

Figure 13.4 Three basic gates for building circuits.

13.3 Advanced Case Study: Gates, Circuits, and Design Patterns

In this section we'll study the design and implementation of a class hierarchy and set of programs for simulating **logic gates** and circuits.⁷ We'll see how a well-designed class hierarchy enables us to model in software the same modular, component-based circuit construction that hardware designers use. We'll use the heuristics for programming with inheritance discussed earlier in the chapter as well as some of the **design patterns** from [GHJ95] that have become part of the standard toolkit for object-oriented programmers.

13.3.1 An Introduction to Gates and Circuits

In both an abstract and concrete sense, current computers are built from gates and chips, sometimes called **digital logic**. Although computers now have specially-designed chips for tasks like playing sound, doing graphics, and reading disks, at some low-level everything can be built from gates that regulate when and how electricity flows through a circuit. We'll use a standard set of three gates to construct different circuits and then use these circuits to construct other circuits. Eventually this process can lead to a working computer. Instead of physically building the gates we'll model them in software. The relationship between mathematical logic and digital logic was first recognized by Claude Shannon (see his biography in Sec. 4.3.3.)

The three gates we'll use are shown in Fig. 13.4. They are the **and-gate**, the **or-gate**, and the **inverter** or **not-gate**. Each of these gates corresponds to a boolean operator with the same name in C++. Traditionally, the behavior of these gates is shown with truth tables identical to those in Table 4.3 for the logical operators. Prog. 13.9 simply creates one of each gate and tests it with all possible inputs to show how gate behavior is the same as the behavior of the logical operators shown in Table 4.3.

⁷This example was motivated by a related example in [AS96]. If you read only one (other) book in computer science, that should be the one. It is simply the best introductory book on computer science and programming there is, though it's not easy reading.

Program 13.9 `gatetester.cpp`

```
#include <iostream>
using namespace std;

#include "gates.h"
#include "wires.h"

// show truth tables for each digital logic gate

int main()
{
    Gate * andg = new AndGate();
    Gate * org  = new OrGate();
    Gate * inv  = new Inverter();

    GateTester::Test(andg);
    GateTester::Test(org);
    GateTester::Test(inv);

    return 0;
}
```

`gatetester.cpp`

O U T P U T

```
prompt> gatetester
testing and (0)
-----
0 0      :      0
1 0      :      0
0 1      :      0
1 1      :      1
-----
testing or (0)
-----
0 0      :      0
1 0      :      1
0 1      :      1
1 1      :      1
-----
testing inv (0)
-----
0        :      1
1        :      0
-----
```

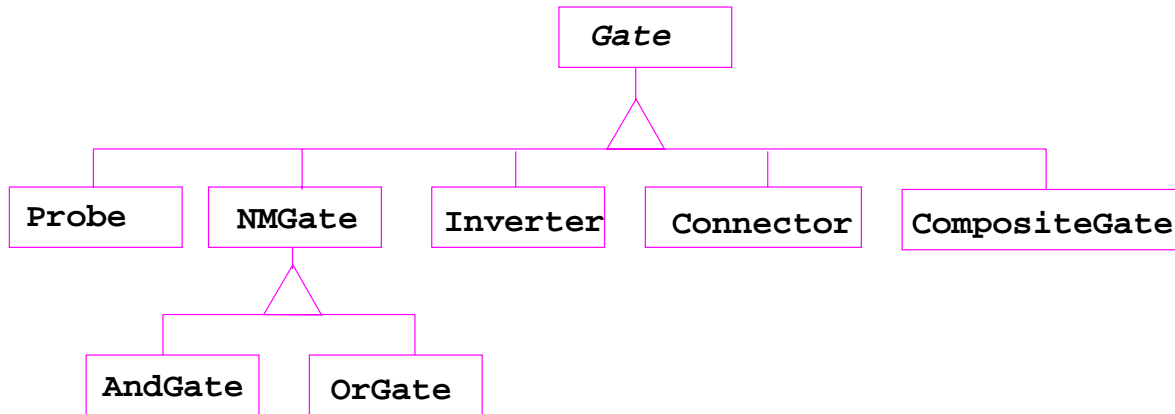


Figure 13.5 Hierarchy of components for building circuits.

The output is displayed using ones and zeros instead of true and false, but one corresponds to true and zero corresponds to false. If you look at *gatetester.cpp* carefully, you'll notice that `new` is called for three different types, but the returned pointer is assigned to variables of the same type: `Gate`. The inheritance hierarchy that enables this assignment is shown in Fig. 13.5. The class `GateTester`, included via `#include "gates.h"`, contains a static method `Test`. We could have made `Test` a free function, but by making it a static function in the `GateTester` class we avoid possible name clashes with other functions named `Test`.⁸ Gates by themselves aren't very interesting, to build circuits we need to connect the gates together using wires. Complex circuits are built by combining gates and wires together. Once a circuit is built, it can become a component in other circuits, acting essentially like a more complex gate.

13.3.2 Wires, Gates, and Probes

Prog. 13.10, *gatemiredemo.cpp* shows another method of gate construction. Wires are created, and then gates are constructed with wires attached to each gate's input(s) and output(s). The gates in Prog. 13.9 were constructed without wires attached to the inputs and outputs, but as we'll see later it's possible to attach wires after a gate has been constructed as well as to construct a gate from wires as shown in Prog. 13.10. All three of the principle logic gates: `and`, `or`, `inverter`, can be given a name when constructed as shown for the `andg` gate pointer.

The gates in *gatemiredemo.cpp* are wired together as shown in Fig. 13.6. An `and`-gate and an `or`-gate are attached so that the output of the `and`-gate feeds into the `or`-gate. In addition, `Probe` objects are attached to the output wires of the gates. As shown in

⁸The C++ `namespace` feature (see Sec. A.2.3) could also be used to avoid name conflicts, but several compilers still don't support namespaces.




Figure 13.6 A simple example using gates, wires, and probes.

Fig. 13.5, a Probe *is-a* Gate. Abstractly, gates are attached to wires (and vice versa), so a probe is similar to an and-gate in this respect. A Probe object prints a message whenever the current on the wire it's monitoring changes, but also prints a message when it's first attached to a wire.

When a wire is constructed, the current on the wire is set to zero/false. The current changes when it's either explicitly changed using `Wire::SetCurrent`, or when a change in current propagates from an input wire to an output wire through a gate. A careful reading of the program and output shows that wires can be printed, and that each wire is numbered in the order in which it's created (a static counter in `wires.h`, Prog. G.15, keeps track of how many wires have been created.) After the circuit is constructed, the probes detect and print changes caused by changes in the circuit.



Program 13.10 `gatewiredemo.cpp`

```
#include <iostream>
using namespace std;

#include "gates.h"
#include "wires.h"

int main()
{
    Wire * in  = new Wire(); // and-gate in
    Wire * in2 = new Wire(); // and-gate in
    Wire * in3 = new Wire(); // or-gate in
    Wire * aout = new Wire(); // and-gate out
    Wire * oout = new Wire(); // or-gate out

    Gate * andg = new AndGate(in,in2,aout,"andgate");
    Gate * org  = new OrGate(in3,andg->OutWire(0),oout);
    cout << "attaching probes" << endl;
    Probe * p = new Probe(aout); // attach to the and-out wire
    Probe * q = new Probe(oout); // attach to the or-out wire

    cout << "set " << *in << " on" << endl;
    in->SetSignal(true);
}
```

```

cout << "set " << *in2 << " on" << endl;
in2->SetSignal(true);
cout << "set " << *in << " off" << endl;
in->SetSignal(false);
cout << "set " << *in3 << " on" << endl;
in3->SetSignal(true);
return 0;
}

```

 gatwiredemo.cpp

After the probes are attached, the current on wire 0, one of the and-gate inputs, is turned on (or set). Since the other and-gate input has no current, no current flows out of the and-gate. When the current to wire 1 is set, the and-gate output (wire 3) becomes set and the probe detects this. Since the and-gate output is one of the or-gate inputs, the or-gate output (wire 4) is also set and the other probe detects this change. The probes continue to detect changes as current is turned off and on as illustrated in the program and output.

```

O U T P U T

prompt> gatwiredemo
attaching probes
(wire 3)      signal= 0
(wire 4)      signal= 0
set (wire 0) on
set (wire 1) on
(wire 4)      signal= 1
(wire 3)      signal= 1
set (wire 0) off
(wire 4)      signal= 0
(wire 3)      signal= 0
set (wire 2) on
(wire 4)      signal= 1

```

13.3.3 Composite Gates and Connectors

Program 13.12 shows two ways of constructing the **xor-gate** illustrated in Fig. 13.7. The output of an xor-gate is set when one of its inputs is set, but not when both are set. The truth-table generated by `GateTester::Test` for an xor-gate is shown in the output. Both methods create a `CompositeGate` object, another of the gates in the hierarchy shown in Fig. 13.5. A `CompositeGate` is a gate as shown in the diagram, but it's a gate made up of other gates. In particular, a composite gate can be formed from the basic gate building-blocks, but also from other composite gates. A collection of connected gates is also known as a **circuit**, so a `CompositeGate` object represents a circuit. The key idea here is that a circuit is also a gate for building other circuits.




Figure 13.7 Building an xor circuit.

Program Tip 13.8: The class `CompositeGate` is a concrete example of the *Composite* design pattern [GHJ95]. There, the pattern is a solution to a problem stated as “you want clients to be able to ignore the difference between compositions of objects and individual objects. Clients will treat all objects ...uniformly.”

A linked-list can also be viewed as a composite. A node is an individual item in a list, but it also represents a complete list since the node provides access to the rest of the list. Just as a node contains information and a pointer to another node, a `CompositeGate` contains information about the gate and many pointers to other gates, the gates that make up the circuit. Rather than dwell on the implementation of the class, we’ll see how it’s used to create complex circuits.

It’s almost simpler to create a new class `XorGate` than to build a composite gate that works like an xor-gate. However, creating a new class requires writing code and recompiling a program. As we’ll see in the final program from this chapter, it’s possible to create a gate-construction language or program that can build new gates while a program is running. The only member function of a class `XorGate` that differs in a substantive way from either `AndGate` or `OrGate` is `XorGate::Act`, the method that determines how a signal propagates through the gate.

Program 13.11 xorgate.cpp

```
class XorGate : public Gate
{
public:
    virtual void Act()
    {
        myOuts[0]->SetSignal(
            (myIns[0]->GetSignal() || myIns[1]->GetSignal()) &&
            !(myIns[0]->GetSignal() && myIns[1]->GetSignal())
        )
    }
};
```

xorgate.cpp

The partial class declaration and definition shown above captures in boolean logic and code exactly the relationship shown in digital logic in Fig. 13.7. The output is set when either input is set, but not when both inputs are set.

Constructing CompositeGate objects Three CompositeGate methods allow a complex gate to be constructed from other gates.

- AddGate adds a gate to a composite gate. Presumably the added gates will be connected in some way (otherwise the composite gate won't be very useful.)
- AddIn adds an input wire to a composite gate. Presumably each input wire is connected to a gate that's part of the composite object. Each call of AddIn adds a new input wire.
- AddOut adds an output wire to a composite gate. As with AddIn, presumably each added output wire is connected to one of the gates added to the composite.

Each of these methods is shown in MakeXOR of Prog. 13.12. Note that each call of AddIn and AddOut adds a wire that is an input (respectively output) of a gate already added to the composite. The input and output wires could be specified first, then the gates added; the net effect is the same.

Using the method Gate::clone. The MakeXOR function also shows the method Gate::clone applied to the AndGate object ag. The method clone is abstract⁹ in Gate so every concrete subclass must provide an implementation. Client programs typically define objects and reference them through pointers of type Gate * more often than by pointers of a specific subclass like AndGate * or CompositeGate *. Since clone is virtual, the object actually cloned returns a copy of itself.

```
void DoStuff(Gate * g)
// post: do something with a copy of g
{
    Gate * copy = g->clone();
    // what kind of gate is copy? we can't tell but
    // we can apply any generic Gate method to copy
}
```

In this example, the object referenced by copy is some kind of gate, and if clone works as expected copy is a duplicate of the gate g passed to the function DoStuff. The Gate::clone method is an example of what's often called a **virtual constructor**. The clone method is used to create objects, like a constructor, but the clone method is virtual so it creates an object whose type isn't known at compile time.

⁹Recall that I use *abstract* rather than the more C++ specific term *pure virtual*.

Program Tip 13.9: The clone method is a concrete example of what's called the *Factory Method* design pattern in [GHJ95]. There, the pattern is a solution to a problem paraphrased as “client code can't anticipate what kind of objects it must create or wants to delegate responsibility of creation to subclasses in a class hierarchy”.

Using Connectors The functions MakeXOR and MakeXOR2 illustrate the differences between calling Connect to connect wires to gate inputs and output (in MakeXOR) and constructing gates from existing wires (MakeXOR2). When gates are constructed without wires attached as they are in MakeXOR, the gate functions InWire and OutWire are used to access input wires and output wires, respectively, for attaching these wires to other wires using connectors. A connector is a gate that simply transfers current from one wire to another as though the wires are joined or soldered together.

As the output shows, the circuit created by MakeXOR uses more wires than the circuit created by MakeXOR2. When gates are constructed without wires in client code, each gate creates its own wires for input and output. Counting the input and output wires for each gate in Fig. 13.7 shows that there are 11 wires: $3 \times (2 \text{ and-gates}) + 3 \times (1 \text{ or-gate}) + 1 \times (2 \text{ inverters})$. The wires for the gate created by MakeXOR2 are explicitly created in the client program. There are fewer wires since, for example, the connections between the inputs of the right-most and-gate (whose output is the circuit's output) and their sources (the outputs of the or-gate and inverter) require only two wires whereas four wires are used by MakeXOR.

Program 13.12 xordemo.cpp

```
#include <iostream>
using namespace std;

#include "gates.h"
#include "wires.h"
#include "tvector.h"

// illustrate connecting wires to gates using Connect

CompositeGate * MakeXOR()
// post: return an xor-gate
{
    CompositeGate * xorg = new CompositeGate(); // holds xor-gate
    Gate * ag = new AndGate(); // build components
    Gate * ag2= ag->clone(); // and gate a different way
    Gate * og = new OrGate();
    Gate * inv = new Inverter();

    Connect(og->InWire(0), ag->InWire(1) ); // wire components
    Connect(og->InWire(1), ag->InWire(0) );
    Connect(ag->OutWire(0), inv->InWire(0));
}
```

```

Connect(inv->OutWire(0), ag2->InWire(1));
Connect(og->OutWire(0), ag2->InWire(0));

xorg->AddGate(ag); xorg->AddGate(ag2); // add gates to xor-circuit
xorg->AddGate(inv); xorg->AddGate(og);

xorg->AddOut(ag2->OutWire(0)); // add inputs/outputs
xorg->AddIn(og->InWire(0)); xorg->AddIn(og->InWire(1));

return xorg;
}

CompositeGate * MakeXOR2()
// post: returns an xor-gate
{
    CompositeGate * xorg = new CompositeGate();
    tvector<Wire *> w(6); // need 6 wires to make circuit
    tvector<Gate *> gates; // holds the gates in the xor-circuit
    int k;
    for(k=0; k < 6; k++)
    { w[k] = new Wire();
    }
    gates.push_back(new OrGate( w[0], w[1], w[2] )); // create wired gates
    gates.push_back(new AndGate(w[0], w[1], w[3] )); // share inputs
    gates.push_back(new Inverter(w[3], w[4] )); // and out->inv in
    gates.push_back(new AndGate(w[2], w[4], w[5] )); // combine or, inv

    for(k=0; k < gates.size();k++) // add gates to xor
    { xorg->AddGate(gates[k]);
    }
    xorg->AddIn(w[0]); xorg->AddIn(w[1]); // add inputs/outputs
    xorg->AddOut(w[5]);

    return xorg;
}

int main()
{
    CompositeGate * g = MakeXOR();
    CompositeGate *g2 = MakeXOR2();
    cout << "circuit has " << g->CountWires() << " wires" << endl;
    GateTester::Test(g);
    cout << "circuit has " << g2->CountWires() << " wires" << endl;
    GateTester::Test(g2);

    return 0;
}

```

xordemo.cpp

The code in `MakeXOR2` exploits the `Gate` class hierarchy by creating a vector of pointers to `Gate *` objects, but creating different kinds of gates for each pointer to reference. A vector of `Gate *` pointers is also used in the private section of the `CompositeGate` class to store the gates used in constructing the composite object. Although the functions `MakeXOR` and `MakeXOR2` create different digital circuits, the circuits are identical from

a logical view point: they compute the same logical operator as shown by the truth tables. The different functions create a `CompositeGate` using the same process.

1. Create an initially empty composite.
2. Construct gates, wire them together and add the gates to the composite.
3. Specify input wires and output wires for the composite.
4. The composite object is finished.

As we've noted, steps two and three can be interchanged, the relative order in which these steps are executed does not affect the final composite gate.

```

O U T P U T

prompt> xordemo
circuit has 11 wires
testing composite: 4 gates, 2 in wires, 1 out wires
-----
0 0      :      0
1 0      :      1
0 1      :      1
1 1      :      0
-----
circuit has 6 wires
testing composite: 4 gates, 2 in wires, 1 out wires
-----
0 0      :      0
1 0      :      1
0 1      :      1
1 1      :      0
-----
```

Pause to Reflect



13.16 Suppose a probe `pin` is added to the input wire `in` as part of Prog. 13.10:

```
Probe * q = new Probe(out); // in original program
Probe * pin = new Probe(in); // added for this exercise
```

As a result of adding this probe three lines of output are added, what are the lines and where do they appear in the output? (hint: one line is printed when the probe is attached.)

13.17 If the `AndGate` instance `andg` in Prog. 13.10 is tested at the end of `main`, the truth table printed is the standard truth table for an and-gate.



This happens even though the output of `andg` is connected to the input of the or-gate `org`. Why? (hint: is the circuit consisting of the and-gate and or-gate combined a `Gate` object?)

```
GateTester::Test( andg );
```

13.18 The probe `p` can be removed from the wire `about` at the end of Prog. 13.10 using a `Wire` member function. What's the function and what call uses it to remove the probe (see `wires.h`, Prog. G.15 for `Wire` methods)?

13.19 Write the function `RemoveProbe` whose header follows. (See `wires.h` and `gates.h` in Howto G.)

```
void RemoveProbe(Probe * p)
// post: p is removed from the wire it monitors/probes
```

13.20 The return type of the function `MakeXOR` is `CompositeGate *` in Prog. 13.12, `xordemo.cpp`. If the return type is changed to `Gate *` an xor-gate is still returned, but the call of `MakeXOR` below fails to compile.

```
CompositeGate * g = MakeXOR();
```

If `g`'s type is changed to `Gate *` the definition of `g` compiles, but then the output statement below fails to compile.

```
cout << "circuit has " << g->CountWires() << " wires" << endl;
```

What's the cause of this behavior (hint: `CountWires` is not a `Gate` method.)

13.21 The circuit diagrammed in Fig. 13.8 shows a circuit that is logically equivalent to an or-gate, but which is constructed from an and-gate and three inverters. Write a function `MakeOR` that returns a `CompositeGate` representing the circuit diagrammed in Fig. 13.8. Draw a similar circuit that's logically equivalent to an and-gate using only inverters and or-gates.




Figure 13.8 Building an or-gate from other basic gates.




Figure 13.9 A disabler circuit.

13.22 The circuit diagrammed in Fig. 13.9 is a **disabler** circuit. The signal on the wire labelled *Control* determines if the signal on the (other) input wire is passed through to the output wire of the disabler. When the control signal is zero (off), the input signal goes through to the output, i.e., the input and the output are the same. When the control signal is set, i.e., true/one, the input signal is stopped, or disabled, and the output wire is false/zero regardless of the value on the input wire.

Write a function `MakeDisabler` that returns a disabler circuit. Construct both gates without wires so that you must use `Connect` to wire the circuit together. How many wires are used in the circuit? (Do this exercise on paper, not necessarily by writing and testing a function.) Implement an alternative version called `MakeDisabler2` which does not use `Connect` so that both gates in the circuit are constructed with wires. How many wires are used in the circuit?

13.23 Write the method `Disabler::Act` that represents the logic of a disabler circuit. Model the function on the version of `XorGate::Act` shown in Sec. 13.3.3.

13.24 The **comparator** circuit shown in Fig. 13.10 determines whether the signal on the wire labeled *R* is less than the signal on the wire labeled *C*, where one/zero are used for true/false. (*continued*)

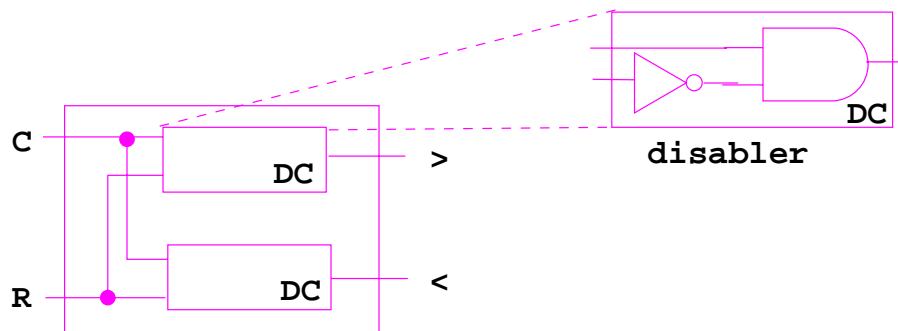


Figure 13.10 A comparator circuit for selecting the larger of two values.

Write a truth table for the circuit by tracing all four possible combinations of zero/one for inputs and labeling the corresponding outputs. Verify that if the signals are the same, the outputs are both zero. If $R < C$ then the lower output wire labeled $<$ is one/true and the upper wire is zero/false. If $R > C$ then the upper output labeled $>$ is one/true and the lower wire is zero/false.

13.25 Assume a function `MakeDisabler` is written that returns a disabler-gate/circuit. Use this function to write a `MakeComparator` function that returns a composite-gate encapsulating a comparator circuit.

13.26 Do you expect the truth tables printed by the two calls of `GateTester::Test` that follow to be the same? Why?

```
void TruthTwice(Gate * g)
{
    Gate * copy = g->clone();
    GateTester::Test(g);
    GateTester::Test(copy);
}
```

13.3.4 Implementation of the wire and Gate classes

The interactions between classes in the `Gate` hierarchy and the class `Wire` are fairly complex. It's not essential to understand these interactions to write simple programs like the ones we studied in previous sections, but a solid understanding of the interactions is needed before you write your own `Gate` subclasses or write more involved programs.

Once we've looked at the classes and their implementations in more detail, we'll be able to make judgments about the overall design of the `Gate/Wire` framework. We'll see that there are some problems in the `Gate` class hierarchy that make it more difficult to add an `XorGate` subclass than it should be. It's not difficult to add such a class, but the process would be considerably more simple with the introduction of a new class encapsulating behavior common to `AndGate`, `OrGate`, and what would be `XorGate`. As we've stressed, software should be grown: the design process does not finish when you have a working program or prototype. Since programs and classes evolve, it makes sense to step back and examine a design and implementation after the initial kinks have been ironed out.

ProgramTip 13.10: Class methods sometimes need to be *re-factored* into other classes, or into new classes that weren't part of an initial design.

Re-factoring means you don't add new functionality, but you re-distribute (to existing classes) or re-assign (to new classes) existing behavior and functionality to make classes and code more re-usable.

As a start towards understanding the design we'll consider the simple code in Prog. 13.13 that creates an or-gate, attaches a probe to the output of the gate, and sets

one of the input gates to true. The interactions and method calls made by all classes for the three lines of code in *gwinteraction.cpp* are shown in Fig. 13.11.

Program 13.13 *gwinteraction.cpp*

```
#include <iostream>
using namespace std;

#include "gates.h"
#include "wires.h"

int main()
{
    Gate * org = new OrGate();
    Probe * p = new Probe(org->OutWire(0));
    org->InWire(0)->SetSignal(true);

    return 0;
}
```

gwinteraction.cpp

OUTPUT

```
prompt> gwinteraction
(wire 2)      signal= 0
(wire 2)      signal= 1
```

Two separate concepts generate almost all the interactions shown in Fig. 13.11. We'll give an overview of each concept, discuss why they're used in the *Wire/Gate* framework, and then provide a more in-depth look at each of them.

1. A *Wire* object can have any number of gates attached to it. Every time the signal on a wire changes, the wire notifies all the attached gates that the signal has changed using the method `Gate::Act`. Each gate responds differently when it's acted on, e.g., probes print a value, or-gates propagate a true value to their output wire if one of their input wires is set, and so on.
2. When a *Gate* is constructed without wires, e.g., as in *gwinteraction.cpp* or in *MakeXOR* as opposed to *MakeXOR2* of Prog. 13.12, *xordemo.cpp*, the gate creates its own wires. Rather than calling `new Wire` directly, a gate requests a wire from a *WireFactory* associated with the entire *Gate* hierarchy by a static instance variable of the *Gate* class.

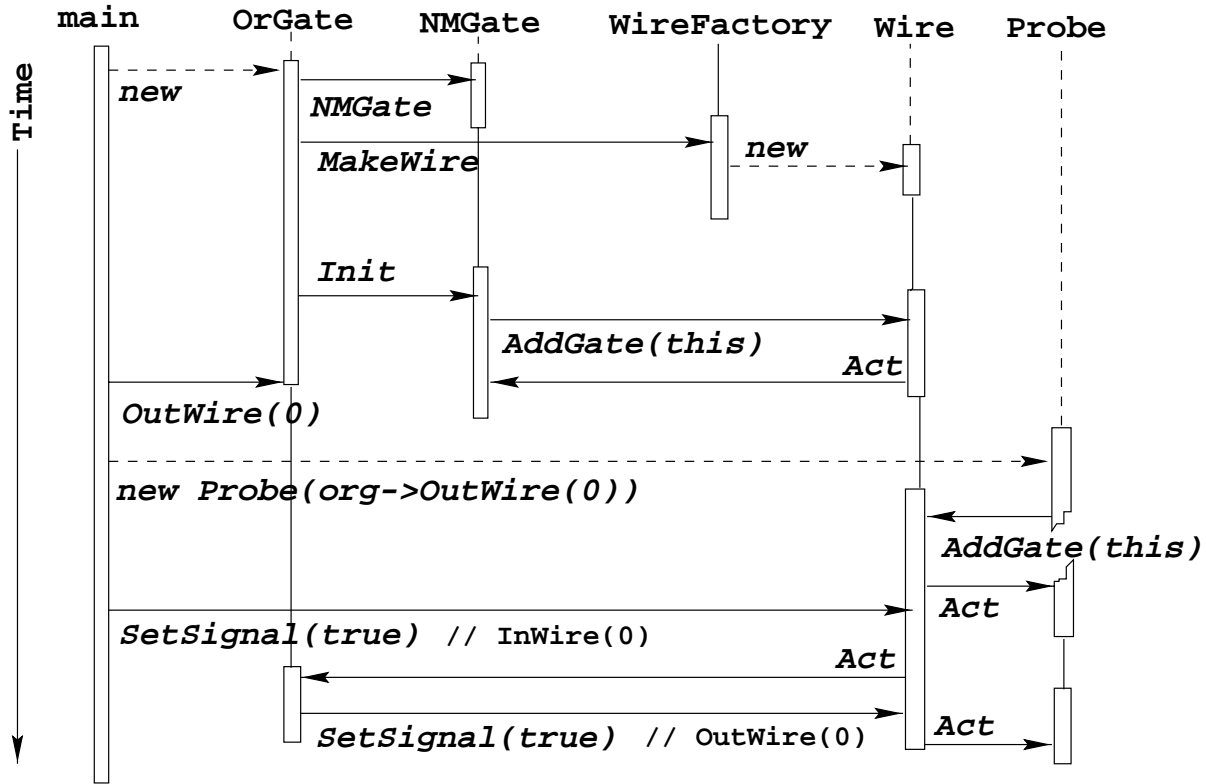


Figure 13.11 Interaction diagram: creating an or-gate with no connected wires, attaching a probe to the output of the gate, and setting the signal on the first of the gate's two inputs.

13.3.5 Gates and Wires: Observers and Observables

Look carefully at the inputs to the xor-gate diagrammed in Fig. 13.7 and the comparator diagrammed in Fig. 13.10. In both cases one wire is attached to the inputs of two different gates. Any change in the wire must propagate a signal through both gates. Suppose a probe is attached to one of the input wires that feeds into more than one gate. Then a change in the wire must notify two gates and a probe, which is really three gates since a probe is-a gate. How are changes in a circuit propagated? In the framework discussed here, a gate attaches itself to a wire, or can be attached by another object to a wire using the method `Wire::AddGate`. For example, a `Probe` instance adds itself during construction to the wire it probes.

```

Probe::Probe(Wire * w)
    : myWire(w)
    // post: probe is attached to w
    
```

```
{
    myWire->AddGate(this);
}
```

It's almost as though each attached gate listens to the wire, waiting for a change. However, a gate doesn't actively listen, it is notified by the wire when the wire's signal changes. The wire notifies all the gates that have been attached to it using the following code.

```
void Wire::SetSignal(bool signal)
// post: notify all attached/listening gates if signal changes
{
    if (signal != mySignal)
    {
        mySignal = signal;
        int k;
        for(k=0; k < myGates.size(); k++)
        {
            myGates[k]->Act( );
        }
    }
}
```



You can look at the code in *wires.h* for details (see Howto G), but the code above is mostly self-explanatory from the names of the instance variables and the syntax of how they're used — for example, `myGates` seems to be a `tvector` object from how it's used. Gates that have been attached using `AddGate` can subsequently be removed using `Wire::RemoveGate`. Gate identity for removal is based on pointer values, so any object added can be removed since the address of an object doesn't change.

Program Tip 13.11: In [GHJ95] the *Observer* pattern is a solution to a problem “when a change to one object requires changing others, and you don't know how many objects need to be changed.” The Observer pattern is sometimes called Observer/Observable or Publish/Subscribe.

In the code above you can see that a wire's gates are notified in the same order in which they are added to the wire. Suppose `Wire` object `w2` notifies the first of the two gates that are (hypothetically) attached to `w2`. Since a gate's `Act` method may set other wires, that will in turn call other `Act` methods, the second gate attached to `w2` may have its `Act` method invoked well after other gates have acted. In one of the modifications in the Exercise section you'll be asked to introduce time into the `Wire/Gate` framework to account for these anomalies.

The Observer pattern is common outside of programming. Volunteer firemen are notified when there's an event they must respond to, but the firemen do not actively phone the fire department to find fires. The firemen correspond to gates in our framework, the fire department is the wire notifying the firemen. Auctions sometimes model the pattern: bidders are notified when a new, higher bid has been made. A bidder actively monitoring new bids doesn't quite fit the model, but a bidder that responds only when notified of a new bid does.

Bjarne Stroustrup (b. 195?)

Bjarne Stroustrup is truly the “father” of C++. He began its design in 1979 and is still involved with both design and implementation of the language. His interests span computer science, history, and literature. In his own words:



... C++ owes as much to novelists and essayists such as Martin A. Hansen, Albert Camus, and George Orwell, who never saw a computer, as it does to computer scientists such as David Gries, Don Knuth, and Roger Needham. Often, when I was tempted to outlaw a feature I personally disliked, I refrained from doing so because I did not think I had the right to force my views on others. In writing about creating software, Stroustrup (p. 693) [Str97] mentions several things to keep in mind, three are ideas we’ve emphasized in this book: (1) There are no “cook-book” methods that can replace intelligence, experience, and good taste in design and programming, (2) Experimentation is essential for all nontrivial software development, and

(3) Design and programming are iterative activities.

Stroustrup notes that it is as difficult to define what a programming language is as to define computer science.

Is a programming language a tool for instructing machines? A means of communicating between programmers? A vehicle for expressing high-level designs? A notation for algorithms? A way of expressing relationships between concepts? A tool for experimentation? A means of controlling computerized devices? My view is that a general-purpose programming language must be all of those to serve its diverse set of users.

For his work in the design of C++, Stroustrup was awarded the 1994 ACM Grace Murray Hopper award, given for fundamental contributions made to computer science by work done before the age of 30. Most of this material is taken from [Str94].

13.3.6 Encapsulating Construction in wireFactory

The simple three-line program in Prog. 13.13, constructs an or-gate without providing wires when the or-gate is constructed. The or-gate makes its own wires, and the program connects a probe to the created output wire. As we saw in the two different functions

MakeXOR and MakeXOR2 of Prog. 13.12, a gate can be created by attaching existing wires to the gate when the gate is constructed, or by creating a gate and then connecting wires to the input/output wires the gate constructs itself. Where do these self-constructed wires come from? The simplest method is to create new wires using `new Wire()` — sample code for the `Inverter` constructor shows this (this isn't the real constructor, which uses a different technique discussed later). An `Inverter` has an input, an output, a name, and a number.

```
Inverter::Inverter(const string& name)
: myIn(new Wire(name)), myOut(new Wire(name))
  myName(name), myNumber(ourCount)
{
    ourCount++;
    myIn->AddGate(this);
}
```

Since an `Inverter` creates the wires using the `new` operator, the class is responsible for deleting the wires in its destructor. This approach tightly couples the `Gate` and `Wire` classes. If a better wire class is designed, or we want to run a circuit simulation using a `LowEnergyWire` class representing a new kind of wire that's a subclass of `Wire`, we'll have to rewrite every gate's constructor to use the new kind of wire. We can't reduce the coupling inherent in the circuit framework because wires and gates do depend on each other, but we can reduce the coupling in how gates create wires. To do this we design a `WireFactory` class. When a client wants a wire, the wire is "ordered" from the factory rather than constructed using `new`. If a new wire class is created, we order wires from a new factory that makes the new kind of wires. Because we use inheritance to model is-a relationships, the new kind of wires can be used in place of the original wires since, for example, a `LowEnergyWire` is-a `Wire`. By isolating wire creation in a `WireFactory`, changing the kinds of wires used by all gates means simply changing the factory, and the factory is created in one place so it can be changed easily. The `Inverter` constructor actually used in `gates.cpp` illustrates how a factory isolates wire construction in one place.

```
Inverter::Inverter(const string& name)
: myIn(ourWireFactory->MakeWire(name)),
  myOut(ourWireFactory->MakeWire(name)),
  myName(name), myNumber(ourCount)
{
    ourCount++;
    myIn->AddGate(this);
}
```

The `my/our` naming convention tells us that `ourWireFactory` is a static instance variable. The factory is shared by every `Gate` object since it's defined as a protected static data member in the abstract `Gate` superclass. This means every `Inverter`, every `AndGate`, and every gate subclass not yet implemented can share the factory.

Program Tip 13.12: Using a factory class to isolate object creation decreases the coupling between the created objects and their collaborating classes. This design pattern is called *Abstract Factory* in [GHJ95]. A factory class is used when “a system should be independent of how its products are created, composed, and represented” or when “a system should be configured with one of multiple families of products”.

Our `WireFactory` class is not abstract, but we’ll explore how to create more than one kind of factory in the exercises by creating an abstract base class from which `WireFactory` derives. The `Gate::clone` method outlined in Program Tip 13.9 as a realization of a factory *method* shares characteristics with the `WireFactory` class that is a factory *class*: both isolate object creation so that clients can use objects without knowing how to create them.

13.3.7 Refactoring: Creating a `BinaryGate` class

When I first designed the `Gate` hierarchy in Fig. 13.5 I anticipated creating classes like `And3Gate`, an and-gate with three inputs that sets its output only when all three inputs are set. I considered an `And3Gate` to be a 3-1-gate, a gate with three inputs and one output. The existing `AndGate` class represents a 2-1-gate while the comparator circuit diagrammed in Fig. 13.10 is a 2-2-gate with two inputs and two outputs. Similarly, the full-adder diagrammed in Fig. 13.13 is a 3-2-gate. Thinking there would be some common behavior in these gates I created a class `NMGate` to model an n-m-gate as I’ve just described. Since a subclass is responsible for calling its superclass constructor, this leads to the constructor below for an `AndGate` instance constructed without wires.

```
AndGate::AndGate(const string& name)
    : NMGate(ourCount, name)
// post: this and-gate is constructed
{
    tvector<Wire *> ins(2), outs(1);
    ins[0] = ourWireFactory->MakeWire(myName);
    ins[1] = ourWireFactory->MakeWire(myName);
    outs[0] = ourWireFactory->MakeWire(myName);
    NMGate::Init(ins, outs);
    ourCount++;
}
```

The `AndGate` constructor creates two input wires, one output wire, and puts these wires into vectors for initializing the parent `NMGate` class. The general class `NMGate` is initialized with vectors of wires for input and output so that it can be used for a 3-2-gate as well as an 8-8-gate. The `OrGate` constructor shows striking similarities to the `AndGate`.

```
OrGate::OrGate(const string& name)
    : NMGate(ourCount, name)
```

```

{
    tvector<Wire *> ins(2), outs(1);
    ins[0] = ourWireFactory->MakeWire(myName);
    ins[1] = ourWireFactory->MakeWire(myName);
    outs[0] = ourWireFactory->MakeWire(myName);
    NMGate::Init(ins,outs);
    ourCount++;
}

```

This duplicated code will be replicated in any new 2-1-gate, e.g., if we implement an `XorGate` class. The `Act` methods of these classes differ because the gates model different logic, and the `clone` methods differ since each gate must return a copy of itself, but the other `AndGate` and `OrGate` methods are the same. Since 2-1-gates are quite common, and we may be implementing more “basic” 2-1-gates in the future, it’s probably a good idea to refactor the behavior in common to the 2-1-gates into a new class `BinaryGate`. The new class derives from `NMGate` and is a parent class to `AndGate` and `OrGate`. The `AndGate` constructor will change as follows.

```

AndGate::AndGate(const string& name)
    : BinaryGate(ourCount,name)
// post: this and-gate is constructed
{
    ourCount++;
}

```

The behavior common to the `AndGate` and `OrGate` constructors has been factored out into the `BinaryGate` constructor. Similarly, all the methods whose behavior is the same in the binary gate subclasses are factored into the new `BinaryGate` superclass.

Pause to Reflect



13.27 The method `Wire::AddGate` is implemented as follows.

```

void Wire::AddGate(Gate * g)
// post: g added to gate collection, g->Act() called
{
    myGates.push_back(g);
    g->Act();
}

```

Identify each call of `g->Act()` whose source is `AddGate` that appears in the interaction diagram of Fig. 13.11. Which of the calls generate(s) output?

13.28 Constructing an `Inverter` and connecting a probe to its output generates the output shown. (*continued*)

```

Gate * inv = new Inverter();
Probe * p = new Probe(inv->OutWire(0));

```

O U T P U T

```
(wire 1)      signal= 1
```

Why is the wire labeled `(wire 1)`, where is wire 0? Draw an interaction diagram like Fig. 13.11 for these two statements. Trace all method calls, particularly the `Gate::Act` calls, and show why the call of `g->Act()` in `Wire::AddGate` shown in the previous exercise is necessary to get the behavior shown in the output — what would the output of the probe be if the call `g->Act()` wasn't included in the method `AddGate`? Why?

- 13.29** The statements below construct a disabler circuit as diagrammed in Fig. 13.9. The circuit isn't formed as a composite, but the gates and wires together make a disabler circuit with a probe attached to the circuit's output wire.

```
Wire * controller= new Wire();
Gate * ag=        new AndGate();
Gate * inv=       new Inverter(controller, ag->InWire(1));
Probe * p=        new Probe(ag->OutWire(0));
ag->InWire(0)->SetSignal(true); // send a signal through
```

Since the `controller` is false/zero when constructed, the signal set should propagate through the disabler. Draw an interaction diagram like Fig. 13.11 for these five statements.

- 13.30** As implemented, the `WireFactory` class cannot recycle used wires, i.e., if a `Gate` is destroyed, the wires it may have ordered from the factory are not re-used. The factory does keep track of all the wires ever allocated/ordered, and cleans the wires up when the factory ceases to exist.

In what function does the current `WireFactory` destroy all the wires allocated during the factory's lifetime? Sketch a design that would allow the factory to recycle wires no longer needed. You'll need to identify how the factory stores recycled wires and how the factory collaborates with the `Gate` classes to get wires back when a gate no longer needs them.

- 13.31** The class `NMGate` is an abstract class because it has at least one abstract/pure virtual function, e.g., `Act`. However, there is an `NMGate` constructor and an `NMGate` class has state: the input and output wires. Why is the class an abstract class, which means it's not possible to create an `NMGate` object, but the class still has a constructor and state? Note that the statement below will not compile for two reasons: the constructor is protected and the class is abstract.

```
Gate * g = new NMGate(); // won't compile
```

13.32 Why is `ourCount++` used in the body of the refactored `AndGate` constructor at the end of Sec. 13.3.7? Why isn't the increment factored into the `BinaryGate` constructor?

13.33 The following statement, added as the last statement in `main` of Prog. 13.12, `xordemo.cpp`, produces the output shown.

```
cout << g2->deepString() << endl;
```

The output shows all the components of the composite gate `g2` created by `MakeXOR2`. The method `deepString` is implemented in each `Gate` subclasses, although it often defaults to the same function as `toString`. Why are the and gates numbered 2 and 3? where are and gates numbered 0 and 1? Draw the circuit for this composite and label every gate and wire with its number.

OUTPUT

```
composite: 4 gates, 2 in wires, 1 out wires
all-in   (wire 11) (wire 12)
all-out  (wire 16)
      or (1)
      in  (wire 11) (wire 12)      out (wire 13)
----
      and (2)
      in  (wire 11) (wire 12)      out (wire 14)
----
      inv (1)
      in  (wire 14)      out (wire 15)
----
      and (3)
      in  (wire 13) (wire 15)      out (wire 16)
----
-----
```

13.34 Instead of refactoring `AndGate` and `OrGate` into a new `BinaryGate` class, suppose a new constructor is added to the `NMGate` class in which the number of inputs and outputs is specified as shown in the following:

Is this a better solution than introducing a new class `BinaryGate`? Why? Write the constructor that takes the number of inputs and outputs as parameters.

```
AndGate::AndGate(const string& name)
    : NMGate(2,1,ourCount,name)
// post: this and-gate is constructed
{
    ourCount++;
}
```

13.3.8 Interactive Circuit Building

Prog. 13.12, *xordemo.cpp*, shows how a composite circuit can be built by creating gates and wires, then wiring them together. In Sec. 13.3.3 we described how to create new class declarations and definitions using an `XorGate` class as an example. Both these methods for creating circuits require writing, compiling, testing, and debugging programs. A different approach is outlined in the run of *circuitbuilder.cpp*. A complete version of this program is not provided, you'll be asked to write in as an exercise. We'll discuss why it's a useful program and some of the design issues that arise in developing it.

A graphical circuit-building program in which the user creates new gates by choosing from a palette of standard gates, uses the mouse to wire gates together, and tests the circuits built, might be the best way of designing and building new circuits. However, a text-base interactive circuit builder is easier to design and implement. Many of the classes and ideas in a text-based program may transfer to a graphics-based program, so we'll view the text-based program as a useful prototype.




Figure 13.12 Building A half-adder circuit.

We'll use the interactive circuit building program to build a *half-adder*, a circuit for adding two one-digit binary numbers diagrammed in Fig. 13.12.¹⁰ We'll use the half-adder to build a *full-adder*, a circuit that basically adds three one-bit numbers, though we'll view the inputs as two numbers and a carry from a previous addition, diagrammed in Fig. 13.13. Full-adders can be wired together easily to form an *n-bit ripple-carry adder* for adding two *n*-bit binary numbers that we'll explore in an exercise.

¹⁰A binary digit is usually called a **bit**, which is almost an acronym for **binary digit**.

Binary, or base 2, numbers are added just like base 10 numbers, but since the only values of a binary digit (or bit) are zero and one, we get Table 13.1 as a description of the half-adder.

Table 13.1 Adding two one-bit numbers

A	B	S	C
0	0	0	0
0	1	1	0
1	0	1	0
1	1	0	1

The output labeled *S* in Fig. 13.12 and Table 13.1 is the sum of two bits. The output labeled *C* is the carry. Since we have $1 + 1 = 10$ in base 2, the last line of the table shows the sum is zero and the carry is one, where the sum is the rightmost or least-significant digit. Similarly in Fig. 13.13 the sum and carry represent adding the three input bits. A table for the full-adder is shown in the output of *circuitbuilder.cpp*.

Before looking at a run of the program we'll outline a list of requirements for an interactive circuit builder. The program doesn't meet all these requirements in the run shown, but you can add features as explored in chapter exercises.

1. The program should allow the user to choose standard gates for building circuits, but the list of gates should grow to include circuits built during the program. In other words, the program may start with only three gates (and, or, inverter), but any circuits built with the program become gates used in building other circuits.
2. The program should be simple to use, commands should correspond to user expectations. First-time users should be able to use the program without much help, but experienced users should be able to use their experience to build circuits quickly.
3. The program should be able to load circuits built by the program. This means the user should be able to save newly constructed circuits and load these circuits in a later run.
4. Connecting gates and wires should not require an in-depth knowledge of the `Gate` and `Wire` classes we've studied. Circuit designers shouldn't need to be experts in object-oriented programming and design to use the program.
5. The program should be flexible enough to adapt to new requirements we expect to receive from users once the program has been reviewed and tested. For example, users make mistakes in building circuits, it would be nice to support *undo* features to change gates and connections already created.

In the run below there is no facility for saving and loading circuits and there is no *undo* command, but attempts are made to meet the other requirements. The program shows an initial collection of the three standard gates available for creating circuits. In the run, the user builds the half-adder diagrammed in Fig. 13.12 by creating gates, printing the composite made from the gates in order to find the name of each wire, then connecting




Figure 13.13 Building a full-adder from half-adders and an or-gate.

the gates and specifying inputs and outputs for the composite gate constructed. After the new circuit is finished, the user types *stop*, the circuit is tested, and the new circuit is added to the list of available gates.

The full-adder diagrammed in Fig. 13.13 is built next using the same process:

- gates are added to the composite: two half-adders and an or-gate
- the composite is printed (the half-adders show up as composites), the wires between the gates are connected, and the inputs and outputs are specified.
- the circuit is finished, tested, and added to the toolkit of available circuits.

I'll point out some of the features shown in the program that try to meet the initial requirements.

Designing for Both Novice and Expert Users. The command *add* which adds gates to the composite being constructed comes in three forms, each illustrated in the run.

- *add and*, the user specifies the gate to add
- *add*, the user doesn't specify a gate, and is prompted for one
- the user presses enter/return when prompted for a gate and list of available gates is printed (see the end of the output)

Minimal Knowledge of Gate and Wire Classes. Every gate displayed is shown with inputs and outputs. The input and output wires are numbered, and users connect wires by using a wire's number rather than typing `and(1)->InWire(0)` which might be used in a program, but shouldn't be demanded from a user.¹¹

¹¹Implementation aside: wire numbers can be used to find wires only because a `WireFactory` supports wire lookup by number.

OUTPUT

```
prompt> circuitbuilder
0.      and
1.      or
2.      inverter
command: add and
command: add and
command: add or
command: add inverter
command: show
current circuit
composite: 4 gates, 0 in wires, 0 out wires
all-in
all-out
      and (1)
      in (wire 8) (wire 9)      out (wire 10)
----
      and (2)
      in (wire 11) (wire 12)    out (wire 13)
----
      or (1)
      in (wire 14) (wire 15)    out (wire 16)
----
      inv (1)
      in (wire 17)      out (wire 18)
-----
connections: none

command: connect 10 17
command: connect 18 12
command: connect 16 11
command: connect 14 8
command: connect 9 15
command: in 14
command: in 9
command: out 13
command: out 10
command: test
output continued →
```

```

O U T P U T

testing composite: 4 gates, 2 in wires, 2 out wires
-----
0 0      :      0 0
1 0      :      1 0
0 1      :      1 0
1 1      :      0 1
-----
command: stop
name for circuit: half
command: add half
command: add half
command: add or
command: show
current circuit
composite: 3 gates, 0 in wires, 0 out wires
all-in
all-out
           composite: 4 gates, 2 in wires, 2 out wires
all-in   (wire 25) (wire 20)
all-out  (wire 24) (wire 21)
output elided/removed
-----
           composite: 4 gates, 2 in wires, 2 out wires
all-in   (wire 36) (wire 31)
all-out  (wire 35) (wire 32)
output elided/removed
-----
           or (4)
           in (wire 41) (wire 42)      out (wire 43)
connections: none
command: connect 24 31
command: connect 21 42
command: connect 32 41
command: in 36
command: in 25
command: in 20
command: out 35
command: out 43
command: test

output continued →

```

```

                                O U T P U T
testing composite: 3 gates, 3 in wires, 2 out wires
-----
0 0 0   :   0 0
1 0 0   :   1 0
0 1 0   :   1 0
1 1 0   :   0 1
0 0 1   :   1 0
1 0 1   :   0 1
0 1 1   :   0 1
1 1 1   :   1 1
-----
command: stop
name for circuit: full
command: add
gate name:
0.      and
1.      or
2.      inverter
3.      half
4.      full

```

13.3.9 simpleMap: Mapping Names to Gates



In my prototype for the interactive builder program I used a structure called a *map*. I'll show the simple version I used in this prototype which is good enough for the prototype, easy to understand, and not fully functional. You don't need the class `SimpleMap` to write the interactive circuit builder, but you'll need to implement something just like it (see *simplemap.h*, Prog. G.17.)

A `SimpleMap` is templated on two classes, one class is the **value** stored in the map, the other is the **key** used to look up a value. In *simplemapdemo.cpp*, `Gate *` values are stored in the map and `int` values are keys to retrieve gate pointers. The same kind of map is used in *circuitbuilder.cpp*, but strings are used to lookup a gate rather than integers. Users are more comfortable typing `add` and than typing `add 0`, where 0 is the index of the and-gate stored in a map.

Program 13.14 `simplemapdemo.cpp`

```

#include <iostream>
using namespace std;

```

```

#include "simplemap.h"
#include "gates.h"

int main()
{
    SimpleMap<int, Gate *> gatemap;
    gatemap.insert(0, new AndGate("map-and-gate"));
    gatemap.insert(1, new OrGate("map-or-gate"));
    gatemap.insert(2, new Inverter("map-not-gate"));

    Gate * g = 0; // get g from map
    SimpleMapIterator<int, Gate*> git(gatemap);
    for(git.Init(); git.HasMore(); git.Next())
    {
        int index = git.Current();
        g = gatemap.getValue(index);
        cout << index << "\t" << *g << "\t" << *(g->clone()) << endl;
    }
    return 0;
}

```

simplemapdemo.cpp

The program shows how a map works as a gate-toolkit. The program retrieves a gate and makes a copy of it using `clone`. The copy could be added to a composite being constructed by the user. When a new circuit is finished it can be easily added to the toolkit using the method `SimpleMap::insert`.

O U T P U T

```

prompt> simplemapdemo
0      and (0) map-and-gate      and (1) map-and-gate
1      or (0) map-or-gate       or (1) map-or-gate
2      inv (0) map-not-gate     inv (1) map-not-gate

```

13.4 Chapter Review

We discussed inheritance, a powerful technique used in object-oriented programming for re-using a common interface. We saw several examples of inheritance hierarchies in which superclasses specified an interface, and subclasses implemented the interface with different behavior, but using a common naming convention. Inheritance allows an object that's an instance of a subclass to be substituted for, or used-as-an instance of, the corresponding superclass. In this book inheritance always models an "is-a" relationship, which ensures that objects can be substituted for other objects up an inheritance hierarchy.

Topics covered include:

- Streams form an inheritance hierarchy. A function with an `istream` parameter can receive many kinds of streams as arguments including `cin`, `ifstream`, and

istringstream objects.

- Prototypes are first-attempts at designing and implementing a program or classes that allow the programmer and the client to get a better idea of where a project is headed.
- Inheritance in C++ requires superclass functions to be declared *virtual* so that subclasses can change or specialize behavior. We use *public* inheritance which models an is-a relationship. Virtual functions are also called polymorphic functions. (Other uses of inheritance are possible in C++, but we use inheritance only with virtual functions and only with public inheritance.)
- Virtual superclass functions are always virtual in subclasses, but the word `virtual` isn't required. It's a good idea to continue to identify functions as virtual even in subclasses, because a subclass may evolve into a superclass.
- Subclasses should call superclass constructors explicitly, otherwise an implicit call will be generated by the compiler.
- An inherited virtual function can be used directly, overridden completely, or overridden while still calling the inherited function using `Super::function` syntax.
- Data and functions declared as *protected* are accessible in subclasses, but not to client programs. Data and functions declared as *private* are not accessible to subclasses except using accessor and mutator functions that might be provided by the superclass. Nevertheless, a subclass contains the private data, but the data isn't directly accessible.
- Abstract base classes contain one pure virtual function, a function identified with the ugly syntax of `= 0`. An abstract base class is an interface, it's not possible to define an object whose type is an abstract base class. It's very common, however, to define objects whose type is `ABC *` where `ABC` is an abstract base class. An abstract base/superclass pointer can reference any object derived from the superclass.
- Flexible software should be extendible, programming in the future tense is a good idea. Using abstract classes that can have some default function definitions, but should have little state, is part of good programming practice.
- Several design patterns were used in designing and implementing a `Gate/Wire` framework for modeling digital circuits. The patterns used include *Composite*, *Factory*, *Abstract Factory*, and *Observer*.
- Programs should be grown rather than built; refactoring classes and functions is part of growing good software.
- A class `SimpleMap` is a usable prototype of the map classes you'll study as you continue with computer science and programming. The map class facilitates the implementation of an interactive circuit building program.



13.5 Exercises

- 13.1** Design a hierarchy of math quiz questions that cover the operations of addition, subtraction, multiplication, and division. You might also consider questions involving

ratios, fractions, or other parts of basic mathematics. Each kind of question should have both easy and hard versions, i.e., addition might require carrying in the hard version. Keep the classes simple to make it possible to write a complete program; assume the user is in fourth or fifth grade.

Design and implement a quiz class that uses the questions you've just designed (and tested). The quiz should use different questions, and the questions should get more difficult if the user does well. If a user isn't doing well, the questions should get simpler. The quiz class should give a quiz to one student, not to two or more students at the same time. Ideally, the quiz class should record a student's scores in a file so that student progress can be tracked over several runs of the program.

13.2 Implement the class `MultipleChoice` shown in Fig. 13.1. You'll need to decide on some format for storing multiple choice questions in a file, and specify a file when a `MultipleChoice` question object is created. Incorporate the new question into *inheritquiz.cpp*, Prog. 13.2, or design a new quiz program that uses several different quiz questions.

13.3 We studied a templated class `LinkSet` designed and implemented in Sec. 12.3.6 (see Progs. 12.11 and 12.12, the interface and implementation, respectively.) New elements were added to the front of the linked list representing the set elements. Design a class like the untemplated version of the set class, `LinkStringSet`, that was developed first. The new class supports only the operations `Add` and `Size`. Call the class `WordList`, it can be used to track the unique words in a textfile as follows.

```
void ReadStream(WordStreamIterator& input, WordList * list)
// post: list contains one copy of each word in input
{
    string word;
    for(input.Init(); input.HasMore(); input.Next())
    {
        word = input.Current();
        ToLower(word);
        StripPunc(word);
        list->Add(word);
    }
    cout << list->Size() << " different words" << endl;
}
```

Make the function `Add` a pure virtual function and make the helper function `FindNode` from `LinkStringSet` virtual and protected rather than private. Then implement three subclasses each of which uses a different technique for maintaining the linked list (you may decide to use doubly linked lists which make the third subclass slightly simpler to implement.)

- A class `AddAtFront` that adds new words to the front of the linked list. This is similar to the class `LinkStringSet`.
- A class `AddAtBack` that adds new words to the end of the linked list (keep a pointer to the last node, or use a circularly linked list.)
- A class `SelfOrg` that adds new nodes at the back, but when a node is found using the virtual, protected `FindNode`, the node is moved closer to the front by one position. The idea is that words that occur many times move closer to

the front of the list so that they'll be found sooner.

Test each subclass using the function `ReadStream` shown above. Time the implementations on several text files. Try to provide reasons for the timings you observe. As a challenge, make two additions to the classes once they work. (1) Add an iterator class to access the elements. The iterator class will need to be a friend of the superclass `WordList`, but friendship is *not* inherited by subclasses. You'll need to be careful in designing the hierarchy and iterator so the iterator works with any subclass. (2) Make the classes templated.

13.4 Program 12.4, *frogwalk3.cpp* in Sec. 12.1.6, shows how to attach an object that monitors two random walkers to each of the walkers. The class `Walker` is being observed by the class `WalkRecorder`, though we didn't use the term *Observer* when we discussed the example in Chap. 12.

Create an inheritance hierarchy for `WalkRecorder` objects that monitor a random walker in different ways. Walkers should accept any number of `WalkRecorders`, rather than just one, by storing a vector of pointers rather than a single pointer to a `WalkRecorder`. Implement at least two different recorders, but try to come up with other recorders that you think are interesting or useful.

- Implement a recorder that works like the original `WalkRecorder` in tracking every location of all the walkers it's recording.
- Implement an `ExtremeRecorder` class that tracks just the locations that are furthest left (lowest) and right (highest) reached by any walker being monitored by the `ExtremeRecorder`. Alternatively, have the recorder keep track of one pair of extremes per walker rather than one pair of extremes for all walkers (this is tricky.)
- Use the graphics package in Howto H and create a class that displays a walker as a moving square on a canvas. Each walker monitored should appear as a different color. Walkers can supply their own colors, or the recorder can associate a color with a walker (it could do this using a `SimpleMap` object or in several other ways.)



13.5 Design a hierarchy of walkers each of which behaves differently. The walkers should wander in two-dimensions, so a walker's location is given by a `Point` object (see *point.h*, Prog. G.10). The superclass for all walkers should be named `Walker`.

Design a `WalkerWorld` class that holds all the walkers. `WalkerWorld::Step` asks the world to ask each of its walkers to take one step, taking a step is a virtual function with different implementations by different `Walker` subclasses. You can consider implementing a hierarchy of `WalkerWorld` classes too, but at first the dimensions of the world in which the walkers roam should be fixed when the world is created. The lower-left corner of the world has location $(0, 0)$; the upper-right corner has location $(\text{maxX}, \text{maxY})$. In a world of size 50×100 the upper-right corner has coordinates $(49, 99)$.

Consider the following different behaviors for step-taking, but you should be imaginative in coming up with new behaviors. A walker should always start in the middle of the world.

- A random walker that steps left, right, up, and down with equal probability. A

walker at the edge of the world, e.g., whose location is $(0,x)$, can't move off the edge, but may have only three directions to move. A walker in the corner of the world has only two choices.

- A walker that walks immediately to the north edge of the world and then hugs the wall circling the world in a clockwise direction.
- A walker that wraps around the edge of the world, e.g., if it chooses to walk left/west from location $(0,y)$ its location becomes $(\max X,y)$.

You'll probably want to add at least one `WalkRecorder` class to monitor the walkers, a graphics class makes for enjoyable viewing.

- 13.6** Function objects were used to pass comparison functions encapsulated as objects to sorting functions, see Sec. 11.3 for details. It's possible to use inheritance rather than templates to enforce the common interface used by the comparison function objects described in Sec. 11.3. Show how the function header below can be used to sort using function objects, although the function is only templated on one parameter (contrast it to the declaration for `InsertSort` in *sortall.h*, Prog. G.14.)



```
template <class Type>
void InsertSort(tvector<Type> & a,
int size, const Comparer & comp);
```

You should show how to define an abstract `Comparer` class, and how to derive subclasses that are used to sort by different criteria.

- 13.7** The circuit constructed by the statements below is self-referential. Draw the circuit and trace the calls of `Gate::Act` through the or-gate, inverter, and probe. What happens if the circuit is programmed? What happens if the or-gate is changed to an and-gate?

```
Gate * org = new OrGate("trouble");
Gate * inv = new Inverter();
Probe * p = new Probe(inv->OutWire(0));

Connect(org->OutWire(0), inv->InWire(0));
Connect(inv->OutWire(0), org->InWire(1))
```

- 13.8** Implement a complete program for interactively building circuits. Invent a circuit-description language you can use to write circuits to files and read them back. You should try to use a `Factory` for creating the gates and circuits used in the program, but you'll need a factory to which you can add new circuits created while the program is running. Using a `SimpleMap` can make the factory implementation easier, but you'll need to think very carefully about how to design the program.
- 13.9** Implement a class `GateFactory` that encapsulates creation of the four standard gate classes: `AndGate`, `OrGate`, `Inverter`, `CompositeGate` as well as a class `XorGate`. The factory class is used like the `WireFactory` class, but for creating gates rather than wires, see the code on the next page.

For example, the code below creates a disabler-circuit, see Fig. 13.9.

```

GateFactory gf;
Gate * cg = gf.MakeComposite();
Gate * ig = gf.MakeInverter();
Gate * ag = gf.MakeAndGate();
// connect wires, add gates, input and output wires, to cg

```

This class enables gates to be created using a factory, but it doesn't force client programs to use the factory. Nor does it stop clients from creating hundreds of factories. The second concern can be addressed using a design pattern called *Singleton*. A singleton class allows only one object to be created. Clients can have multiple pointers to the object, but there's only one object. The class `Singleton` in `singleton.h` illustrates how to do this.

Program 13.15 `singleton.h`

```

#ifndef _SINGLETON_H
#define _SINGLETON_H

// demo code for a singleton implementation

class Singleton
{
public:
    static Singleton * GetInstance();
    // methods here for Singleton behavior
private:
    static Singleton * ourSingleton;
    Singleton(); // constructor
};

Singleton * Singleton::ourSingleton = 0;

Singleton * Singleton::GetInstance()
{
    if (ourSingleton == 0)
    {
        ourSingleton = new Singleton(); // ok to construct
    }
    return ourSingleton;
}

Singleton::Singleton()
{
    // nothing to construct in this simple example
}

#endif

```

`singleton.h`

Show by example how client code uses a singleton object. Assume there's a void method `Singleton::DoIt()` and write code to call it. Explain how client programs are prevented from creating `Singleton` objects and how the class limits itself to creating one object. Then modify either your `GateFactory` or the existing `WireFactory` class to be singletons.

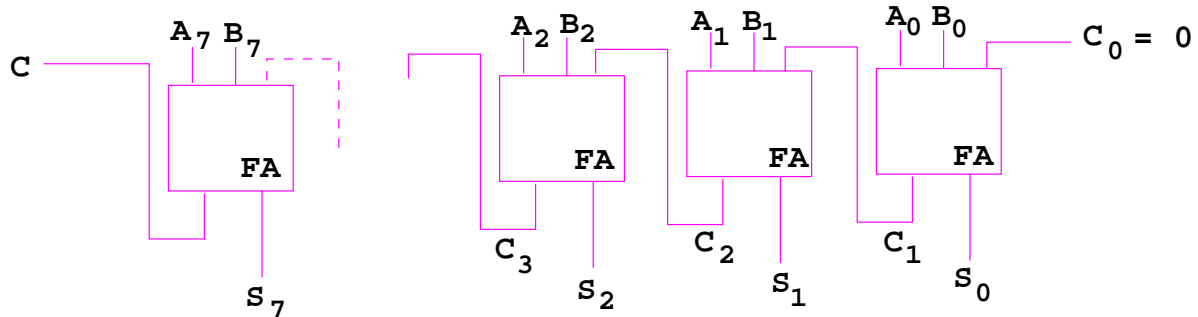


Figure 13.14 A ripple-carry adder for 8-bit numbers.

13.10 The circuit in Fig. 13.14 is an 8-bit ripple-carry adder, a concrete version of the more general n -bit ripple-carry adder. The circuit adds two 8-bit numbers represented by A and B , where $A = A_7A_6A_5A_4A_3A_2A_1A_0$, and A_0 is the least-significant bit. The largest 8-bit value is 1111111 which is 255_{10} (base 10). Each box labeled FA is a full-adder, see Fig. 13.13 for details. This ripple-adder is a 17-9-gate circuit, with 17 inputs: 8 bits for A , 8 bits for B , and the initial carry-in value, and 9 outputs: 8 bits for the sum and a final carry-out.

Write an English description for how the ripple-carry adder works. Note that the initial carry-in C_0 is set to zero. Other carries ripple through the circuit, hence the name. Then write a function `RippleAdder` to create and return a composite-gate representing an n -bit ripple-carry adder where n is a parameter to the function. Assume you have a function `FullAdder`. To test the function you'll need to implement the `FullAdder` function which in turn will require implementing a `HalfAdder` function.

13.11 In real circuits, electricity does not travel through a circuit instantaneously, but is delayed by the gates encountered. Different gates have different built-in delays, and the delays of the built-in gates affect circuits built up from these gates.

For example, we'll assume a delay of 3 time-units for an and-gate, 5 units for an or-gate, and 2 units for an inverter (you'll be able to change these values in the program you write.) Assume a disabler-circuit as diagrammed in Fig. 13.9 has the input to the and-gate from the outside on, the input to the inverter off, so that the output signal is on. If the inverter-input signal is set to true, the circuit's output will change to false 5 time-units later. There will be a 2-unit delay for the inverter followed by a 3-unit delay for the and-gate.

Develop a new class called `TimedGate` that acts like a gate, but delays acting for a set amount of time. This is a non-trivial design, so you'll need to think very carefully about how to incorporate delays into the circuit system. Assume you'll be using only `TimedGates`, not mixing them with regular gates. One way to start is shown in the following:.

```

class TimedGate : public Gate
{
public:
    // substitute me for g
    TimedGate(Gate * g, int delay);

    virtual int InCount() const
        {return myGate->InCount();}
    virtual int OutCount() const
        {return myGate->OutCount();}
    virtual Wire * InWire(int n) const
        {return myGate->InWire(n);}
    virtual Wire * OutWire(int n) const
        {return myGate->OutWire(n);}
    virtual string toString() const;
        // can use g's toString

    virtual Gate * clone()
    {
        return new TimedGate(myGate->clone(), myDelay);
    }
    virtual void Act(); // act with delay

protected:
    Gate * myGate;
    int myDelay;
};

```

This class can be used as-a gate. It forwards most requests directly to the gate it encapsulates as shown. The constructor and the `TimedGate::Act` function require careful thought.

A `TimedGate` object must remove the `Gate g` it encapsulates from the wires connected to `g`'s inputs. Then the `TimedGate` object substitutes itself for the the inputs. All this happens at construction.

In addition, you'll need to define some kind of structure that stores timed events so that they happen in the correct order. In my program I used a static `EventSimulator` object that all `TimedGates` can access. Events are put into the simulator, and arranged to occur in the proper order. Again, you'll need to think very carefully about how to do this.

13.12 The circuit in Fig. 13.15 is designed to control an elevator. It's a simple circuit designed to direct the elevator up or down, which are the circuit's outputs. The inputs are the current floor and the requested floor. The diagram shows a circuit for an elevator in a four-story building. The current floor is specified by the binary number C_1C_0 , so that 00 is the first floor¹², 01 is the second floor, 10 is the third floor, and 11 is the fourth floor. The digit C_1 is the most significant digit. Similarly, R_1R_0 is a binary

¹²This is a book on C++, so floors are numbered beginning with zero.

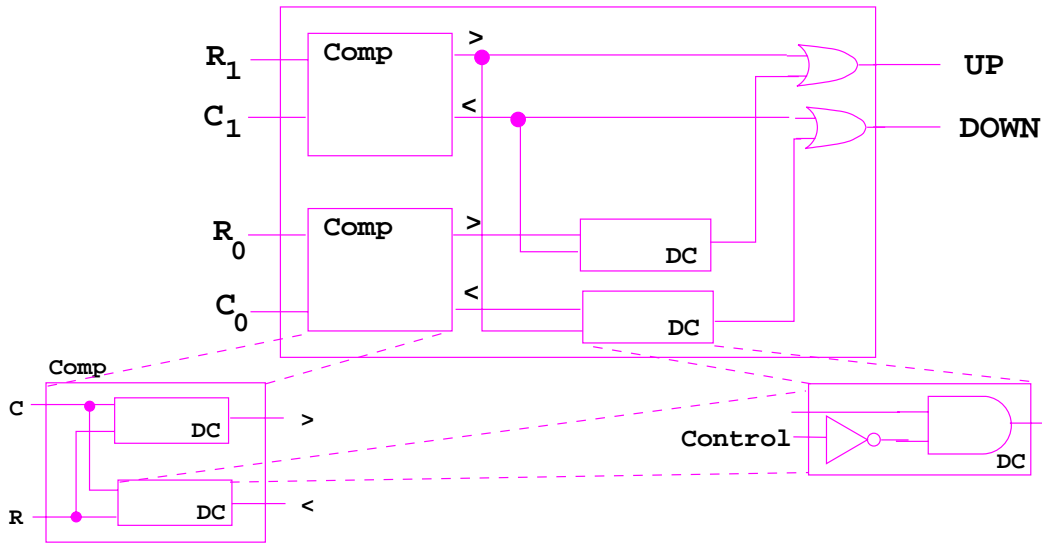


Figure 13.15 A circuit for choosing which direction an elevator travels. The inputs labeled C are the current floor (2 bits) and the R inputs are for the requesting floor.

representation of the requested floor.

The purpose of the circuit is to direct the elevator up when the requested floor is greater than the current floor, and down when the requested floor is less than the current floor. Write an English description of why the circuit works. Be attentive to the order of inputs to the comparator gates and see Fig 13.10 and Fig. 13.9 and the associated descriptions.

Write a truth table by hand for a 4-input 2-output circuit, or build the circuit with a program and have `GateTester::Test` print the truth table for the circuit. Try to generalize the circuit to a building with 2^N floors rather than four floors.

Bibliography

- [AA85] Donald J. Albers and G.L. Alexanderson. *Mathematical People*. Birkhäuser, 1985.
- [ACM87] ACM. *Turing Award Lectures: The First Twenty Years 1966-1985*. ACM Press, 1987.
- [AS96] Harold Abelson and Gerald Jay Sussman. *Structure and Interpretation of Computer Programs*. MIT Press, McGraw Hill Book Company, second edition 1996.
- [Asp90] William Aspray. *Computing Before Computers*. Iowa State University Press, 1990.
- [Aus98] Matthew H. Austern *Generic Programming and the STL* Addison-Wesley, 1998.
- [Ben86] Jon Bentley. *Programming Pearls*. Addison-Wesley, 1986.
- [Ben88] Jon Bentley. *More Programming Pearls*. Addison-Wesley, 1988.
- [Ble90] Guy E. Blelloch. *Vector Models for Data-Parallel Computing*. MIT Press, 1990.
- [Boo91] Grady Booch. *Object Oriented Design with Applications*. Benjamin Cummings, 1991.
- [Boo94] Grady Booch. *Object Oriented Design and Analysis with Applications*. Benjamin Cummings, second edition, 1994.
- [BRE71] I. Barrodale, F.D. Roberts, and B.L. Ehle. *Elementary Computer Applications in Science Engineering and Business*. John Wiley & Sons Inc., 1971.
- [Coo87] Doug Cooper. *Condensed Pascal*. W.W. Norton, 1987.
- [DR90] Nachum Dershowitz and Edward M. Reingold. Calendrical calculations. *Software-Practice and Experience*, 20(9):899–928, September 1990.
- [(ed91] Allen B. Tucker (ed.). *Computing Curricula 1991 Report of the ACM/IEEE-CS Joint Curriculum Task Force*. ACM Press, 1991.
- [EL94] Susan Epstein and Joanne Luciano, editors. *Grace Hopper Celebration of Women in Computing*. Computing Research Association, 1994. Hopper-Book@cra.org.

- [Emm93] Michele Emmer, editor. *The Visual Mind: Art and Mathematics*. MIT Press, 1993.
- [G95] Denise W. Güler. Pioneering women in computer science. *Communications of the ACM*, 38(1):45–54, January 1995.
- [Gar95] Simson Garfinkel. *PGP: Pretty Good Privacy*. O'Reilly & Associates, 1995.
- [GHJ95] Erich Gamma and Richard Helm and Ralph Johnson and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Programming*. Addison-Wesley, 1995.
- [Gol93] Herman H. Goldstine. *The Computer from Pascal to von Neumann*. Princeton University Press, 1993.
- [Har92] David Harel. *Algorithmics, The Spirit of Computing*. Addison-Wesley, second edition, 1992.
- [Hoa89] C.A.R. Hoare. *Essays in Computing Science*. Prentice-Hall, 1989. (editor) C.B. Jones.
- [Hod83] Andrew Hodges. *Alan Turing: The Enigma*. Simon & Schuster, 1983.
- [Hor92] John Horgan. Claude e. shannon. *IEEE Spectrum*, April 1992.
- [JW89] William Strunk Jr. and E.B. White. *The Elements of Style*. MacMillan Publishing Co., third edition, 1989.
- [Knu97] Donald E. Knuth. *The Art of Computer Programming*, volume 1 Fundamental Algorithms. Addison-Wesley, third edition, 1997.
- [Knu98a] Donald E. Knuth. *The Art of Computer Programming*, volume 2 Seminumerical Algorithms. Addison-Wesley, third edition, 1998.
- [Knu98b] Donald E. Knuth. *The Art of Computer Programming*, volume 3 Sorting and Searching. Addison-Wesley, third edition 1998.
- [KR78] Brian W. Kernighan and Dennis Ritchie. *The C Programming Language*. Prentice-Hall, 1978.
- [KR96] Samuel N. Kamin and Edward M. Reingold. *Programming with class: A C++ Introduction to Computer Science*. McGraw-Hill, 1996.
- [Mac92] Norman Macrae. *John von Neumann*. Pantheon Books, 1992.
- [McC79] Pamela McCorduck. *Machines Who Think*. W.H. Freeman and Company, 1979.
- [McC93] Steve McConnell. *Code Complete*. Microsoft Press, 1993.
- [Neu95] Peter G. Neumann. *Computer Related Risks*. Addison Wesley, 1995.

- [Pat96] Richard E. Pattis. *Get A-Life: Advice for the Beginning C++ Object-Oriented Programmer*. Turing TarPit Press, 1996.
- [Per87] Alan Perlis. The synthesis of algorithmic systems. In *ACM Turing Award Lectures: The First Twenty Years*. ACM Press, 1987.
- [PL90] Przemyslaw Prusinkiewicz and Aristid Lindenmayer. *The Algorithmic Beauty of Plants*. Springer-Verlag, 1990.
- [RDC93] Edward M. Reingold, Nachum Dershowitz, and Stewart M. Clamen. Calendrical calculations, ii: Three historical calendars. *Software-Practice and Experience*, 23(4):383–404, April 1993.
- [Rie96] Arthur Riel. *Object-Oriented Design Heuristics*. Addison-Wesley, 1996.
- [Rob94] Eric S. Roberts. Loop exits and structured programming: Reopening the debate, 1994. submitted for publication.
- [Sla87] Robert Slater. *Portraits in Silicon*. MIT Press, 1987.
- [Str87] Bjarne Stroustrup. *The C++ Programming Language*. Addison Wesley, 1987.
- [Str94] Bjarne Stroustrup. *The Design and Evolution of C++*. Addison-Wesley, 1994.
- [Str97] Bjarne Stroustrup. *The C++ Programming Language*. Addison Wesley, third edition, 1997.
- [Mey96] Scott Meyers. *More Effective C++*. Addison-Wesley, 1996.
- [Wei94] Mark Allen Weiss. *Data Structures and Algorithm Analysis in C++*. Benjamin Cummings, 1994.
- [Wil56] M.V. Wilkes. *Automatic Digital Computers*. John Wiley & Sons, Inc., 1956.
- [Wil87] Maurice V. Wilkes. Computers then and now. In *ACM Turing Award Lectures: The First Twenty Years*, pages 197–205. ACM Press, 1987.
- [Wil95] Maurice V. Wilkes. *Computing Perspectives*. Morgan Kaufmann, 1995.
- [Wir87] Niklaus Wirth. From programming language design to compiler construction. In *ACM Turing Award Lectures: The First Twenty Years*. ACM Press, 1987.