

Claude Kirchner (Ed.)

ccc1-690

Rewriting Techniques and Applications

5th International Conference, RTA-93
Montreal, Canada, June 16-18, 1993
Proceedings

BIBLIOTHEQUE DU CERIST

Springer-Verlag

Berlin Heidelberg New York
London Paris Tokyo
Hong Kong Barcelona
Budapest

Series Editors

Gerhard Goos
Universität Karlsruhe
Postfach 69 80
Vincenz-Priessnitz-Straße 1
W-7500 Karlsruhe, FRG

Juris Harmanis
Cornell University
Department of Computer Science
4130 Upson Hall
Ithaca, NY 14853, USA

Volume Editor

Claude Kirchner
INRIA Lorraine and CRIN
615 Rue du Jardin Botanique, F-54602 Villers les Nancy Cedex, France

CR Subject Classification (1991): D.3, F.3.2, F.4, I.1, I.2.2-3

ISBN 3-540-56868-9 Springer-Verlag Berlin Heidelberg New York
ISBN 0-387-56868-9 Springer-Verlag New York Berlin Heidelberg

6258

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable for prosecution under the German Copyright Law.

© Springer-Verlag Berlin Heidelberg 1993
Printed in Germany

Typesetting: Camera ready by author
Printing and binding: Druckhaus Beitz, Hemsbach/Bergstr.
45/3140-543210 - Printed on acid-free paper

Preface

This volume contains the proceedings of RTA-93, the Fifth International Conference on Rewriting Techniques and Applications held June 16-18, 1993, in Montreal, Canada.

There were 91 submissions to RTA-93 authored by researchers from countries including Canada, France, Germany, Italy, India, Japan, the Netherlands, the People's Republic of China, Russia, Spain, United Kingdom, and the United States of America. Papers covered many topics: term rewriting; termination; graph rewriting; constraint solving; semantic unification, disunification and combination; higher-order logics and theorem proving, with several papers on distributed theorem proving, theorem proving with constraints, and completion.

Each submission was reviewed by at least three program committee members or their outside referees. All the members of the program committee met on February 1993 in Nancy and selected 29 papers and 6 system descriptions demonstrated during the conference and documented in this volume.

As for the proceedings of the previous conference, I welcomed the idea of presenting in the proceedings a list of open problems in the field and an update of the previous list of such open problems, showing altogether the strong activity of the term rewriting community in the large.

Three invited speakers gave a talk on their recent works related to the topics of RTA. Sergei Adian presented his work on algorithmic problems for groups and semigroups, Leo Bachmair the impact of rewriting techniques on theorem proving and Jean Gallier a general method for proving properties of typed lambda terms.

I am very grateful to the program committee for their efforts and cooperation in deciding the program and other related matters to RTA-93; to Mitsuhiro Okada for taking great care of the local arrangements for the conference; to the invited speakers Sergei Adian, Leo Bachmair and Jean Gallier, and lastly to Marian Vittek for doing everything that needed to be done to facilitate my task in organizing the program committee.

RTA-93 was sponsored by INRIA (France), the Centre de Recherche en Informatique de Nancy (France), Concordia University (Canada), the Center for Pattern Recognition and Machine Intelligence, Montreal (Canada), the Natural Science and Engineering Research Council (Canada), le Fonds pour la Formation de Chercheurs et l'Aide à la Recherche (Quebec) and the National Science Foundation (USA), and was held under the auspices of the European Association for Theoretical Computer Science.

Nancy, April 1993

Claude Kirchner
Chair, RTA-93

Program committee

Hubert Comon (Orsay)
Bruno Courcelle (Bordeaux)
Harald Ganzinger (Saarbrücken)
Tieh Hsiang (Stony Brook)
Claude Kirchner (Nancy)
Jan Willem Klop (Amsterdam)
Klaus Madlener (Kaiserslautern)
Paliath Narendran (Albany)
Mike O'Donnell (Chicago)
Mitsuhiro Okada (Montreal)
Leszek Pacholski (Wrocław)
Michael Rusinowitch (Nancy)
Mark Stickel (Menlo Park)

Organizing committee

Ronald Book (Santa Barbara)
Nachum Dershowitz (Urbana)
Jean Gallier (Philadelphia)
Deepak Kapur (Albany)
Claude Kirchner (Nancy)
Klaus Madlener (Kaiserslautern)
Pierre Lescanne (Nancy)
David Plaisted (Chapel Hill)

Local arrangements

Mitsuhiro Okada (Montreal)

Referees

A. Arnold	P. Audebaud	J. Avenhaus
H.-J. Bürckert	L. Bachmair	S. Bailey
D. Basin	H. Baumeister	A. Bockmayr
F. de Boer	M.P. Bonacina	A.M. Borzyszkowski
A. Boudet	W. Bousdira	P. Casteran
W. Charatonik	T. Chen	C. Choffrut
E.A. Cichon	E. Contejean	T. Deiss
B. Delsart	J. Denzinger	E. Domenjoud
D. Dougherty	F. Fages	M. Falaschi
D. Fehrer	M. Fernández	M.C.F. Ferreira
R. Fettig	A. Geser	R. Gilleron
G. Gonthier	B. Gramlich	S. Hölldobler
M. Hanus	D. Hofbauer	J. Hong
M. Huber	U. Hustadt	P. Jacquet
P. Johann	J.-P. Jouannaud	T. Jurdziński
X. Kühler	R. Kennaway	D. Kesner
A. Kisielewicz	H.-J. Kreowski	D. Krob
M. Kutylowski	S. Lange	D. Lugiez
C. Lynch	F. Müller	C. Marché
R. McCloskey	W. McCune	R. McNaughton
A. Middeldorp	B. Mu	R. Nieuwenhuis
T. Nipkow	D. Niwiński	V. van Oostrom
E. Orłowska	F. Otto	M. Parigot
M. Piotrów	D. Plaisted	E. Poll
L. Puel	D. Rémy	C.R. Ramakrishnan
P. Rao	S.A. Rebelsky	B. Reinert
P. Rety	M.M. Richter	C. Ringeissen
W. Sadfi	G. Salzer	A. Sattler-Klein
M. Schmidt-Schauss	K. Schulz	H. Seidl
D.J. Sherman	A. Skowron	G. Smolka
W. Snyder	R. Socher-Ambrosius	Z. Splawski
J. Steinbach	J.-M. Steyaert	J. Stuber
S. Tison	J. Tiuryn	X. Toenne
Y. Toyama	P. Urzyczyn	S. Vorobyov
U. Waldmann	I. Walukiewicz	R. Wichagen
C.P. Wirth	H. Zantema	H. Zhang

Table of Contents

INVITED TALK:

Rewrite Techniques in Theorem Proving

L. Bachmair (University of New York at Stony Brook) 1

Redundancy Criteria for Constrained Completion

C. Lynch and W. Snyder (Boston University) 2

Bi-rewriting, a Term Rewriting Technique for Monotonic Order Relations

J. Levy and J. Agusti (CSIC, Blanes) 17

A Case Study of Completion Modulo Distributivity and Abelian Groups

H. Zhang (The University of Iowa City) 32

A Semantic Approach to Order-Sorted Rewriting

A. Werner (University of Karlsruhe) 47

Distributing Equational Theorem Proving

J. Avenhaus and J. Denzinger (University of Kaiserslautern) 62

On the Correctness of a Distributed Memory Gröbner Basis Algorithm

S. Chakrabarti and K. Yelick (University of California at Berkeley) . . 77

Improving Transformation Systems for General E -Unification

M. Moser (Technical University of Munich) 92

Equational and Membership Constraints for Infinite Trees

J. Niehren (DFKI, Saarbrücken), A. Podelski (DEC, Paris) and R. Treinen (DFKI, Saarbrücken) 106

Regular Path Expressions in Feature Logic

R. Backofen (DFKI, Saarbrücken) 121

INVITED TALK:

Proving Properties of Typed Lambda Terms: Realizability, Covers, and Sheaves

J. Gallier (University of Pennsylvania at Philadelphia) 136

Some Lambda Calculi with Categorical Sums and Products

D.J. Dougherty (Wesleyan University) 137

Paths, Computations and Labels in the λ -Calculus

A. Asperti (University of Bologna) and C. Laneve (INRIA Sophia-Antipolis) 152

Confluence and Superdevelopments

F. van Raamsdonk (CWI, Amsterdam) 168

Relating Graph and Term Rewriting via Böhm Models <i>Z.M. Ariola (University of Oregon at Eugene)</i>	183
Topics in Termination <i>N. Dershowitz and C. Hoot (University of Illinois at Urbana)</i>	198
Total Termination of Term Rewriting <i>M.C.F. Ferreira and H. Zantema (University of Utrecht)</i>	213
Simple Termination is Difficult <i>A. Middeldorp (University of Tsukuba) and B. Gramlich (University of Kaiserslautern)</i>	228
Optimal Normalization in Orthogonal Term Rewriting Systems <i>Z. Khasidashvili (INRIA Rocquencourt)</i>	243
A Graph Reduction Approach to Incremental Term Rewriting <i>J. Field (IBM T.J. Watson Research Center)</i>	259
Generating Tables for Bottom-up Matching <i>E. Lippe (Software Engineering Research Centre, Utrecht)</i>	274
INVITED TALK: On Some Algorithmic Problems for Groups and Monoids <i>S. I. Adian (Steklov Mathematical Institute, Moscow)</i>	289
Combination Techniques and Decision Problems for Disunification <i>F. Baader (DFKI, Saarbrücken) and K. Schulz (University of Munich)</i>	301
The Negation Elimination from Syntactic Equational Formula is Decidable <i>M. Tajine (University Louis Pasteur, Strasbourg)</i>	316
Encompassment Properties and Automata with Constraints <i>A.-C. Caron and J.-L. Coquide and M. Dauchet (University of Lille)</i>	328
Recursively Defined Tree Transductions <i>J.-C. Raoult (IRISA, Rennes)</i>	343
AC-Complement Problems: Satisfiability and Negation Elimination <i>M. Fernández (LRI, Orsay)</i>	358
A Precedence-Based Total AC-Compatible Ordering <i>A. Rubio and R. Nieuwenhuis (University of Barcelona)</i>	374
Extension of the Associative Path Ordering to a Chain of Associative Commutative Symbols <i>C. Delor and L. Puel (LRI, Orsay)</i>	389
Polynomial Time Termination and Constraint Satisfaction Tests <i>D.A. Plaisted (University of North Carolina, Chapel Hill)</i>	405

Linear Interpretations by Counting Patterns	
<i>U. Martin (University of St Andrews, Fife)</i>	421
Some Undecidable Termination Problems for Semi-Thue Systems	
<i>G. Sénizergues (LABRI, Bordeaux)</i>	434
SYSTEM DESCRIPTIONS	
Saturation of First-Order (Constrained) Clauses with the Saturate System	
<i>P. Nivela and R. Nieuwenhuis (University of Barcelona)</i>	436
MERILL: An Equational Reasoning System in Standard ML	
<i>B. Matthews (University of Glasgow)</i>	441
Reduce the Redex \rightarrow ReDuX	
<i>R. Bündgen (University of Tübingen)</i>	446
AGG — An Implementation of Algebraic Graph Rewriting	
<i>M. Löwe and M. Beyer (Technical University of Berlin)</i>	451
Smaran: A Congruence-Closure Based System for Equational Computations	
<i>R. M. Verma (University of Houston)</i>	457
LAMBDA LG: Higher Order Algebraic Specification Language	
<i>Y. Gui and M. Okada (Concordia University, Montreal)</i>	462
OPEN PROBLEMS	
More Problems in Rewriting	
<i>N. Dershowitz (University of Illinois at Urbana), J.-P. Jouannaud (LRI, Orsay) and J.W. Klop (CWI, Amsterdam)</i>	468
Authors Index	488

Intimate Computing and the Memory Prosthesis: A Challenge for Computer Systems Research?

(Abstract)

Michael G. Lamming

Rank Xerox EuroPARC

Cambridge, England

At EuroPARC we are trying to build a human *memory prosthesis* — a portable device to help individuals remember things. It will automatically capture and organise predefined classes of information and provide easy ways to recall it when needed, perhaps without even being asked. We call this device a memory prosthesis because it augments normal human memory. It differs from most other information systems in that it focuses on helping the user recall things *they once knew*. Our objective for the memory prosthesis is to assist users with everyday memory problems. Target tasks for the memory aid include: recalling names of people, places, and procedures, finding files, papers and notes, in whatever medium they are expressed, and remembering to perform tasks.

The memory prosthesis is an example of a new class of interactive system we envisage will be made possible by forthcoming advances in micro-electronics. Using cellular radio and infrared technology computers are able to communicate with each other without wires. This new development heralds the dawn of mobile computing. At present radio transceivers are large and power hungry, so much so that the machines to which the transceivers are attached are fairly large. We are looking a short while into the future when mobile computers will be somewhat smaller, indeed small enough to be worn rather than carried — perhaps resembling a watch or piece of jewellery. We look to a time when people don't have to remember to take their computer with them, they *wear* it and take it everywhere.

Such systems will have several fundamental capabilities not previously available on such a wide scale. They will dynamically connect and communicate, not only with each other, but with office equipment, domestic appliances and much of the other business and consumer electronic equipment that surrounds us.

The wireless communication technology used by these systems will be cellular — perhaps based upon the new digital cellular telephone standards. The low-power requirements of a tiny wearable computer will limit the range to a few meters and so communication cells will be small. The consequence of relying on small cells for communication is simple yet profound, mobile computers will know where they are. To find out their location they simply ask the nearest non-mobile object.

So to summarise: computers will be small enough to wear and take everywhere; they will be embedded in domestic appliances, office and consumer equipment; they will talk to each other using cellular wireless communications; and they will know where they are.

Taken together these facilities provide us with another view of mobile computing. Popular views of mobile computing regard it as a tool providing access to information and computation whilst the owner is away from his or her desk. We view it the other way round. Our computers can now gain continuous access to *us* and our immediate environment, wherever we are. In consequence, our personal computer will be able to find out much more about us, and like any other personal assistant, the more it knows the more useful it can be. To distinguish this style of system from *personal computing*, we have coined a new phrase: *intimate computing*.

Carrying a computer around *everywhere* offers almost limitless opportunities to capture useful information. Wherever we go, whatever we do, our tiny computer can *automatically* liaise with the equipment we use to do our work, with the portable computers belonging to the people we meet, and with the devices embedded in the building where we work, to construct a detailed personal cross-reference to much of the information with which we come into contact. Indeed, one of the most likely down-sides for intimate computing is the ease with which we may drown in the incoming tide of unstructured data — *unless it is filtered and organised automatically too*.

Most personal information systems, paper-based or computer-based, require some help from the user to construct a useful database. Typically the *user* has to recognise that an item of information might be required in the future; he or she must then make the effort to *capture* it; and lastly, and perhaps most importantly, he or she has to *organise* the information in a manner that makes it easy to find it again. But to do this, the user must be able to predict the situation in which the information will be needed and think up some indexing terms which he or she guesses might plausibly spring to mind the next time the information is sought. A common problem is to guess incorrectly! For example, Mary may choose to file a useful journal article by author or title, yet subsequently only manage to recall that it was the one her boss gave to her.

This example highlights a well established feature of the human memory system — people are particularly good at recalling activities from their own lives. Psychologists call this mechanism *episodic* or *autobiographical memory*. Experiments have shown that humans are not particularly good at remembering the *time* of an episode in their life, but they are much better at remembering *where* the episode occurred, *who* they were with, or *what* they were doing. We call this the *context*.

On the other hand, computers are excellent at recording the *exact* time an item of information was created, stored, communicated or processed in some way. For example, if Mary chooses to write a note about the journal article on her portable computer, the computer will almost certainly timestamp the note for her. In fact almost every computer transaction is timestamped in some way already. Electronic files are timestamped, telephone call-times are recorded for billing, faxes have the arrival time printed on them, and even each frame of a video sequence contains a time code. Moreover, computers are very good at searching through large bodies of data for items with a particular

timestamp. So if we can give a computer system an exact timestamp it won't take very long to find all the items that are tagged with the same date and time. Yet as we have stressed already, context is fairly easy for humans to remember while exact timestamps are not. If only the context that gave rise to an item of useful information could be used by the computer to find the same item later on...

Previous work at EuroPARC has shown how this might be achieved with mobile technology. For example, Newman and his colleagues have demonstrated a technique called *episode recognition* [2]. Location data obtained from Active Badges can be used to construct automatically, a diary of an individual's life expressed in terms of their location and encounters with other members of staff and visitors. Experiments have shown that these chronicles are a powerful aid to recall, and can be used both to index, and retrieve other less memorable data collected automatically at about the same time. As a result it has been suggested that a more comprehensive diary containing richer descriptions of the user's activities might provide a useful indexing mechanism for navigating through a huge database of personal information [1]. We now believe it is possible to design a computer system in which imprecise informal yet personal memories we have for past events can be used as keys to recover detailed information about the event itself.

Clearly our primary motivation for building this system is to provide more effective support for human memory. But in doing so, we are encountering all sorts of technical problems for which we have no convenient solution. Nevertheless, our programme of work proceeds in anticipation of acceptable solutions becoming available shortly. By trying to build this demanding application we hope to create another small focus for research in computer science and engineering and highlight some of the technical challenges that lie ahead for all of us.

References

- [1] Lamming, M. G., & Newman, W. M. (1992). *Activity-based Information Retrieval: Technology in Support of Personal Memory*. In F. H. Vogt (Ed.), *Information Processing '92. Proceedings of the 12th World Computer Congress, Vol. III* pp. pp 68-81. Madrid: Elsevier Science Publishers (North-Holland).
- [2] Newman, W., Eldridge, M., & Lamming, M. (1991). *Pepys: Generating Autobiographies by Automatic Tracking*. In *Proceedings of the second European conference on computer supported cooperative work*. Amsterdam.

Active Programming Strategies in Reuse

Mary Beth Rosson and John M. Carroll

IBM T. J. Watson Research Center
Yorktown Heights, New York 10598, USA

Abstract. In order to capitalize on the potential for software reuse in object-oriented programming, we must better understand the processes involved in software reuse. Our work addresses this need, analyzing four experienced Smalltalk programmers as they enhanced applications by reusing new classes. These were *active* programmers: rather than suspending programming activity to reflect on how to use the new components, they began work immediately, recruiting code from example usage contexts and relying heavily on the system debugger to guide them in applying the borrowed context. We discuss the implications of these findings for reuse documentation, programming instruction and tools to support reuse.

1 Introduction

A key attraction of object-oriented programming languages is the potential they offer for the reuse of software components. A well-designed object class defines a tightly encapsulated bundle of state and behavior that can be “plugged into” a target application to fill some functional need — hence the popular metaphor of a “software IC” [4,5]. And while most of this potential has been asserted rather than demonstrated, empirical evidence documenting the advantages of an object-oriented language for code reuse is beginning to emerge [17]. At this point, however, we know very little about the *process* of component reuse and thus how we might best support reuse activities.

A programmer attempting to recruit existing software components for his or her current project must carry out two basic tasks. First, the candidate component(s) must be identified. This may be trivial in cases where the component was self-generated or is already familiar to the programmer (see, e.g., [6,16]). However, much of the missed potential in software reuse arises in situations where the programmer knows little or nothing about the component in advance. As component libraries increase in size, the difficulty of locating novel functionality increases commensurately. Not surprisingly, researchers have begun to apply a variety of classification and information-retrieval techniques to address the difficult problem of locating unknown functionality within large class libraries [12,21].

Once a candidate component has been identified, the programmer must incorporate the component into the ongoing project. Again, if the component is self-generated or already familiar, this process is simplified: the programmer already knows what it does and how it is used, and merely must apply this knowledge to the new situation. But for unfamiliar components, the programmer must engage in at

least some form of analysis, determining what the component does and how it can contribute to current needs, and then designing and implementing the code needed to extract the desired functionality [2, 10]. Researchers are only beginning to explore how one might document code intended for reuse (see, e.g., [14]). But from the perspective of a programmer considering reuse, one requirement is clear: understanding how to use a component must take less time and effort than (re)building the component itself. Indeed, given programmers' general preference for self-generated code, the cost of reusing a component should be considerably less than that of creating it.

This paper seeks to elaborate the requirements for reuse documentation and tool support through analysis of experts carrying out a reuse task. We observed Smalltalk programmers enhancing an application through the reuse of classes we provided. Most generally, our goal was to characterize the strategies and concerns of the programmers as they attempted to reuse the novel classes — by understanding what does and does not work well in the current reuse situation, we can begin to reason about possible modifications or enhancements. More specifically, however, we were interested in the role that examples might play in documenting reusable components. We have been researching example-based programming environments for learning and for reuse [3,13,20,22], and this empirical setting provides an opportunity to examine experts' natural strategies for finding and applying example information.

2 The Reuse Situation

Four experienced Smalltalk programmers participated in the study. All had been programming in Smalltalk/V® PM [8] for over two years, and had over 10 years of general programming experience. All had worked on user interface development in Smalltalk, largely on building components for advanced user interfaces (e.g., multimedia objects, direct manipulation techniques, visual programming).

Each programmer completed two reuse projects, in two separate sessions. The reuse situation approximated the application prototyping activities these programmers carry out in their normal work environment, in that both projects involved an enhancement to the user interface of an already-written interactive application. The applications were simple but non-trivial examples of Smalltalk projects; in debriefing sessions after the experiment, all of the programmers judged that these were representative reuse programming tasks. The order of the projects was counterbalanced — one project served as the first project for two of the programmers, the other as the first for the other two. During their second sessions, programmers were introduced to the Reuse View Matcher [22] and were allowed to use this tool while completing the project. Due to space limitations, this paper will not discuss the second set of sessions involving the Reuse View Matcher.

The programmers were read brief instructions at the beginning of each session, describing the application they were to enhance, and identifying the class they were to reuse in making this enhancement. They were told that they were not expected to spend "more than a couple of hours" on the project and that they should not worry

if they did not complete it in this amount of time. Finally, the programmers were asked to "think aloud" while they worked, to vocalize their plans and concerns as they worked as much as possible without interfering with their activities [9].

After hearing these instructions, the programmers were given an extended introduction (approximately 20 minutes) to the application to be enhanced; this involved going over a hierarchical view of the major application classes, a design diagram of application objects and their connections, descriptions of typical interaction scenarios, as well as a comprehensive walk-through of the code. The intent was to familiarize them with the application enough so that their problem-solving efforts would focus on the reuse of the new class rather than on understanding how the existing application worked. No information other than the name was provided about the class to be reused.

During the reuse task, programmers worked at their own pace in a standard Smalltalk/V PM environment. The experimenters took notes and made videotapes of the programming activity on the display, occasionally prompting the programmer to comment on a plan or concern. All projects were completed within one and a half to two hours.

2.1 The Color-Mixer Project

One of the projects consisted of an enhancement to a color-mixer. The color-mixer converts rgb values input by the user to create custom colors; these colors are stored in and retrieved from a database of named colors. The original application has three buttons for red, green and blue (see Figure 1); clicking one of these buttons brings up a dialog box in which the user types an integer to manipulate a color component. The color being edited is displayed as a "swatch", and is flanked by the list of saved colors. Users can select colors from the list, as well as adding and deleting colors.

Because everything in Smalltalk is an object, and because objects typically inherit a good deal of their functionality, it is difficult to characterize the "size" of applications. However, the most important objects in the color-mixer are instances of six classes (see Figure 1): ColorMixer, ColorMixWindow, ButtonPane, ListPane, GraphPane and Dictionary. The last four classes in the list are components of the standard library. The number of methods in these six classes ranges from six to 54, with an additional 118 to 338 inherited methods.

The programmer's task was to replace the button-dialog box input style with horizontal sliders. No information was provided concerning the appearance or functionality of the slider, only that they were to use the new class HorizSliderPane. A typical solution involves the editing of the existing `openOn:` method (this is the method that creates and initializes the windows and subpanes, and the button creation code must be replaced with analogous code for the sliders), and the addition of four new methods (to handle activity in each of the sliders, and to draw any given slider).

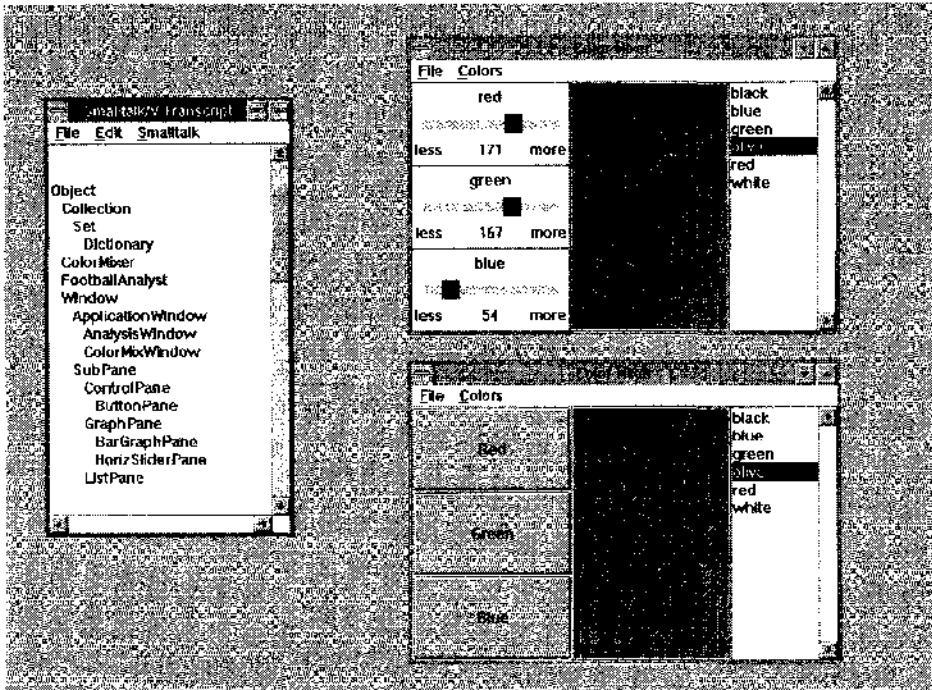


Figure 1. The Color-mixer Project: On the left is a listing of the major classes involved in the color-mixer and football analyst applications; indentation in the list signifies superclass-subclass relationships. In the upper right is the original color-mixer; beneath it is the application enhanced to use sliders as input devices.

The class library included an example application already making use of HorizSliderPane. The example usage was a football analysis program, in which five sliders are used to manipulate defensive player characteristics (e.g., speed, age, height), and the predicted consequences of the characteristics (e.g., sacks, interceptions, tackles) are graphed in a separate pane. This application uses five main classes (FootballAnalyst, HorizSliderPane, BarGraphPane, AnalysisWindow, and Dictionary; only Dictionary is part of the standard library; see Figure 1); method count ranges from five to 33, with from 118 to 363 inherited methods. Because one of our research goals was to examine experts' strategies for discovering and employing example usage information, the programmers were not told of the example application in advance.

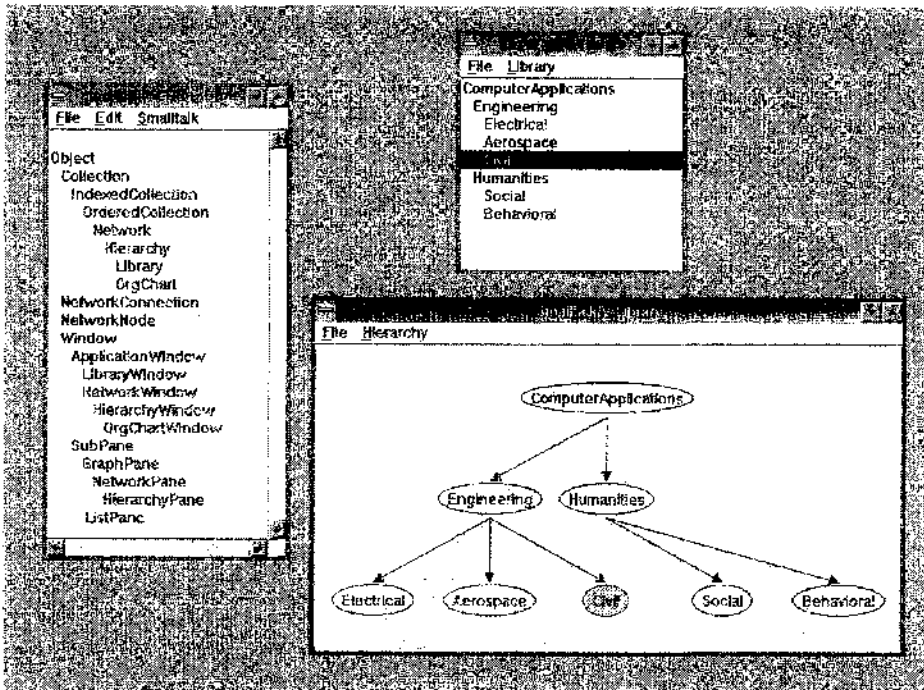


Figure 2. The Library Project: On the left is a listing of the major classes involved in the library and organization chart applications; indentation in the list signifies superclass-subclass relationships. In the upper right is the original library application; beneath it is the application enhanced to use a graphical hierarchy.

2.2 The Library Project

The second project consisted of enhancements to a library acquisitions application. This application manages a hierarchical collection of book categories (e.g., Computer Applications broken into Electrical Engineering, Aerospace Engineering, etc.); categories are annotated with information about acquisitions (e.g., number of books, titles). Hierarchical structure is conveyed via an indented list (see Figure 2), and users manipulate the categories by selecting a list item and making menu selections. In this way, they can add and delete categories, rename categories, and browse and edit the acquisitions information.

The library project uses five main classes (Library, NetworkNode, ListPane, NetworkConnection and LibraryWindow; only ListPane is part of the standard hierarchy, and the Library class inherits from two novel superclasses, Network and Hierarchy; see Figure 2). The method count for these five classes ranges from 4 to 54, with from 118 to 319 inherited methods.

Programmers were asked to enhance this project by using the new class HierarchyPane; again, they were told nothing of the appearance or functionality of

the target class. An instance of this class is able to graph a hierarchical network of nodes (see Figure 2). It also can identify nodes or connections selected via a mouse click. Finally, the subpane allows users to name nodes by typing directly onto the graphed elements.

HierarchyPane differs from HorizSliderPane, in that much of its functionality is inherited from its superclass NetworkPane. Further, it was designed to work in concert with a number of other novel classes (HierarchyWindow, Node, NetworkNode and NetworkConnection), whereas HorizSliderPane is a relatively "standalone" component. A typical solution for reusing HierarchyPane in the library application involves creation of a new LibraryWindow class as a subclass of HierarchyWindow (thereby inheriting the ability to draw, select, and name nodes in the graph), and the updating of five methods from the original LibraryWindow class (the methods for adding, removing and showing acquisitions for a selected category, the method defining the menu, and the `openOn:` method).

As for the color-mixer project, the class hierarchy included an example usage of HierarchyPane — an organization chart, in which the nodes correspond to employees, and in which employees of various job descriptions (e.g., staff member, secretary, visitor) can be added to the hierarchy, given names, reassigned, and given project descriptions. The example uses seven main classes (OrgChart, OrgChartWindow, HierarchyPane, Node, NetworkNode, NetworkConnection, and TextField; none of these are part of the standard hierarchy, and both OrgChartWindow and HierarchyPane inherit from novel superclasses; see Figure 2). The method count for these classes ranges from 1 to 37, with inherited methods ranging from 118 to 442. Programmers were not told in advance about the HierarchyPane usage example.

3 Reuse of Uses

In most discussions of component reuse in object-oriented systems, the focus has been on the class or classes reused. Design methodologies attempt to articulate characteristics of reusable classes [15,18] and tool builders develop techniques for classifying and retrieving useful classes [12,21]. The dominant metaphor is "construction" — the programmer finds parts that can be reused, modifies them as necessary and connects them together (see, e.g., [4,5]).

Our observations suggest that this focus on components may be over-simplified. To develop the knowledge needed to reuse the components directly, the programmers would have had to stop work on their overarching goal — enhancing the project they had been given — and spend time analyzing and reflecting on the target class. These programmers were too focussed on their end goal to engage in protracted analysis. Instead, they made active use of all resources available in the environment, and began programming immediately. This led them to reuse the components only indirectly, through the reuse of "uses". That is, the main entity participating in the reuse programming was not the target class but rather the *example application of that class*. The programming consisted of finding and reusing the patterns of component

reuse reified in the example application. As one programmer put it, on discovering the example application, "so there's a solution in the system!"

The extensive reuse of the example occurred despite mixed feelings expressed by the programmers. There was a sense that this wasn't the "right" way to reuse a class, that it was somehow cheating or taking the easy way out. One programmer said that he would look at the example only if all else failed, but then immediately began to work with it. Another viewed the example as a mixed blessing, because although it offered information on how to use the target class, it now required analysis itself: "Whenever you provide help, you provide trouble, now I have to understand this!" However, when probed about these feelings at the close of the experiment, the programmers indicated that the strategy of borrowing heavily from examples is one they use frequently in prototyping Smalltalk applications, and that their reservations were due to a perceived demand to use more conservative methods in this experimental situation.

Smalltalk provides explicit support for the identification and reuse of example usage context through its "senders" query which returns a list of methods in which a target message is sent. An experienced programmer can browse this list and make reasonable guesses as to which other classes if any are already using the class of interest; if motivated, they can then explore these other classes to discover why and how the target class is being used. All of the programmers made early and repeated use of the senders query; further, they showed an ability to discriminate among the various messages defined for the target class, asking for senders only on the more important methods (e.g., a method providing the contents for the subpane): "AnalysisWindow seems to be figuring prominently as a sender of interesting messages".

3.1 Reusing Pieces of an Example

The most common reuse of the example applications consisted of borrowing code used as the interface to the target class, both blocks of code copied out of methods and entire methods. For instance, all of the programmers borrowed code from the example applications' `openon:` methods; by convention this is a message sent to a window which instantiates the various subpanes, defining their graphical and behavioral characteristics. The instantiation of subpanes in Smalltalk/V is often complex, and typically includes the definition of events that the subpane will handle. Thus copying an instantiation code snippet (8-15 lines of code) can save considerable time in working out exactly how a new kind of subpane needs to be initialized.

Sometimes the borrowed code was not directly reusable itself, but rather was used more as a functional specification. In working out slider event handling for the color-mixer project, the programmers copied over the `sliderActive:` method from the football program. This method does three things: first, the affected slider processes the mouse activity; second, the relevant player characteristic is updated; and third, predicted player performance is graphed. Only the first of these events maps directly to (and thus could be reused in) the color-mixer project. Nonetheless,

the programmers were able to understand the code in `sliderActive:` as a specification of what they needed to do in their own version: process slider activity, re-set the model data (in this case, the color settings), and display the results (the new color swatch).

On a few occasions, the borrowed code came from work the programmers had just completed themselves (as in the "new code reuse" situation described in [6]). For example, both programmers working on the color-mixer first developed the code for one slider, then worked from that code to implement the other two. In these cases, the programmers knew exactly what needed to be changed, and the "programming" consisted simply of the physical edits.

In general, the copy/edit strategy worked quite well (see also [16]). It reduced the amount of typing required of the programmer, and helped to insure that the details of the code (e.g., placement of line separators) would be correct. More importantly, it removed the burden of analyzing the target class enough to generate the correct protocol for a particular usage situation, enabling a rapid programming progress. For many parts of the borrowed protocol (e.g., the event definitions in the `openOn:` method), the programmers knew what parts of the code needed to be edited and how to do this.

However, the copy/edit strategy did lead to some problems stemming from the novel parts of the target class' protocol, in that the programmers were now able to copy and "use" protocol that they didn't fully understand. A good example comes from one programmer's work on the color-mixer. In the football analyst example, each slider is instantiated with a different starting value. Because the slider instantiation code was copied from the football `openOn:`, instantiation of the `value` variable also became part of the color-mixer `openOn`. The `value` attribute is not generic to subpanes, so the programmer did not know off-hand whether it was prerequisite to slider functioning, and if so, what a reasonable starting value would be for the color-mixer. The programmer did not know enough about the protocol for sliders to answer these questions, so he simply made a guess. Later on, this guess caused problems, as the initial positions of the sliders did not match the starting color (white). Subsequently, the programmer solved the problem not by going back and correcting the initialization code, but rather by adding code at a later point that simulated the selection of white in the color list pane.

In some cases there was a conflict between the component interface suggested by the example, and the current design of the project. In the football program, the activity of all the sliders is handled by a single method `sliderActive:`. Modeling on the example, one of the programmers began by copying over the method and modifying it to refer to color-mixer objects. However, in the course of doing this, he recognized that there would be a problem in discriminating among the different slider instances. Despite the suggestion by the football example that multiple sliders could be managed by one method, he decided to change his approach and work from the more familiar model of the buttons used by the original user interface. Noting that three separate methods had been written to handle button activity, he developed an analogous set of three slider activity methods.

The Smalltalk environment is very supportive of the copying/editing of example usage code. Programmers can open as many code browsers as they like, and can freely select and paste text among them. In this study, the programmers almost always had at least two browsers open (one for the example and one for the project) and often used more when the code involved a number of embedded messages. In this way, they could preserve their top-level context while going off to answer a question or to find additional relevant code in other classes or methods.

3.2 Reusing an Application Framework

All of the programmers' initial efforts to reuse the example application involved bringing methods or pieces of methods *from* the example application *into* the project. However, the two programmers working on the library project ultimately decided to create a new kind of library window, one that was a sibling of `OrgChartWindow` (i.e., had `HierarchyWindow` as a superclass, this was in fact the solution requiring least programming effort). In doing this, they were deciding to *inherit* rather than borrow from the example usage context. After this decision, their activity shifted, as they began bringing code from the original library window into the new window. This was in marked contrast to the programmers working on the color-mixer project, who appeared to never even consider inheriting functionality from the football example.

The decision to subclass reflects a desire to reuse more than just the snippets of code involving the target class; in this case, the programmers elected to adopt the entire application context of the example. In Smalltalk/V PM, this context is normally managed by a window; the window communicates with the underlying application objects (e.g., a hierarchical collection of employees) and with the subpanes used to display application information. Thus reuse of the context can be accomplished by subclassing the application window; reuse of this sort is often referred to as reuse of an "application framework" [7]. Framework reuse brings along the component of interest "for free" in some sense, in that it is already a component of the framework, and the example window already has the code needed to interface between the component and other application objects.

Deciding to reuse the example's application framework had a remarkable effect on the programmers' reuse efforts. What had at first been a rather complex process of tracking down individual methods and instance variables distributed across `NetworkWindow`, `HierarchyWindow` and `OrgChartWindow`, and copying and editing methods or pieces of methods, now became a straightforward process of copying over and updating the menu functions from the original `LibraryWindow` class. One of the programmers spent over an hour reaching the decision to subclass; once he did, he was rather frustrated at the thought of throwing away all the work he had done so far, but even so was able to complete the project in fifteen minutes.

The problems of tracking down functionality distributed throughout an inheritance hierarchy have been noted before; Taenzer, Ganti and Podar [23] refer to this as the "yoyo" problem. The Smalltalk/V class hierarchy browser offers little support for dealing with hierarchically distributed function, as programmers must

navigate from superclass to superclass in search of methods. Taenzler et al. [23] point to this problem as an argument against reuse via inheritance, suggesting that understanding how to subclass an extensive hierarchy requires much more distributed code analysis than simply reusing a component. Our situation offers a new twist on considering whether to reuse functionality directly or through inheritance: when a component has already been incorporated into a rich application framework, programmers may find that indirect inheritance of the component's functionality (i.e., through subclassing the framework) will simplify enormously the task of reusing the component.

Several factors seemed to contribute to the programmers' decision to reuse the application framework for graphical hierarchies. One was simply the difficulties in tracking down, borrowing and integrating function. There seemed to be a sense that the process was more complicated than it should be, e.g., "I should probably be trying to inherit some of this...". When asked later, one of the programmers indicated that it was his realization of how many of his borrowed methods were inherited from superclasses of `OrgChartWindow` that made him decide to move the library window. For the other programmer, a critical incident was his effort to compile a key method (the one allowing selection of nodes in the graph), and discovering a instance variable of the example window that had no analog in the library application. Up to that point, he had seemed willing to work with the complexity of tracking down and borrowing example protocol, but adding a new (and mysterious) piece of state information was too much.

Another factor may have been the similarity between the example usage and the project. On first discovering the `HierarchyWindow` class, one programmer tried a simple experiment while voicing his belief that it would never work: he tried opening a `HierarchyWindow` "on" the library object (an instance of `Library`, part of the `Collection` hierarchy). To his (and our!) surprise, this experiment was successful. Of course, the `LibraryWindow` functionality was not present, but at least the book collection was displayed in a nice graphical hierarchy. This experiment may seem extreme, in that it has a rather low probability of pay-off. However, it was simple to do, and it provided the programmer with considerable insight into the example application that he was able to apply to his later efforts.

The subclassing strategy did simplify the reuse programming project. However, it also introduced some rather subtle problems. There was considerable overlap in the functionality of the example and of the library (e.g., both had facilities for adding and removing elements in a hierarchy, for renaming these elements). One of the programmers, having decided to subclass, wanted to inherit as much functionality as he could. So, when updating the menu selections, rather than copying over the methods from the original library window and editing them to work in this new context, he first tried simply inheriting the methods defined in the superclasses. On the surface, this strategy seemed to work — he was now able to add and delete library categories and rename them. He never realized that the underlying library structure was not being manipulated correctly (the relationships among categories weren't being specified). It may be that programmers following a subclassing strat-

egy are more likely to satisfice, accepting generic inherited functionality that is almost but not quite right simply because it is there and is already working.

4 The Reuse Programming Process

The programmers were opportunistic in the objects of reuse — extensive recruitment of the example contexts reduced considerably the amount they needed to learn about the target class. But they were also opportunistic in how they went about doing the reuse task. They spent little time in deliberated analysis of the example, in understanding how it was going to help or interfere with their enhancement efforts. Rather they began using the code of the example immediately to make progress on their goal. These were *active users* of Smalltalk [1]: as has often been observed for human problem-solving [11] the process we observed was very locally driven, with specific features of the environment and the evolving solution determining each succeeding step.

4.1 Getting Something to Work With

An early goal for all of the programmers was to get an instance of the target class up and running, so that they could see what it looked like. One of the programmers working on the library project was able to use the organization chart example to do this. After discovering the example, he immediately took on the goal of starting it up. He found an `OrgChart` class method `fromUITData`, the name of which signalled to him that it was a special “set-up” method, and that he could use it to create an appropriate `OrgChart` object and start up the application. By doing this, he was able to see what a `HierarchyPane` looked like, as well as to experiment with the interaction techniques it supported.

With respect to programming activities, the focus of initial efforts for all of the programmers was on modifying the project's `openOn:` method to include the new class: “I want to get one of these things as a subpane”. However, while there was some browsing of the target class methods to see how to do this, the browsing tended to yield inferences about class functionality rather than usage protocol; as we noted earlier, the programmers seemed to resist carrying out an analysis of the target class comprehensive enough to allow them to write code to instantiate it for their project. Instead, they sometimes looked for clues in the code they were replacing. Thus the two programmers working on the color-mixer examined the code used to create the buttons, thinking about how they might modify it for sliders (e.g., what events a slider might handle in contrast to a button).

One programmer working on the color-mixer tried to take advantage of other code in the `openOn:` method as well. Noting that `HorizSliderPane` is a subclass of `GraphPane`, he examined the code instantiating the color swatch (an instance of `GraphPane`), thinking that he might be able to build a slider definition from it. This led to a variety of problems, as he began to hypothesize that the slider functionality was somehow built from the scroll bars present in every subpane, and that the pro-

tool controlling these scroll bars for GraphPanels must be critical in creating sliders. This was certainly a reasonable hypothesis on functional grounds, but in fact was quite misleading.

The programmers seemed to feel that successfully instantiating the target class within the project context was a momentous event. It appeared that this was considered to be the major hurdle of the project, and now they could get on with business as usual, adding the remainder of the component's functionality (i.e., its event handling). One explanation for this is that the programmers could "see their end goal in sight" — a new and improved view of their project data. But another equally important factor is that by instantiating the new component as part of the project, the programmers could now rely much more on the environment to guide their programming. In a Smalltalk application, objects are created and code references are established only when the application is run, making the code alone inherently ambiguous and mental simulation of it difficult. In contrast, if the programmer is able to start up an application, all ambiguities in the code are resolved, and the programmer can use Smalltalk's sophisticated interactive debugging tools to analyze and modify the code.

4.2 Debugging into Existence

We have seen that the programmers relied heavily on code already in the environment in attacking the reuse projects. But they also relied heavily on the tools of the environment to locate and make sense of the relevant code. In particular, they repeatedly started up the application they were working on, and looked to see where it "broke" to plan their next move.

Smalltalk is particularly supportive of this debugging-centered style of program construction. The language is non-typed and compiled incrementally, which permits rapid and repeated experimentation with the code used to run an application. The debugger and inspector tools support such experimentation directly, providing flexible access to and manipulation of the runtime context for an application (i.e., objects and their state, messages in progress).

In some cases, the programmers knew something of the steps they would need to take, but used the debugger to help them in carrying these out. Thus, once they had copied the instantiation code from the example application's `openon:`, they knew that certain modifications would be necessary: instance variable names needed to be changed, the menu name needed to be changed, the project would need a drawing method, etc. Some of the programmers even carried out some anticipatory activity, perhaps creating a method that they knew they would need, but that they also knew was not yet functional. However, for the most part, they relied on the system to detect the absence of methods or the inappropriate states of objects. In a typical scenario, the programmer would start up the project application, receive a "message not understood" error, return to the example in search of a method with that name, copy the method, perhaps making a few changes, try again and see how far it got, make some changes and try again. This sort of cycle might be repeated

many times, but the programmers seemed comfortable with it, and seemed confident that they were making progress.

In other cases, the debugger was used to untangle more subtle problems. So, for example, the superclass `HierarchyWindow` uses the `network` instance variable to point to the main application object, whereas the original `LibraryWindow` class uses `library`. A thorough analysis of the example would have revealed the relevant mapping between these two variables. However, the two programmers working on this project simply borrowed the example code as-is and used the debugger to ascertain what role the `network` variable was playing and how to provide this information within their project.

The compiler was used in this opportunistic fashion as well. When dealing with complex pieces of borrowed code, the programmers often would attempt to compile the code before they had completed editing it. The system would flag variable names not defined for the class (e.g., the `HierarchyWindow` code refers to `graphPane`, while the `LibraryWindow` uses `pane`), and the programmers would then replace the unknown name with the name of the analogous variable. This minimized the amount to which they needed to read through and analyze the unfamiliar code.

5 Summary and Implications

Our observations describe a process of component reuse in which the component is reused only indirectly, through the reuse of its “uses” — bits of protocol or even entire application frameworks. The programmers we studied pursued this style of reuse piecemeal and opportunistically; they focused initially on getting a runnable albeit skeletal result which they could exercise and improve incrementally, relying heavily on interactive debugging. We have characterized these as “active” programming strategies, an orientation in which programmers directly and immediately enlist and transform their software materials in favor of withdrawing from such activity to analyze and plan.

5.1 Scope of Active Reuse

This work was exploratory empirical research in its scope and scale. It addressed a particular programming situation, application prototyping, which may differ significantly from other situations. However, at least some of our observations are consistent with studies of other reuse situations. Lange and Moher [16] observed that an experienced programmer extending a library of software components was quite likely to use existing components with related functionality as templates or models for the new components. Detienne [6] found that programmers designing and implementing new applications sometimes reused their own code as they worked. Interestingly, the programmers in this study chose not to borrow code from other applications, perhaps because the other applications available were only peripherally related to the problems being solved.

Further research is needed to assess the generality of the more specific strategies we observed. All four programmers relied extensively on the system tools to organize their work, using multiple browsers to maintain their context across different parts of the hierarchy, and using the debugger and inspector to track down and modify missing or inappropriate pieces of borrowed code. It is not clear though what the boundary conditions for such an approach might be — it may be that they are only likely to occur in a tool-rich interpreted environment like Smalltalk.

Some strategies were unique to a particular programmer. For example, only one programmer made the effort to “run” the example application before borrowing code from it. He felt that this gave him a chance to preview the functionality he would be incorporating; it may be that across a wider variety of reuse projects, perhaps involving more complex components, such a strategy would be more prevalent. In another case, one programmer experimented with opening a graphical HierarchyWindow “on” his application data. The success of this experiment conveyed a great deal to him about what the graphical network framework expected in terms of data structures. It is important to understand the generality of such techniques and strategies.

5.2 Consequences of Active Reuse

Beyond the question of generality, we can ask about the consequences of the active programming strategies we observed. For example, two of the programmers did not produce a perfectly correct result, and it is not clear whether or how their problems would have been detected and corrected given unlimited time, or given instructions that emphasized the accuracy of the result. Indeed, the active programming we observed may be inadvisable from a software engineering perspective, if the small errors or inefficiencies introduced by reliance on example code are very difficult to unearth subsequently. Further research is needed to determine what if any strategies experts have developed for minimizing this downside inherent in reuse by example.

It is important not to lose sight of the main benefit of this style of software reuse: these active strategies reflect a creative and effective resolution of the inherent tension between the need to distance oneself from one's own project to study someone else's code, and continuing to make concrete progress toward a desired result. Elsewhere we have characterized such a tension as the “production paradox” [1], wherein users are too focussed on the product they are creating to acquire the skills that will facilitate its creation. In this Smalltalk reuse setting, the programmers' borrowing of example code allowed them to quickly incorporate at least some approximation of the new functionality into their own project; they could then work within their own project context to “learn” the minimum necessary for successful reuse.

5.3 Training and Tools for Active Reuse

Our work has a variety of implications for how objected-oriented programming should be conceptualized, taught and supported. Most generally, it suggests the desirability of a broader view of component reuse: the pluggable “software IC” metaphor [4,5] is not the only way reuse has been conceptualized, but it is a dominant image in talking and thinking about reuse. Both of our target objects (the slider and the graphical hierarchy) could be used as pluggable components; the slider, in particular, is an interface widget and eminently pluggable. However, all four programming projects described here reused the target classes through use of some or all of their example usage contexts. This suggests a more situational view of reuse in which pluggable, context-free reuse is the simple and ideal case.

The programmers we studied invented the strategies we observed or learned them informally from colleagues. As we noted, they occasionally expressed some embarrassment at their own reluctance to fully analyze code they wanted to reuse and their predilection for “stealing” usage protocol. If these practices survive — indeed emerge from — the natural selection pressures of professional programming, we should at least consider that perhaps they should be the topic of instruction in (Smalltalk) programming.

This implication for instruction entrains a related implication for the documentation of software components. Our four programmers were able to find example uses of the target classes, but in many situations this would not be true, and hence an example-oriented reuse strategy would be thwarted. Of course, imagining example-based documentation on a large scale raises many consequent issues. Who will build the examples? One resource is the test programs built in the course of development, and often discarded afterward. Delivering these along with software components would provide some support for the example-oriented strategy at virtually no cost. Another question is what makes a good example. There is a literature on concept formation in cognitive psychology that addresses the issue of how examples are abstracted in comprehension [19]. It is an interesting and open question whether and how similar characteristics bear on reuse.

Finally, this work embodies three themes for tool support: the sequence of activities in reuse programming, recruitment of example usage code, and the use of the system debugger. Our four programmers seemed to follow a loose script: first they instantiated the component in the project context, then they successively elaborated it function by function. Throughout this process they made extensive use of example usage contexts and of the debugger. An obvious implication is to provide tools that more explicitly integrate and coordinate the information needed at each point along the way. Thus tool support might guide reuse activities through a reuse script (for example, a list of target class behaviors to instantiate in the project context), using this script to coordinate the programmer's work with the example usage code, the project code, and the interactive debugging facilities.