

Preface

This book is primarily intended to be a text for the programming component in an introductory two semester computer science course (some materials are a little advanced and might postponed to later semesters). This intent shows in terms of references to "students", "assignments" and "later studies in computing". However, the book should be equally suited to an individual who wants to learn how to program their own personal computer. If used as a text for an introductory computer science course, the contents of this book would normally be supplemented by a variety of other materials covering everything from "ethics" to "Turing machines".

Intended readership

It is assumed that you are "computer literate". In your school or college, or at work, you will have used word processing packages, and possibly other packages like spreadsheets and data bases. Although most schools do provide at a limited introduction to programming (usually in Pascal or possibly a dialect of Basic), this text does not rely on such past experience.

Assumed knowledge

It is expected that you will use one of the modern "Integrated Development Environments" (IDE) on a personal computer. Examples of such environments include the Borland environment for Intel PCs and the Symantec environment for Macintosh/PowerPC machines. These environments are supplied with comprehensive reference manuals . These manuals give details like how to start programs; they list and explain menu options; they include illustrations showing window layouts with details of all controls. These kinds of reference material are not duplicated in this book.

Programming environment

Content

Part I Introduction

The first section of the book contains some introductory material on computer hardware, operating systems and programming languages. Most of the material should be largely familiar to anyone who has done a "computer literacy" subject at school, or who reads popular computing magazines. The aim of this section is to provide background context knowledge so that you will understand what is happening when you do something like "compile a program" or "link with a library". Even if you feel confident that you know this material, you should skim through this section as it does introduce some terminology employed later in the text.

Hardware	Chapter 1 provides a brief overview of computer hardware giving a fairly standard presentation of computer architecture along with a few details of how data are input and output when being processed by a program.
Machine and assembly languages	Chapter 2 looks at programs in the machine, describing the bit patterns that represent instructions and explaining how these bit patterns get into memory. The chapter includes short sections on assembly language programming and system's software such as assemblers and loaders. The basic structural components of assembly language programs – sequences of instructions, conditional branches to build loops and make selections, and subroutine call and return – are illustrated.
Operating systems	Operating systems are briefly reviewed in Chapter 3. The treatment is partly historical, explaining the development of operating systems. Terms such as "time sharing", "virtual memory", and "multiprogramming" are introduced along with some details of the systems where these ideas first evolved and the reasons why these features were added to operating systems.
High level languages	Chapter 4 has again a somewhat historical viewpoint; it presents the evolution of high level programming languages. Most commonly used programming languages are closely related relying on the same "imperative procedural" programming constructs. The "imperative procedural languages" share the same view of programs – really the same view as in assembly language. Programs calculate results from data values given as inputs; the input data, intermediate values, and final result values are represented by <i>variables</i> . <i>Expressions</i> describe how to combine data values; the most common expressions simply define arithmetic manipulations. Variables are updated by assignment of the values calculated by executing instructions implementing the expressions. In addition to <i>assignment statements</i> , these languages have <i>iterative statements</i> (for loop constructs), <i>selection statements</i> , and subroutine call mechanisms. Although there are major differences in style and appearance, diverse languages like FORTRAN, COBOL, and C++ have underlying similarities resulting from their common imperative procedural model. (There are languages, like Prolog, that do embody quite different models.)
Algol family of languages	Languages like Pascal, Modula2, C, C++, Simula, Eiffel really all belong to the same linguistic family; their roots are in a language called Algol-60. They have diverged to varying extents and, to some degree, have each become specialized for particular kinds of application. Nevertheless, they remain closely in structure.
Recursive functions	In addition to touching on the development of programming languages, Chapter 4 also provides a simple introductory example illustrating the idea of recursion.
C++ development environment	Chapter 5 provides a general picture of the kind of "Integrated Development Environment" that you will be using to create and run your programs. The treatment is non-specific; all the environments include: <ul style="list-style-type: none"> • an editor (for entry of program code and of data files), • a compiler (for translating code), • function libraries (compiled code for standard useful components like code to calculate the square root of a number, or code to translate between a string of printed digits and a number in the form that the computer can manipulate),

- a linker (puts together a complete program by combining your code with components from the libraries),
 - a "make system" (this organizes the compilation and linking processes for you; this make system may be fully automatic and largely hidden in the "program" or "project" management facilities provided by the IDE),
- and
- a visual debugger (a supplied program that you may use to control the execution of your own program; the debugger will allow you to do things like arrange for your program to stop at places where you want to check on the values of chosen data elements or where you want observe exactly what is happening as each individual statement is executed).

Of course, the exact layout of an editor window, the compiler options, the libraries etc all do vary a little between different environments.

Chapter 5 also contains an introduction to input and output. Any interesting program will produce some results that are to be output; most require some input data. The C++ language does not define input and output commands; instead it relies on function libraries. There are different input/output libraries that one may chose to use; the most common is the `iostream` library. The `iostream` library uses some of the more advanced features of C++ (such as object oriented design and "operator overloading"); but it can be used long before you get to really understand how its code works. Essentially, the `iostream` library provides two "objects" – `cout` and `cin`. The `cout` object "owns" the connections to the output device (normally the computer's screen) and provides all services related to output (such as copying character strings so that a program can get its output displayed on the screen, and converting numbers into sequences of printable/displayable digits). Similarly, the `cin` object "owns" the input device (usually this is the keyboard) and takes responsibility for translating keyed characters into values suable by a program. Simple use of the `cin` and `cout` objects is illustrated with some small examples.

C++ input and output

Part II Simple programs

The second part introduces the fundamental programming constructs of imperative procedural languages.

Programs are built up from a *main* routine (main program) that utilizes "procedures" (or "functions", or "subroutines" – the names vary); these procedures may be specifically written for a particular application or may come from libraries. But all routines, (main and all its subroutines) are constructed in the same way. So we start by looking at a number of simple programs that just have a single main routine that makes limited use of standard functions from the input/output and mathematical libraries.

The constructs for building an individual routine include:

- variable definitions,

- simple (arithmetic) expressions that combine the values of variables to obtain new results,
- assignment statements that change variables,
- sequences of assignment (and input and output) statements that perform processing tasks,
- declarations of constants,
- iterative constructs for building loops
- selection statements for making choices.

Sequence Chapter 6 deals with the basics. It focuses on programs that read data values, perform a few computational steps, and print the results. Such programs involve solely a sequence of variable definitions, input statements, assignment statements, and output statements. The data values used may be integers, real numbers, or individual printable characters. These basics are covered in Chapter 6.

Iteration Chapter 7 adds loop constructs. Usually, a calculation will involve something more complex than a fixed sequence of steps. Instead, you may need to do the same (or similar) processing steps for each element in some collection, or you may need to repeat a particular operation many times.

"Boolean" variables, variables that represent the values true and false, and boolean expressions are also covered. This is a somewhat untidy area of C/C++ at the moment; the proposed standard for C++ will improve things.

C++ has a minor alternative to the normal while loop in its "do ... while" construct. This is covered briefly. Then "for loops" are introduced. Two other control statements used with loops – "break" and "continue" – are noted but illustrations of their use is deferred to later more complex examples.

Selection Selection statements, switch and if, are covered in Chapter 8.

File i/o Chapter 9 introduces the use of files. Most realistic programs require reasonable amounts of input data (e.g. the names and marks for each of several hundred students enrolled in a course). It is very tiresome to have to type in such data "live" (make a mistake and you may have to do it all again). Even when testing programs, one often has to type in a substantial amount of data; retyping of data for each test run is unacceptable. It is worth learning quite early on how to use file i/o.

Part III Functions and data aggregates

The third part of this text focuses mainly on functions., but also introduces more complex forms of data such as arrays and structs.

Functions Functions have a number of important roles:

- they allow a processing step, e.g. validated input of a data value, to be coded once and used from many points in a program rather than having the code repeated at each point,

- they capture and preserve expertise; experts can develop functions that implement particular algorithms, e.g. a sort algorithm; once developed, the function can be added to a library from where it can be copied and used in numerous programs;
- they serve as a basis for design.

It is rare for one to be working with individual data elements as done in the examples in Part II. Usually, data values form related groups. Arrays represent a simple case of structured data; arrays group sets of data elements that are identical in type. The data elements making up an array are most often simple built in types so you can have arrays of reals, or arrays of integers, or arrays of characters; you can also have arrays of programmer defined structured types. The compiler knows that the only "structure" possessed by an array is that all the data elements belong together (allowing you to write code to do things like calculate how much memory space is used by an array), but at the same time the individual data elements in the array should be capable of being separately referenced so the compiler supports a mechanism that allows you to do things like ask for the "1st", "2nd", ..., "199th" element.

Arrays

Application specific groupings of data are defined using "structs" (also known as "structures" or "records"). A struct declaration describes the data that belong together. Thus, in a program to handle sales records one might have a grouping "customer" (that groups data values of name, address, amount owed, date last contacted, ...), while in a word processor program you might have a grouping "format_record" (that groups data values of text font, style, size, ...). Usually, the individual data members that make a struct are of different types; a customer record might contain an array of characters for the name, and several different integers or reals (phone number, amount owed, etc). You do get structs where the data members are of similar types, e.g. a struct to represent a point in three dimensions would have three integer (or real) number data members to represent the different coordinates. Once you have defined the form of a structure, the compiler allows you have variables that are instances of that struct and to do things like assign values to complete structs or individual data members within structs. The major role of structs is this grouping of related data.

Structures or "structs"

Chapter 10 introduces functions. It covers definition of functions and result types. The working a function calls, including recursive functions, is illustrated. Minor features like "inline" and "default arguments" are explained.

The arrays are covered in Chapter 11. The first examples use one dimensional arrays; these include "strings" – the arrays of characters used to represent textual data. Multidimensional arrays are then covered. Example programs illustrating the use of arrays include an encryption program, and a simple version of an "image analysis" program.

Chapter 12 consists of lots of examples using functions and arrays. Some of these examples are complete programs; one program provides a simple model of heat diffusion, another is an implementation of the game "hangman". Other examples build up little "libraries" of useful functions; these include a "cursor graphics" package for

Lots of examples

both Symantec and Borland environments. This package is used in a number of the other examples to provide simple visual outputs.

Functions for standard algorithms

Chapter 13 looks at some problems for which "standard" functions have been developed. These problems involve things like sorting or searching through arrays of integers. The methods for solving such tasks (the algorithms) are now well known; functions implementing these algorithms are available in libraries.

Tools

We diverge slightly with Chapter 14 which looks at two tools that help in the development of programs. The more important of these two tools is a "code coverage tool". A "code coverage tool" helps a programmer check that all parts of a program's code have been executed during test runs; this is an important part of the process of getting a program ready for practical use. Unfortunately, neither of the IDEs includes a code coverage tool; the examples use output from the "tcov" tool on Unix. The second tool is a profiler. Although less useful than code coverage tools, profilers are more widely available and are included in both IDEs. A programmer can use a profiler to determine where a program spends most time. This can be used to focus any work that is to be done on speeding up the program. Most of the programs that you will write initially spend their time largely in the system's input/output libraries so there won't often be much scope for speed improvements.

Design and documentation

Chapter 15 contains some suggestions for how you might go about designing, and documenting your programs. The design approach is the one underlying all the examples using functions that have been presented in Chapters 10...13. The design process is a form of "divide and conquer". You start by considering the overall task the program is to perform and break it down into a separate functions that deal with particular aspects like data input, computational steps, and results display. You then consider each function in turn, analysing it and breaking it down into functions. This "top-down functional decomposition" process is repeated until the functions are sufficiently simple that you can see how to express their processing in terms of the sequence, iterative loop, and selection constructs introduced in Part II. Top down functional decomposition is a limited approach. It works well for programs that basically "read data, process data, print results" and where the data are simple, e.g. an array of numbers like the heat diffusion example in Chapter 12. This design approach gets relegated to playing a minor support role in Part IV where more complex programs are considered.

Enum, struct, and union

Chapter 16 presents enumerated types (enums), structures (structs), and unions. These constructs allow a programmer to define new data types to supplement the integers, reals, and characters that are built in to the language.

Enumerated types are useful when you want to have a data element that must have one of a restricted number of values. For example, you want a "colour" variable that can be 'red', 'blue', or 'yellow'; or a "text_style" variable that can be 'plain', 'bold', or 'italic'. You can just use integers for such data; the advantage of enumerated types is that they give you a mechanism for telling the compiler about the semantics of these "colour" and "text_style" variables. If you define an enumerated "colour" type, the

compiler can check that colour variables are used correctly; the compiler will verify that only colour values like 'red' get assigned to a colour variable.

Technically, in C++ there isn't much difference between a struct and a simple class. In this text, a distinction is made. Structs are used only to achieve a grouping of related data elements. The examples in Chapter 16 illustrate the definition of structs and the ways in which code can manipulate both complete variables of struct types and the individual data members within a struct variable.

You should regard "unions" as a "read-only" feature of C++. For completeness, they are covered in section 16.3 using examples taken from the Xlib graphics library used on Unix. You will need to use unions when working with some C libraries like Xlib, but it may be years before you get to write code where it would be useful to define new unions.

Chapter 17 has a few examples illustrating the use of structs. One example illustrates the idea of a "file of records", something commonly required in simple business data processing applications.

Using structs

Next, bit-level operations are introduced. Of course all data are just bits. But usually we think about data as reals, or integers, or characters and forget the bit patterns. But there are a few situations where you actually want to use bits. Chapter 18 presents the bit manipulation facilities of C++ and has some illustrative examples like "hashing" and the use of a bit-vector used to summarise properties of a more elaborate data record. The main example in this chapter is a simple information retrieval system that has a file with the texts of newspaper stories and a separate index file that helps you find stories that use particular keywords.

Bits and pieces

Part IV A touch of class

The programming approach developed in Part III is fine for straightforward programs that have a single set of data that must be transformed in some way. Some scientific, engineering, and statistical calculations do have just arrays of numbers for data and, essentially, they simply calculate the value of some function of their input data. Most programs are more complex.

There are two complications. Firstly, the data structures become more elaborate; and, secondly, you start encountering problems where you have no idea how many data structures you require until the program is actually running.

The more elaborate structures involve more than just lots more data fields. There are typically constraints on the values in different data fields. For instance, you might want a data structure that represents a point in two dimensional space with its position recorded using both cartesian (x, y) and polar (r, theta) coordinates. If you change a point's x-coordinate, you really should update its (r, theta) coordinates to match; there is this constraint, they should define the same place in two-dimensional space. Somehow, you need to be able to shield the data values from arbitrary changes. Instead of allowing direct access to the data in a structure, you provide an alternative functional

Protecting data from arbitrary change

style of interface that permits changes but arranges that these are done in such a way that all constraints are still satisfied.

Design problems and top down functional decomposition

Large programs contain many different types of data and have to perform many functions. Consider for example, a simple system for maintaining hospital records. This would need data structures representing patients, wards, operating theatre schedules, laboratory tests. These data would be accessed by admissions staff, clerks calculating bills, laboratory staff recording results of blood tests, surgeons wanting lists of their day's tasks. It is difficult to identify the "function" of the program that could serve as the top for the kind of top-down functional decomposition approach illustrated in Part III. Even if you can come up with a design, frequently you will find that it is not very satisfactory. Everything seems interlinked. All the different functions scramble around the same data structures and arrays. Change something in the section used by the pathology lab, and you will find that somehow this alters the way the admissions system works.

Designing around the data

Often it is better to try a different decomposition of the problem, a decomposition that focuses on the data. So you identify important elements in the program: patient, theatre, ward. For each you identify the data owned; a patient owns a name, a ward number, a list of results of lab. tests, You also identify the tasks performed; a patient record can handle requests to note the result of another lab. test, report what ward and bed it is associated with, note an addition to its bill and so forth. By considering the data, you can break the overall problem down into more manageable parts. These separate parts are largely independent and can be developed in isolation. Once the parts are working, you make up the overall program by describing interactions, e.g. you can describe how the "laboratory" gives a "patient" a new "lab.-result" to note, and also informs "finance" to arrange additional charging.

Classes

Classes provide the basis for protection of data, and the packaging together of functionality and data. They also serve as the basis for the design of most the programs that you will write in future.

Dynamic data

Sometimes you know how many data elements your programs must manipulate, or at least you can identify a reasonable upper bound. Such knowledge simplifies design; you can plan out arrays to hold the data elements required. But there are situations where you don't know how many data elements you have and where guessing a maximum is inappropriate. All the modern languages in the Algol family allow you to create data structures "dynamically" as and when the program needs them. Often, these dynamic structures are used to build up program models of complex networks.

Beginners class

Chapter 19 introduces C++ classes. The first section covers the syntax, using simple examples like Points and Customers. The following two sections present more elaborate examples. First there is class `Bitmap`. This is a more systematic, more complete version of the idea of a bitmap (set) record structure such as that used to identify use of keywords in the information retrieval example given in Chapter 18. The second class is `Number`. Not your everyday integer, these `Numbers` have hundreds of digits. But you can do all the standard arithmetic operations using them.

Chapter 20 looks at dynamic data and "pointers". Dynamic data are data structures created at run time. Obviously in order to use these data, a program has to know where they've been located in memory. That is the role of pointers; pointers hold addresses of dynamic structures.

Dynamic data and pointers

Chapter 20 starts by explaining how it all works, providing a model of the "heap" (the area of memory used for these data). Next, there is a general introduction to pointers and ways of using pointers. There are a few simple examples illustrating pointers with "strings" (character arrays). Then there is a larger example illustrating the use of dynamic data in a little simulation of aircraft landing at an airport. Aircraft are objects that are created dynamically. The simulation allows aircraft movements to be controlled. If controlled correctly, aircraft land at the airport and, then, the corresponding dynamic data structures are destroyed.

Pointers aren't used solely with dynamic data structures. Chapter 20 explains some other uses. The older C language made much more use of pointers than does C++; you are cautioned against continuing with some of the older C style usages, such as use of pointers rather than array indexing.

The final section of Chapter 20 illustrates how to use pointers and dynamic data structures to build a very simple "network" of interlinked data elements.

Very often programs need to maintain collections of data. For example, a "patient" needs a collection of some kind to store the results of all the lab. tests performed; a "theatre" needs a list of the patients due for surgery that day; the air-traffic control example needed a collection of "aircraft" in the controlled air space.

Collection classes

Collections get used in different ways. You may want to represent a "queue" – first come first served. Alternatively you might have a priority queue, the tasks in the queue have priorities and an urgent task added later may jump to the front of the queue. You might just want a list; the order is arbitrary, you remove things from the list as you choose. Or you might have a collection where the stored items have "keys" (unique identifiers like "driver licence numbers") and you need to find an item given its key.

Just as some algorithms have been standardized and become standard functions in function libraries, so too have some mechanisms for storing collections of data. A collection class will define a set of operations for adding and removing items that are to be stored. It will also define some data structure that it uses internally to keep track of the stored items. Collection classes can be standardized and put into "class libraries". These libraries facilitate reuse, just as did the simpler function libraries.

Chapter 21 introduces a variety of simple collection classes and shows how they are implemented. The examples include queue, priority queue, dynamic array, list, and binary search tree. This introduction is fairly brief. You may well study such structures in more detail in a later "Data Structures and Algorithms" course.

Chapter 22 contains a couple of simple examples that illustrate the "world of interacting objects" that is characteristic of programs designed using classes.

World of Interacting Objects

Chapter 23 expands on the initial treatment of C++ classes. Additional topics covered include: shared class properties, friends, iterators, operator functions, resource manager classes with destructors, and an introduction to "inheritance". "Inheritance" is

Intermediate class

again motivated mainly by considerations of program design; it is shown how you can identify similarities among classes, represent these similarities using an inheritance class hierarchy, and then exploit this hierarchy to simplify the overall design of a program.

Two more trees

Chapter 24 presents two more "collection classes". They are both more sophisticated versions of the binary search tree illustrated in Chapter 21. The first AVL tree performs exactly the same task as earlier the binary tree class, it just performs better. Class BTree shows how a collection of data that is too large to fit into memory can be held on disk but still permit rapid retrieval of information. You may wish to ignore details of the algorithms used for these two examples because they are more complex than others in this text. But, you will probably find class AVL and class BTree quite useful building blocks for later programs.

Templates and Exceptions

Chapters 25 and 26 are both short. Chapter 25 gives a brief introduction to C++'s template facilities. Templates are a mechanism for making code more general purpose. You can write a "template" for a function like a "sort" function; this "template sort" describes how to sort any data for which assignment and the > (greater than) and == (equals) operations are defined. This saves copying and editing of standard code to make minor changes needed to accommodate different data. You can also define "template classes". The most common "template classes" are other "collection classes". There are some advantages in using templates for such classes, though they do tend to complicate syntax and some compilers still don't handle templates particularly well. The advantages relate to better type security (i.e. less chance of run-time errors due to programmer confusion as to the nature of the data being used).

Exceptions, Chapter 26, are intended as an error handling mechanism. The examples used in this book deal with errors, e.g. unreadable input files, by printing an apology and quitting. This isn't the best approach (for example, it wouldn't do for the program that controls Space Shuttle launches). A better mechanism is for the code that encounters an insurmountable problem to stop and return an error report to the code that called it. The calling code can handle the error report and, possibly, can implement a different strategy (e.g. ask the user to select another input file).

Supermarket

Chapter 27 is again just an example. Once again it is a simulation. This one makes limited use of inheritance and employs standard collection classes.

Design and documentation

Chapter 28 looks again at design and documentation, now from the perspective of "object based" programs.

Part V Object oriented programming

Finally section five provides a brief introduction to "object-oriented" (OO) programming.

Polymorphism

The introduction to inheritance in Chapter 23 showed you could exploit similarities to factor details out of code. The example used there was the "document" that owned a collection of "things" such as "text paragraphs, and "circuit components". When it

needed to save its data to file, the document used a loop where it got a pointer to the next "thing" in its collection and told that "thing" to write itself to the file. The document object didn't even know the range of types of "thing" that it could have in its collection; it knew only that all "things" could write themselves to files. The pointer it used to reference a thing was a "polymorphic pointer" – it pointed to "things" that had different shapes (some shaped like text paragraphs, others like circuit components).

Chapter 29 develops a complete program that illustrates this style of design and use of polymorphism. The example is a form of "dungeons and dragons" game. The dungeon is inhabited by numerous creatures drawn from different races. All creatures have essentially the same behaviour (they seek to end the game by eliminating the human player), but each different kind (class) of creature goes about this task in a distinct way. The game program works with a collection of creatures. At appropriate times, the creatures get their chance to proceed – the code is simple "for each creature in creature list do creature.run". Here the reference variable creature is polymorphic, as we move along the list the variable creature may reference first a ghost, then a patrol.

Polymorphism maybe great fun but in itself it isn't the reason why object oriented methodologies are increasing in popularity. So why OO?

Really, the reason is that OO represents another step in the progression towards higher levels of design reuse.

Function libraries allow you to reuse algorithms. You don't write a sine function, you use sin() from the math library; you don't write a sort routine, you use qsort() from stdlib. Design by top down functional decomposition allows you to consider the overall role of a program and break down, over and over again, until you have functions that either exist in a library or that you can write. You build up a program from a set of reusable (and special purpose) functions that rummage through global data structures.

Class libraries and object based programs allow you to reuse abstract data types. You don't have to implement code to interrogate the operating system and build up a representation of today's date, you use an instance of class Date from the class library. You don't implement your own search structure, you use an instance of class AVL or class BTree from some library. Your approach to design changes. Your first step is to identify a partitioning of the overall program system into separate components. You then identify their classes, determining what data each instance of a class must own, and what these objects do. Finally, you build up your program as a unique system of interacting objects (each jealously guarding its own data).

When inheritance was introduced in Chapter 23, it was shown that this was a way of representing and exploiting similarities. Now, many application programs have substantial similarities in their behaviour. Just think about the way the different word processor, spreadsheet, diagram drawing and other programs work on your personal computer. They all have the same arrangements of "window", "menus", "undo/redo" editing operations and so forth. Obviously, they all have a lot of rather similar code. The data drawn in the windows, and the specific menu commands differ, but all the code for picking up a mouse click, recognizing it as a menu selection, finding which menu option must be similar in all the programs.

Reuse! Reuse!

Reuse with functions

Reusable classes and object based design

Can we reuse more?

Reusable patterns of interaction?

If you could abstract out this similar code, you could greatly decrease the cost of developing new applications. The new applications could reuse not just individual classes, but groups of classes whose patterns of interactions are already defined at least in part.

Increasingly, this is being done. Once patterns of interaction, e.g. undo/redo an edit step, have been identified as being similar in many programs, an attempt is made to identify how this pattern can be represented in an abstract form. This leads to the identification of the classes involved and a particular sequence of interaction among them. Thus you might determine that in response to a menu request a Window object creates a Command object to hold the information needed to change (backwards and forwards) the data held by a Document object. You can define these as partially implemented abstract classes with the calls needed for Undo/Redo already encoded.

If you need to implement a similar editing operation in your own program, you can create classes that "inherit" from the partially implemented abstract classes. Your classes are just the same as those defined, except yours hold some data. Although overall patterns of interactions will be defined in the abstract classes, a few operations will have been left undefined. These are the operations that actually create and change data. All you need do is implement those few operations. You reuse the overall pattern of interactions.

Framework class library

The classes needed to build a particular kind of a program can be provided as a "framework". So, nowadays, you frequently get to program using a "framework class library".

Chapter 30 provides something of the flavour of an OO library by building a little framework. It isn't much, just a framework for programs that involve filling in and searching records that are displayed on the screen. For display, it uses the cursor graphics functions from Chapter 12. For storage, it uses the BTree class from Chapter 24.

Chapter 31 introduces the OWL (Object Windows Library) and TCL (THINK Class Library) framework class libraries. These frameworks (and similar frameworks like MFC, Microsoft Foundation Classes) are provided as part of the integrated development environments. They represent the reusable models for "Windows" and "Macintosh" programs. They embody many of the standard patterns that have now been identified as reusable structures for large scale programs.

Examples

You can't learn to program by learning some rules that define the syntax of a programming language and reading a few isolated code fragments.

You need to see a programming language in action. You have to see complete programs developed using the language constructs, as covered at a particular stage of your course, in order to understand how programs can be created.

It is not sufficient just to read the code of a final program. You need to see how the programmer got from a problem specification to that code. You have to learn how to "design" programs and how to record your design decisions.

Design is an iterative process. You look at the problem. You isolate some aspects that can be considered further in isolation. You then work on each of these aspects, either finding ways of decomposing them into still simpler isolatable aspects or identifying how the work required could be expressed in terms of code. *Design*

Different problems suit different design styles. The earlier examples in this text are mostly programs that have the form "read some data, apply a function that computes something from those data, output the result". These problems are well suited to a design style known as top down functional decomposition. The aspect that you focus on is the main "function" of the program and you break this function down into simpler auxiliary functions. You specify what these functions do. You identify what data they need to exchange or share. Finally, you write some code.

Examples illustrating this approach are given in Chapters 12 and 17. They include things like a "Hangman" game, Conway's Life program, a program that keeps "customer records" in a file on disk. Each example has a fairly detailed discussion of design as well as code. *Examples of functional design*

The later examples are more suited to "object based" design. You start by identifying "objects" that will be used in your program, and identify how these objects interact. You focus then on the different kinds, or classes, of objects that you need; trying to characterise each in terms of what such objects own and what they do. You elaborate on your initial ideas of how these objects interact. Your ideas about the classes objects get defined more formally. You elaborate the descriptions of the functions that define their behaviours. Finally, you write some code. You may not have to write all the code. As you developed your program design you may have discovered opportunities to reuse existing classes.

Chapters 20 and 22 contain examples, they include things like simulations of aircraft landing at airports and a system for maintaining files of reference cards (a kind of "Filofax" system). Again these examples, particularly those in Chapter 22, contain fairly lengthy discussions of design. *Examples of object based design*

The final examples use "object oriented design". These extend on the earlier object based examples showing how more advanced language features can provide a cleaner overall design based on exploiting similarities that may exist among the different kinds of objects needed in a program. The example in Chapter 27 illustrates a simulation that makes relatively limited use of object oriented techniques. More elaborate examples are given in Chapters 29 and 30. *Examples of object oriented design*

These examples should "bring you up to speed". You won't quite be ready to exploit the full power of your computer, its development system, and the "framework class libraries" that it includes. But you will have got far enough to understand the principles of those frameworks so that you can proceed on with one of the books that covers programming for your chosen system.

Many of the examples don't include the full code. The code, and code for "useful components" like the collection classes, is available. (You need to be able to use the "ftp" file exchange program on the Internet. You should find the code at the site named `ftp.cs.uow.edu.au`; the directories containing the code will be in `/pub/oop/ABC`). While you can simply fetch the code, it is better if you implement the examples for yourself using the code given and the design outlines for any components for which no code is given.

Note on filenames in examples

The development environments differ in their rules regarding filenames. In this text, it is assumed that there are no restrictions on name length; if your system restricts filenames, you will have to change those names that are too long.

The environments use suffixes on filenames to distinguish content. For example, a file for data may be supposed to have a name that ends with the suffix `.dat`. Two different kinds of files are used hold program code. "Header" files (which contain just things like lists of function names and constants) typically have names that end with the suffix `.h`. Different systems use different suffixes for the files containing the main code. Code files may be supposed to end with the suffix `.cp`, or `.cc`, or `.cpp`, or `.CC` or something else. The text uses the suffix `.cp`; you will have to change this to the suffix appropriate to your system.

Table of Contents

Part I Introduction

1 Computer Hardware

- 1.1 CPU and instructions
- 1.2 Memory AND DATA
- 1.3 Bus
- 1.4 Peripherals
 - 1.4.1 Disks and tapes
 - 1.4.2 Other I/O devices

2 Programs: Instructions in the Computer

- 2.1 Programming with bits!
- 2.2 Loaders
- 2.3 Assemblers
- 2.4 Coding in assembly language
- 2.5 From Assembly Language to "High Level" languages

3 Operating Systems

- 3.1 Origins of "Operating Systems"
- 3.2 Development of Operating Systems
 - 3.2.1 Multiprogramming
 - 3.2.2 Timesharing
 - 3.2.3 File Management
 - 3.2.4 Virtual Memory
 - 3.2.5 "Spooling"
 - 3.2.6 Late 1960s early 1970s system
- 3.3 Networking
- 3.4 Modern systems
 - 3.4.1 UNIX
 - 3.4.2 Macintosh OS

4 Why have "high-level" languages?

- 4.1 Limitations of Assembly Language and Origins of High Level Languages
- 4.2 High level languages constructs
 - 4.2.1 Statements
 - 4.2.2 Data types and structures
- 4.3 Evolution of high level languages
- 4.4 FORTRAN
- 4.5 BASIC
- 4.6 Lisp
- 4.7 COBOL
- 4.8 ALGOL
- 4.9 The "Algol Family"
 - 4.9.1 AlgolW, Pascal, and Modula2
 - 4.9.2 ADA

- 4.9.3 BCPL, C, and C++
- 4.9.4 Simula, SmallTalk, and Eiffel

5 C++ development environment

- 5.1 Integrated Development Environment
- 5.2 C++ input and output
- 5.3 A simple example program in C++
 - 5.3.1 Design
 - 5.3.2 Implementation

Part II Simple Programs

6 Sequence

- 6.1 Overall structure and main() Function
- 6.2 Comments
- 6.3 Variable definitions
- 6.4 Statements
- 6.5 Examples
 - 6.5.1 "Exchange rates"
 - 6.5.2 pH
- 6.6 Naming rules for variables
- 6.7 Constants
- 6.7.1 Example program with some const definitions.
- 6.8 Initialization of Variables
- 6.9 Converting data values

7 Iteration

- 7.1 While loops
- 7.2 Examples
 - 7.2.1 Modelling the decay of CFC gases
 - 7.2.2 Newton's method for finding a square root
 - 7.2.3 Tabulating function values
- 7.3 Blocks
- 7.4 "Boolean" variables and expressions
 - 7.4.1 True and False
 - 7.4.2 Expressions using "AND"s and "OR"s
- 7.5 Short forms: C/C++ abbreviations
- 7.6 Do ... While
- 7.7 For loop
- 7.8 Break and Continue statements

8 Selection

- 8.1 Making choices
- 8.2 A realistic program: Desk Calculator
- 8.3 IF...
- 8.4 Terminating a program

- 8.5 Example Programs
- 8.5.1 Calculating some simple statistics
- 8.5.2 Newton's method for roots of polynomials
- 8.6 What is the Largest? What is the Smallest?

9 Simple use of files

- 9.1 Dealing with more data
- 9.2 Defining filestream objects
- 9.3 Using input and output filestreams
- 9.4 Stream states
- 9.5 Options when opening filestreams
- 9.6 When to stop reading data?
- 9.7 More Formatting options
- 9.8 Example

Part III Functions and Data Aggregates

10 Functions

- 10.1 Form of a function Definition
- 10.2 Result types
- 10.3 Function declarations
- 10.4 Default argument values
- 10.5 Type safe linkage, Overloading, and Name Mangling
- 10.6 How functions work
- 10.7 Inline functions
- 10.8 A recursive function
- 10.9 Examples of simple functions
- 10.9.1 GetIntegerInRange
- 10.9.2 Newton's square root algorithm as a function
- 10.10 The rand() function
- 10.11 Examples
- 10.11.1 π -Canon
- 10.11.2 Function plotter

11 Arrays

- 11.1 Defining one dimensional arrays
- 11.2 Initializing one dimensional arrays
- 11.3 Simple Examples Using One-Dimensional arrays
- 11.3.1 Histogram
- 11.3.2 Plotting once more
- 11.4 Arrays as arguments to functions
- 11.5 Strings: Arrays of characters
- 11.6 Multi-dimensional arrays
- 11.6.1 Definition and initialization
- 11.6.2 Accessing individual elements
- 11.6.3 As arguments to functions

11.7	Arrays of Fixed length Strings
11.8	Examples using arrays
11.8.1	Letter Counts
11.8.2	Simple encryption
11.8.3	Simple image processing
12	Programs with functions and arrays
12.1	Curses
12.2	Heat diffusion
12.3	Menu Selection
12.4	Pick the keyword
12.5	Hangman
12.6	Life
13	Standard algorithms
13.1	Finding the Right Element: Binary Search
13.1.1	An iterative binary search routine
13.1.2	Isn't it a recursive problem?
13.1.3	What was the cost?
13.2	Establishing Order
13.3	A simples sort of sort
13.4	Quicksort: a better sort of sort
13.4.1	The algorithm
13.4.2	An implementation
13.4.3	An enhanced implementation
13.5	Algorithms, Functions, and Subroutine Libraries
14	Tools
14.1	The "Code Coverage" Tool
14.2	The Profiler
15	Design and documentation : 1
15.1	Top down functional decomposition
15.2	Documenting a design
16	Enum, Struct, and Union
16.1	Enumerated types
16.2	Structs
16.3	Unions
17	Examples using structs
17.1	Reordering the class list again
17.2	Points and Rectangles
17.3	File of Records
18	Bits and pieces
18.1	Bit manipulations

- 18.2 Making a Hash of it
- 18.2.1 Example hashing function for a character string
- 18.2.2 A simple "hash table"
- 18.2.3 Example: identifying the commonly used words
- 18.3 Coding "Property Vectors"
- 18.4 Pieces (Bit fields)

Part IV

19 Beginners' Class

- 19.1 Class declarations and definitions
- 19.1.1 Form of a class declaration
- 19.1.2 Defining the member functions
- 19.1.3 Using class instances
- 19.1.4 Initialization and "constructor" functions
- 19.1.5 const member functions
- 19.1.6 inline member functions
- 19.2 Example: Bitmaps (Sets)
- 19.3 Numbers – a more complex form of data
- 19.4 A glance at the "iostream" classes

20 Dynamic data and pointers

- 20.1 The "Heap"
- 20.2 Pointers
- 20.2.1 Some "pointer" basics
- 20.2.2 Using pointers
- 20.2.3 Strings and hash tables revisited
- 20.3 Example: "Air Traffic Controller"
- 20.4 & : the "address of" operator
- 20.5 Pointers and Arrays
- 20.6 Building Networks

21 Collections of data

- 21.1 Class Queue
- 21.2 Class PriorityQueue
- 21.3 Class DynamicArray
- 21.4 Class List
- 21.5 Class BinaryTree
- 21.6 Collection class Libraries

22 A World of Interacting Objects

- 22.1 RefCards
- 22.1.1 Design
- 22.1.2 Implementation
- 22.2 InfoStore
- 22.2.1 Initial design outline for InfoStore

- 22.2.2 Design and Implementation of the Vocabulary class
- 22.2.3 Other classes in the InfoStore program
- 22.2.4 Final class design for the InfoStore program

23 Intermediate class

- 23.1 Shared Class Properties
- 23.2 Friends
- 23.3 Iterators
 - 23.3.1 ListIterator
 - 23.3.2 TreeIterator
- 23.4 Operator functions
 - 23.4.1 Defining operator functions
 - 23.4.2 Operator functions and the iostream library
- 23.5 Resource Manager Classes and Destructors
 - 23.5.1 Resource management
 - 23.5.2 Destructor functions
 - 23.5.3 The assignment operator and copy constructors
- 23.6 Inheritance
 - 23.6.1 Discovering similarities among prototype classes
 - 23.6.2 Defining Class Hierarchies in C++
 - 23.6.3 But how does it work?!
 - 23.6.4 Multiple Inheritance
 - 23.6.5 Using Inheritance

24 Two more "trees"

- 24.1 AVL Trees
 - 24.1.1 What's wrong with binary trees?
 - 24.1.2 Keeping your balance
 - 24.1.3 An implementation
 - 24.1.4 Testing!
- 24.2 BTree
 - 24.2.1 Multiway trees
 - 24.2.2 A tree on a disk?
 - 24.2.3 BTree: search, insertion, and deletion operations
 - 24.2.4 An implementation
 - 24.2.5 Testing

25 Templates

- 25.1 A general function and its specializations
- 25.2 Language Extensions for Templates
 - 25.2.1 Template declaration
 - 25.2.2 Template instantiation
- 25.3 Specialized Instantiations of Template Code
- 25.4 A Template Version of QuickSort
- 25.5 The Template Class "Bounded Array"
- 25.6 The Template Class Queue

- 26 Exceptions**
- 26.1 C++'s Exception Mechanism
- 26.2 Example: Exceptions for class Number
- 26.3 Identifying the Exceptions that you intend to throw

- 27 Example: Supermarket**
- 27.1 Background and Program Specification
- 27.2 Design
- 27.2.1 Design preliminaries
- 27.2.2 Scenarios: identifying objects, their classes, their responsibilities, their data
- 27.2.3 Filling out the definitions of the classes for a partial implementation
- 27.3 A partial implementation
- 27.4 Finalising the Design
- 27.4.1 Histograms
- 27.4.2 Simulating the checkouts and their queues
- 27.4.3 Organizing the checkouts

- 28 Design and documentation: 2**
- 28.1 Object-Based Design
- 28.2 Documenting a design

Part V

- 29 The Power of Inheritance and Polymorphism**
- 29.1 The "Dungeon" Game
- 29.2 Design
- 29.2.1 Preliminaries
- 29.2.2 WindowRep and Window classes
- 29.2.3 DungeonItem hierarchy
- 29.3 An implementation
- 29.3.1 Windows classes
- 29.3.2 Class Dungeon
- 29.3.3 DungeonItems

- 30 Reusable designs**
- 30.1 The RecordFile Framework: Concepts
- 30.2 The Framework classes: Overview
- 30.3 The Command Handler Classes
- 30.3.1 Class declarations
- 30.3.2 Interactions
- 30.3.3 Implementation Code
- 30.4 Collection Classes and their Adapters
- 30.5 Class Record
- 30.6 The Windows Class Hierarchy
- 30.6.1 Class Responsibilities
- 30.6.2 Implementation Code

- 30.6.3 Using the concrete classes from a framework
- 30.7 Organizational Details

31 Frameworks for understanding

- 31.1 "Resources" and "Macros"
- 31.2 Architects, Experts, and Wizards
- 31.3 Graphics
- 31.4 Persistent Data
- 31.5 The Event Handlers
- 31.6 The Same Patterns in the code

Figures

Part I Introduction

- 1.1 Schematic diagram of major parts of a simple computer.
- 1.2 Principal components of a CPU.
- 1.3 CPU registers.
- 1.4 Simplified representation of an instruction inside a computer.
- 1.5 Memory organized in words with integer addresses.
- 1.6 Some of the 256 possible bit patterns encodable in a single byte and the (printable) characters usually encoded by these patterns.
- 1.7 Representing integers in two byte (16-bit) bit patterns in accord with the "two's complement notation scheme".
- 1.8 Bus
- 1.9 Bits are stored in concentric "tracks" on disk.
- 1.10 Tracks divided into storage blocks.
- 1.11 Read/write head assemble and multiplatter disks.
- 1.12 Disk controller.
- 1.13 A disk data transfer using direct memory access.
- 1.14 Simple file directory and file allocation scheme.
- 1.15 Controller for a simple i/o device such as a keyboard.
- 1.16 Input from a simple character-oriented device like a keyboard.

- 2.1 Executing a program.
- 2.2 View into the machine.
- 2.3 Simple layout for an instruction word.
- 2.4 "Assembling" an instruction.
- 2.5 Assembly language source deck.
- 2.6 Generating a symbol table in "pass 1" of the assembly process.
- 2.7 Generating code in "pass 2" of the assembly process.
- 2.8 Sharing a piece of code as a subroutine.
- 2.9 Organization of program in memory of an early computer.

- 3.1 Part of main memory is reserved for "system's" code and data.
- 3.2 System's code, library code, and user's code.
- 3.3 Operating the computer.
- 3.4 "Mainframe" computer system of early 1960s.
- 3.5 Early multiprogrammed system.
- 3.6 Drum storage.
- 3.7 MIT's pioneering "time-sharing" computer system.
- 3.8 Program "overlay" scheme.
- 3.9 O.S. of late 1960s early 1970s.
- 3.10 Early networked system, such as the airline seat reservation systems.

- 4.1 An array, or matrix, of places where an engineer wants to analyze properties of a steel beam that gets represented as an array of data values in a program.
- 4.2 Linking a program from separately compiled and library files.
- 4.3 Simple arrangement of code and data areas in memory (as used by early FORTRAN compilers and similar systems).
- 4.4 A "recursive" problem solving process.
- 4.5 Stack during execution of recursive routine.
- 4.6 Algol style compilation and loading processes.
- 4.7 Compiling, assembling, and link-loading a C program.

- 5.1 Illustration of typical editing and project windows of an example integrated development environment.
- 5.2 Display from a debugger.

Part II Simple Programs

- 6.1 Screen dump of program run through the debugger.

- 8.1 Principles of root finding!

- 9.1 Programs using file i/o to supplement standard input and output.
- 9.2 Ends of files.

Part III Functions and data aggregates

- 10.1 Stack during function call.
- 10.2 A "stack backtrace" for a recursive function.
- 10.3 Estimating π .
- 10.4 Character based plotting of a function.

- 11.1 Stack with when calling `double mean(double data_array[], int numitems)`.

- 12.1 Picture produced using the cursor graphics package.
- 12.2 Creating a program from multiple files
- 12.3 Curses model for an output window.
- 12.4 The heat diffusion experiment.
- 12.5 Temperature plots for the heat diffusion example.
- 12.6 The curses based implementation of the Hangman game.
- 12.7 Configurations for Conway's Life.

- 13.1 Illustrating selection sort.
- 13.2 A bureaucracy of clerks "Quicksorts" some data. The circled numbers identify the order in which groups of data elements are processed.
- 13.3 Graphs – networks of nodes and edges.
- 13.4 The bipartite matching (or "high school prom") problem. Find a way of matching elements from the two sets using the indicated links.
- 13.5 Maximal flow problem: what is the maximum flow from source to sink through the "pipes" indicated.
- 13.6 The "package wrapping" (convex hull) problem.

- 15.1 From program specification to initial functional decomposition.
- 15.2 "Call graph" for a program.
- 15.3 Further examples of "call graphs".
- 15.4 "Flowchart" representation, an alternative to a pseudo-code outline of a function.

- 16.1 XEvents – an example of a union from the Xlib library.

- 17.1 A file of "customer records".

- 18.1 A simplified form of hash table.
- 18.2 Simple information retrieval system.
- 18.3 Bit maps and chosen bit numbering.
- 18.4 Finding and counting the common bits.

Part IV A Touch of Class

- 19.1 A simple preliminary design diagram for a class.
- 19.2 Large integers represented as arrays.
- 19.3 "Overflow" may occur with any fixed size number representation.
- 19.4 Mechanism of addition with carry.
- 19.5 Shift and simple product sub-operations in multiplication.

- 20.1 Program "segments" in memory: code, static data, stack, and "heap".
- 20.2 "Blocks" allocated on the heap.
- 20.3 Pointer assignment.
- 20.4 Representing an array of initialized character pointers.
- 20.5 Hash table using pointers to strings.
- 20.6 The "Air Traffic Control" problem.
- 20.7 Building a "list structure" by threading together "list nodes" in the heap.

- 21.1 Queue represented as a "circular buffer" and illustrative Append (put) and First (get) actions.
- 21.2 "Tree" of partially ordered values and their mapping into an array.
- 21.3 Addition of data element disrupts the partial ordering.
- 21.4 Restoration of the partial ordering.
- 21.5 Removal of top priority item followed by restoration of partial ordering.
- 21.6 A "Dynamic Array".
- 21.7 The first append operation changing a List object.
- 21.8 Adding an element at the tail of an existing list.
- 21.9 Removing an item from the middle of a list.
- 21.10 A "binary search tree".
- 21.11 Inserting a new record.
- 21.12 Removing a "leaf node".
- 21.13 Removing a "vine" node.
- 21.14 Removing a "internal" node and promoting a successor.
- 21.15 Printed representation of binary tree.

- 22.1 First idea for classes for RefCards example.
- 22.2 Object interactions when loading from file (activities resulting from execution of `UserInteraction::Initialize()` for an UI object).
- 22.3 Object interactions resulting from `UserInteraction::Terminate()`.
- 22.4 Object interactions resulting from a User's "add card" command.
- 22.5 Object interactions resulting from a User's "show card" or "change card" commands.
- 22.6 Object interactions resulting from a User's "delete card" command.
- 22.7 Object interactions resulting from a User's "view" command.
- 22.8 Illustration of direct file dependencies in the RefCard program.
- 22.9 First idea for classes for InfoStore example.
- 22.10 Object interactions while loading a vocabulary file.
- 22.11 Object interactions while saving to a vocabulary file.
- 22.12 Revised model for classes.
- 22.13 Object interactions while adding articles to information store.
- 22.14 Final design for class InfoStore.
- 22.15 Final design for class IndexEntry.

- 23.1 Tree and tree iterator.
- 23.2 Illustration of groupings involved in concatenated use of `ostream& operator<<()` functions.
- 23.3 Resource manager class with memory leak.
- 23.4 Assignment leading to sharing of resources.
- 23.5 Similarities among classes.
- 23.6 Virtual tables.

- 24.1 Rearranging a tree to maintain perfect balance.
- 24.2 Example trees: AVL and not quite AVL.
- 24.3 Effects of some additions to an AVL tree.
- 24.4 Manoeuvres to rebalance a tree after a node is added.
- 24.5 Rebalancing a tree after a deletion.
- 24.6 Structure of a node for a "2-3" multiway tree.
- 24.7 An example "2-3" tree.
- 24.8 Splitting a full leaf node to accommodate another inserted key.
- 24.9 Another insertion, another split, and its ramifications.
- 24.10 The tree grows a new root.
- 24.11 Features of a BTree.
- 24.12 Mapping a binary tree onto records in a disk file.
- 24.13 Mapping a file onto disk blocks.
- 24.14 BTree nodes and the BTree file.
- 24.15 Simple insertion into a partially filled node.
- 24.16 Splitting a BTreeNode.
- 24.17 Reorganizing a pair of BTreeNodes after a "split".
- 24.18 Splitting leading to a new root node.
- 24.19 Moving data from a sibling to restore BTree property of a "deficient" BTreeNode.
- 24.20 Combining BTreeNodes after removing a key.
- 24.21 Removal of a key may lead to a change of root node.

- 25.1 Illustration of conceptual mechanism for template instantiation.

- 27.1 Scenario for creation of Customer objects.
- 27.2 Scenario for deletion of Customer objects.
- 27.3 Scenarios for creation and deletion of Checkout objects.
- 27.4 A hierarchy of "Activity" classes.
- 27.5 Interactions with Customer::Run().
- 27.6 Module structure and header dependencies for partial implementation of "Supermarket" example.
- 27.7 Design diagram for class Checkout.
- 27.8 Final design diagram for class Manager.
- 27.9 Final design diagram for class Shop.

Part V Object Oriented Programming

- 29.1 "Dungeon" game's display.
- 29.2 Class WindowRep and the Windows class hierarchy.
- 29.3 NumberItem and EditText Windows
- 29.4 DungeonItem class hierarchy.

- 29.5 A digitized line.
- 29.6 Some of the object interactions in `Dungeon::Run`.
- 29.7 Object interactions during `Player::Run`.
- 29.8 Module structure for Dungeon game example.

- 30.1 A "record" as handled by the "RecordFile Framework".
- 30.2 The menu of commands for record manipulation.
- 30.3 Record from another "RecordFile" program.
- 30.4 Class hierarchy for "RecordFile" Framework.
- 30.5 Design diagram for class `Document`.
- 30.6 Framework defined interactions for an `Application` object handling an "Open" command.
- 30.7 Opening a document whose contents are memory resident (interactions initiated by a call to `ArrayDoc::OpenOldFile()`).
- 30.8 Some of the interactions resulting from a call to `Document::DoNewRecord()`.
- 30.9 Example trace of specific object interactions while a `RecordWindow` is "posed modally".
- 30.10 Module structure for a program using the RecordFile framework.

- 31.1 Example of "network" data structure with pointers.