



Advanced CORBA® Programming with C++
Michi Henning
Steve Vinoski
Publisher: Addison Wesley
First Edition February 12, 1999
ISBN: 0-201-37927-9, 1120 pages

Book for “Rain Manager”, IT-SC



Enjoy the life together.

Review

Here is the CORBA book that every C++ software engineer has been waiting for. *Advanced CORBA® Programming with C++* provides designers and developers with the tools required to understand CORBA technology at the architectural, design, and source code levels. This book offers hands-on explanations for building efficient applications, as well as lucid examples that provide practical advice on avoiding costly mistakes. With this book as a guide, programmers will find the support they need to successfully undertake industrial-strength CORBA development projects.

The content is systematically arranged and presented so the book may be used as both a tutorial and a reference. The rich example programs in this definitive text show CORBA developers how to write clearer code that is more maintainable, portable, and efficient. The authors' detailed coverage of the IDL-to-C++ mapping moves beyond the mechanics of the APIs to discuss topics such as potential pitfalls and efficiency. An in-depth presentation of the new Portable Object Adapter (POA) explains how to take advantage of its numerous features to create scalable and high-performance servers. In addition, detailed discussion of advanced topics, such as garbage collection and multithreading, provides developers with the knowledge they need to write commercial applications.

Other highlights:

In-depth coverage of IDL, including common idioms and design trade-offs

Complete and detailed explanations of the Life Cycle, Naming, Trading, and Event Services

Discussion of IIOP and implementation repositories

Insight into the dynamic aspects of CORBA, such as dynamic typing and the new DynAny interfaces

Advice on selecting appropriate application architectures and designs

Detailed, portable, and vendor-independent source code

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book and Addison-Wesley was aware of the trademark claim, the designations have been printed in initial caps or all caps.

The authors and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein. The publisher offers discounts on this book when ordered in quantity for special sales. For more information, please contact:

Corporate, Government, and Special Sales

Addison Wesley Longman, Inc.

One Jacob Way

Reading, Massachusetts 01867

(781) 944-3700

Library of Congress Catalog-in-Publication Data

Henning, Michi

Advanced CORBA® Programming with C++ / Michi Henning, Steve Vinoski.

p. cm. — (Addison-Wesley professional computing series)

Includes bibliographical references and index.

ISBN 0-201-37927-9

1. C++ (Computer program language) 2. CORBA (Computer architecture)

I. Vinoski, Steve. II. Title. III. Series.

QA76.73.C153 H4581999

005.13'3—dc21

98-49077

CIP

Copyright © 1999 by Addison Wesley Longman, Inc.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form, or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior consent of the publisher. Printed in the United States of America.

Published simultaneously in Canada.

Text printed on recycled and acid-free paper.

4 5 6 7 8 9 10—CRS—0302010099

Third printing, August 1999

Dedication

To Anni and Harry, for setting me on the path.

To Jocelyn, for letting me follow it.

—Michi

To Cindy, my wife and best friend—for your sacrifices, your support, your patience, and your love.

—Steve

[Preface](#)

[Prerequisites](#)

[Scope of this Book](#)

[Acknowledgments](#)

[Michi's Acknowledgments](#)

[Steve's Acknowledgments](#)

[1. Introduction](#)

[1.1 Introduction](#)

[1.2 Organization of the Book](#)

[1.3 CORBA Version](#)

[1.4 Typographical Conventions](#)

[1.5 Source Code Examples](#)

[1.6 Vendor Dependencies](#)

[1.7 Contacting the Authors](#)

[I: Introduction to CORBA](#)

[2. An Overview of CORBA](#)

[2.1 Introduction](#)

[2.2 The Object Management Group](#)

[2.3 Concepts and Terminology](#)

[2.4 CORBA Features](#)

[2.5 Request Invocation](#)

[2.6 General CORBA Application Development](#)

[2.7 Summary](#)

[3. A Minimal CORBA Application](#)

[3.1 Chapter Overview](#)

[3.2 Writing and Compiling an IDL Definition](#)

[3.3 Writing and Compiling a Server](#)

[3.4 Writing and Compiling a Client](#)

[3.5 Running Client and Server](#)

[3.6 Summary](#)

[II: Core CORBA](#)

[4. The OMG Interface Definition Language](#)

[4.1 Chapter Overview](#)

[4.2 Introduction](#)

[4.3 Compilation](#)

[4.4 Source Files](#)

[4.5 Lexical Rules](#)

[4.6 Basic IDL Types](#)

[4.7 User-Defined Types](#)

[4.8 Interfaces and Operations](#)

[4.9 User Exceptions](#)

[4.10 System Exceptions](#)

[4.11 System Exceptions or User Exceptions?](#)

[4.12 Oneway Operations](#)

[4.13 Contexts](#)

[4.14 Attributes](#)

- [4.15 Modules](#)
- [4.16 Forward Declarations](#)
- [4.17 Inheritance](#)
- [4.18 Names and Scoping](#)
- [4.19 Repository Identifiers and pragma Directives](#)
- [4.20 Standard Include Files](#)
- [4.21 Recent IDL Extensions](#)
- [4.22 Summary](#)

[5. IDL for a Climate Control System](#)

- [5.1 Chapter Overview](#)
- [5.2 The Climate Control System](#)
- [5.3 IDL for the Climate Control System](#)
- [5.4 The Complete Specification](#)

[6. Basic IDL-to-C++ Mapping](#)

- [6.1 Chapter Overview](#)
- [6.2 Introduction](#)
- [6.3 Mapping for Identifiers](#)
- [6.4 Mapping for Modules](#)
- [6.5 The CORBA Module](#)
- [6.6 Mapping for Basic Types](#)
- [6.7 Mapping for Constants](#)
- [6.8 Mapping for Enumerated Types](#)
- [6.9 Variable-Length Types and _var Types](#)
- [6.10 The String_var Wrapper Class](#)
- [6.11 Mapping for Wide Strings](#)
- [6.12 Mapping for Fixed-Point Types](#)
- [6.14 Mapping for Sequences](#)
- [6.15 Mapping for Arrays](#)
- [6.16 Mapping for Unions](#)
- [6.17 Mapping for Recursive Structures and Unions](#)
- [6.18 Mapping for Type Definitions](#)
- [6.19 User-Defined Types and _var Classes](#)
- [6.20 Summary](#)

[7. Client-Side C++ Mapping](#)

- [7.1 Chapter Overview](#)
- [7.2 Introduction](#)
- [7.3 Mapping for Interfaces](#)
- [7.4 Object Reference Types](#)
- [7.5 Life Cycle of Object References](#)
- [7.6 Semantics of _ptr References](#)
- [7.7 Pseudo-Objects](#)
- [7.8 ORB Initialization](#)
- [7.9 Initial References](#)
- [7.10 Stringified References](#)
- [7.11 The Object Pseudo-Interface](#)
- [7.12 _var References](#)
- [7.13 Mapping for Operations and Attributes](#)
- [7.14 Parameter Passing Rules](#)
- [7.15 Mapping for Exceptions](#)
- [7.16 Mapping for Contexts](#)

[7.17 Summary](#)

[8. Developing a Client for the Climate Control System](#)

[8.1 Chapter Overview](#)

[8.2 Introduction](#)

[8.3 Overall Client Structure](#)

[8.4 Included Files](#)

[8.5 Helper Functions](#)

[8.6 The main Program](#)

[8.7 The Complete Client Code](#)

[8.8 Summary](#)

[9. Server-Side C++ Mapping](#)

[9.1 Chapter Overview](#)

[9.2 Introduction](#)

[9.3 Mapping for Interfaces](#)

[9.4 Servant Classes](#)

[9.5 Object Incarnation](#)

[9.6 Server main](#)

[9.7 Parameter Passing Rules](#)

[9.8 Raising Exceptions](#)

[9.9 Tie Classes](#)

[9.10 Summary](#)

[10. Developing a Server for the Climate Control System](#)

[10.1 Chapter Overview](#)

[10.2 Introduction](#)

[10.3 The Instrument Control Protocol API](#)

[10.4 Designing the Thermometer Servant Class](#)

[10.5 Implementing the Thermometer Servant Class](#)

[10.6 Designing the Thermostat Servant Class](#)

[10.7 Implementing the Thermostat Servant Class](#)

[10.8 Designing the Controller Servant Class](#)

[10.9 Implementing the Controller Servant Class](#)

[10.10 Implementing the Server main Function](#)

[10.11 The Complete Server Code](#)

[10.12 Summary](#)

[11. The Portable Object Adapter](#)

[11.1 Chapter Overview](#)

[11.2 Introduction](#)

[11.3 POA Fundamentals](#)

[11.4 POA Policies](#)

[11.5 POA Creation](#)

[11.6 Servant IDL Type](#)

[11.7 Object Creation and Activation](#)

[11.8 Reference, ObjectId , and Servant](#)

[11.9 Object Deactivation](#)

[11.10 Request Flow Control](#)

[11.11 ORB Event Handling](#)

[11.12 POA Activation](#)

[11.13 POA Destruction](#)

[11.14 Applying POA Policies](#)

[11.15 Summary](#)

[12. Object Life Cycle](#)

[12.1 Chapter Overview](#)

[12.2 Introduction](#)

[12.3 Object Factories](#)

[12.4 Destroying, Copying, and Moving Objects](#)

[12.5 A Critique of the Life Cycle Service](#)

[12.6 The Evictor Pattern](#)

[12.7 Garbage Collection of Servants](#)

[12.8 Garbage Collection of CORBA Objects](#)

[12.9 Summary](#)

[III: CORBA Mechanisms](#)

[13. GIOP, IIOP, and IORs](#)

[13.1 Chapter Overview](#)

[13.2 An Overview of GIOP](#)

[13.3 Common Data Representation](#)

[13.4 GIOP Message Formats](#)

[13.5 GIOP Connection Management](#)

[13.6 Detecting Disorderly Shutdown](#)

[13.7 An Overview of IIOP](#)

[13.8 Structure of an IOR](#)

[13.9 Bidirectional IIOP](#)

[13.10 Summary](#)

[14. Implementation Repositories and Binding](#)

[14.1 Chapter Overview](#)

[14.2 Binding Modes](#)

[14.3 Direct Binding](#)

[14.4 Indirect Binding via an Implementation Repository](#)

[14.5 Migration, Reliability, Performance, and Scalability](#)

[14.6 Activation Modes](#)

[14.7 Race Conditions](#)

[14.8 Security Considerations](#)

[14.9 Summary](#)

[VI: Dynamic CORBA](#)

[15. C++ Mapping for Type any](#)

[15.1 Chapter Overview](#)

[15.2 Introduction](#)

[15.3 Type any C++ Mapping](#)

[15.4 Pitfalls in Type Definitions](#)

[15.5 Summary](#)

[16. Type Codes](#)

[16.1 Chapter Overview](#)

[16.2 Introduction](#)

[16.3 The TypeCode Pseudo-Object](#)

[16.4 C++ Mapping for the TypeCode Pseudo-Object](#)

[16.5 Type Code Comparisons](#)

- [16.6 Type Code Constants](#)
- [16.7 Type Code Comparison for Type any](#)
- [16.8 Creating Type Codes Dynamically](#)
- [16.9 Summary](#)

[17. Type DynAny](#)

- [17.1 Chapter Overview](#)
- [17.2 Introduction](#)
- [17.3 The DynAny Interface](#)
- [17.4 C++ Mapping for DynAny](#)
- [17.5 Using DynAny for Generic Display](#)
- [17.6 Obtaining Type Information](#)
- [17.7 Summary](#)

[V: CORBAservices](#)

[18. The OMG Naming Service](#)

- [18.1 Chapter Overview](#)
- [18.2 Introduction](#)
- [18.3 Basic Concepts](#)
- [18.4 Structure of the Naming Service IDL](#)
- [18.5 Semantics of Names](#)
- [18.6 Naming Context IDL](#)
- [18.7 Iterators](#)
- [18.8 Pitfalls in the Naming Service](#)
- [18.9 The Names Library](#)
- [18.10 Naming Service Tools](#)
- [18.11 What to Advertise](#)
- [18.12 When to Advertise](#)
- [18.13 Federated Naming](#)
- [18.14 Adding Naming to the Climate Control System](#)
- [18.15 Summary](#)

[19. The OMG Trading Service](#)

- [19.1 Chapter Overview](#)
- [19.2 Introduction](#)
- [19.3 Trading Concepts and Terminology](#)
- [19.4 IDL Overview](#)
- [19.5 The Service Type Repository](#)
- [19.6 The Trader Interfaces](#)
- [19.7 Exporting Service Offers](#)
- [19.8 Withdrawing Service Offers](#)
- [19.9 Modifying Service Offers](#)
- [19.10 The Trader Constraint Language](#)
- [19.11 Importing Service Offers](#)
- [19.12 Bulk Withdrawal](#)
- [19.13 The Admin Interface](#)
- [19.14 Inspecting Service Offers](#)
- [19.15 Exporting Dynamic Properties](#)
- [19.16 Trader Federation](#)
- [19.17 Trader Tools](#)
- [19.18 Architectural Considerations](#)
- [19.19 What to Advertise](#)

- [19.20 Avoiding Duplicate Service Offers](#)
- [19.21 Adding Trading to the Climate Control System](#)
- [19.22 Summary](#)

[20. The OMG Event Service](#)

- [20.1 Chapter Overview](#)
- [20.2 Introduction](#)
- [20.3 Distributed Callbacks](#)
- [20.4 Event Service Basics](#)
- [20.5 Event Service Interfaces](#)
- [20.6 Implementing Consumers and Suppliers](#)
- [20.7 Choosing an Event Model](#)
- [20.8 Event Service Limitations](#)
- [20.9 Summary](#)

[VI: Power CORBA](#)

[21. Multithreaded Applications](#)

- [21.1 Chapter Overview](#)
- [21.2 Introduction](#)
- [21.3 Motivation for Multithreaded Programs](#)
- [21.4 Fundamentals of Multithreaded Servers](#)
- [21.5 Multithreading Strategies](#)
- [21.6 Implementing a Multithreaded Server](#)
- [21.7 Servant Activators and the Evictor Pattern](#)
- [21.8 Summary](#)

[22. Performance, Scalability, and Maintainability](#)

- [22.1 Chapter Overview](#)
- [22.2 Introduction](#)
- [22.3 Reducing Messaging Overhead](#)
- [22.4 Optimizing Server Implementations](#)
- [22.5 Federating Services](#)
- [22.6 Improving Physical Design](#)
- [22.7 Summary](#)

[A. Source Code for the ICP Simulator](#)

- [A.1 Overview](#)
- [A.2 Transient Simulator Code](#)
- [A.3 Persistent Simulator Code](#)

[B. CORBA Resources](#)

- [B.1 World Wide Web](#)
- [B.2 Newsgroups](#)
- [B.3 Mailing Lists](#)
- [B.4 Magazines](#)

[Bibliography](#)

Preface

For years, both of us have been (and still are) teaching CORBA programming with C++ to software engineers all over the world. One of the most frequently asked questions in our courses is, "Where can I find a book that covers all this?" Although many books have been written about CORBA, most of them focus on high-level concepts and do not address the needs of software engineers. Even though CORBA is conceptually simple, the devil lies in the detail. Or, more bluntly, books focusing on high-level concepts are of little use when you must find out why your program is dumping core.

To be sure, there are resources available about CORBA, such as newsgroups, Web pages, and the Object Management Group (OMG) specifications. However, none of them really meets the needs of a programmer who must get the code to work (and preferably by yesterday). We wrote this book so that there would finally be a tutorial and reference that covers CORBA programming with C++ at the level of detail required for real-life software development. (And, of course, we wrote it so that we would have a good answer for our students.)

Writing such a book is a tall order. Explaining the CORBA specification and APIs is one thing, and it's a necessary part of the book. However, knowing the various APIs will not, by itself, make you a competent programmer (only a knowledgeable one). To be competent, you need not only knowledge of the mechanics of the platform but also an understanding of how the different features interact. You must combine them effectively to end up with an application that performs and scales well and is maintainable, extensible, portable, and deployable.

To help you become competent (as opposed to merely knowledgeable), we go beyond the basics in a number of ways. For one thing, we provide advice as to what we consider good (and bad) design, and we make no attempt to hide problems with CORBA (which, like any other complex software system, has its share of wrinkles). Second, we go beyond the APIs by explaining some of CORBA's internal mechanisms. Even though you can use an ORB without knowing what goes on under the hood, it is useful to understand these mechanisms because they have a profound influence on how well (or how poorly) an application will perform. Third, we devote considerable space to a discussion of the merits of various design decisions; typically, when a design provides a gain in one area it also involves a loss in another. Understanding these trade-offs is crucial to building successful applications. And fourth, where appropriate, we make recommendations so that you are not left without guidance.

Inevitably, our approach required us to make value judgments, and, just as inevitably, a number of people will disagree with at least some of the recommendations we make. Whether you agree or disagree with us, you should still profit from our approach: if you agree, you can stick to the advice we give; if you disagree, the discussion will have at least encouraged you to think about the topic and form your own opinion. Either way,

you are better off than you would be with a book that just dumps the facts on you without providing the deeper insight required to use them.

Prerequisites

This book is not a beginner's book, in the sense that we do not devote much space to explaining the structure of the OMG or the specification adoption process. We also do not provide a high-level overview of the architectural goals of CORBA or all its services and facilities (see [31] for a high-level overview). Instead, we assume that you want to know how to write real CORBA applications with C++. Despite the lack of overview material, you should be able to follow the material even if you have never seen CORBA before. If you have experience in network programming or have used another RPC platform, you will find it easy to pick things up as you go.

Much of this book consists of source code, so we expect you to be literate in C++. However, you do not need to be a C++ guru to follow the code. We have avoided obscure or little-understood features of C++, preferring clarity to cleverness. If you understand inheritance, virtual functions, operator overloading, and templates (not necessarily in intricate detail), you will have no problems. Some of the source code uses the Standard Template Library (STL), which is now part of the ISO/IEC C++ Standard. We have limited ourselves to very simple uses of this library, so you should be able to understand the source code even if you have never seen STL code before.

If you have never written threaded code, you will find the chapter on writing threaded servers tough going. Unfortunately, there was not enough room to provide an introduction to programming with threads. However, the Bibliography lists a number of excellent books on the topic.

Despite our best efforts to show realistic and working source code, we had to make a number of compromises to keep code examples understandable and of manageable size. When we demonstrate a particular feature, we often use straight-line code, whereas in a realistic application the code would better be encapsulated in a class or helper function. We have also minimized error handling to avoid obscuring the flow of control with lots of exception handlers. We chose this approach for didactic purposes; it does not imply that the code pretends to reflect best possible engineering practice. (The Bibliography lists a number of excellent books that cover source code design in great detail.)

Scope of this Book

OMG members are continually improving CORBA and adding new features. As a result, available ORB implementations conform to different revision levels of the specification. This book covers CORBA 2.3. (At the time of writing, CORBA 2.3 is being finalized by the OMG.) Throughout the text, we indicate new features that may not yet be available in your ORB implementation; this allows you to restrict yourself to an earlier feature set for maximum portability.

Despite its size, our main regret is that this book is too short. Ever-increasing page counts and ever-closer deadlines forced us to drop chapters on the Dynamic Invocation Interface (DII), the Dynamic Skeleton Interface (DSI), and the Interface Repository (IFR). Fortunately, the vast majority of applications do not need those features, so dropping these chapters is not much of a loss. If your application happens to require the dynamic interfaces, the background we provide here will enable you to easily pick up what you need from the CORBA specification.

Another feature notable by its absence is Objects-By-Value (OBV). We chose not to cover OBV because it is too new for anyone to have any substantial engineering experience with it. In addition, at the time of writing, there are still a number of technical wrinkles to be ironed out and we expect the OBV specification to undergo further changes before it settles down.

Size and time limitations also meant that we could not cover every possible CORBA service. For example, we did not cover the Transaction Service or Security Service because each of them would require a book of its own. Rather than being complete, we have restricted ourselves to those services that are most essential for building applications: the Naming, Trading, and Event Services. We cover those services in more detail than any other publication we are aware of.

An important part of this book is the presentation of the Portable Object Adapter (POA), which was added in CORBA 2.2. The POA provides the server-side source code portability that was missing from the (now deprecated) Basic Object Adapter. The POA also provides a number of features that are essential for building high-performance and scalable applications. We have therefore paid particular attention to showing you how to use the POA effectively in your designs.

Overall, we believe this book offers the most comprehensive coverage to date of CORBA programming with C++. We have arranged the material so that you can use the book both as a tutorial and as a reference. Our hope is that after the first reading, you will have this book open at your side when you are sitting at your terminal. If so, we will have achieved our goal of creating a book that is used by real engineers to build real applications.

Acknowledgments

As with any book, the authors are only part of the story, and this is the place to thank the large number of people who have contributed to making this book possible. At Addison Wesley Longman, Mike Hendrickson and our editor, Deborah Lafferty, believed us when we told them that this book needed to be written. Without their faith in us, you would not be reading this. Brian Kernighan reviewed several drafts and made us redo the job where necessary. His clarity of thought and critical eye have greatly improved this book. John Fuller and Genevieve Rajewski, our production editors, put up with all our naive questions and enabled two amateurs to take a book to camera-ready stage. Our copy editor, Betsy Hardinger, edited every page in this book with meticulous attention to detail. Her efforts taught us more about clarity of style than we thought possible.

Particular thanks go to Colm Bergin, Jonathan Biggar, Bart Hanlon, Jishnu Mukerji, and Doug Schmidt, our expert reviewers. They read the entire manuscript and spotted many problems that would have otherwise gone unnoticed. Their comments kept us honest throughout. Alan Shalloway reviewed the book from the perspective of a newcomer and made valuable suggestions on how to improve the presentation of some of the more difficult topics.

Todd Goldman and Tim Gill from Hewlett-Packard gave us permission to draw on earlier ORB training material written by Michi. John Vinoski and Dan Rabideau of Green Bay Engraving take credit for designing the Möbius strip on the cover.

We are grateful to Steve's employer, IONA Technologies, for allowing us to use the next generation of their Orbix product (called "ART") to develop and test our code examples. Their generosity provided us with the opportunity to make sure that our examples were correct and functional. The fact that ART conforms to CORBA 2.3 allowed us to target the most recent version of the CORBA specification available as of this writing.

We also would like to thank the many contributors to `comp.object.corba` and the `corba-dev` mailing list. The discussions there have influenced much of the content of this book. A number of people have provided feedback, corrections, and constructive criticism since the first printing of this book. Rather than list them all here (and have to keep updating this Preface for each new printing), we have placed a list of everyone who contributed at <http://www.awl.com/cseng/titles/0-201-37927-9>. Our thanks go to all these people for helping to make this a better book.

Michi's Acknowledgments

I would like to thank my former employer, DSTC Pty Ltd, for providing me with an environment that was conducive to writing. Joachim Achtzehnter, Martin Chilvers, Wil Evers, Ted McFadden, and Michael Neville reviewed parts of the manuscript and made valuable suggestions for improvement. Particular thanks go to David Jackson, who read all my drafts and made sure that loose ends were not allowed to remain hanging. Finally, I would like to thank my wife, Jocelyn, and our son, Tyson, for their love and

encouragement. Without their support and patience, this book would have never been written.

Steve's Acknowledgments

I would like to thank my employer, IONA Technologies, for supporting my efforts to write this book, which occasionally kept me away from the office. In particular, I would like to thank Barry Morris for his support and encouragement, Stephen Keating for taking up the slack when I had to miss work because of all-night writing sessions, and the whole IONA Boston product development team for their patience and support.

I would also like to thank Bart Hanlon, who not only reviewed this book but also was my manager at my former employer, for continually encouraging me for several years to tackle this project and for teaching me a lot about tackling projects in general. In the technical realm, I have learned from many people over the course of my career, but I owe much to John Morris, Craig Bardenheuer, Denis deRuijter, Dale LaBossiere, Tom Moreau, and Bob Kukura, who at one time or another greatly influenced my education in the realms of distributed systems and engineering in general. I would also like to thank my *C++ Report* co-columnist, Doug Schmidt, a true technical visionary whose work in object-oriented network programming, C++ frameworks, and CORBA has paved the way for books such as this one. He not only helped review this book, but also agreed to let me use material from our columns in writing it.

Finally, without the support of my family, I would have never been able to even consider writing this book. I'd like to thank my wife, Cindy, and our children, Ryan and Erin, for putting up with my extremely long hours and days of writing and working. Thanks also to my parents, Ed and Dooley, who have always supported me with their seemingly limitless patience and love. I also owe my brother, John, a special thanks for his wonderful artwork on our book cover.

Michi Henning and Steve Vinoski

October 1998

Chapter 1. Introduction

- 1.1 Introduction
- 1.2 Organization of the Book
- 1.3 CORBA Version
- 1.4 Typographical Conventions
- 1.5 Source Code Examples
- 1.6 Vendor Dependencies
- 1.7 Contacting the Authors

1.1 Introduction

CORBA (Common Object Request Broker Architecture) is now well established in the mainstream of software development and has found phenomenal industry acceptance. CORBA is supported on almost every combination of hardware and operating system in existence. It is available from a large number of vendors (even as freeware), supports a large number of programming languages, and is now being used to create mission-critical applications in industries as diverse as health care, telecommunications, banking, and manufacturing. The increasing popularity of CORBA has created a corresponding increase in demand for software engineers who are competent in the technology.

Naturally, CORBA has had to evolve and grow (sometimes painfully) to reach its current levels of popularity and deployment. When the first version of CORBA was published in 1991, it specified how to use it only in C programs. This was a result of building CORBA from proven technology. At that time, most production-quality distributed systems were written in C.

By 1991, object-oriented (OO) languages such as Smalltalk, C++, and Eiffel had been in use for years. Not surprisingly, many developers thought it strange that a language-independent distributed OO system such as CORBA could be programmed only using C, a non-OO, procedural language. To correct this short-coming, several development groups at companies such as Hewlett-Packard, Sun Microsystems, HyperDesk Corporation, and IONA Technologies started developing their own proprietary mappings of CORBA to the C++ language. These proprietary mappings of CORBA to C++, all invented independently, differed in many ways. As most C++ programmers know, C++ is a multiparadigm language that supports varied approaches to application development, including structured programming, data abstraction, OO programming, and generic programming. The proprietary C++ mappings reflected this diversity; each of them mapped different CORBA data types and interfaces into different (sometimes very different) C++ types and classes. The mapping differences reflected not only the varied backgrounds of the developers but also the ways they intended to use CORBA to build systems as diverse as software integration middleware, operating systems, and even desktop tool kits.

When the Object Management Group (OMG) issued a Request For Proposals (RFP) for a standard mapping of CORBA to C++, these developers and other groups submitted their mappings to the standardization process. As is common for OMG RFP submissions, the submitting groups joined forces to try to reach consensus and arrive at a single C++ mapping specification that would draw from the strengths of all the submitted mappings. The process of producing a single standard C++ mapping for CORBA took approximately 18 months, lasting from the spring of 1993 until the fall of 1994. For technical reasons, such as the richness of C++ and its support for diverse programming styles, the consensus-building process was not an easy one. At one point, because of the competitive spirit and political nature of some of the parties involved (both characteristics are inevitable in any industry standards group), the C++ mapping standardization effort fell apart completely. However, the need for a standard C++ mapping eventually overcame all obstacles, and the standardization was completed in the fall of 1994.

The C++ mapping was first published with CORBA 2.0. Since its adoption, the mapping has been revised several times to fix flaws and to introduce minor new functionality. Despite this, the mapping has remained surprisingly stable and portable even while the C++ language was undergoing its own standardization process. The standard C++ mapping removed a major obstacle to broad acceptance of CORBA because it created source code portability, at least for the client side. The server side still suffered from portability problems until CORBA 2.2.

CORBA 2.0 also removed another major obstacle by providing the Internet Inter-ORB Protocol (IIOP). IIOP guarantees that system components developed for different vendors' ORBs can interoperate with one another, whereas before CORBA 2.0, different system components could communicate only if all of them used the same vendor's ORB.

The C++ mapping and IIOP were key features that initiated CORBA's move into the mainstream and made it a viable technology for many commercial companies. This increased popularity of CORBA also meant an increased demand for extensions and bug fixes. As a result, the specification has been revised three times since the publication of CORBA 2.0. CORBA 2.1 was largely a cleanup release that addressed a number of defects. CORBA 2.2 added one major new feature: the Portable Object Adapter (POA). The POA, together with an update to the C++ mapping, removed the server-side portability problems that existed to that point. CORBA 2.3, the most recent release, as of this writing, fixed many minor bugs and added one major new feature, Objects-By-Value.

The OMG has now grown to more than 800 members, making it the world's largest industry consortium, and CORBA has become the world's most popular and widely used middleware platform. In our estimation, C++ is the dominant implementation language for CORBA (although Java is making some inroads for client development). Demand for CORBA-literate C++ programmers continuously outstrips supply, and it seems likely that CORBA will remain the dominant middleware technology for at least several more years. This book is all about making you CORBA-literate and giving you the information you need to be able to write production-quality CORBA-based systems.

1.2 Organization of the Book

The book is divided into six parts and two appendices.

Part I, Introduction to CORBA, provides an overview of CORBA and presents the source code for a minimal CORBA application. After reading this part, you will know the basic architecture and concepts of CORBA, understand its object and request dispatch model, and know the basic steps required to build a CORBA application.

Part II, Core CORBA, covers the core of CORBA with C++: the Interface Definition Language (IDL), the rules for mapping IDL into C++, how to use the POA, and how to support object life cycle operations. This part introduces the case study we use throughout this book; following each major section, we apply the material presented there to the case study so that you can see how the various features and Application Programming Interfaces (APIs) are used for a realistic application. After reading this part, you will be able to create sophisticated CORBA applications that exploit many CORBA features.

Part III, CORBA Mechanisms, presents an overview of the CORBA networking protocols and shows the mechanisms that underpin CORBA's object model, such as location transparency and protocol independence. After reading this part, you will have a good idea of what goes on beneath the hood of an ORB and how design choices made by various vendors influence a particular ORB's scalability, performance, and flexibility.

Part IV, Dynamic CORBA, covers dynamic aspects of CORBA: type `any`, type codes, and type `DynAny`. After reading this part, you will know how you can use these CORBA features to deal with values whose types are not known at compile time. This knowledge is essential for building generic applications, such as browsers or protocol bridges.

Part V, CORBAServices, presents the most important CORBA services, namely the Naming, Trading, and Event Services. Almost all applications use one or more of these services. The Naming and Trading Services allow applications to locate objects of interest, whereas the Event Service provides asynchronous communication so that clients and servers can be decoupled from each other. After reading this part, you will understand the purpose of these services and you will be aware of the architectural consequences and trade-offs implied by their use.

Part VI, Power CORBA, discusses how to develop multithreaded servers and presents a number of architectural and design issues that are important for building high-performance applications.

Appendix A shows the source code for an instrument control protocol simulator that you can use if you want to experiment with the source code in this book.

Appendix B contains a list of useful resources you can use to get more information about various aspects of CORBA.