Developing A Reflective Model of Collaborative Systems

Paul Dourish

Rank Xerox Research Centre—Cambridge Laboratory (EuroPARC)
61 Regent Street, Cambridge, CB2 1AB, United Kingdom
and
Dept. of Computer Science, University College, London

dourish@europarc.xerox.com

Abstract

Recent years have seen a shift in perception of the nature of HCI and interactive systems. As interface work has increasingly become a focus of attention for the social sciences, we have expanded our appreciation of the importance of issues such as work practice, adaptation and evolution in interactive systems.

The reorientation in our view of interactive systems has been accompanied by a call for a new model of design centred around user needs and participation. This paper argues that a new process of design is not enough, and that the new view necessitates a similar reorientation in the *structure* of the systems we build. It outlines some requirements for systems which support a deeper conception of interaction, and argues that the traditional system design techniques are not suited to creating such systems. Finally, using examples from ongoing work in the design of an open toolkit for collaborative applications, it illustrates how the principles of computational reflection and metaobject protocols can lead us towards a new model based on open abstraction which holds great promise in addressing these issues.

Categories and Subject Descriptors: C.2.4 [**Computer-Communication Networks**]: Distributed Systems—*distributed applications*; D.2.2 [**Software Engineering**]: Tools and Techniques—*user interfaces*; D.2.10 [**Software Engineering**]: Design—*methodologies*; H.1.10 [**Models and Principles**]: General.

General Terms: Design**.**

Addition Key Words and Phrases**:** Collaborative applications, system architecture, computational reflection, metaobject protocol, open implementations.

1 Introduction

The last ten years or so have seen a remarkable shift in perspectives on the design, evaluation and use of interactive systems. The field of HCI has moved from being a relatively minor component of software engineering to being the focus of attention for researchers from a variety of disciplines, including psychology and social science. Studies and investigations from these perspectives have led to a gradual evolution in our conception of "the interface" and of computer-based work in general. As a result, HCI has increasingly come to concern itself not just with the *mechanism* of the interface, but with a range of related issues concerning the context in which interactive systems are used.

1.1 Studies of Work at the Interface

To ground discussion of this new view of interactive systems development, I will discuss three areas of research which have informed it: the customisation of interactive systems; their embedding within a social organisation; and the co-adaptation of systems and work practices.

1.1.1 Customisation

Customisation and adaptation of computer systems have been studied in a variety of contexts. Trigg *et al* [1987] studied adaptation in the Notecards hypertext system. They described four aspects of adaptability which could allow a tool to be used in different application areas by users with different working styles. These were: *flexibility* (providing generic, reusable objects and behaviours); *parameterisability* (offering a range of alternative behaviours that users could select); *integrability* (linking with other applications in the environment); and *tailorability* (allowing users to make changes to the system itself). Their work showed how an adaptable system could be applied widely, essentially serving as an infrastructure within which a variety of information management applications could be generated.

MacLean *et al* [1990] were also concerned with customisation in the Buttons systems. Buttons are graphical on-screen objects which encapsulate behaviour; they can be incorporated into on-line documents and sent through electronic mail. Buttons support multiple levels of customisation. At the simplest level, their name, placement or graphical appearance can be customised. More advanced users can change explicit parameters to customise them for particular needs; and more radical changes in their behaviour can be made by modifying the Lisp code that they contain. Buttons were designed around these levels of tailorability to flatten the learning curve for interactive systems. Establishing a "tailoring culture", in which customisations are encouraged and shared, was also crucial to the investigation.

Sharing customisations (and customisation activity) has been a focus of attention for other researchers. Mackay [1990b; 1991] studied the patterns of sharing customisations (in the form of configuration files, either in whole or in part) in technical organisations. She found that a remarkable amount of customisation was performed (or shared) in groups rather than being a private activity—even when the software being customised was single-user. Nardi and Miller [1991] showed that similar collaborative activity goes on in the notionally "single-user" activity of programming and debugging spreadsheets.

Customisation is clearly a widespread and important phenomenon. The studies highlight the importance of designing systems with an understanding of how such flexibility can be harnessed and exploited.

1.1.2 Social Aspects of System Use

Ethnographic studies of work practice and technology have pointed to the strong social elements present in apparently individual use of technology. Heath and Luff [1991], studying the activities in a control room of London's underground transport system, observed a range of complex practices employed by the controllers. In particular, they showed how controllers would *peripherally monitor* each other's activities in order to coordinate their own actions with those of their colleagues. Heath and Luff further observed that individuals would quietly offer commentaries on their own activities, specifically so that these might be used by colleagues in their own work. Others uses of peripheral monitoring and anticipation in similar settings have been observed by Filippi and Theureau [1993] and Suchman [1992].

Harper *et al* [1991] studied a different setting—in this case, the control rooms of air traffic control centres. Again, they observed that apparently individualistic activity is frequently designed to facilitate coordination between individuals. Their investigations pointed to the role that physical artefacts can play in this process. The "flight strips"—used by air traffic controllers to record details of flights currently in their managed air space—are notionally designed to record information for an individual. However, practices emerge around such artefacts in which they become the focus of coordination activity.

Suchman [1987] discussed these issues more generally, concentrating in particular on the situatedness of technology and the use of technology. Technological interactions are embedded within social and organisational situations, and their influence must be taken into account when attempting to analyse or predict aspects of system interaction. These studies all call into question the extent to which interactions and activity can be predicted from an external viewpoint; and yet such predictions are at the basis of the traditional system development approach.

1.1.3 Coadaptation of Systems and Practices

The third aspect, drawing together elements from the previous two, concerns the longer-term mutual evolution of systems and work practices. Mackay [1990a] presented a variety of case studies illustrating this pattern, some of which were discussed above. Perhaps of most interest here is one study on Information Lens [Malone et al. 1988; Mackay 1989]. Lens is a system designed to filter and sort electronic mail. Investigations of use revealed that users adapted—or even *subverted*—features of Information Lens to match it to their working styles. This in turn led to a change in the development strategy, which incorporated and developed the mechanisms that the users had created. This is a spiralling process, based on the mutual interaction of system development and patterns of use—in Mackay's terms, a *coadaptive* phenomenon.

Echoes of this unanticipated interaction between technology and working practice are to be found in Sproull and Kiesler's [1991] work on the organisational impacts of electronic communication. Such technologies are often construed as "electronic office memos", and subject to the same rules and procedures. However, the experience is typically that aspects of electronic mail—in particular, rapid turn-around and the emergence of self-organising interest groups—results in a very different pattern of use, often much to the surprise of the organisation. As patterns become better established, they affect the further development of the technology, and so the coadaptive spiral continues.

1.2 A Revised View of Interactive Systems

There is a set of general issues underlying the particular concerns raised in all these studies. Together, they form part of wider reappraisal of the nature of computer-based work. They encourage us to look beyond widget design in understanding interactive systems. For instance, focussing on the development of the notion of user interface, Grudin [1993] has shown how it extends beyond the computer itself, incorporating the wider social and organisation issues arising from the use of computer systems. Elsewhere [Grudin 1990] he has placed these developments within a historical framework of changing concerns in interface design. At the micro-level, these sentiments are echoed by Bowers and Rodden [1993], who found the same factors at work in a very specific setting—the installation of a large network of CSCW systems in a particular observation.

Studies of customisation, of the social nature of computer-based work, and of the coevolution of systems and work practices, all address issues in the organisation of computer-based work practices. They highlight the strong relationship between these practices and the social organisation of work generally. Computer-based work is highly socially organised. As a result, these studies place an emphasis on variability in use; they point out the complexity of the relationship between the general pattern of use and the details of particular activities.

It is not surprising, then, that studies such as these should result in a reorientation of our view of the fundamental nature of interactive systems; a reorientation which takes into account this expanded understanding of computer usage. Indeed, such a reorientation has been taking place. It moves away from a view of systems as fixed, "black-box" artifacts which could be objectively studied and assessed, towards a new view of "systems-in-use" which acknowledges the influence of these other elements. In the new view, systems are *situated* within particular organisations and practices; they are *dynamic*, placing greater importance in the study of patterns of use and the cycle of software adoption; and they *evolve*, with specific working practices and behaviours emerging around the interactive system, while at the same time the system itself is tailored to particular working styles. In other words, systems must be able to support the variability exposed by investigations like those discussed above.

This shift in perspectives has largely addressed interactive systems *as they appear to the user*, in particular settings. Indeed, it has been part of a general trend towards participative or user-centred approaches to system building. However, it raises other important issues, particularly for the designers and implementors of interactive systems; issues which reach below the interface.

In this paper, I want to concentrate on the *structure* of the systems we create. I believe that the implications of our reconception of HCI forces us to reconsider our notions of system-building. This means changes not only to the *process* of design, but also—critically—to the *artefacts* of design. In particular, I will show how fundamental mechanisms that we use in constructing software systems do not support the design of interactive systems which change and evolve. Drawing from work done in other areas of systems development, I will show how the principles of *computational reflection* lead to a more open model of systems development, with revisability as a key element. Using examples from ongoing work in the design of an open CSCW toolkit, I will illustrate how this approach tackles a number of existing flexibility problems.

Our starting point for this investigation, then, is a question—what does the shift in our view of interactive systems imply for the nature of the systems that we design?

2 Implications for Interactive Systems Design

Our increasing appreciation of the issues discussed above has been accompanied by a call for a form of design which is more sensitive to them—a form of design oriented specifically around user needs and user involvement (see, for example, [Norman and Draper 1986; Ehn 1988;Bødker and Grønbæk 1991]). Following on from this, Grudin [1991] has highlighted conflicts between the "user-centred" approach to design and the traditional systems development process.

Taking this as a starting point, I want to explore how we can construct systems which support the sorts of practices outlined above. Principally, this involves looking at technical implications of the new approach for interactive systems, and then working towards an architectural model which addresses them. So if we generalise some aspects of the "systems-in-use" perspective, we can identify two sets of consequences for their design, concerned with the cycle of design, and creating evolving systems.

2.1 The Cycle of Design

The first of these is a reconsideration of the cycle of design. In the traditional "waterfall" model of software engineering [Royce 1970], the "design" of the system happens at a fixed point—after specification and before implementation. More recent models, such as those of Boehm [1988] or Booch [1991], adopt a more iterative, prototype-based approach. Here, "design" is distributed more evenly through the development process, not concentrated in one place.

Despite their differences, there is at least one point of fundamental agreement among all these models. They state that, at some point, a product is *delivered* to a user community; at which point, for that revision of the software, the design process is over.

This is an assumption that we must reconsider in the light of the "systems-in-use" model. When we take this perspective, we are forced to ask questions such as: "When does design happen?"; "Who does the design?"; and "When does it stop?". When we look at an interactive system as an evolving artifact in use, it follows that the process of design does *not* end with the delivery of the system to some community of users. Instead, it continues as they use and adapt the system. This leads to a second and more focussed set of concerns for system developers, looking at how systems are structured, constructed and delivered.

2.2 Creating Evolving Systems

The developer of an interactive system must not only be concerned with the traditional issues of system design, but also with the issues of providing a system which is amenable to evolution and adaptation. We can focus on three particular aspects of this problem:

1. *Open infrastructures*. From the perspective of systems-in-use, we begin to see delivered systems as not being closed and static, but rather as infrastructures for further specialisation, refinement and end-user design. They provide a framework within which users can change and adapt the basic system to their own patterns of usage. The system developer is concerned with appropriate openness within the system and ensuring that it lends itself to these adaptations. (The nature of "appropriateness" will be considered in more depth later; for the moment, we can consider "appropriate" as being "lending itself to appropriation".) In addition, extending the model of customisation, we must consider the ways in which the system can lend itself to customisation of *function* (the "semantics" of manipulating information in an interactive system) as well as *presentation* (surface-level issues of views and interaction).

2. *Dynamic and reactive systems*. When we think of user behaviour changing over time, then we must consider how the interactive system will support and respond to these changes. From this point of view, systems need to be designed to dynamically react to patterns of use and activity. The system's response must be situated in the same sense as is the user's activity. User's needs are highly dependent on many contextual factors, such as patterns of activity and changes in configuration. Similarly, system behaviour should react to contexts of use.

3. *Adaptive and evolving systems*. In addition to the "immediate" view of system reactivity, the developer must also be concerned with the longer-term view of the evolution of the system. Research on customisation and coadaptivity shows that this evolution has its roots in the social aspects of work and is enabled, in part, through the sharing of customisations. This implies that it is important to address issues such as the nature of the customisation mechanisms, and the means by which they can be distributed, shared and themselves evolve over time—support for Maclean et al.'s "tailoring culture".

   "Tailoring" in this context does not just imply the ability to make changes. It is also crucial that changes and adaptations be separated from core functionality in a principled way, while maintaining the "reach" of tailorability into that core. This separation is crucial if tailorings and adaptations are to be transportable, not just from person to person, but also across software releases. As software product cycles shrink, it is essential that users can rely on this. Without such a separation, a user is unlikely to make the investment of time and effort which tailoring involves, since the effects will be lost with the next software release.

The issue for system designers, then, is to develop a set of techniques for constructing software systems which enable the distribution of the design phase throughout the whole lifecycle of a system, and which support software adaptation and evolution. This is a significant departure from the traditional focus of software design, a focus on developing fixed algorithms which manipulate models of the application domain. Instead, it focuses on the way in which such models might be constructed and manipulated—a meta-level problem, one step removed from the application domain itself.

3 Reflection and Open Abstraction

This view has developed from ongoing work in the design of systems for Computer-Supported Cooperative Work (CSCW). CSCW systems, by their nature, have very strong requirements for flexibility and openness. Customisation may be performed not only by users but by groups as a whole, and even a single group might employ a wide variety of working styles in the course of their work. These factors bring the system developer face-to-face with the issues of reactivity and adaptability. These problems are magnified for developers of *generic toolkits*. These are used to generate a variety of CSCW applications, which may embody different models of collaboration. As a result, a major goal of my work has also been to provide the application developer with sufficient flexibility to create a range of application styles. Reflecting the shift in design focus outlined above, the emphasis here is on a framework in which mechanisms and interactional styles can be *created*, rather than the traditional approach of providing a selection of mechanisms from which particular components can be *selected*.

The systems approach that I am developing is based on the principles of *computational reflection* ([Smith 1982; Maes 1987]), and in particular the *metaobject protocol* [Kiczales et al. 1991]. This approach is a crucial stepping-stone towards the goals of flexible design which were outlined above. It provides a way of incorporating the flexibility we need not only into the design process, but into the *artefacts of design* themselves; it is fundamentally about systems which are open to explicit change and adaptation. It's worth taking some time, then, to look at the mechanics of reflection, and the way it establishes a link between generic models of system action and the performance of that action.

In the rest of this section, I will outline the reflective approach, show how it has been developed into the metaobject protocol and illustrate how such a protocol can be used in a particular case (representation in a programming language). I will show how it can be further generalised into the notion of an *open implementation*, and then return to look at the design issues in interactive systems.

3.1 Computational Reflection and Metaobject Protocols

Computational reflection is the principle that a computational system can embody, within itself, a model of its own behaviour (a self-representation) which is *causally connected* to that behaviour. Causal connection implies that the representation not only describes, but also controls, the behaviour of the system. First, this results in systems which can examine their own behaviour through examination of the model; the system can "reason" about its own activity. Second, such systems can make changes to the model and hence change their own behaviour. Essentially, in addition to the traditional "base-level" computation which concerns the system's application domain, reflection enables "meta-level" computation, which concerns the system's own manipulation and execution of base-level concepts.

This principle was originally demonstrated as part of the execution model of 3-Lisp, a reflective dialect of the Lisp programming language. 3-Lisp's reflective facilities were realised by giving the language explicit access to its own interpreter (the program controlling its behaviour) [des Rivières and Smith 1984]. By looking at the interpreter structures, 3-Lisp programs could examine their own execution states. For instance, a program could look at the function call sequence recorded in the interpreter's data structures, and so ask questions like, "how was this function arrived at?". Further, by making changes to those same structures, programs could alter future behaviour; for instance, modifications to the processing of binding structures would allow programs to change the ways in which values were associated with variables. Since the interpreter structures represent a program's execution, providing access gave 3-Lisp programs the facilities to reason about and control their own behaviour.

More recently, the principles embodied in 3-Lisp's reflective model have been combined with the techniques of object-oriented programming to yield the metaobject protocol. The metaobject protocol embodies a reflective self-representation in the structures of object-oriented programming. The self-representation in a metaobject protocol-based system is less explicit than that of 3-Lisp. Much of the representation is encoded in the object-oriented structure. Access to the representation, and manipulation of it, is provided through the object system, using the standard techniques of object-oriented programming (subclassing, specialisation, overloading, *etc.*). (An example, to illustrate the mechanics of modifying system behaviour in a metaobject protocol, is provided below.)

The first full metaobject protocol (or MOP) was developed within the definition of the Common Lisp Object System (CLOS) [Bobrow et al. 1988; Bobrow et al. 1993[1]]. The CLOS MOP creates a reflective object system, using its own object mechanisms to create an object-oriented representation of its behaviour[2]. The reflective model can be changed through standard object-oriented techniques, just like any other object-oriented definition. This allows application programmers to adapt the object system semantics for particular needs; perhaps for efficiency on particular platforms, compatibility with other systems, or specialised behaviours which enable the development of specialised applications.

3.2 Using The CLOS MOP

To explain the CLOS MOP a little more concretely, I will present an example in which an application programmer can revise design decisions in the implemented language. This particular example focuses on the issue of “instance representation” in the implementation, but it is illustrative of the general approach.

In an object-oriented programming language such as CLOS, the implementor must design a representation for instances (objects[3]), which will record such properties as the identity of the instance and the values of its slots A simple and obvious mechanism might be to allocate enough memory for all of the instance slots, define that to be the size of an instance of that particular class, and then compile slot references in programs into the appropriately-valued offsets into the instance body.

This is shown in figure 1a; the positions of the slots (`x` and `y`) are well-known, and so references to `y` can always be compiled as references to the address of the object plus the size of slot `x`. This representation is efficient for many applications which might be built with the programming language, and it lets the compiler generate fast code for slot access.

However, there are application programs which are not well served by this representation. For instance, consider a knowledge-based application. The application programmer might wish to define a class referring to people. The class might have very many slots (many hundreds or more), referring to various properties that individuals might have. However, any given instance of that class would perhaps define and use only a few of them. In this case, the simple representation outlined above—which allocates memory for slots *whether or not* they hold values—would be inappropriate. Instead, the application programmer would prefer a mechanism which only allocates space when the slot was assigned a value—perhaps one based on a lookup table for each instance. Figure 1b illustrates this alternative representation, where each instance maintains pointers only to the slots which actually have values.

In traditional languages, an implementation’s choice of representation scheme is hidden and fixed. Since it is hidden, the programmer becomes aware of it only indirectly—*e.g.* through its effect on perfor-

1. Since then, metaobject protocols have been incorporated into related languages, such as EuLisp [Bretthauer et al. 1992] and Dylan [Shalit 1992]
2. In other words, CLOS is not only reflective, but also *meta-circular* (defined in itself).
3. Throughout this paper, I will use CLOS terminology for object-oriented concepts. Individual objects are *instances* of *classes*. Each instance contains a number of defined variables called *slots*. Class-specific behaviours are defined as *methods*, grouped together into *generic functions* which define the patterns of interaction.

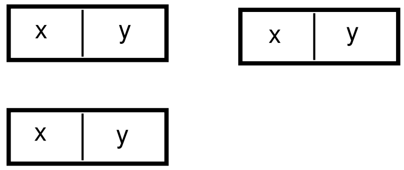

Figure 1a: In a packed representation, each data value is held within the object at a fixed and well–known position.

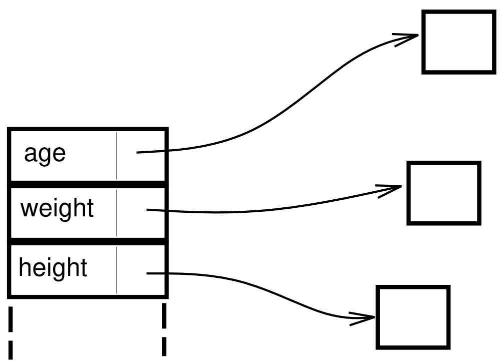

Figure 1b: In a sparse representation, each object contains a table which associates assigned slots with pointers to the data they contain.

mance in particular cases. Since it is fixed, the programmer cannot use high-level information about the program's behaviour to inform implementation decisions. Instead he or she has to write code so that it suits the decisions already made by the implementor—what Kiczales [1992] refers to as "*coding between the lines*".

In a MOP-based language, however, the application programmer can *revise* the decisions of the language implementor—in this case, by changing the representation model used. This is done through traditional object-oriented techniques; in fact, the programming of the CLOS MOP is performed in CLOS. For this example, the mechanism would be roughly as follows:

1. There is a class, called `standard-class`, of which normal user-defined classes are instances. `Standard-class` is known as the *metaclass* of such classes—a metaclass is the class of a class.

2. The metaobject protocol defines generic functions for classes, including the functions implementing instance allocation and slot lookup. Methods for these operations specialise on `standard-class`, and hence are applied to its instances (normal classes).

3. A new metaclass, embodying the new instance representation such as the sparse table-driven approach, is defined as a subclass of `standard-class`. Call this `sparse-class`.

4. Since `sparse-class` is a subclass of `standard-class`, it inherits the same methods for instance allocation and slot lookup as `standard-class`. However, we can define new, more specific methods which will override the existing ones and apply only to instances of `sparse-class`.

5. The applications programmer can now define new methods for `allocate-instance` (the generic function for creating instance representations) and `slot-value-using-class` (for looking up slot

values) specifically for instances of `sparse-class`. These implement the new, sparse slot representation.

6. The programmer can now create classes (like `person`) which have `sparse-class` as their metaclass. These classes will inherit their class-like behaviour from `sparse-class` and hence use the sparse, table-driven instance representation scheme.

So, in this example, the programmer has used knowledge of the requirements of a specific application to revise implementation decisions. The revision was performed by using the system's reflective model of its own behaviour. New, specialised behaviours were associated with a subclass of one of the system's internal classes. Since the MOP guarantees the use of generic functions to implement internal behaviours, this approach can be used to redirect the object system's behaviour; and since it exploits the polymorphism of object-oriented programming techniques, the new behaviour applies only to the objects we specify (those whose metaclass is `sparse-class`). Without this facility, it would have been necessary to write the application's code in a convoluted way, to live with inefficient performance, or to abandon this particular object system altogether. So there are benefits both to the application developer—who can tailor the system to the needs of a particular application—and to the language implementor—who can design a language with much wider applicability than traditional ones.

By defining CLOS's behaviour in terms of the metaobject protocol, the developers of the language had the means to make their language open and adaptable. They avoided a traditional problem within language and toolkit design—a premature commitment forced upon system designers in making implementation decisions which limit the choices open to the later users of that system. Instead, the system specifies default behaviours, the base-level behaviours of the object system (or whatever); but it also provides the mechanism by which those behaviours can be revised to make them more appropriate in particular circumstances.

The result, of course, is that the designers of a MOP-based system have a much less specific idea of how their system will be used. Through default behaviours, they specify a particular system, which should be generally useful; but, through the generic behaviours of the metaobject protocol, they define a framework within which users can create their own customised systems. This is a complex, two-level design task.

3.3 Extending the MOP Approach

The examples of reflective systems given above have concentrated on programming languages (3-Lisp and CLOS). Indeed the primary use of reflective techniques to date has been to provide flexible semantics for programming languages. However, we have seen that the essence of the reflective approach is closely related to the problems of openness and adaptability which were discussed for interactive systems design. So, is it possible that we could adapt reflective techniques for use in other areas?

There are examples which show that we can. One starting point would be Silica [Rao 1991], a reflective window system which forms the basis of the Common Lisp Interface Manager (CLIM). Silica provides a basic window system infrastructure for generating graphical, windowing applications. However, it also provides a set of metalevel abstractions which can be used by application developers to *reach into* the windowing infrastructure and tailor it to the needs of specific applications, much as we saw with CLOS. In Silica, the metalevel interface is defined in terms of *metaobjects* (windows and window components) and the *contracts* between them. Contracts manage the relationships between system components; they deal with issues such as geometry management, window "stacking", refreshing and so forth. The metalevel interface is used to avoid the same sorts of problems which motivated the CLOS metaobject protocol—that is, situations in which implementation decisions in the (window system) infrastructure limit its applicability to particular situations.

Silica represents a new departure in our discussion of reflective systems. Unlike 3-Lisp or CLOS, which are general-purpose programming languages, Silica cannot be defined *in its own terms*—a window system does not provide a language which can be used to construct window systems. CLOS and 3-Lisp are *metacircular*, whereas Silica is not. So Rao's notion of *implementational reflection*, as embodied in Silica, shows how we can apply reflective techniques to a much wider range of application areas than we have seen so far.

Some more recent work, arising from the metaobject protocol experiences, has opened out these notions into a more general means of providing a system's clients with control over the abstractions that they use. Kiczales [1992] presents metaobject protocols as one technique which can be used to realise *open implementations*—system implementations which augment traditional abstraction barriers with *modification interfaces*, allowing higher-level users to "reach in" and make appropriate changes. It is an approach which can be applied to a wide range of problems arising in areas where infrastructures must support a wide range of applications. Kiczales also introduces the complementary notion of *open behaviour*, in which it is the semantics, rather than the implementation, which are open to change from the higher levels. An example of this, again in the CLOS domain, is PCLOS [Paepcke 1988]. PCLOS exploits metalevel hooks into the mechanisms by which objects are created, accessed and destroyed, and so creates a *persistent* version of the language which maps objects onto long-term storage in a database. Here, it is the *behaviour* of the system which has been extended, rather than the implementation for particular applications or platforms.

These notions are very general. While they originate in work on programming language design, they point to the application of reflective techniques in a much wider range of software application areas—including the domain of interactive systems.

4 A Reflective Toolkit for CSCW Design

The approach to system architecture that I have presented here results from my current work on the design of Prospero, a flexible toolkit for CSCW systems. Individual CSCW applications need to be flexible along various dimensions. First, they must be *statically* flexible, such as in terms of customisation to particular individual or group practices or working styles (explored in more detail by Greenberg [1991]). Second, they must be *dynamically* flexible, in response to changes in group behaviour in the course of specific collaborations or even specific collaborative sessions. Third, they must be *implementationally* flexible, as infrastructural and interoperative requirements change. At the same time, a toolkit needs to provide developers with sufficient flexibility to generate applications for a wide range of groups, applications and usage settings.

Just as most programming languages fix the implementations of the language's abstractions, existing CSCW toolkits such as GroupKit [Roseman and Greenberg 1992] or MMConf [Crowley et al. 1990] are forced, through their structure, to impose fixed models of their abstractions, such as distributed data management. This follows directly from the traditional structuring techniques in software development, which hide implementation details behind abstraction barriers, out of reach of the applications developer. There is clearly value to this approach. In particular, it isolates the applications developer from toolkit concerns, focussing attention on those areas specific to the application. The cost is that the range of applications which can be developed within the toolkit is greatly restricted, since the implementation decisions within the toolkit constrain the kinds of interactions which can be supported in applications. The very isolation that toolkits provide prevents the developer from using high-level information about the application to inform lower-level decisions where appropriate.

My current work uses reflective techniques to address these problems. Structuring a CSCW toolkit around a metaobject protocol allows us to tackle two issues. First, as well as providing default behaviours which specify the natural behaviour of the system, it gives programmers the opportunity to *specialise* and *refine* the generic framework in the toolkit. This means that the toolkit can be used to provide customised support for particular situations and applications. Second, the self-representation is present in the applications at run-time. This allows appropriately-written programs to respond dynamically and adaptively as they are used.

This section outlines the way in which the reflective approach is applied in the CSCW domain. Taking a very high-level description of generic application behaviour, it shows that, in areas of concern for CSCW application developers, a variety of strategies can be supported within a single metalevel framework.

4.1 Using Reflection in CSCW Design

To apply reflection to the design of a CSCW toolkit, we must "open up" the implementation. This involves specifying the generic behaviours which underlie the system's operation, and the generic entities on which these behaviours act. Providing explicit access to these generic behaviours allows the toolkit user (*i.e.* the programmer) to specialise them for particular situations. These generic behaviours can be broken down into sub-protocols, or specific areas of responsibility.

Clearly, there are a huge number of potential areas of responsibility within the toolkit. One of the major issues in MOP design is the identification of a particular set of concerns that the design should address in order to create a system which is flexible but manageable. We do this by looking at the particular experiences of developers building both applications and toolkits for cooperative systems. We can see a number of issues which are embedded in the design of the infrastructure and yet have strong implications for the kinds of applications which can be supported— areas which are candidates for this approach.

Currently, my work concentrates on three main areas: the management of user data distributed across time and space; mechanisms for managing conflict in user interactions; and control over the linkage between the components of multiple users' interfaces. For each area, the approach is fundamentally the same, and involves specifying generic behaviour. This is defined in terms of generic function invocations on metaobjects, or explicit representations of the system's behaviour. These generic behaviours can be specialised by application developers through incremental modifications to the representations and the actions over them.

Before looking at the use of reflective techniques in these areas, it's necessary to lay down some high-level structure which relates them. A full description of the approach used in Prospero is beyond the scope of this paper, and so a simplified account is presented here; the interested reader is referred to other treatments (*e.g.* [Dourish 1994]).

Consider a system which operates in terms of generic edit operations applied to shared objects. The most general layer of functionality is provided by the function[1]:

```
(edit-object object user editop) -> state-marker
```

`Object` is a local reference to a globally shared object in the collaborative workspace; `user` is a representation of the user performing the operation; and `editop` is an encoding of the operation being performed. The generic function `edit-object` applies an edit operation to an object, and returns `state-marker` which describes the new state. It is implemented in terms of a number of lower-level generic functions:

```
(find-object object) -> shobject
(lock-object shobject user editop) -> lockid
(apply-edit shobject user editop) -> change-marker
(propagate change-marker lockid) -> state-marker
```

These functions perform the component operations of `edit-object`: mapping from local objects, pre-

1. Functions are given here in the format used by the Lisp programming language; the first term is the function name, and subsequent terms name arguments to the function.

sented within the interface, to object components of the shared workspace; obtaining access to those objects; applying changes; and then propagating those changes more widely and releasing the lock. We use change-markers and state-markers as encapsulations of the state of the system at various points. Change-markers record edits made which have not yet been committed; state-markers checkpoint global status. The model presented by this protocol uses these for synchronisation, as it presents a view of edit changes being performed locally; however, as long as it is true to this generic model, implementations may behave differently.

Having set up this general framework, we can now investigate how system variability in the areas of data distribution, conflict management and interface linkage can be managed.

4.1.1 Data Distribution

The issue of data distribution has been a bone of contention within the CSCW implementation community for some time. The term "data distribution" covers mechanisms by which the system manages the user data storage and manipulation. This data may be replicated or distributed across multiple computers, but the system must present a view of a single, coherent data store. Systems such as MMConf [Crowley et al. 1990] take a *fully replicated* approach in which each participant in a conference has a private copy of the data. Others, such as Rapport [Ahuja et al. 1990], use *centralised* architectures, which concentrate data at a single point in the network. Greenberg et al. [1992] have argued in favour of *hybrid* systems which combine these approaches. Each of these solutions—replicated, centralised or hybrid—makes some trade-off between efficiency and complexity when they are the only approach taken by a particular toolkit.

It seems clear that there can be no solution which is appropriate in every case. Not only are there occasions where any of the centralised, replicated or hybrid approaches are appropriate, but further, there are times when we might need *others*. For instance, consider *disconnected* systems in which some interfaces involved in the collaboration are not permanently connected to the others, or situations in which network latency is high and intra-group interaction low. Here it might be useful to adopt *migratory* mechanisms, which allow data objects to move from one node to another in the network. Migration isn't a separate strategy in itself; it can be combined with any of the other three basic techniques. Other approaches can be posited which will be particularly appropriate for other situations, and we need to be able to express this variability in the toolkit.

Perhaps more importantly, the data distribution approach adopted by a toolkit or application can have important consequences for the appearance, functionality and usability of the application. The use of a centralised data store, for instance, can negatively affect the response time of the system; while the use of a replicated approach has implications for the maintenance of data consistency. This is at odds with the traditional view that such factors as data distribution are sufficiently "low-level" that they can be safely encapsulated and hidden behind an abstraction barrier[1].

Rather than make these decisions up-front, we can take the reflective approach. Within a toolkit, we not only provide some default mechanism for managing data within a multi-user system, but we also give access to the mechanism by which data distribution is accomplished. This allows programmers, who may find the default behaviour inappropriate in their case (*e.g.* because of the network topology they are using), to "reach in" to the toolkit and provide new mechanisms to be used in their applications.

The sample subprotocol outlined above manages data distribution largely through the `find-object` and `propagate` mechanisms, which isolate the location and distribution details from the details of actually making changes to the objects. Using `find-object`, we can encode new mechanisms for mapping between interface objects and the underlying shared data. For instance, in a centralised system, `find-ob-`

1. The interaction between the distributed data management and the issue of synchrony of interaction is also critical, but merits a longer investigation than can be presented here.

`ject` will always return a pointer to the central object store, and `propagate` will return the locally changed object to the server. On the other hand, in a fully replicated system, the shared object reference is always local, and more complicated methods on `propagate` will allow changes to be synchronised appropriately. It is important to note, though, that this approach does not merely provide a switch between these two modes. Instead, it provides a framework in which *new solutions* can be devised. The generality of `find-object` and `propagate` allows many alternatives, including hybrid and migratory systems, to be created.

Since the representations are available at run-time, rather than simply when the system is defined, other opportunities present themselves. We can amend the data distribution mechanism using a dynamic model, which would allow distribution strategies to be changed in the course of an on-going collaborative session. This allows an application to adapt to the needs of the group as they arise. For instance, consider two users sharing a "scrawl"-style whiteboard application, connected on the same ethernet segment. Since their connection has fairly low round-trip packet times and high data integrity, the system requirements for data management are fairly minimal; a centralised approach is probably entirely adequate. However, things change if a third user joins their conference from some distance, connected via a much slower dial-up line. In this situation, a centralised approach is no longer appropriate, since the bandwidth of the link to the third user is not sufficient to support a network interaction with a data server for each action at the interface. The system must switch, at run-time, from one algorithm to another—from a centralised to a replicated data representation. A reflective approach provides the potential for multiple behaviours within the same generic framework, thus supporting this form of dynamic adaptation. If distribution is associated with an object through a mixin[1] class, then changing the class of the object will result in the dynamic switch to a different behaviour.

The reflective approach provides a framework within which new mechanisms can be defined, and the means to attach use mechanisms selectively in different parts of the system. This gives three principal benefits which would not be available with traditional solutions. First, the application developer is no longer constrained by decisions within the toolkit, but is free to adapt the toolkit mechanisms to the needs of the application. Second, the toolkit developer no longer needs to "second guess" the specific needs of the developer, or restrict applicability of the toolkit to a subset of potential applications. Third, using the reflective model at run-time allows applications to respond dynamically to their environments and the requirements of particular situations, all within a single coherent framework. We will see this pattern of benefits repeat itself in other areas where we apply computational reflection to toolkit design.

4.1.2 Conflict Management

An important area of concern for collaborative applications is the management, or avoidance, of conflicts within the shared workspace. A conflict might occur, for instance, when two users apply a change to the same object at once. Various techniques have been employed to deal with this sort of situation, including floor control, exclusion and locking. Some systems, such as ShrEdit [McGuffin and Olson 1992], "lock" regions of the shared workspace, preventing simultaneous updates since only one user can hold a lock on a region at any given time. Others, such as GROVE [Ellis and Gibbs 1989], use an algorithm which "fixes up" conflicts afterwards, in effect imposing a post-hoc serialisation on the changes that users make.

The essence of conflict management strategies is that the system be able to provide guarantees that users' changes to the data will not lead to a loss of synchronisation or data integrity. Prospero exploits explicit representations of such guarantees; this approach is explored in more detail in [Dourish 1994]. In this ac-

1. A *mixin* class is one which can be added to other classes in order to bring some new behaviour, orthogonal to that defined in the base classes.

count, however, I will focus solely on systems providing rapid access to a single thread of control.

Even if we choose a simple approach such as locking, then we have to consider the impact that particular locking mechanisms, defined within the toolkit, might have on higher-level usage issues. For instance, in a collaborative system supporting free-form sketching or brainstorming, the emphasis is probably on unencumbered access to the shared work surface. If each user had to explicitly request and explicitly relinquish locks on the data or control of the floor, the overhead would be too high and the progress of work would be severely disrupted. A looser form of control would be needed. On the other hand, looser control would be inappropriate in systems where data integrity must be rigorously maintained and controlled. In a collaborative software engineering application, or a multi-user CAD system which generates control instructions for a milling machine, data errors due to unchecked conflicts could be potentially disastrous, and a much stronger and more explicit form of locking would be required. Neither approach satisfies the needs of a generic toolkit.

Addressing these problems in a reflective toolkit, we attempt to provide a metalevel interface which defines the generic operations involved in requesting, obtaining and releasing locks. In the simple protocol outlined at the start of section 4, we focus on the call to `lock-object`, and the implicit `release-lock` called from `propagate`. As before, the protocol itself does not *embody* a locking policy. Instead, it deals with a procedure by which locks are obtained, and a facility for creating and installing new mechanisms. The generic function specifies that, as well as the object to be locked, the function arguments include the user requesting the lock and the type of operation to be performed. The system can take this information into account when selecting the locking mechanism. So, different locking strategies may apply to different users or activities, and for different sorts of objects within the same system. The implementor can rely on the object system's generic dispatch mechanism to dynamically select the appropriate locking implementation. So the programmer can not only tailor locking strategies to particular applications, but can also build systems in which the locking mechanisms used rely on specific details of the user or object involved.

The basic mechanism is sufficiently open that a wide range of locking strategies can be defined. Not only will it allow the implementation of standard strong and weak locks, but also multi-way locks (held by multiple people at once), tickle locks (which, when idle, may be implicitly reassigned to other users when they perform an operation) and so on. Indeed, we can reproduce schemes such as GROVE's dOPT algorithm, in which explicit locks are not used at all, by making `lock-object` to construct an appropriate "state vector". This will be distributed by the call to `propagate`, so that other nodes can use this information to resolve ambiguities arising out of conflicting or mis-ordered operations. In this case, we regard the dOPT state vector as an implicit "lock", in the sense that it is an object which will allow conflict resolution. In other words, the same basic mechanism can be used to encode a form of conflict management which is hardly "lock-based" at all.

4.1.3 Interface Linkage

One of the most obvious differences between CSCW systems is the level at which they "link" interface features. Linkage determines the level of control that users have over the way their own interfaces appear, without affecting other users of the same collaborative application. The grossest level of linkage is screen-replication, as used for example in Timbuktu [Faralon 1987]. Screen linkage means that all users see exactly the same thing on their screens. Shared X systems [Garfinkel et al. 1989] link interfaces at the level of *windows*; users share the contents of a window, while their screens may show other, independent applications and window placement can vary from person to person. Many explicit multi-user tools such as ShrEdit are much looser, and will replicate only the data[1]. Here, users may have different views of the data, and may be provided with individual edit cursors. Within this class of systems, there are further differences in what each

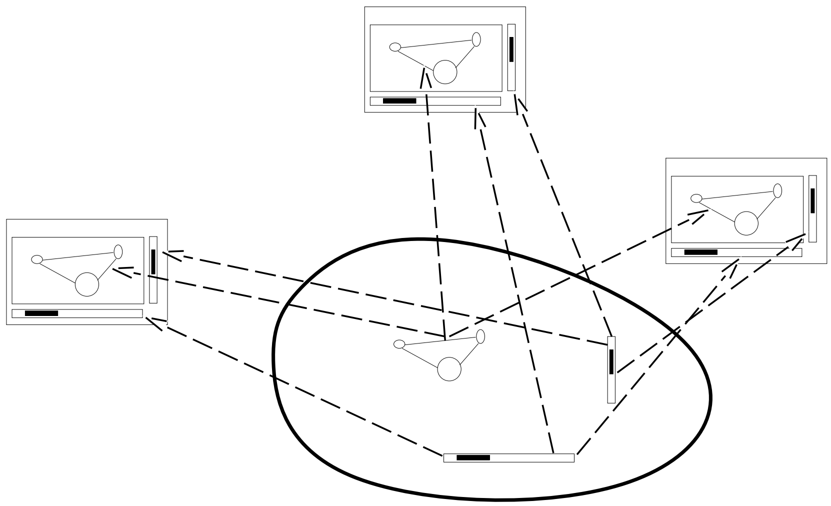

Figure 2: Not only data objects, but also interface components, can be part of the shared workspace.

user can see of the other's interfaces.

While many systems separate users and isolate their interfaces, research on groups interacting through synchronously shared systems has shown how low-level cues can be used by collaborators to create an awareness of the activity and progress of the group as a whole [Dourish and Bellotti 1992]. Recent work, such as that of Dewan and Choudhary [1991] or Haake and Wilson [1992], has looked at the provision of switchable linkage states, in which users can choose how much their interfaces will mirror each other's. A similar two-mode switching facility was available in rIBIS [Rein and Ellis 1991].

Once again, we can see the requirements for flexibility within applications and toolkits, and that this flexibility can have a dynamic component. The non-dynamic aspect is the same now-familiar toolkit-level problem; that different applications require different linkage strategies, and so a generic toolkit must be able to support a range of linkage options. Dewan's work with Suite, or Haake and Wilson's with SEPIA, tackle just this problem, as well as addressing the dynamic problem of switching between these different linkage modes during a collaborative session. However, both systems provide flexible linking through specific "modes", which pre-define linkage strategies. This only addresses part of the problem. What if different users or groups require different linkage strategies, or some situation calls for a strategy which has not been predefined? There are no options while remaining within these frameworks. While clearly the goal here is to provide flexibility, it is flexibility within the constraints of the set of needs anticipated by the toolkit designer. While these mechanisms are *parameterised*, they are not *open*.

The approach taken in Prospero is to make aspects of interface components into shared objects, exactly like workspace data objects. This is illustrated in figure 2. In this way, components such as menus, button states and cursor positions, are subject to the mechanisms outlined earlier for dynamic distribution algorithms. The same mechanisms which maintain consistency between workspace objects can be used to synchronise interfaces.

The result is that the same flexibility which can be applied to distributed data management also applies

1. Note that our concern here is with replication of interface features, rather than with the underlying data representations discussed earlier. So, in these systems, it is only the data which is guaranteed to be consistently replicated *between interfaces*.

to interface management. Interface components can be explicitly shared, causing that aspect of interfaces to be linked; or they can be separated and broadcast, allowing each individual control while being able to see other's states; or they can be private, in which case other users cannot see them. In addition, another property they have in common with shared data objects is that they can be moved into and out of the shared workspace in the course of a collaboration. This means that the linked aspects of the users' interfaces can be dynamically controlled and adjusted as the collaboration continues

5 Reflection and Adaptive Computation

In this paper, I have mainly been concerned with the use of computational reflection and related techniques as ways of opening up implementations and providing flexibility. The techniques I have discussed have been developed primarily in the domain of programming language design. Recently, similar issues have come to prominence in a number of other areas. Various radical solutions have been adopted which go beyond the traditional separation of "mechanism and policy"; and often, these address issues very similar to those discussed here within the context of CSCW.

The design of communication protocols on data networks has traditionally been based on a "layered" approach. This form of design is exemplified by the seven-layer ISO protocol stack [Zimmerman 1980]. End-to-end communication requirements are broken down into different areas of responsibility, such as data representation conversion, direct host-to-host communication, inter-network communication and so on. Each component is encapsulated in a layer and, on a given machine, each layer interacts only with the layers directly above and below it. More recently, however, the need to handle interactive multimedia traffic, and issues arising in the design of protocols for gigabit networks, have resulted in a breakdown in this model. In its place, a flatter approach is emerging in which more of the traffic management is controlled *directly by the application*, rather than being hidden in the network software. O'Malley and Peterson [1992] have described a model in which the application can compose kernel-internal microprotocols into larger units optimised for their particular requirements. Clark and Tennenhouse [1990] propose the concept of *Integrated Layer Processing* as a mechanism to avoid the inter-layer inefficiencies which emerge when the infrastructure is examined from the point of view of particular end-to-end systems.

In the area of operating systems, much functionality which has usually been in the domain of the system itself is being opened up to external control. One of the most obvious examples is Mach's external pager [Rashid et al. 1987], which allows user programs to involve themselves in aspects of the virtual memory system's operation. Similarly, Anderson et al. [1992] describe "scheduler activations" as a means to avoid trade-offs in the implementation of threads, which are traditionally a completely opaque abstraction over an implementation based either in the operating system kernel or in a user library. Scheduler activations provide a finer grain of control and are explicitly designed to allow application-specific customisation.

This trend is repeated in many other areas too, such as interprocessor communication [Felten 1992] or even microprocessor design [Athanas and Silverman 1993]. The same principle is at work in all of these examples. They are all based on an understanding that traditional closed abstractions are not always appropriate for high-level systems design in general, and in particular for the design of infrastructural components. The various solutions are oriented around a *downward flow of information*, from the higher levels (applications) to the lower levels (toolkits and infrastructures), in order to support better interaction between the two. This downward flow—from application requirements to the details of system support—mirrors the problems that Prospero addresses with reflective techniques.

Reflection achieves this by opening up the underlying implementation and allowing the applications programmer to explore alternative implementations and behaviours within the metalevel framework. This corresponds with what Kay [1993] has characterised as *late-binding systems*—those in which design and implementation decisions which affect observable behaviour are delayed until they can be resolved with as much context as possible. While the techniques discussed here have been derived from work on programming language design, late-binding is useful and important in interactive systems, especially in CSCW systems where contextual factors play such a large part in the interaction.

Prospero is a toolkit for CSCW applications under development, based on these principles. It concentrates primarily on issues in data distribution, conflict management and interface linkage. The examples in this paper have been drawn from this work in progress. The system provides default behaviours which can be used to construct applications in the usual way, where they are appropriate. More importantly, though, it also provides a metalevel framework which can be used to revise implementation decisions, to extend the structure to cover new areas, and to make the toolkit more appropriate for a range of applications.

6 Summary

The primary focus of this paper has been on models of implementing interactive systems. I have argued that recent years have seen a fundamental reorientation in our view of interactive systems and their use. In turn, this forces a reorientation in our view of system design and structure. In particular, appreciation of the need for (and use of) customisation facilities, the role of work practice and situation in system use, and the co-adaptive nature of system use and user behaviour lead us to a model of systems design which emphasises openness, dynamic behavior and evolution of systems and practices.

The move away from static systems leads us to reconsider the architectures which underlie interactive systems. It is not enough to simply change the process of design; instead, we need to change the nature of the artefacts themselves. By drawing on the principles and techniques of computational reflection, derived originally from research into programming language semantics, I have outlined a model of interactive system design which is oriented specifically towards these new goals of flexibility and adaptation. In particular, this model is currently being used as the implementational basis of a toolkit for CSCW design, and I have outlined how this toolkit tackles a number of current problems in CSCW toolkits which must be used in a wide range of different circumstances and situations.

This work is currently ongoing. It is hoped that the reflective toolkit for CSCW can provide insights into the general application of notions of open implementation and behaviour to a range of current problems in interactive system design.

Acknowledgments

The ideas expressed in this paper would never have seen the light of day without the contributions and encouragement of many people. I would particularly like to thank Hal Abelson, Bob Anderson, Victoria Bellotti, Danny Bobrow, Jon Crowcroft, Gregor Kiczales and Wendy Mackay for fruitful and enlightening discussions. I am very grateful to Annette Adler, Jon Crowcroft, Laura Dekker, Marge Eldridge, Lorna Goulden, Jonathan Grudin, Gillian Ritchie, Lisa Tweedie and the anonymous reviewers for careful readings and valuable comments on earlier drafts of this paper.

References

S. Ahuja, J. Ensor, and S. Lucco, *"A Comparison of Application Sharing Mechanisms In Real-time Desktop Conferencing Systems"*, Proc. Conference on Office Information Systems COIS'90, Boston, Mass., April 1990.

Thomas Anderson, Brian Bershad, Edward Lazowska and Henry Levy, "*Scheduler Activations: Effective Kernel Support for the User-Level Management of Parallelism*", ACM Transactions on Computer Systems, 10(1), pp. 53–79, February 1992.

Peter Athanas and Harvey Silverman, "*Processor Reconfiguration Through Instruction-Set Metamorphosis*", IEEE Computer, March 1993.

Daniel Bobrow, Linda Demichiel, Richard Gabriel, Sonya Keene, Gregor Kiczales and David Moon, "*Common Lisp Object System Specification*", X3J13 Document 88-002R, June 1988.

Daniel Bobrow, Richard Gabriel and Jon L White, "*CLOS in Context: The Shape of the Design Space*", in Paepcke (ed.), "Object-Oriented Programming: The CLOS Perspective", MIT Press, Cambridge, Mass., 1993.

Suzanne Bødker and Kaj Grønbæk, "*Cooperative Prototyping: Users and Designers in Mutual Activity*", Intl. Journal of Man-Machine Studies, 34(3), pp. 453–479, 1991.

Barry Boehm, "*A Spiral Model of Software Development and Enhancement*", IEEE Computer, pp. 61–72, May 1988.

Grady Booch, "*Object Oriented Design*", Benjamin/Cummings, Redwood City, Ca., 1991.

John Bowers and Tom Rodden, *"Exploding the Interface: Experiences of a CSCW Network"*, Proc. InterCHI'93, Amsterdam, Netherlands, April 1993.

Harry Bretthauer, Harley Davis, Jürgen Kopp and Keith Playford, "*Balancing the EuLisp Metaobject Protocol*", in Proc. IMSA'92 Workshop on Reflection and Metalevel Architectures, Tokyo, Japan, November 1992.

David Clark and David Tennenhouse, "*Architectural Considerations for a New Generation of Protocols*", ACM SIGCOMM Communications Review, 20(4), pp. 200–208, 1990.

Terry Crowley, Paul Milazzo, Ellie Baker, Harry Forsdick and Raymond Tomlinson, "*MMConf: An Infrastructure for Building Shared Multimedia Applications*", in Proc. ACM Conference on Computer-Supported Cooperative Work CSCW'90, Los Angeles, October, 1990.

Prasun Dewan and Rajiv Choudhary, "*Flexible User Interface Coupling In a Collaborative System*", in Proc. ACM Conf. Human Factors in Computing Systems CHI'91, New Orleans, Louisiana, April 1991.

Paul Dourish, "*A Divergence-Based Model of Synchrony and Distribution in Collaborative Systems*", EuroPARC Technical Report EPC-94-102, Rank Xerox EuroPARC, Cambridge, UK, 1994.

Paul Dourish and Victoria Bellotti, "*Awareness and Coordination in Shared Workspaces*", in Proc. ACM Conference on Computer-Supported Cooperative Work CSCW'92, Toronto, Canada, November 1992.

Pelle Ehn, "*Work-Oriented Design of Computer Artifacts*", Arbetslivscentrum, Stockholm, 1988.

Clarence Ellis and Simon Gibbs, "*Concurrency Control in Groupware Systems*", Proc. ACM Conference on Management of Data SIGMOD'89, Seattle, Washington, 1989.

Farallon Computing, "*Timbuktu: The next best thing to being there*", 1987.

Edward Felten, "*The Case for Application-Specific Communication Protocols*", Technical Report TR-02-03-11, Department of Computer Science, University of Washington, Seattle, Washington, 1992

Geneviève Filippi and Jacques Theureau, *"Analyzing Cooperative Work in an Urban Traffic Control Room for the Design of a Coordination Support System"*, Proc. Third European Conference on Computer Supported Cooperative Work ECSCW'93, Milano, Italy, September 1993.

Dan Garfinkel, Phil Gust, Mike Lemon and Steve Lowder, "*The SharedX Multi-User Interface User's Guide, Version 2.0*", Software Technology Lab Report STL-TM-89-07, Hewlett-Packard Laboratories, Palo Alto, Ca., 1989.

Saul Greenberg, "*Personalisable Groupware: Accommodating Individual Roles and Group Differences*", in Proc. European Conference on Computer-Supported Cooperative Work ECSCW'91, Amsterdam, Netherlands, September 1991.

Saul Greenberg, Mark Roseman, Dave Webster and Ralph Bohnet, "*Human and Technical Factors of Distributed Group Drawing Tools*", Interacting with Computers 4(3), pp. 364-392, 1992.

Jonathan Grudin, "*The Computer Reaches Out: The Historical Continuity of Interface Design*", in Proc. ACM Conference on Human Factors in Computing Systems CHI'90, Seattle, Washington, April 1990.

Jonathan Grudin, "*Obstacles to User Involvement in Software Product Development, with Implications for CSCW*", Intl. Journal of Man-Machine Studies, 34, pp. 435-452, 1991.

Jonathan Grudin, *"Interface: An Evolving Concept"*, Communications of the ACM, 36(4), pp. 110–119, April 1993.

Jörg Haake and Brian Wilson, "*Supporting Collaborative Writing of Hyperdocuments*", in Proc. ACM Conference on Computer-Supported Cooperative Work CSCW'92, Toronto, Canada, November 1992.

Richard Harper, John Hughes and Dan Shapiro, "*Harmonious Working and CSCW: Computer Technology and Air Traffic Control*", in Bowers and Benford (Eds.) "Studies in Computer Supported Cooperative Work", North-Holland, Amsterdam, 1991.

Christian Heath and Paul Luff, "*Collaborative Activity and Technological Design: Task Coordination in London Underground Control Rooms*", in Proc. European Conference on Computer-Supported Cooperative Work ECSCW'91, Amsterdam, Netherlands, September 1991.

Alan Kay, "*The Early History of Smalltalk*", in Proc. ACM Conf. History of Programming Languages HOPL-II, Cambridge, Mass, published in SIGPLAN Notices 28(3), March, 1993.

Gregor Kiczales, Jim des Rivières and Danny Bobrow, "*The Art of the Metaobject Protocol*", MIT Press, Cambridge, Mass., 1991.

Gregor Kiczales, "*Towards a New Model of Abstraction in Software Enginering*", in Proc. IMSA'92 Workshop on Reflection and Metalevel Architectures, Tokyo, Japan, Nov. 4-7, 1992

Wendy Mackay, "*How Do Experienced Information Lens Users Use Rules?*", in Proc. ACM Conference on Human Factors in Computing Systems CHI'89, Austin, Texas, 1989.

Wendy Mackay, "*Users and Customisable Software: A Co-Adaptive Phenomenon*", PhD thesis, Sloan School of Management, MIT, Cambridge, Mass., 1990 (a).

Wendy Mackay, "*Patterns of Sharing Customisable Software*", in Proc. ACM Conference on Computer-Supported Cooperative Work CSCWW'90, Los Angeles, Ca., October 1990 (b).

Wendy Mackay, "*Triggers and Barriers to Customising Software*", in Proc. ACM Conference on Human Factors in Computing Systems CHI'91, New Orleans, Louisiana, April 1991.

Lola McGuffin and Gary Olson, "*ShrEdit: A Shared Electronic Workspace*", CSMIL Technical Report, Cognitive Science and Machine Intelligence Laboratory, University of Michigan, 1992.

Allan MacLean, Kathleen Carter, Thomas Moran and Lennart Lovstrand, "*User-Tailorable Systems: Pressing the Issues with Buttons*", in Proc. ACM Conference on Human Factors in Computing Systems CHI'90, Seattle, Washington, April 1990.

Pattie Maes, "*Computational Reflection*", Technical Report 87.2, Artificial Intelligence Laboratory, Vrije Universiteit, Brussels, Belgium, 1987.

Tom Malone, Kevin Crowston, Ramana Rao, David Rosenblitt and Stu Card, "*Semi-Structured Messages Are Surprisingly Useful for Computer-Supported Coordination*", ACM Transactions on Office Information Systems, 5(2), pp. 115–131, 1987.

Bonnie Nardi and James Miller, "*Twinkling Lights and Nested Loops: Distributed Problem Solving and Spreadsheet Development*", in Greenberg (ed.), "Computer-Supported Cooperative Work and Groupware", Academic Press, 1991.

Don Norman and Steve Draper (eds.), "*User-Centered Systems Design*", Lawrence Erlbaum Associates, 1986.

Sean O'Malley and Larry Peterson, "*A Dynamic Network Architecture*", ACM Transactions on Computer Systems, 10(2), May 1992.

Andreas Paepcke, "*PCLOS: A Flexible Implementaton of CLOS Persistance*", Proc. European Conference on Object-Oriented Programming ECOOP'88, Springer-Verlag, 1988.

John Patterson, Ralph Hill, Stephen Rohall and Scott Meeks, "*Rendezvous: An Architecture for Synchronous Multi-*

User Applications", in Proc. ACM Conference on Computer-Supported Cooperative Work CSCW'90, Los Angeles, Ca., October 1990.

Ramana Rao, "*Implementational Reflection in Silica*", in Proc. European Conference on Object-Oriented Programming ECOOP'91, Geneva, Swizerland, 1991.

Richard Rashid, Avadis Tevanian, Michael Young, David Golub, Robert Baron, David Black, William Bokosky and Jonathan Chew, "*Machine-Independent Virtual Memory Management for Paged Uniprocessor and Multiprocessor Architectures*", in Proc. Conf. Architectural Support for Programming Languages and Operating Systems, Palo Alto, Ca., October 1987.

Gail Rein and Clarence Ellis, "*rIBIS: A Real-Time Group Hypertext System*", Intl. Journal of Man-Machine Studies, 34, pp. 349-367, 1991.

Jim des Rivières and Brian Smith, "*The Implementation of Procedurally Reflective Languages*", Technical Report ISL-4, Xerox PARC, Palo Alto, California, June 1984.

Mark Roseman and Saul Greenberg, "*GroupKit: A Groupware Toolkit for Building Real-Time Conferencing Applications*", Proc. ACM Conference on Computer-Supported Cooperative Work CSCW'92, Toronto, Canada, November 1993.

W. W. Royce, "*Managing the Development of Large Software Systems*", Proc. WESTCON, Ca., 1970.

Andrew Shalit, "*Dylan: An Object-Oriented Dynamic Language*", Apple Computer, Cupertino, Ca., 1992.

Brian Smith, "*Reflection and Semantics in a Procedural Language*", Report MIT-TR-272, MIT Laboratory for Computer Science, Cambridge, Mass., 1982.

Lee Sproull and Sara Kiesler, "*Connections: New Ways of Working in the Networked Organisation*", MIT Press, Cambridge, Mass., 1991.

Lucy Suchman, "*Plans and Situated Actions*", Cambridge University Press, Cambridge, UK, 1987.

Lucy Suchman, *"Technologies of Accountability: Of Lizards and Aeroplanes"*, in Button (ed.), "Technology in Working Order: Studies of Work, Interaction and Technology", Routledge, 1992.

Randy Trigg, Tom Moran and Frank Halasz, "*Adaptability and Tailorability in NoteCards*", in Proc. Interact'87, Stuttgart, Germany, 1987.

H. Zimmerman, *"OSI Reference Model—The ISO Model of Architecture for Open Systems Interconnection"*, IEEE Transactions on Communications, 28(4), pp. 425–432, April 1980.