Other analysis that we have done [97] concerning the search/replace facility shows that Ten is indeed less usable because the correspondence between result and display templates for the searching task is not as the user anticipates. In general, the move to greater screen real estate has meant a sacrifice in visibility of correspondence detail that facilitates predictability.

## 8.6 Conclusions

In this chapter, we have initiated an assessment of multiagent architectures with the formal agent model. Having introduced two well-known multiagent architectures, MVC and PAC, we identified a major difference between them in terms of the abstract property of correspondence which we desire to map down to a concrete form in some architecture. PAC advocates the isolation of correspondence information between *System* and *Interface* agents, whereas MVC does not. We presented local correspondence as a more task-related form of the abstract version of correspondence presented in Chapter 7. The formal definition led directly to a more precise interpretation of the duties of the control agent in a PAC compound agent, which we isolated in terms of its correspondence role and its hierarchical role.

In the case study of *Spy vs. Ten*, we showed how a formal development of the *System* agent architecture can lead to a semi-formal heuristic for the design of potential graphical interfaces to support properties such as predictability and synthesis, in which the user tries to recognize structures of the *Interface* and relate them to structures of the *System*.

---

<sup>2</sup>This information was obtained through private communication with David Duce from Rutherford Labs.

# Chapter 9

## Conclusions

### 9.1 Summary of the thesis

Recalling the dual focus of attention in this thesis—human-computer interaction and software engineering—we can summarize the results which have been presented.

From an HCI perspective, we have provided informal, semi-formal and fully formal mechanisms for the design and analysis of interactive systems. Informally, we have defined an interaction framework which describes the major components in an interactive system, the user, the system and the interface, and we used this framework both to contextualize a great portion of HCI research and to motivate the more formal treatment of interactive properties. At the more formal level, we have provided a constructive bridge between the abstract and general computational models of interaction and the heuristic design of interactive systems. A formal model of the agent has been used as the building block for interactive system description. We have provided the beginnings of a methodology for incorporating psychological knowledge of users by demonstrating the link between task analysis and agent description. The formulation of interactive properties on agents provides constraints for the design of an interactive system and a means for evaluating existing systems to understand user difficulties. We have provided several examples of how a formal or semi-formal analysis can be used to explain scenarios of user behaviour.

From a software engineering perspective, the agent model is a formal model of an object-oriented specification language. It is the first specification language to our knowledge with a compositional semantics that incorporates existing model-oriented techniques for a description of an entity's internal state behaviour along with existing process algebra techniques for a description of the external event behaviour. The agent language answers criticisms of previous model-oriented techniques, such as Z, by providing a mechanism for describing objects. Beyond that, we have also incorporated a means for specifying the communication between objects,



a facility which is not treated as generally in other object-oriented formalisms.

It is interesting to note that the requirements for our agent model which lead to the development of the hybrid notation arose directly from our need to express a particular class of interactive properties. Thus we have shown how due consideration of nonfunctional requirements in software engineering can increase the benefits of formal methods as well as influence the development of better formal techniques.

## 9.2 Contributions of this thesis

Our goal in this thesis was to provide a scientific approach to the development of more usable interactive systems. We outline now the contributions of this thesis toward that goal.

Formality rarely provides its own context; it is often necessary to provide informal motivation for the formalism. This motivation is usually culled from common sense about the world around us. The interaction framework is such a common sense description of interaction in the everyday world. Its description is free of psychological or computational jargon. The framework has served two major purposes in our own work. It has provided the context for our research, enabling us to compare and classify previous HCI research as well as our own. It has also motivated the development of the agent formalisms used in this thesis to describe relevant interactive properties.

We consider interactive systems as a collection of communicating stimulus-response agents. A simple model of an agent was described which was consistent with the interaction framework and allowed for the precise description of several interactive properties. These properties were motivated by an informal description of the translations made between the major components in an interactive system—the *User*, *System*, *Input* and *Output*. We provided an account of how those translations affect the overall goal-directed behaviour of the human user. Some formal descriptions were aimed at describing the ease and coverage of translations in the framework which are themselves seen as implicit agent descriptions based on their stimulus-response behaviour. The correspondence property was described as a relationship between two separate agents that would eventually be embodied by a third translation agent. Other more abstract properties, such as predictability, synchronicity and consistency, classified agents by the relationship between their input history and their state or response history.

The problem with the simple agent model was that its near black-box description did not allow for a constructive discussion of interactive properties that relate the goals or results of the interaction with the immediate and visible information that the user sees. At the abstract level of the simple agent, these properties can only be discussed by use of mappings from the state space to the respective result and display spaces, as is done in the red-PIE model and Sufrin and He's model

of interactive processes. In designing a system, however, we need to be able to construct the agents from the result and display information of a task analysis.

So motivated by the desire to express interactive properties more constructively, we refined the agent model. The refinement was intended to support a hybrid view of an agent, using complementary descriptions of its internal, state-based behaviour and its external, event-based behaviour. The state of the agent was given as an attribute-value mapping and the two complementary views were made consistent by a communication description which linked the operations of the internal description with the events of the external description that occur on input, output, synchronized and internal channels. We also defined two composition operators, corresponding to the synchronous combination of independent agents (those with no attributes in common) and the interleaved combination of dependent agents (those with attributes in common).

The description of the refined agent model was produced in the Z notation, but we felt that Z did not provide a natural means for building the description of an agent, i.e., Z was not a suitable design notation. We, therefore, provided a language for agents which made clear the internal, external and communication descriptions. One criterion for developing this language was that it be familiar to those with experience in other formal notations, and so it was made to look similar to a model-oriented notation, such as Z or VDM, for the internal description, and similar to a process algebra notation, such as CSP or CCS, for the external description.

The refined agent model and its associated language were then used to show how interactive properties relating result and display information could replace the previous more abstract properties on agents. Templates have been previously used in abstract models as a means of modelling the focus of attention of a user. We gave a constructive definition of templates for restricting the state of an agent to those attributes which are relevant to a particular task. The identification of result and display templates as task-dependent descriptions derived from task analysis links agent design in a more user-centred interactive system development method. Reformulations of interactive properties related input to task-specific result and display templates yielded versions of predictability, consistency, etc., that were more relevant to the user's understanding of the interaction.

The final contribution of this thesis was to initiate the formal description of multiagent architectures which have been previously used as heuristic guidelines for interactive system development. Architectures are realistic platforms for the description of implementations, and it is necessary to show how the properties derived in the abstract can be traced down to this more concrete level. We discussed two known multiagent architectures, MVC and PAC, and highlighted a difference between them. PAC provides for the explicit description of correspondence between the *System* (model or abstraction) and *Interface* (view or presentation) agents, whereas MVC does not. We then showed how this feature of the control agent in

PAC is a manifestation of the abstract correspondence property expressed between agents. A more extensive case study then showed how the agent description of the *System* leads to a natural agent description of the *Interface*, which was then used to analyze the actual graphical interfaces to existing interactive systems to explain confirmed user confusion or irritation.

### 9.3 Future work

Though the contributions of this thesis as stated above are significant, another real contribution of this work is the research agenda that it motivates. We will summarize the main aims of this agenda and then comment on improvements to the current work which are necessary based on our experience using agents.

The agent model is intended as the formal system modelling notation for the ESPRIT Basic Research Action project 3066 (AMODEUS) [13]. Within the project, system modelling fits into a larger research scheme with very courageous aims. The AMODEUS project is interested in assimilating different HCI modelling techniques in order to see how they can be coordinated toward a more effective design practice. The modellers fall into three main categories. There are psychologists whose main research is in the development of models of the user. There are computer scientists whose main interest is in the development of models for system description. These two modelling domains are linked by a third domain concerned with how practicing interactive systems developers document the rationale behind their design decisions.

As briefly discussed in Chapters 7 and 8, the main exercises within the project are based on separate analyses of scenarios of interaction between user and system. These scenarios are used as a means of eliciting information comparing and contrasting the utility of the different modelling domains. Earlier versions of the agent model and notation have been used with relative success by the system modellers, and some of those examples have been provided within this thesis. Many more scenarios have been investigated [72, 8, 42], and further are planned. Some of the case studies address issues in computer-supported cooperative work (CSCW) and highly interactive display-based systems. Continued application of the agent model to these scenarios in case studies specifically geared to test its ability to capture relevant information concerning interactive behaviour will undoubtedly lead to further refinements of the model and an increased confidence in its utility.

We realize that the formal agent model cannot alone address all of the important aspects of interaction, not even from the restricted viewpoint of system modelling. For example, the agent model does not provide a way of determining the relevant agents in the system from some requirements specification or scenario description. Therefore, a number of system modellers are working on ways of applying design heuristics in the form of an expert system design tool to help a designer determine the agents and their communications connections before embarking on a formal

description [70]. We also plan to investigate the possibility of applying the agent model more within the user modelling domain as a way of expressing the results of psychological theory in a language of design.

Refinement from specification toward implementation within the agent model should also be possible. Sufrin and He demonstrated how a refinement ordering can be defined on their model of interactive processes, leading to a definition of operational and data refinement in terms of the result and display behaviour. With slight modifications, this definition of the refinement ordering can be mimicked in the agent model, with the additional benefit that refinement can be restricted to task, since the result and display information in the agent model is defined in terms of task. The agent language could be augmented to wide spectrum language so that refinement can remain within the agent model and proceed from high-level specification to implementation, as described by Morgan's refinement calculus on Dijkstra's guarded command language [116, 115].

We mentioned in Chapter 5 that the inclusion of the simplest traces model of a process algebra was intended as an example of how more complicated process algebras could be incorporated. We have not investigated this point very thoroughly, and it would be valuable to see what advantages actually arise from the use of other process algebra models. We suspect that the inclusion of more sophisticated models, such as the failures-divergence model of CSP, would not automatically introduce the ability to express liveness properties because of the constraint within the model that both internal and external descriptions determine overall behaviour. Whereas the external description of an agent in a failures-divergence semantics may be shown to satisfy some liveness properties—guaranteeing that some desirable behaviour will happen—the corresponding internal description may prohibit the agent from engaging in the desired property.

An interesting comment on the development of the agent model is that it was directly influenced by two different perspectives on its behaviour, the internal and the external. The need for a hybrid notation arose when it was realized that the two different perspectives corresponded to different natural notations. Hence, a model-oriented language more naturally expressed internal behaviour and a process algebra notation more naturally expressed external behaviour. Research at York is now investigating a classification of perspectives on system models in order to determine how hybrid notations can be adapted to express the different perspectives [16].

We first described the behaviour of an agent as stimulus-response. In moving toward a concrete notation that corresponded to both Z and CSP, we lost the ability to simply express the connection between a state transition, the stimuli that triggered it and the responses that resulted from it. We admit that this connection is much more naturally expressed within standard Z by the use of decorations ? and ! to mark input and output respectively. However, the primitive communication

operators in Z do not allow the arbitrary communication between the output of one operation and the input to other operations, and so we abandoned the use of Z exclusively. We are not familiar with any attempts in the Z community which have solved this communication problem. However, it may be possible to define in Object-Z a communication component to the object class, allowing for a calculus of objects. We hope to pursue this point by using the agent model as an example of how communication can be incorporated into a model-oriented notation.

In summary, we see the work of this thesis as a solid justification for the use of formal methods in system development to promote non-functional as well as functional requirements. Detailed examination of the non-functional requirements related to usability have shown that they can be addressed scientifically to answer relevant research questions. In addition, we point out that the application of formal methods to areas such as HCI can indicate how to improve the existing formal notations themselves.

# Bibliography

- [1] G. Abowd, A. Dix, and M. Harrison. State of the art: Formal aspects of user interfaces. Internal report, Human-Computer Interaction Group, Department of Computer Science, University of York. Presented at Eurographics'90. May 1990.
- [2] G. Abowd, A. Dix, and M. Harrison. Formalising user recognisable structures of graphics packages. In *Proceedings of the Eurographics Workshop on Formal Methods in Computer Graphics*, Marina di Carrara, Italy, June 1991.
- [3] G. D. Abowd. Properties of a graphical interface within a formal interactive system architecture. In *Proceedings of the Eurographics Workshop on Formal Methods in Computer Graphics*, Marina di Carrara, Italy, June 1991.
- [4] G. D. Abowd and R. Beale. A framework for the analysis and design of interactive systems. Technical Report YCS (156), University of York, Department of Computer Science, 1991.
- [5] G. D. Abowd and R. Beale. Users, systems and interfaces: a unifying framework for interaction. In D. Diaper and N. Hammond, editors, *HCI'91: Usability Now*. British Computer Society Special Interest Group on Human-Computer Interaction, Cambridge, 1991.
- [6] G. D. Abowd, J. Bowen, A. Dix, M. Harrison, and R. Took. User interface languages: A survey of existing methods. Technical Report PRG-TR-5-89, Oxford University Computing Laboratory Programming Research Group, October 1989. Also published as internal report 2487-TN-PRG-1008 Issue 1.0 for ESPRIT project 2487 (REDO).
- [7] G. D. Abowd and M. D. Harrison. On a constructive approach to applying formal methods in HCI. Technical Report YCS (151), University of York, Department of Computer Science, December 1990.
- [8] G. D. Abowd and M. D. Harrison. Design scenarios for M1.5. Working paper RP1/WP10, Esprit BRA project 3066 (AMODEUS), March 1991.

- [9] G. D. Abowd, M. D. Harrison, and C. R. Roast. Modelling predictability in interactive systems. Working paper RP1/WP9, Esprit BRA project 3066 (AMODEUS), September 1990.
- [10] H. Alexander. Executable specifications as an aid to dialogue design. In H. J. Bullinger and B. Shackel, editors, *Human-Computer Interaction — INTERACT'87*, pages 739–744. North Holland, 1987.
- [11] H. Alexander. Formally-based techniques for designing human-computer dialogues. In D. Diaper and R. Winder, editors, *People and Computers III*, pages 201–214. Cambridge University Press, 1987.
- [12] H. Alexander. *Formally-Based Tools and Techniques for Human-Computer Dialogues*. Ellis Horwood Ltd., 1987.
- [13] AMODEUS consortium. Assimilating models of designers, users and systems. Esprit Basic Research Action 3066, Technical Annex, 1989.
- [14] S. O. Anderson. Proving properties of interactive systems. Technical report, Department of Computer Science, Heriot-Watt University, 1985.
- [15] S. O. Anderson. Proving properties of interactive systems. In M. D. Harrison and A. F. Monk, editors, *People and Computers: Designing for Usability*. Cambridge University Press, 1986.
- [16] P. Andrews. First year report on doctoral research. Department of Computer Science, University of York, June 1991.
- [17] P. Barnard. Interacting cognitive subsystems: A psycholinguistic approach to short-term memory. In A. Ellis, editor, *Progress in the psychology of language*, volume 2, chapter 6. Lawrence Erlbaum, 1985.
- [18] P. Barnard. Cognitive resources and the learning of human-computer dialogs. In J. M. Carroll, editor, *Interfacing Thought: Cognitive Aspects of Human-Computer Interaction*, pages 112–158. MIT Press, 1987.
- [19] P. Barnard and M. Harrison. Towards a framework of modelling interactions. Project deliverable D3, Esprit BRA project 3066 (AMODEUS), September 1990.
- [20] P. Barnard and M. D. Harrison. Integrating cognitive and system models in human computer interaction. In A. G. Sutcliffe and L. A. Macauley, editors, *People and Computers V*. Cambridge University Press, 1989.

- [21] P. Barnard, M. Wilson, and A. MacLean. Approximate modelling of cognitive activity with an expert system: A theory-based strategy for developing an interactive design tool. *The Computer Journal*, 31(5):445-456, 1988.
- [22] P. J. Barnard and M. D. Harrison. Towards a framework for modelling human-computer interaction. Working paper RP3/WP5, Esprit BRA project 3066 (AMODEUS), May 1991. Submitted to Esprit Conference 1991.
- [23] L. Bass, R. Little, R. Pellegrino, S. Reed, R. Seacord, S. Sheppard, and M. R. Szczur. The arch model: Seeheim revisited. User Interface Developer's Workshop Report, April 1991.
- [24] R. Bird and P. Wadler. *Introduction to Functional Programming*. Prentice-Hall International, 1988.
- [25] R. Bornat. Imperative languages in distributed computing. In D. A. Duce, editor, *Distributed Computing Systems Programme*, IEE Digital Electronics and Computing Services, 1984.
- [26] N. Boudriga, F. Elloumi, and A. Mili. The lattice of specifications: Applications to a specification methodology. Internal report from Tunis University. 1990.
- [27] J. P. Bowen, R. B. Gimson, and S. Topp-Jørgensen. Specifying system implementations in Z. Technical Monograph PRG-63, Oxford University Computing Laboratory Programming Research Group, February 1988.
- [28] S. D. Brookes. *A Model for Communicating Sequential Processes*. D.Phil. thesis, Oxford University, 1983. Also published as Carnegie-Mellon Technical Report CMU-CS-83-149.
- [29] S. Burbeck. How to use model-view-controller (MVC). Technical report, Softsmarts, Inc., 1987.
- [30] K. Butler, J. Bennett, P. Polson, and J. Karat. Report on the workshop on analytical models: Predicting the complexity of human-computer interaction. *SIGCHI Bulletin*, 20(4):63-79, April 1989.
- [31] S. K. Card, J. D. Mackinlay, and G. G. Robertson. The design space of input devices. In J. C. Chew and J. Whiteside, editors, *Empowering People, Proceedings of CHI'90 Conference*, pages 117-124, 1990.
- [32] S. K. Card, T. P. Moran, and A. Newell. *The Psychology of Human Computer Interaction*. Lawrence Erlbaum, 1983.



- [33] S. K. Card, G. G. Robertson, and J. D. Mackinlay. The information visualizer, an information workspace. In S. P. Robertson, G. M. Olson, and J. S. Olson, editors, *Reaching Through Technology, Proceedings of the CHI'91 Conference*, pages 181-188. ACM Press, 1991.
- [34] D. Carrington, D. Duke, R. Duke, P. King, G. Rose, and G. Smith. Formal specification in Object-Z: Introduction and case studies. Technical Report 105, Key Centre for Software Technology, Dept. of Computer Science, University of Queensland, July 1989.
- [35] J. M. Carroll. Infinite detail and emulation in an ontologically minimized HCI. In J. C. Chew and J. Whiteside, editors, *Empowering People, Proceedings of CHI'90 Conference*, pages 321-327, 1990.
- [36] J. M. Carroll and M. B. Rosson. Usability specifications as a tool in iterative development. In R. Hartson, editor, *Advances in Human-Computer Interaction* Ablex, 1984.
- [37] J. Coenen, W.-P. de Roever, and J. Zwiers. Assertional data reification proofs: Survey and perspective. In *Fourth Refinement Workshop of the BCS FACS Special Interest Group*, January 1991.
- [38] J. Coutaz. PAC, an object oriented model for dialog design. In H. J. Bullinger and B. Shackel, editors, *Human-Computer Interaction -- INTERACT'87*, pages 431-436. North-Holland, Amsterdam, 1987.
- [39] J. Coutaz. *Interface Homme-Ordinateur: Conception et Réalisation*. PhD thesis, University of Grenoble, Laboratoire de Génie Informatique, December 1988.
- [40] J. Coutaz. Architecture models for interactive software. In *Proc. of ECOOP'89*, 1989.
- [41] J. Coutaz. Architecture models for interactive software: Failures and trends. In *Engineering for Human-Computer Interaction*, 1989. Proceedings of the IFIP WG 2.7 conference.
- [42] J. Coutaz, L. Nigay, M. Harrison, and G. Abowd. Design scenarios for ml.5. Working paper RP2/WP11, Esprit BRA project 3066 (AMODEUS), March 1991.
- [43] B. J. Cox. *Object-Oriented Programming: an Evolutionary Approach*. Addison-Wesley, 1986.

- [44] M. Curry, A. Monk, and P. Wright. Obstacles to the use of formal notations in software design practice. In University of York, Department of Computer Science Technical Report, September 1990.
- [45] J. Davies. *Specification and Proof: Real-Time Systems*. D.Phil. thesis, Oxford University, 1991.
- [46] J. Davies and S. Schneider. An introduction to timed CSP. Technical Monograph PRG-75, Oxford University Computing Laboratory Programming Research Group, August 1989.
- [47] A. Dix. Nondeterminism as a paradigm for understanding the user interface. In M. D. Harrison and H. W. Thimbleby, editors, *Formal methods in Human-Computer Interaction*, Cambridge Series on Human-Computer Interaction, chapter 4. Cambridge University Press, 1990.
- [48] A. J. Dix. *Formal Methods and Interactive Systems: Principles and Practice*. D.Phil. thesis, University of York, 1987.
- [49] A. J. Dix. *Formal Methods for Interactive Systems*. Academic Press, 1991.
- [50] A. J. Dix and M. D. Harrison. Principles and interaction models for window managers. In M. D. Harrison and A. F. Monk, editors, *People and Computers: Designing for usability*, pages 352-366. Cambridge University Press, 1986.
- [51] A. J. Dix and M. D. Harrison. Interactive systems design and formal development are incompatible? In J. A. McDermid, editor, *Proceedings 1988 Refinement Workshop*. Butterworth Scientific, 1989.
- [52] A. J. Dix, M. D. Harrison, and E. E. Miranda. Using principles to design features of a small programming environment. In I. Sommerville, editor, *Software Engineering Environments*, pages 135-150. Peter Peregrinus, 1986.
- [53] A. J. Dix, M. D. Harrison, C. Runciman, and H. W. Thimbleby. Interaction models and the principled design of interactive systems. In H. Nichols and D. S. Simpson, editors, *European Software Engineering Conference*, pages 127-135. Springer Lecture Notes, 1987.
- [54] A. J. Dix and C. Runciman. Abstract models of interactive systems. In P. Johnson and S. Cook, editors, *People and Computers: Designing the interface*, pages 13-22. Cambridge University Press, 1985.
- [55] D. Duke and R. Duke. A history model for classes in Object-Z. Technical Report 120, Key Centre for Software Technology, Dept. of Computer Science, University of Queensland, November 1989.

- [56] R. Duke, P. King, G. Rose, and G. Smith. The Object-Z specification language: Version 1. Technical Report 91-1, Key Centre for Software Technology, Dept. of Computer Science, University of Queensland, May 1991.
- [57] J. Finlay, A. Green, P. Barnard, and M. Harrison. Linking user and system models: an interaction structure. Working paper RP3/WP2, Esprit BRA project 3066 (AMODEUS), January 1991.
- [58] FOCUS. Foundations of opto-electronic computer systems, action 3180. In *Synopses of Basic Research: actions, working groups and networks of excellence, Volume 8*, pages 192-195, January 1991.
- [59] J. Foley and A. van Dam. *Fundamentals of Interactive Computer Graphics*. Addison Wesley, 1984.
- [60] J. Goguen, J. Thatcher, and E. Wagner. An initial algebra approach to the specification, correctness and implementation of abstract data types. In Yeh, editor, *Current Trends in Programming Methodology*, volume IV: Data Structuring, pages 80-149. Prentice-Hall, 1978.
- [61] A. Goldberg. *Smalltalk-80: The Interactive Programming Environment*. Addison-Wesley, 1983.
- [62] M. Green. The University of Alberta user interface management system. *ACM Conference of the Special Interest Group for Graphics*, 19(3):205-214, July 1985.
- [63] T. R. G. Green, F. Schiele, and S. J. Payne. Formalisable models of user knowledge in human-computer interaction. In G. C. van der Veer, T. R. G. Green, J.-M. Hoc, and D. Murray, editors, *Working with Computers: Theory versus Outcome*, Computers and People Series, chapter 1. Academic Press, 1988.
- [64] J. Grudin. The case against user interface consistency. *Communications of the ACM*, 4(3):245-264, 1989.
- [65] J. Guttag and J. Horning. The algebraic specification of abstract data types. *Acta Informatica*, 10:27-52, 1978.
- [66] M. Hanlon and J. Newton. FOREST case study: The MVC. Technical report, Advanced Techniques Group, Data Logic Limited, Queens House, Greenhill Way, Harrow, October 1990. 1990.
- [67] D. Harel. Statecharts: a visual formalism for complex systems. *Science of Computer Programming*, 8(3):231-274, June 1987.

- [68] D. Harel. On visual formalisms. *Communications of the ACM*, 31(5):514-530, May 1988.
- [69] M. Harrison and G. Abowd. Formal methods in human-computer interaction: a tutorial. Technical Report YCS 155, University of York, Department of Computer Science, 1991. Also a tutorial presentation at CHI'91 conference.
- [70] M. Harrison, J. Coutaz, L. Nigay, J. E. Finlay, and G. D. Abowd. Agent architectures and the application of transformation rules in interactive system development. Project deliverable D2, Esprit BRA project 3066 (AMODEUS), September 1990.
- [71] M. D. Harrison. Modelling user structures within system specifications. In Colloquium on Formal Methods in HCI: III, sponsored by IEE Professional Group C5 (Man-Machine Interaction). December 1989.
- [72] M. D. Harrison, J. Coutaz, J. E. Finlay, G. D. Ahowd, and L. Nigay. Interaction analysis techniques from a system modelling viewpoint. Project deliverable D1, Esprit BRA project 3066 (AMODEUS), September 1990.
- [73] M. D. Harrison and A. J. Dix. A state model of direct manipulation. In M. D. Harrison and H. W. Thimbleby, editors, *Formal Methods in Human Computer Interaction*, pages 129-151. Cambridge University Press, 1990.
- [74] M. D. Harrison, C. R. Roast, and P. C. Wright. Complementary methods for the iterative design of interactive systems. In G. Salvendy and M. Smith, editors, *Designing and Using Human-Computer Interfaces and Knowledge Based Systems*, pages 651-658. Elsevier Scientific, 1989.
- [75] M. D. Harrison and H. W. Thimbleby, editors. *Formal Methods in Human Computer Interaction*. Cambridge University Press, 1990.
- [76] J. He. Process refinement. In J. A. McDermid, editor, *The Theory and Practice of Refinement*. Butterworth Scientific, 1989. Proceedings of 1988 York Refinement Workshop.
- [77] J. He. A trace-state based approach to process specification and design. Oxford University Computing Laboratory Programming Research Group, 1989.
- [78] J. He, C. Hoare, and J. Sanders. Data refinement refined. In *European Symposium on Programming*, volume 213 of *Lecture Notes in Computer Science*. Springer-Verlag, 1986.
- [79] S. Hekmatponr and D. Ince. Evolutionary prototyping and the human-computer interface. In H.-J. Bullinger and B. Shackel, editors, *Proceedings of INTERACT'87*, pages 479-484. North-Holland, 1987.

- [80] M. Hennessy. *Algebraic Theory of Processes*. MIT Press, Cambridge, Massachusetts, 1988.
- [81] C. A. R. Hoare. Communicating sequential processes. *Communications of the ACM*, 21(8):666-77, 1978.
- [82] C. A. R. Hoare. *Communicating Sequential Processes*. Prentice-Hall International, London, 1985.
- [83] A. Howes and S. J. Payne. Display-based competence: towards user models for menu-driven interfaces. University of Lancaster.
- [84] J. Hughes. Specifying a visual file system in Z. In *Formal Methods in HCI: III*. IEE Professional Group C5 (Man-machine interaction), 1989.
- [85] E. L. Hutchins, J. D. Hollan, and D. A. Norman. Direct manipulation interfaces. *Human-Computer Interaction*, 1:331-338, 1985.
- [86] INMOS Limited. *occam Programming Manual*. Prentice-Hall International, London, 1984.
- [87] ISO 8807. Information processing systems—open systems interconnection — LOTOS—a formal technique based on the temporal ordering of observational behaviour. Technical report, ISO Standards Authority, 1988.
- [88] R. J. K. Jacob. Survey and examples of specification techniques for user-computer interfaces. Technical report, Naval Research Laboratory, Washington, D.C., 1983.
- [89] R. J. K. Jacob. Using formal specifications in the design of a human-computer interface. *Communications of the ACM*, 26(4):259-264, 1983.
- [90] C. B. Jones. *Systematic Software Development Using VDM*. Prentice-Hall International, 1986.
- [91] G. Jones. *Programming in occam*. Prentice-Hall International, 1986.
- [92] M. B. Josephs. A state-based approach to communicating processes. *Distributed Computing*, 3, 1988.
- [93] D. E. Kieras and P. G. Polson. An approach to the formal analysis of user complexity. *International Journal of Man-Machine Studies*, 22:365-394, 1985.
- [94] C. Knowles. Can cognitive complexity theory (CCT) produce an adequate measure of system usability? In D. M. Jones and R. Winder, editors, *People and Computers IV*, pages 291-307. Cambridge University Press, 1988.

- [95] D. Kozen. Results on the propositional  $\mu$ -calculus. *Theoretical Computer Science*, 23, 1983.
- [96] G. E. Krasner and S. T. Pope. A cookbook for using the model-view-controller user interface paradigm in Smalltalk-80. *JOOP*, 1(3), August 1988.
- [97] S. Krening. The structure of interaction. Working paper RP3/WP6, Esprit BRA project 3066 (AMODEUS), June 1991.
- [98] J. Laird, A. Newell, and P. Rosenbloom. SOAR: an architecture for general intelligence. *Artificial Intelligence*, 33:1-64, 1987.
- [99] L. Lamport. A simple approach to specifying concurrent systems. *Communications of the ACM*, 32(1):32-45, January 1989.
- [100] K. Lano.  $Z^{++}$ , an object-orientated extension to Z. In *Proceedings of the fifth annual Z User Meeting*, December 1990.
- [101] C. H. Lewis. A research agenda for the uieties in human-computer interaction. *Human-Computer Interaction*, 5(2-3):125-143, 1990.
- [102] B. Liskov and S. Zilles. Programming with abstract data types. *ACM SIG-PLAN Notices*, 9(4):50-59, April 1974.
- [103] R. Macdonald. Z usage and abuse. Report 910003, Royal Signals and Radar Establishment, Malvern, Worcestershire, February 1991.
- [104] Z. Manna. *A Mathematical Theory of Computation*. McGraw Hill, 1974.
- [105] L. S. Marshall. *A Formal Description Method for User Interfaces*. PhD thesis, University of Manchester, United Kingdom, 1986. Also published as technical report UMCS-87-1-2.
- [106] M. Martins. Formal specification of highly interactive text editors. Technical Report RAL-87-093, Rutherford Appleton Laboratory, November 1987.
- [107] A. Mével and T. Guéguen. *Smalltalk-80*. Macmillan Education, 1987.
- [108] F. Mili and A. Mili. Relational heuristics for programming: Advances and perspectives. Internal report from Oakland University and Tunis, November 1989.
- [109] A. Milner. *A Calculus of Communicating Systems*, volume 92 of *Lecture Notes in Computer Science*. Springer-Verlag, 1980.
- [110] R. Milner. *Communication and Concurrency*. Prentice-Hall International, UK, Ltd., London, 1989.

- [111] F. Moller and C. Tofts. A temporal calculus of communicating systems. LFCS Report Series ECS-LFCS-89-104, LFCS, Department of Computer Science, University of Edinburgh, December 1989.
- [112] A. Monk, editor. *Fundamentals of Human-Computer Interaction*. Computers and People Series. Academic Press, 1984.
- [113] A. F. Monk and A. J. Dix. Refining early design decisions with a black-box model. In D. Diaper and R. Winder, editors, *People and Computers III, HCI'87*, pages 147-158. Cambridge University Press, 1987.
- [114] T. P. Moran. The command language grammar: a representation for the user interface of interactive systems. *International Journal of Man Machine Systems*, 15, 1981.
- [115] C. Morgan. *Programming from Specifications*. Prentice-Hall International, 1990.
- [116] C. Morgan, K. Robinson, and P. Gardiner. On the refinement calculus. Technical Monograph PRG-70, Oxford University Computing Laboratory Programming Research Group, October 1988.
- [117] C. C. Morgan and B. A. Sufrin. Specification of the Unix filing system. *IEEE Transactions on Software Engineering*, 10(2):128-142, March 1984.
- [118] K. Myers and N. Hammond. Definition of scenarios for M1 workshop. Internal report IR6, Esprit BRA project 3066 (AMODEUS), July 1990.
- [119] A. Newell and H. Simon. *Human Problem Solving*. Prentice-Hall, 1972.
- [120] W. Newman. A system for interactive graphical programming. In *SJCC 1968*, pages 47-54, Washington D.C., 1968. Thompson Books.
- [121] L. Nigay. Modelisation des architectures logicielles des systems interactifs. Master's thesis, Institut National Polytechnique de Grenoble, Lahoratoire de Génie informatique, June 1990.
- [122] D. Norman. Some observations on mental models. In D. Gentner and A. Stevens, editors, *Mental Models*, pages 7-14. Erlbaum, 1983.
- [123] D. A. Norman. Four stages of user activities. In B. Shackel, editor, *Human-Computer Interaction--INTERACT'84*, pages 507-511. Elsevier Science Publishers, 1984.
- [124] D. A. Norman. Cognitive engineering. In D. A. Norman and S. Draper, editors, *User-Centered System Design*, pages 31-62. Erlbaum, 1986.

- [125] D. A. Norman. *The Psychology of Everyday Things*. Basic Books, 1988.
- [126] D. L. Parnas. A technique for software module specification with examples. *Communications of the ACM*, 15(5):330-, May 1972.
- [127] S. J. Payne. Task-action grammars. In B. Shackel, editor, *Human-Computer Interaction—INTERACT'84*, pages 527–532. Elsevier Science Publishers, 1984.
- [128] S. J. Payne and T. R. G. Green. Task-action grammars: a model of mental representation of task languages. *Human-Computer Interaction*. 2(2):93–133, 1986.
- [129] G. Pfaff and P. ten Hagen, editors. *Seeheim Workshop on User Interface Management Systems*. Berlin, 1985. Springer-Verlag.
- [130] G. D. Plotkin. A structural approach to operational semantics. Technical Report DAIMI FN-19, Computer Science Department, Aarhus University, 1981.
- [131] V. Pratt. A decidable mu-calculus. In *Proceedings of the 22nd IEEE Conference on the Foundations of Computer Science*, 1981.
- [132] G. Reed. *A Uniform Mathematical Theory for Real-Time Distributed Computing*. D.Phil. thesis, Oxford University, 1988.
- [133] P. Reisner. Formal grammar and human factors design of an interactive graphics system. *IEEE Transactions on Software Engineering*, SE-7(2):229–240, 1981.
- [134] P. Reisner. Analytic tools for human factors of software. Research Report RJ 3808 (43605), IBM Research Laboratory, San Jose, 1983.
- [135] P. Reisner. Formal grammar as a tool for analysing ease of use: some fundamental concepts. In J. C. Thomas and M. L. Schneider, editors, *Human Factors in Computer Systems*, pages 53–78. Ablex, 1983.
- [136] P. Reisner. What is inconsistency. In D. Diaper, D. Gilmore, G. Cockton, and B. Shackel, editors, *Human-Computer Interaction—INTERACT'90*, pages 175–181. Elsevier Science Publishers, 1990.
- [137] C. R. Roast and P. C. Wright. Incorporating the user's perspective into a system model. Technical Report YCS 148, University of York, Department of Computer Science, 1990.



- [138] G. G. Robertson, J. D. Mackinlay, and S. K. Card. Cone trees: Animated 3D visualizations of hierarchical information. In S. P. Robertson, G. M. Olson, and J. S. Olson, editors, *Reaching Through Technology, Proceedings of the CHI'91 Conference*, pages 189-194. ACM Press, 1991.
- [139] A. Roscoe. *Denotational Semantics for Occam*, volume 197 of *Lecture Notes in Computer Science*, pages 306-329. Springer-Verlag, 1985. July 1984 Seminar on Concurrency.
- [140] A. Roscoe and C. Hoare. The laws of occam programming. Technical Monograph PRG-53, Oxford University Computing Laboratory Programming Research Group, February 1986.
- [141] M. Ryan. Developments of MAL in forest research. WP1 (Logic and Language) Deliverable NFR/WP1.1/IC/R/001/A, Forest Research Project, October 1990.
- [142] F. Schiele and T. Green. HCI formalisms and cognitive psychology: the case of task-action grammars. In M. D. Harrison and H. W. Thimbleby, editors, *Formal methods in Human-Computer Interaction*, Cambridge Series on Human-Computer Interaction, chapter 2. Cambridge University Press, 1990.
- [143] R. L. Schnapp. *SuperPaint 1.1*. Silicon Beach Software, 1988. User's manual for Macintosh software package.
- [144] S. A. Schneider. *Correctness and Communication in Real-Time Systems*. D.Phil. thesis, Oxford University, 1989.
- [145] S. Schuman, D. Pitt, and P. Byers. Object-oriented process specification. Technical report, Department of Mathematics, University of Surrey, 1990.
- [146] S. A. Schuman and D. Pitt. Object-oriented subsystem specification. In L. Meertens, editor, *Program Specification and Transformation*. Elsevier Science Publishers, 1987.
- [147] B. D. Sharratt. The incorporation of early interface evaluation into command language grammar specifications. In D. Diaper and R. Winder, editors, *People and Computers III—Proceedings of HCI'87*. Cambridge University Press, 1987.
- [148] B. D. Sharratt. Top-down interactive systems design: some lessons learnt from using command language grammar specifications. In H.-J. Bullinger and B. Shackel, editors, *Proceedings of INTERACT '87*, pages 395-399. North-Holland, 1987.

- [149] J. L. Sibert, R. Belliardi, and A. Kamran. Some thoughts on the interface between UIMS and application programs. In *User Interface Management Systems*, pages 183–189. Springer-Verlag, 1985.
- [150] J. L. Sibert, W. D. Hurley, and T. W. Bleser. An object-oriented user interface management system. In *ACM SIGGRAPH'86*, 1986.
- [151] T. Simon. Analysing the scope of cognitive models in human-computer interaction: a trade-off approach. In D. M. Jones and R. Winder, editors, *People and Computers IV*, pages 79–93. Cambridge University Press, 1988.
- [152] J. M. Spivey. *The Z Notation: A Reference Manual*. Prentice-Hall International, 1988.
- [153] J. M. Spivey. *Understanding Z, a Specification Language and its Semantics*. Cambridge University Press, 1989.
- [154] C. Stanley-Smith and T. Cahill. UNIFORM: A language geared to system independence. Project Document UL-TN-1002, Esprit project REDO. 1989.
- [155] C. Stirling and D. Walker. A general tableau technique for verifying temporal properties of concurrent programs. In *Semantics for Concurrency*, Workshops in Computing. Springer-Verlag, 1990. Extended abstract.
- [156] J. E. Stoy. *Denotational Semantics: The Scott-Strachey Approach to Programming Language Theory*. MIT Press, 1977.
- [157] L. Suchman. *Plans and Situated Actions: The Problem of Human Machine Interaction*. Cambridge University Press, 1987.
- [158] B. Sufirin and J. He. Specification, refinement and analysis of interactive processes. In M. D. Harrison and H. W. Thimbleby, editors, *Formal methods in Human-Computer Interaction*, Cambridge Series on Human-Computer Interaction, chapter 6. Cambridge University Press, 1990.
- [159] A. Sutcliffe. Some experiences in integrating specification of human-computer interaction within a structured system development method. In D. Jones and R. Winder, editors, *People and Computers IV*, pages 145–160. Cambridge University Press, 1988.
- [160] H. Thimbleby. Generative user-engineering principles for user interface design. In B. Shackel, editor, *Human-Computer Interaction—INTERACT'84*, pages 661–666. Elsevier Science Publishers, 1984.
- [161] H. Thimbleby. *User Interface Design*. Adison Wesley, 1990.

- [162] R. K. Took. Surface interaction: A paradigm model for separating application and interface. In *Proceedings of CHI '90*, 1990.
- [163] R. K. Took. *Surface Interaction: An Architecture and Formal Model for Separating Application and Interface*. D.Phil. thesis, University of York, 1990.
- [164] P. H. J. van Eijk, C. A. Vissers, and M. Diaz, editors. *The Formal Description Technique LOTOS*. North-Holland, Amsterdam, 1989.
- [165] P. Walsh, K. Lim, J. Long, and M. Carver. JSD and the design of user interface software. *Ergonomics*, 1989.
- [166] J. Whiteside and D. Wixon. Discussion: Improving human-computer interaction—a quest for cognitive science. In J. Carroll, editor, *Interfacing thought: cognitive aspects of human-computer interaction*. MIT press, 1987.
- [167] P. J. Whysall and J. A. McDermid. An approach to object oriented specification using Z. In *Proceedings of the fifth annual Z User Meeting*, December 1990.
- [168] P. J. Whysall and J. A. McDermid. Object oriented specification and refinement. In *Fourth Refinement Workshop of the BCS FACS Special Interest Group*, January 1991.
- [169] T. Winograd and F. Flores. *Understanding Computers and Cognition*. Addison-Wesley, 1987.
- [170] J. Woodcock and C. Morgan. Refinement of state-based concurrent systems. In D. Bjørner, C. Hoare, and H. Langmaack, editors, *VDM'90: VDM and Z—Formal Methods in Software Development*, volume 428 of *Lecture Notes in Computer Science*. Springer-Verlag, 1990.
- [171] J. C. Woodcock. Mathematics as a management tool: Proof rules for promotion. In *Proc. 6th Annual CSR Conference on Large Software Systems*, September 1989.
- [172] P. C. Wright and C. R. Roast. Abstraction and generalisation in the analysis of usability. In University of York, Department of Computer Science Technical Report, September 1990.
- [173] R. Young and J. Whittington. Using a knowledge analysis to predict conceptual errors in text-editor usage. In J. Chew and J. Whiteside, editors, *CHI'90 Conference Proceedings*, pages 91-97. Addison Wesley, 1990.

- 
- [174] R. M. Young and P. Barnard. The use of scenarios in human-computer interaction research: turbocharging the tortoise of cumulative science. In J. Carroll and P. Tanner, editors, *CHI and GI Conference Proceedings: Human Factors in Computing Systems and Graphic Interface*, pages 291–296. ACM, 1987.
- [175] R. M. Young, P. Barnard, T. Simon, and J. Whittington. How would your favourite user model cope with these scenarios? *SIGCHI Bulletin*, 20(4):51–55, 1989.

# Appendix A

## Use of the Z Notation

We make extensive use of the Z notation in this thesis. For the most part, we have adhered to the standard Z notation, as given by Spivey's standard Z reference manual [152]. In this appendix, we discuss some stylistic conventions that we have adopted, and we provide definitions of operators used in the body of the thesis but not defined there nor in the standard Z reference manual.

### Function definitions

When defining a function by predicates describing its effect on arguments in its domain, it is technically necessary that these axioms be bound by some universal quantifier over the domain elements. For example, in the Z reference manual, the projection functions on ordered pairs, are defined as below.

$[X, Y]$
$first : X \times Y \rightarrow X$
$second : X \times Y \rightarrow Y$
$\forall x : X; y : Y$
<ul style="list-style-type: none"><li><math>first(x, y) = x</math></li><li><math>second(x, y) = y</math></li></ul>

We have decided that in most cases the type information of the function is enough to allow the elimination of the universal quantification, leading to a slightly less cluttered presentation that is no less understandable. The projection functions would be rewritten as below.

$\begin{array}{l} \text{first} : X \times Y \rightarrow X \\ \text{second} : X \times Y \rightarrow Y \end{array}$
$\begin{array}{l} \text{first}(x, y) = x \\ \text{second}(x, y) = y \end{array}$

## Formatting quantifications

Many quantifications with bound variables—universal, existential, set comprehensions, mu-expressions and lambda expressions—are rather lengthy and it is always a challenge to present with maximal clarity and minimal bracketing. In this thesis, we have tried to present a consistent format to these quantifications. Where one would normally write such a universal quantification as

$$\forall \text{Decls} \mid P \bullet Q$$

we now write

$$\begin{array}{l} \forall \text{Decls} \\ \mid P \\ \bullet Q \end{array}$$

allowing indentation to replace the need for some bracketing.

As mentioned by Macdonald [103], it is common in Z usage to replace existential quantification by a **where** clause to increase readability. Hence, the expression

$$\begin{array}{l} \exists \text{Decls} \\ \bullet Q \end{array}$$

can be replaced by

$$\begin{array}{l} Q \\ \text{where Decls} \end{array}$$

## Operations on sequences

We will make use of some operations on sequences. The generic relations  $\text{—}$  (“is a prefix of”),  $\text{—}$  (“is a suffix of”) and  $\text{=}$  (“is contained in”) represent useful operations.  $\text{prefixes}$  is a function which generates the set of all prefixes of a given

sequence. The prefix relation `prefix_closed` is true if its argument is a set of sequences which is closed under the prefix relation, that is, the prefixes of every sequence in the set are also in the set.

[X]	
$\neg$	$\text{seq } X \leftrightarrow \text{seq } X$
$\neg$	$\text{seq } X \leftrightarrow \text{seq } X$
$\rightleftharpoons$	$\text{seq } X \leftrightarrow \text{seq } X$
	$\text{prefixes} : \text{seq } X \rightarrow \mathbf{P} \text{seq } X$
	$\text{prefix\_closed } \_ : \mathbf{P} \text{seq } X$
$xs \text{ --- } ys$	$\Leftrightarrow \exists zs : \text{seq } X \bullet zs \hat{\ } xs = ys$
$xs \text{ --- } ys$	$\Leftrightarrow \exists zs : \text{seq } X \bullet zs \hat{\ } xs = ys$
$xs \text{ --- } ys$	$\Leftrightarrow \exists zs, zs' : \text{seq } X \bullet zs \hat{\ } xs \hat{\ } zs' = ys$
$\text{prefixes } ys$	$= \{ zs : \text{seq } X \mid xs \text{ --- } ys \}$
$\text{prefix\_closed } XS$	$\Leftrightarrow \forall zs' : XS \bullet \text{prefixes } zs' \subseteq XS$

Technically, the definition of these operations as generic mandates that their use be indexed by the base set, so we would have to say  $xs \text{ --- } [Y] ys$  to indicate that  $xs$  is a prefix of  $ys$ , where both are sequences from the set  $Y$ . In practice, we will not indicate the base set when it is clear from the context.

## §/ (Distributed sequential composition)

The function `§/` sequentially composes a sequence of homogeneous relations (relations of the form  $X \leftrightarrow X$ ) to obtain one relation.

[X]	
$\text{§/}$	$\text{seq}(X \leftrightarrow X) \rightarrow (X \leftrightarrow X)$
$\text{§/}\langle \rangle$	$= \text{id } X$
$\text{§/}\langle R \rangle$	$= R$
$\text{§/}\langle \langle R \rangle \hat{\ } rs \rangle$	$= R \text{ §/ } rs$

## Sequence filtering

As described in Sufrin and He [158], we can extend the notion of domain and range restriction to act on sequences. These operators will behave like the familiar filtering operations in functional programming and used for trace semantics. The expression  $N \overset{\text{seq}}{\triangleleft} s$  will give a sequence consisting of the resequencing of the restriction of the domain of  $s$  to the set of natural numbers in  $N$ . Similarly, the expression  $s \overset{\text{seq}}{\triangleright} S$

will give a sequence consisting of the resequencing of the restriction of the range of  $s$  to the set of values in  $S$ .

$[X]$
$resequence : (\mathbf{N} \rightarrow X) \rightarrow \text{seq } X$ $sort : \mathbf{P} \mathbf{N} \rightarrow \text{seq } \mathbf{N}$ $increasing : \mathbf{P}(\text{seq } \mathbf{N})$ $\triangleleft^{\text{seq}} : (\mathbf{P} \mathbf{N} \times \text{seq } X) \rightarrow \text{seq } X$ $\triangleright^{\text{seq}} : (\text{seq } X \times \mathbf{P} X) \rightarrow \text{seq } X$
$\forall zs : \text{seq } X; N : \mathbf{P} \mathbf{N}; S : \mathbf{P} X \bullet$ $\quad N \triangleleft^{\text{seq}} s = resequence(N \triangleleft s)$ $\quad \wedge s \triangleright^{\text{seq}} S = resequence(s \triangleright S)$ $\forall f : \mathbf{N} \rightarrow X \bullet resequence(f) = sort(\text{dom } f) \ddagger f$ $\forall N : \mathbf{P} \mathbf{N}; s : \text{seq } X \bullet$ $\quad \text{ran } sort(N) = N$ $\quad \wedge sort(N) \in increasing$ $\quad \wedge s \in increasing \Leftrightarrow \forall i, j : \text{dom } s \bullet i < j \Rightarrow s(i) < s(j)$

For example,

$$\{1, 5, 3, 11\} \triangleleft^{\text{seq}} \langle a, b, c, d, e, f \rangle = \langle a, c, e \rangle$$

$$\langle a, b, c, d, e, f \rangle \triangleright^{\text{seq}} \{c, g, a\} = \langle a, c \rangle.$$

Some specific filtering on traces as suggested by Hoare [82] will be notationally convenient. We first describe channel filtering, which gives the sequence of messages communicated along a channel. We write  $t \downarrow CS$  to represent the restriction of trace  $t$  to events on channels in the set  $CS$ .

$$\left| \begin{array}{l} \downarrow : (\text{seq } Event \times \mathbf{P} \text{ ChannelID}) \text{seq } Event \\ \hline t \downarrow CS = t \ddagger \text{mess} \triangleright^{\text{seq}} C \end{array} \right.$$

## Folding

We define general folding operations *foldl* and *foldr* with their standard definition from functional programming [24].



$[X, Y]$
$foldl : ((X \times Y) \rightarrow X) \rightarrow X \rightarrow seq Y \rightarrow X$ $foldl1 : ((X \times X) \rightarrow X) \rightarrow seq X \rightarrow X$
$\forall \_ op \_ : (X \times Y) \rightarrow X; a : X; y : Y; ys : seq Y$ <ul style="list-style-type: none"> <li>• <math>( foldl op a \langle \rangle ) = a</math></li> <li>• <math>\wedge foldl op a y \hat{\ } ys = foldl op (a op y) ys</math></li> </ul>
$\forall \_ op \_ : (X \times X) \rightarrow X; xs : seq_1 X$ <ul style="list-style-type: none"> <li>• <math>foldl1 op xs = foldl op (head xs) (tail xs)</math></li> </ul>

Similar definitions of the folding operations *foldr* and *foldr1* can be defined.

## Interleaving

Interleaving of sequences is used to define some of the trace semantics for the agent model. We here define interleaving recursively as done by Hoare[82, p. 56].

$[X]$
$interleaves : seq X \leftrightarrow (seq X \times seq X)$
$\langle \rangle interleaves (t, t') \Leftrightarrow (t = \langle \rangle) \wedge (t' = \langle \rangle)$
$s interleaves (t, t') \Leftrightarrow s interleaves (t', t)$
$((x) \hat{\ } s) interleaves (t, t') \Leftrightarrow$ $( (t \neq \langle \rangle \wedge head(t) = x \wedge s interleaves(tail(t), t'))$ $\vee (t' \neq \langle \rangle \wedge head(t') = x \wedge s interleaves(t, tail(t'))))$

## Equivalence relations

An equivalence relation on a type  $X$  is relation which is reflexive, symmetric and transitive. We define the generic predicate  $equivalence[X]$  to be true when its argument is an equivalence relation on the type  $X$ .

$[X]$
$equivalence : (X \leftrightarrow X) \rightarrow Boolean$
$\forall \sim : X \leftrightarrow X \bullet equivalence \sim \Leftrightarrow$ $\forall p, q, r : X \bullet ( p \sim p$ $\wedge p \sim q \Rightarrow q \sim p$ $\wedge (p \sim q \wedge q \sim r) \Rightarrow p \sim r)$

As usual, in practice, we will leave out the base type indication when we use equivalence when it is clear from the context.

The set of equivalence classes derived from an equivalence relation  $\sim$  on the type  $X$  is represented by  $classes(\sim)$ .

$[X]$
$classes : (X \leftrightarrow X) \leftrightarrow PP X$
$dom\ classes = \{\sim : X \leftrightarrow X \mid \text{equivalence } \sim\}$
$\forall \sim : X \leftrightarrow X \mid \sim \in dom\ classes \bullet$ $classes(\sim) = \{ x : X \bullet \{ x' : X \mid x' \sim x \} \}$

# Appendix B

## Some theorems on the refined agent model

In this appendix, we provide the complete proofs for some theorems expressed in the body of this thesis. To assist in these proofs, we will need the following result.

Recall the definition of *compatible*, which relates states which agree over their common attributes.

$$\left| \begin{array}{l} \text{compatible} : \text{State} \leftrightarrow \text{State} \\ \hline (s_1, s_2) \in \text{compatible} \Leftrightarrow s_1 \oplus s_2 = s_1 \cup s_2 \end{array} \right.$$

The relation *compatible* is reflexive and commutative. A less common but important property of *compatible* is that if two states are compatible, and a third state is compatible with the union of the first two, then the third state is compatible with each of the other states. We state and prove this result as Lemma B.1.

### Lemma B.1

$$\begin{array}{l} \forall s, t, u : \text{State} \\ | \quad (s, t) \in \text{compatible} \\ \quad \wedge (s \cup t, u) \in \text{compatible} \\ \bullet \quad (t, u) \in \text{compatible} \\ \quad \wedge (s, u) \in \text{compatible} \end{array}$$

### PROOF OF LEMMA B.1:

Under the assumptions, we can pick  $s, t, u : \text{State}$  satisfying

$$\begin{array}{ll} (s \cup t) \oplus u = (s \cup t) \cup u & [\text{assumption and defn. of compatible}] \\ \Leftrightarrow (s \oplus t) \oplus u = (s \cup t) \cup u & [(s, t) \in \text{compatible}] \\ \Leftrightarrow s \oplus (t \oplus u) = s \cup (t \cup u) & [\text{accoc. of } \oplus \text{ and } \cup] \end{array}$$

By case analysis on  $A \subseteq \mathcal{A}$ , we can satisfy the inference

$$\left( \begin{array}{l} s \oplus t = s \cup t \\ \wedge s \oplus (t \oplus u) = s \cup (t \cup u) \end{array} \right) \Rightarrow (t, u) \in \text{compatible}$$

**Case 1:**  $A = \text{dom } s \cap \text{dom } t$

Since  $(s, t) \in \text{compatible}$ ,  $s$  and  $t$  agree on  $A$ , so

$$A \triangleleft (s \oplus (t \oplus u)) = A \triangleleft (t \oplus u)$$

$$A \triangleleft (s \cup (t \cup u)) = A \triangleleft (t \cup u)$$

Therefore,  $(t, u) \in \text{compatible}$  in this case.

**Case 2:**  $A = \text{dom } t - \text{dom } s$

In this case, we satisfy

$$A \triangleleft (s \oplus (t \oplus u)) = A \triangleleft (t \oplus u)$$

$$A \triangleleft (s \cup (t \cup u)) = A \triangleleft (t \cup u)$$

and so  $(t, u) \in \text{compatible}$  in this case.

**Case 3:**  $A = \text{dom } s - \text{dom } t$

In this case, we satisfy

$$A \triangleleft (t \oplus u) = A \triangleleft u$$

$$A \triangleleft (t \cup u) = A \triangleleft u,$$

and so  $(t, u) \in \text{compatible}$  in this case.

**Case 4:**  $A = \mathcal{A} - (\text{dom } s \cup \text{dom } t)$

In this case, we have

$$A \triangleleft (t \oplus u) = A \triangleleft u$$

$$A \triangleleft (t \cup u) = A \triangleleft u,$$

and so  $(t, u) \in \text{compatible}$  in this case.

Since these cases are exhaustive, we conclude that  $(t, u) \in \text{compatible}$ . By similar arguments, we can show that  $(s, u) \in \text{compatible}$ .

◇ END OF PROOF OF LEMMA B.1

**Theorem 5.1**

- $$\forall S1, S2, S3 : \text{Stateset}$$
- $$\begin{array}{l} | (S1, S2) \in \text{dom } \textit{join} \\ \wedge (\textit{join}(S1, S2), S3) \in \text{dom } \textit{join} \\ \bullet \textit{join}(\textit{join}(S1, S2), S3) = \textit{join}(S1, \textit{join}(S2, S3)) \end{array}$$

**PROOF OF THEOREM 5.1:**

The hypothesis of this theorem ensures that  $\textit{join}(\textit{join}(S1, S2), S3)$  exists. The following predicates are true under this hypothesis.

- $$\begin{array}{l} \forall a : S1.\textit{attrs} \cap S2.\textit{attrs} \\ \bullet S1.\textit{type}(a) = S2.\textit{type}(a) \\ \exists s_1 : S1.\textit{states}; s_2 : S2.\textit{states} \\ \bullet (s_1, s_2) \in \textit{compatible} \\ \forall a : \textit{join}(S1, S2).\textit{attrs} \cap S3.\textit{attrs} \\ \bullet \textit{join}(S1, S2).\textit{type}(a) = S3.\textit{type}(a) \\ \exists s_{12} : \textit{join}(S1, S2).\textit{states}; s_3 : S3.\textit{states} \\ \bullet (s_{12}, s_3) \in \textit{compatible} \end{array}$$

Under these assumptions, we must show that  $\textit{join}(S1, \textit{join}(S2, S3))$  also exists and is identical. The proof, therefore, is split into two parts.

First, to show that  $(S1, \textit{join}(S2, S3)) \in \text{dom } \textit{join}$ , we must satisfy the following four predicates.

$$\forall a : S2.\textit{attrs} \cap S3.\textit{attrs} \tag{B.1}$$

- $S2.\textit{type}(a) = S3.\textit{type}(a)$

$$\exists s_2 : S2.\textit{states}; s_3 : S3.\textit{states} \tag{B.2}$$

- $(s_2, s_3) \in \textit{compatible}$

$$\forall a : S1.\textit{attrs} \cap \textit{join}(S2, S3).\textit{attrs} \tag{B.3}$$

- $S1.\textit{type}(a) = \textit{join}(S2, S3).\textit{type}(a)$

$$\exists s_1 : S1.\textit{states}; s_{23} : \textit{join}(S2, S3).\textit{states} \tag{B.4}$$

- $(s_1, s_{23}) \in \textit{compatible}$

To prove predicate B.1, we begin with the assumption

- $$\begin{array}{l} \forall a : \textit{join}(S1, S2).\textit{attrs} \cap S3.\textit{attrs} \\ \bullet \textit{join}(S1, S2).\textit{type}(a) = S3.\textit{type}(a) \end{array}$$

By the definition of  $\textit{join}$ , this is equivalent to

- $$\begin{array}{l} \forall a : \textit{join}(S1, S2).\textit{attrs} \cap S3.\textit{attrs} \\ \bullet (S1.\textit{type} \cup S2.\textit{type})(a) = S3.\textit{type}(a) \end{array}$$

By the first condition on the definition of *join*, we know that  $S1.type$  and  $S2.type$  agree on common attributes, so the above predicate implies

$$\begin{aligned} \forall a : S2.attrs \cap S3.attrs \\ \bullet S2.type(a) = S3.type(a) \end{aligned}$$

as desired.

To prove predicate B.2, we begin with the assumption

$$\begin{aligned} \exists s_{12} : join(S1, S2).states; s_3 : S3.states \\ \bullet (s_{12}, s_3) \in compatible \end{aligned}$$

By the definition of  $join(S1, S2)$ , this implies

$$\begin{aligned} \exists s_1 : S1.states; s_2 : S2.states; s_3 : S3.states \\ \bullet (s_1, s_2) \in compatible \\ \wedge (s_1 \cup s_2, s_3) \in compatible \end{aligned}$$

By Lemma B.1, this implies

$$\begin{aligned} \exists s_2 : S2.states; s_3 : S3.states \\ \bullet (s_2, s_3) \in compatible \end{aligned}$$

as desired.

To prove predicate B.3 we must satisfy  $\forall a : S1.attrs \cap join(S2, S3).attrs$  By

$$\bullet S1.type(a) = join(S2, S3).type(a)$$

the definition of  $join(S2, S3)$ , we know

$$\begin{aligned} join(S2, S3).attrs &= S2.attrs \cup S3.attrs \\ join(S2, S3).type &= S2.type \cup S3.type \end{aligned}$$

We have two cases to investigate.

**Case 1:**  $a \in S2.attrs$

Since  $(S1, S2) \in \text{dom } join$ , we are done because of the assumption

$$\begin{aligned} \forall a : S1.attrs \cap S2.attrs \\ \bullet S1.type(a) = S2.type(a) \end{aligned}$$

**Case 2:**  $a \in (S3.attrs - S2.attrs)$

Since  $(join(S1, S2), S3) \in \text{dom } join$ , we know

$$\begin{aligned} \forall a : join(S1, S2).attrs \cap S3.attrs \\ \bullet join(S1, S2).type(a) = S3.type(a) \end{aligned}$$

and by the definition of *join* and the assumption that  $a \notin S2.attrs$ , we have

$$\begin{aligned} & \forall a : S1.attrs \cap S3.attrs \\ & \bullet S1.type(a) = S3.type(a) \end{aligned}$$

as desired.

To prove predicate B.4, we can use the assumptions to satisfy

$$\begin{aligned} & \exists s_1 : S1.states; s_2 : S2.states; s_3 : S3 \\ & \bullet (s_1, s_2) \in compatible \\ & \quad \wedge (s_1 \cup s_2, s_3) \in compatible \end{aligned}$$

which by Lemma B.1 and the definition of *join* gives us

$$\begin{aligned} & \exists s_1 : S1.states; s_{23} : join(S2, S3).states \\ & \bullet (s_1, s_{23}) \in compatible \end{aligned}$$

as desired.

We have now shown that  $join(S2, S3)$  and  $join(S1, join(S2, S3))$  are defined. We can now show the equality required by the theorem. Equality of the attribute sets and the type function relies on the associativity of  $\cup$ . To show equality of the state sets, we begin with

$$\begin{aligned} join(join(S1, S2), S3).states &= \{ s_{12} : join(S1, S2).states; s_3 : S3.states \\ & \quad | (s_{12}, s_3) \in compatible \\ & \bullet s_{12} \cup s_3 \} \end{aligned}$$

By the definition of *join* and Lemma B.1, this is equivalent to

$$\begin{aligned} join(join(S1, S2), S3).states &= \{ s_1.states; s_2 : S2.states; s_3 : S3.states \\ & \quad | (s_1, s_2) \in compatible \\ & \quad \wedge (s_2, s_3) \in compatible \\ & \quad \wedge (s_1 \cup s_2, s_3) \in compatible \\ & \quad \wedge (s_2, s_2 \cup s_3) \in compatible \\ & \bullet s_1 \cup s_2 \cup s_3 \} \end{aligned}$$

which is equivalent to  $join(S1, join(S2, S3)).states$ , as desired.

◇ END OF PROOF OF THEOREM 5.1

# Appendix C

## Detailed semantics for the agent language

In this appendix we provide a more detailed semantics for the agent language than provided in Chapter 5. The semantics is given in parts. We first define the overall semantic operator which takes whole agent expressions and maps them into a system of agents. The agent language mirrors the development of the agent model, and so we have separate sections of the language which individually treat the description of internal, communication and external specifications. Each section will have its own semantic operator which maps expressions of its language to specifications in the agent model.

### Agents in the interactive system

An interactive system is a mapping from agent identifiers to the set of agents in *Agent*. We introduce a set of possible agent identifiers.

$[AgentID]$

$IntSys == AgentID \leftrightarrow Agent$

The system semantic function,  $S[-]$ , takes an existing interactive system and an agent language expression and produces a new interactive system. The agent language description represents either the synchronization of existing agents, the interleaving of an existing agent with a 3-part description of a new agent (internal, external, communication specification), or a completely new 3-part description of an agent. The following is a BNF-like description of the agent language syntax.



Square brackets are used to indicate an item which is optional.

```

AgExp ::= agent AgentID           - synchronization
           synchronizes AgentIDList
           [with 3PartSpec]
           endagent AgentID
         | agent AgentID           - interleaving
           interleaves AgentIDList
           [with 3PartSpec]
           endagent AgentID
         | agent AgentID           - 3-part specification
           3PartSpec
           endagent AgentID

```

The system semantic function  $\mathcal{S}[\dots]$  is defined structurally over the elements in *AgExp*. For synchronized combination, the expression

```

agent A1
synchronizes AS
with Spec
endagent A1

```

maps the (fresh) agent identifier *A1* to the synchronous composition of the agents indicated by the sequence of (distinct) agent identifiers *AS*, if such a composition is allowed by *composeall<sub>sync</sub>*. This may then be interleaved with the agent defined by the 3-part specification *Spec*, according to the semantic operator  $\mathcal{A}g[\dots]$  discussed later.

$$\begin{array}{l}
\mathcal{S}[-] : (IntSys \times AgExp) \leftrightarrow IntSys \\
\forall A1 : AgentID; AS : seq_1 AgentID; \rho : IntSys \\
| ( A1 \notin \text{dom } \rho \\
\quad \wedge (AS \ ; \ \rho) \in \text{dom } composeall_{sync} ) \\
\bullet \mathcal{S} \left[ \begin{array}{l} \text{agent } A1 \\ \text{synchronizes } AS \\ \text{endagent } A1 \end{array} \right] = \\
\rho \oplus \{ A1 \mapsto composeall_{sync}(AS \ ; \ \rho) \} \\
\forall A1 : AgentID; AS : seq_1 AgentID; Spec : 3PartSpec \rho : IntSys \\
| ( A1 \notin \text{dom } \rho \\
\quad \wedge AS \ ; \ \rho \in \text{dom } composeall_{sync} \\
\quad \wedge (composeall_{sync}(AS \ ; \ \rho), Ag[ Spec ]) \in \text{dom } compose_{int} ) \\
\bullet \mathcal{S} \left[ \begin{array}{l} \text{agent } A1 \\ \text{synchronizes } AS \\ \text{with } Spec \\ \text{endagent } A1 \end{array} \right] = \\
\rho \oplus \{ A1 \mapsto compose_{int}(composeall_{sync}(AS \ ; \ \rho), Ag[ Spec ]) \}
\end{array}$$

Note that because of the associativity of  $compose_{sync}$  (Theorem 5.5, the order of the agent identifiers in  $AS$  does not matter.

For interleaved combination, the expression

**agent**  $A1$   
**interleaves**  $AS$   
**with**  $Spec$   
**endagent**  $A1$

maps the fresh identifier  $A1$  to the interleaved product of the known agent definitions in  $AS$  and the 3-part specification  $Spec$ , if given.

$$\begin{array}{l}
\forall A1 : AgentID; AS : seq_1 AgentID; \rho : IntSys \\
| ( A1 \notin \text{dom } \rho \\
\quad \wedge AS \sharp \rho \in \text{dom } \text{composeall}_{int}) \\
\bullet S \left[ \begin{array}{l} \text{agent } A1 \\ \text{interleaves } AS \\ \text{with } Spec \\ \text{endagent } A1 \end{array} \right] = \rho \oplus \{A1 \mapsto \text{composeall}_{int}(AS \sharp \rho)\} \\
\forall A1 : AgentID; AS : seq_1 AgentID; Spec : 3PartSpec; \rho : IntSys \\
| ( A1 \notin \text{dom } \rho \\
\quad \wedge (AS \sharp \rho) \sim Ag[Spec] \in \text{dom } \text{composeall}_{int}) \\
\bullet S \left[ \begin{array}{l} \text{agent } A1 \\ \text{interleaves } AS \\ \text{with } Spec \\ \text{endagent } A1 \end{array} \right] = \\
\rho \oplus \{A1 \mapsto \text{composeall}_{int}((AS \sharp \rho) \sim Ag[Spec])\}
\end{array}$$

Note that because of the associativity of  $\text{compose}_{int}$  (Theorem 5.4, the order of the agent identifiers in  $AS$  does not matter.

A stand alone specification of an agent,

```

agent A1
Spec
endagent A1

```

maps the (fresh) identifier  $A1$  to the agent  $Ag[Spec]$ .

$$\begin{array}{l}
\forall A1 : AgentID; Spec : 3PartSpec; \rho : IntSys \\
| ( A1 \notin \text{dom } \rho \\
\quad \wedge Spec \in \text{dom } Ag[-]) \\
\bullet S \left[ \begin{array}{l} \text{agent } A1 \\ Spec \\ \text{endagent } A1 \end{array} \right] = \rho \oplus \{A1 \mapsto Ag[Spec]\}
\end{array}$$

The three part specification of an agent is given by an internal, external and communication language.

```

3PartSpec ::= internal IExp
             communication CExp
             external EExp

```

The agent semantic operator,  $Ag[-]$ , is defined in terms of semantic operators for each of the sublanguages. The domain of  $Ag[-]$  is the set of combinations of internal, communication and external expressions which yield a valid agent description when they are mapped to their respective specifications in the model.

$$\begin{array}{l}
\mathit{Ag}[_] : 3\text{PartSpec} \leftrightarrow \text{Agent} \\
\mathcal{I}[_] : \text{IExp} \rightarrow \text{InternalSpec} \\
\mathcal{C}[_] : \text{CExp} \rightarrow \text{Communication} \\
\mathcal{E}[_] : \text{EExp} \rightarrow \text{ExternalSpec} \\
\hline
\text{dom } \mathit{Ag}[_] = \{ \text{IE} : \text{IExp}; \text{CE} : \text{CExp}; \text{EE} : \text{EExp}; \text{Agent} \\
\quad | \quad ( \theta \text{InternalSpec} = \mathcal{I}[_] \text{IE} ] \\
\quad \quad \wedge \theta \text{Communication} = \mathcal{C}[_] \text{CE} ] \\
\quad \quad \wedge \theta \text{ExternalSpec} = \mathcal{E}[_] \text{EE} ] ) \\
\quad \bullet \left( \begin{array}{l} \text{internal IE} \\ \text{communication CE} \\ \text{external EE} \end{array} \right) \} \\
\mathit{Ag} \left[ \begin{array}{l} \text{internal IE} \\ \text{communication CE} \\ \text{external EE} \end{array} \right] = \mu \text{Agent} \\
\quad \bullet \left( \theta \text{InternalSpec} = \mathcal{I}[_] \text{IE} ] \right. \\
\quad \quad \wedge \theta \text{Communication} = \mathcal{C}[_] \text{CE} ] \\
\quad \quad \left. \wedge \theta \text{ExternalSpec} = \mathcal{E}[_] \text{EE} ] \right)
\end{array}$$

In the following subsections, we will outline the definition of the semantic operators for each of these sublanguages.

## A language for internal specifications

The internal specification language is given by the following BNF-like description.

*ISpecLanguage* ::= **types** *Decls*  
**attributes** *AVmap*  
**invariant** *AVPred*  
**operations** *OpList*

The semantic function  $\mathcal{I}[_]$  can be described in terms of smaller semantic functions on the different parts of *ISpecLanguage*.

*Decls* is a list of type declarations and constructions local to the agent definition. In the declaration, we can build up type definitions from basic sets, exactly as is done in Z. For example, we could write the type definitions for the window.

[*ICON*]  
 $x_{\max}, y_{\max} : \mathbb{N}$   
 $\text{PIXEL} == 0 \dots x_{\max} \times 0 \dots y_{\max}$

There is a semantic function which produces the sets in the valuespace  $\mathcal{V}$  similar to the *Carrier* function in Spivey's semantics. We will also allow the definition of global types outside the definition of any agent, as is allowed in Z.

*AVmap* is a list of the attributes for which each state of the agent has a value. The *AVmap* is written much the same way that components declarations are written in a standard Z schema, so we could write expressions like the following in the *attributes* section of an agent.

```

icon : ICON
iconpos, winpos, winsize : PIXEL
status : open | closed

```

These state declarations are mapped by the semantic function `type[ _ ]` to obtain the *type* information for the agent.

*AVPred* is a predicate on the values of the attributes which must be satisfied by all states of the agent. We allow the same kinds of predicate expressions as is possible in Z. The semantic function `state[ _ ]` takes the *attributes* and the *invariants* information and yields the set of possible states for the agent.

*OpList* is a list describing the possible transitions for the agent. The state transitions are given pre- and postcondition semantics, along with an explicit framing conditions which lists the attributes that can be changed by the operation. Each operation definition gives a family of state transitions indexed by a message identifier along with its typed parameters. The window moving operation for an open window would be defined as follows.

```

operations
  move_open(newpos : PIXEL)
  changes winpos
  pre status = open
  post winpos' = newpos

```

The BNF description of *OpList* is given below. Square brackets are used to indicate an optional occurrence of an item.

*[MessageID]*

```

OpList := MessageID(TypedParams)
  [changes (AttribList)]
  [pre (BeforePred)]
  [post (AfterPred)]
  [OpList]

```

The operation has a name taken from the set *MessageID*. A typed parameter list, of the same form as used to declare the state attributes, lists the generic arguments for the operation. A message is formed by instantiating each argument with an

actual value. For example, a message generated from the definition above would be *move\_open*((0,0)). The **changes** clause provides an explicit framing condition for the operation definition. The precondition is a predicate on the state attributes and the arguments. The postcondition is a predicate on the before and after value of state attributes and the arguments. The values of attributes before and after the state transition are distinguished by undashed and dashed attribute identifier names, as is the convention in Z. Together, the framing, pre- and postconditions define a state transition.

The semantic function  $\text{state}[\_]$  takes a list of operations in *OpList* and creates the *operations* function for the agent. Each single operation defined as above yields a family of state transition relations, indexed by the possible messages.

$$\mid \text{state}[\_]: \text{OpList} \rightarrow (\text{Message} \leftrightarrow (\text{State} \leftrightarrow \text{State}))$$

## A language for communication specification

The communication specification is straightforward, since it simply lists the input and output channels for the agent along with the messages that can be passed along those channels. Synchronized channels arise from synchronous combination of existing agents, it is not possible to explicitly declare a channel as synchronized in the notation. The burden rests with the designer to ensure channel fidelity, i.e., the specification cannot have the same agent using the same channel as input and output, nor can a synchronized channel be an input or output channel for another agent. The internal messages are also declared explicitly as associated to the  $\tau$  channel. A BNF description of the syntax for communication is given below.

$$\begin{aligned} \text{CSpecLanguage} & ::= \{ \text{inputs } \text{TypedChannellist} \\ & \quad \mid \{ \text{outputs } \text{TypedChannellist} \\ & \quad \mid \{ \tau : \text{Messagelist} \} \\ \text{TypedChannellist} & ::= \text{ChannelID} : \text{Messagelist} [\text{TypedChannellist}] \\ \text{Messagelist} & ::= \text{MessageID}(\text{TypedParameters}) [, \text{Messagelist}] \end{aligned}$$

As with the definition of **operations**  $\_$ , the set of messages is generated by instantiating all of the parameters in the *TypedParameters* list with values according to their types. The semantic function  $\mathcal{C}[\_]$  takes an element in *CSpecLanguage* and produces the communications specification.

$$\mid \mathcal{C}[\_]: \text{CSpecLanguage} \leftrightarrow \text{Communication}$$

## A language for external specifications

The external specification language is developed in two parts as described in Section 5.3. In this section we will give the semantics for the constructive language

for generating external specifications of agents. First, we recall the definition of an external specification as given in Chapter 5.

$\begin{array}{l} \text{ExternalSpec} \\ \text{alphabet} : \mathbf{P Event} \\ \text{traces} : \mathbf{P seq alphabet} \\ \hline \langle \rangle \in \text{traces} \\ \text{prefix\_closed}(\text{traces}) \end{array}$
---

The syntax for the constructive language follows very closely the syntax for CSP given by Hoare [82].

$ConEEzp ::= \text{stop}\langle\langle \mathbf{P Event} \rangle\rangle$	- deadlock
$\quad   \text{run}\langle\langle \mathbf{P Event} \rangle\rangle$	- the total behaviour
$\quad   \text{skip}\langle\langle \mathbf{P Event} \rangle\rangle$	- successful termination
$\quad   \langle\langle \text{Event} \rangle\rangle \rightarrow ConEEzp$	- prefix composition
$\quad   ConEEzp \square ConEEzp$	- choice composition
$\quad   ConEEzp ; ConEEzp$	- sequential composition
$\quad   ConEEzp \parallel ConEEzp$	- synchronous parallel composition
$\quad   ConEEzp \# ConEEzp$	- asynchronous parallel composition
$\quad   f(ConEEzp)$	- process relabelling
$\quad   \mu X : \langle\langle \mathbf{P Event} \rangle\rangle \bullet F(X)$	- guarded recursion

We can now give the definition of each of these syntactic constructions in terms of the external specification model. We will give two versions of the semantics for each construct. The first will be a direct denotational definition in terms of the model for external specifications. These definitions are very close to those given by Hoare. The mapping from the syntactic domain  $ConEEzp$  to the semantic domain  $ExternalSpec$  is given by the function  $\mathcal{E}_{con}[-]$ , which can be further broken down into functions for the alphabet,  $\alpha[-]$  and trace set,  $\mathcal{T}[-]$ .

$\mathcal{E}_{con}[-] : ConEEzp \rightarrow ExternalSpec$
$\alpha[-] : ConEEzp \rightarrow \mathbf{P Event}$
$\mathcal{T}[-] : ConEEzp \rightarrow \mathbf{P seq Event}$
$(\mathcal{E}_{con}[ES]).\text{alphabet} = \alpha[ES]$
$(\mathcal{E}_{con}[ES]).\text{traces} = \mathcal{T}[ES]$

The second version of semantics will be an operational semantics which may make it easier to understand the functionality of the different constructs. This operational semantics is given by inference rules in the manner suggested by Plotkin [130] and is the standard semantics usually given for a process algebra in the CCS

family [109, 110]. We will not extend this operational semantics to show that it matches some intuitive equational theory on expressions in *ConEEzp*, as our purpose for presenting it is only to increase understanding.

### Deadlock

$\text{stop}_A$  represents deadlock in which participation in any events from the alphabet is prohibited.

$$\left\{ \begin{array}{l} \text{stop}_- : \mathbf{P}_1 \text{Event} \rightarrow \text{ConEEzp} \\ \alpha[\text{stop}_A] = A \\ \mathcal{T}[\text{stop}_A] = \{ \langle \rangle \} \end{array} \right.$$

There is no inference rule to describe the action of  $\text{stop}_A$  since it has no action.

### Total behaviour

Given a set of events  $A$ , the expression  $\text{run}_A$  denotes the external specification with alphabet  $A$  which can participate in any sequence of events from  $A$ .

$$\left\{ \begin{array}{l} \text{run}_- : \mathbf{P}_1 \text{Event} \rightarrow \text{ConEEzp} \\ \alpha[\text{run}_A] = A \\ \mathcal{T}[\text{run}_A] = \text{seq } A \end{array} \right.$$

Operationally,  $\text{run}_A$  participates in any event in  $A$  and then continues to behave as  $\text{run}_A$ .

$$\frac{e \in A}{\text{run}_A \xrightarrow{e} \text{run}_A}$$

### Successful termination

We introduce a primitive construct,  $\text{skip}_A$  to represent the successful termination of an external specification with alphabet  $A$ . To do this, we must also introduce a special internal event,  $\surd$  (read “tick”) which signals the termination. This event carries no message of significance, so we label it the *null* message. It occurs along the internal channel  $\tau$ . Furthermore, we constrain the external specifications set so that the  $\surd$  event can only appear at the end of a trace.



$null : Message$
$\checkmark : Event$
$skip\_ : P_1 Event \rightarrow ConEEExp$
$\checkmark.channel = \tau$
$\checkmark.message = null$
$\forall ES : ExternalSpec; t : ES.traces$
• $\checkmark \notin \text{ran}(front(t))$
$\alpha[skip_A] = A \cup \{ \checkmark \}$
$\mathcal{T}[skip_A] = \{ \langle \rangle, \langle \checkmark \rangle \}$

The operational semantics of  $skip_A$  is given by one simple inference rule.  $skip_A$  can participate in the event  $\checkmark$  after which it deadlocks.

$$\frac{}{skip_A \xrightarrow{\checkmark} stop_A}$$

### Prefix composition

$e \rightarrow P$  first engages in the event  $e$  and then behaves like  $P$ . The constraint on this construction is that  $e$  must already be in the alphabet of  $P$ .

$\dots \rightarrow \dots : (Event \times ConEEExp) \rightarrow ConEEExp$
$\text{dom}(\dots \rightarrow \dots) = \{ e : Event; P : ConEEExp$
$e \in \alpha[P]$
• $(e, P) \}$
$\alpha[e \rightarrow P] = \alpha[P]$
$\mathcal{T}[e \rightarrow P] = \{ \langle \rangle \} \cup \{ t : \mathcal{T}[P] \bullet \langle e \rangle \sim t \}$

The operational semantics for prefix composition is also covered by one simple inference rule.

$$\frac{e \in \alpha[P]}{e \rightarrow P \xrightarrow{e} P}$$

### Choice composition

Choice between two external specifications, written  $P \square Q$ , indicates that the behaviour can either proceed as described by  $P$  or by  $Q$ . The choice is made by

the first event in which  $P \square Q$  participates. Hoare refers to this choice operator as deterministic (or external) choice and distinguishes it from a nondeterministic choice,  $P \sqcap Q$ . In deterministic choice, the choice can be resolved (externally) by the environment which interacts with  $P \square Q$ , whereas in nondeterministic choice the environment can have no effect. In the traces model, no distinction can be made between these processes, and so we only describe external choice.

$$\left| \begin{array}{l} \_ \square \_ : (ConEEExp \times ConEEExp) \rightarrow ConEEExp \\ \alpha[ P \square Q ] = \alpha[ P ] \cup \alpha[ Q ] \\ \mathcal{T}[ P \square Q ] = \mathcal{T}[ P ] \cup \mathcal{T}[ Q ] \end{array} \right.$$

The operational semantics is given by two inference rules, indicating that  $P \square Q$  can proceed if one of  $P$  or  $Q$  or both can proceed. In the case where both can proceed, the choice is nondeterministic.

$$\frac{P \rightsquigarrow P'}{P \square Q \rightsquigarrow P'}$$

$$\frac{Q \rightsquigarrow Q'}{P \square Q \rightsquigarrow Q'}$$

### Sequential composition

$P ; Q$  behaves like  $P$  until successful termination, marked by participation in the special event  $\surd$ . After successful termination, it behaves like  $Q$ . Unlike CSP, we do not hide participation in the event  $\surd$  in the definition of the external specification.

$$\left| \begin{array}{l} \_ ; \_ : (ConEEExp \times ConEEExp) \rightarrow ConEEExp \\ \alpha[ P ; Q ] = \alpha[ P ] \cup \alpha[ Q ] \\ \mathcal{T}[ P ; Q ] = \mathcal{T}[ P ] \\ \cup \\ \{ t : \mathcal{T}[ P ], t' : \mathcal{T}[ Q ] \mid last(t) = \surd \bullet t \wedge t' \} \end{array} \right.$$

Operationally, if  $P$  can successfully terminate by engaging in the event  $\surd$ , then  $P ; Q$  can proceed to behave as  $Q$ .

$$\frac{P \rightsquigarrow \mathbf{stop}_A}{P ; Q \rightsquigarrow Q}$$

### Synchronous parallel composition

There are two versions of parallel composition which we allow. The first, written as  $P \parallel Q$  demands that  $P$  and  $Q$  synchronize on participation in events of their common alphabet. In the model description, the alphabets are combined. The traces of the synchronized combination are those which when filtered by the alphabet of  $P$  (respectively,  $Q$ ) are a legal trace of  $P$  (respectively,  $Q$ ).

$$\left| \begin{array}{l} \_ \_ : (ConEEzp \times ConEEzp) \rightarrow ConEEzp \\ \alpha[P \parallel Q] = \alpha[P] \cup \alpha[Q] \\ \mathcal{T}[P \parallel Q] = \{ t : \text{seq } \alpha[P \parallel Q] \\ \quad | ( t \stackrel{\text{seq}}{\triangleright} \alpha[P] \in \mathcal{T}[P] \\ \quad \quad \wedge t \stackrel{\text{seq}}{\triangleright} \alpha[Q] \in \mathcal{T}[Q] ) \\ \quad \bullet t \} \end{array} \right.$$

The operational semantics for  $P \parallel Q$  indicates that both  $P$  and  $Q$  can independently participate in events private to themselves, but both must evolve simultaneously on events which they share. Three inference rules sum this behaviour up.

$$\frac{P \rightsquigarrow P'}{P \parallel Q \rightsquigarrow P' \parallel Q} \quad [ e \in \text{alphabet}P - \text{alphabet}Q ]$$

$$\frac{Q \rightsquigarrow Q'}{P \parallel Q \rightsquigarrow P \parallel Q'} \quad [ e \in Q.\text{alphabet} - P.\text{alphabet} ]$$

$$\frac{P \rightsquigarrow P' \quad Q \rightsquigarrow Q'}{P \parallel Q \rightsquigarrow P' \parallel Q'} \quad [ e \in P.\text{alphabet} \cap Q.\text{alphabet} ]$$

### Asynchronous parallel composition

The other version of parallel composition is asynchronous, which we write as  $P \parallel\!\!\! \parallel Q$ . Both  $P$  and  $Q$  can evolve independently, regardless of whether they share events. The alphabet is again the union of the component alphabets. Traces are obtained by interleaving traces of the components.

$$\left| \begin{array}{l} \_ \_ : (ConEEzp \times ConEEzp) \rightarrow ConEEzp \\ \alpha[P \parallel\!\!\! \parallel Q] = \alpha[P] \cup \alpha[Q] \\ \mathcal{T}[P \parallel\!\!\! \parallel Q] = \{ s : \text{seq } \alpha[P \parallel\!\!\! \parallel Q] \\ \quad | \exists t : \mathcal{T}[P]; t' : \mathcal{T}[Q] \bullet s \text{ interleaves } (t, t') \\ \quad \bullet s \} \end{array} \right.$$

The relation *interleaves* is defined in Appendix A.

Operationally,  $P \parallel Q$  is described by two rules similar to the first two rules of  $P \parallel Q$ .

$$\frac{P \xrightarrow{\sim} P'}{P \parallel Q \xrightarrow{\sim} P' \parallel Q}$$

$$\frac{Q \xrightarrow{\sim} Q'}{P \parallel Q \xrightarrow{\sim} P \parallel Q'}$$

### Process relabelling

Sometimes it will be convenient to identify an external specification with a previous one with an appropriate change in the event names. The restriction on such a relabelling is that the mapping to new event names be injective which is total when restricted to the alphabet of the original external specification, so that each old event name is mapped to a unique new event name. In addition,  $\surd$  must be mapped to itself. The new external specification then behaves exactly as the old, with the new event names substituted for the old.

$$\left. \begin{array}{l} \_(-) : ((Event \leftrightarrow Event) \times ConEEExp) \leftrightarrow ConEEExp \\ \text{dom } \_(-) = \{ f : Event \leftrightarrow Event; ES : ConEEExp \\ \quad | ( f \in \alpha[ ES ] \leftrightarrow Event \\ \quad \quad \wedge f(\surd) = \surd ) \\ \quad \bullet (f, ES) \} \\ \alpha[ f(P) ] = f(\alpha[ P ]) \\ \mathcal{T}[ f(P) ] = \{ t : \mathcal{T}[ P ] \bullet t ; f \} \end{array} \right|$$

Note that in the definition of  $\mathcal{T}[ f(P) ]$ , we use the symbol  $;$  to represent the standard Z forward functional composition, not sequential composition of external specifications.

### Recursion

We will allow a simple form of recursion which is uniquely defined using a partial order semantics (see Hoare [82, Section 2.8] and Stoy [156]). The type *ConEEExp* can be considered a complete partial order using the following ordering relation.

$$\left. \begin{array}{l} \_ \sqsubseteq \_ : ConEEExp \leftrightarrow ConEEExp \\ ES \sqsubseteq ES' \Leftrightarrow ( \alpha[ ES ] = \alpha[ ES' ] \\ \quad \wedge \mathcal{T}[ ES ] \subseteq \mathcal{T}[ ES' ] ) \end{array} \right|$$

In this partial order, we can define the least upper bound (*lub*) of two constructed external specifications with the same alphabet by forming the trace set from the union of the individual trace sets. We can extend this to obtain lubs for *chains* of specifications as well. Therefore, *ConEEzp* is a complete partial order (with respect to the alphabet  $A$ ), on which all guarded continuous functions  $F : \text{ConEEzp} \rightarrow \text{ConEEzp}$  have a unique fixed point solution to the recursive equation  $\mu X : A \bullet F(X)$ , which is the lub of the infinite chain of successive applications of  $F$  to  $\text{stop}_A$ , the bottom element in the complete partial order. We have purposely designed *ConEEzp* so that all of the constructors are guarded and continuous, so the recursive equation will always have a unique solution.

$$\begin{array}{l}
 \text{chain} : \mathbf{P}(\text{seq } \text{ConEEzp}) \\
 \text{lub} : \text{seq } \text{ConEEzp} \leftrightarrow \text{ConEEzp} \\
 \mu X : \_ \bullet \_ (X) : (\mathbf{P}_1 \text{Event} \times (\text{ConEEzp} \rightarrow \text{ConEEzp})) \rightarrow \text{ConEEzp} \\
 \hline
 \text{ESS} \in \text{chain} \Leftrightarrow \forall i : 1..(\# \text{chain} - 1) \bullet \text{ESS}(i) \sqsubseteq \text{ESS}(i + 1) \\
 \text{dom } \text{lub} = \text{chain} \\
 \alpha[ \text{lub}(\text{ESS}) ] = \alpha[ \text{ESS}(1) ] \\
 \mathcal{T}[ \text{lub}(\text{ESS}) ] = \bigcup_{1.. \# \text{ESS}} \mathcal{T}[ \text{ESS} ] \\
 \alpha[ \mu X : A \bullet F(X) ] = A \\
 \mathcal{T}[ \mu X : A \bullet F(X) ] = \mathcal{T}[ \text{lub}(\langle F^0(\text{stop}_A), F^1(\text{stop}_A), \dots \rangle) ]
 \end{array}$$

Though the definition of recursion depends on the alphabet, in practice it is not indicated when the context of use makes clear what alphabet is intended.