

# A Discrete EVent system Simulator

Jim Nutaro

September 7, 2011



# Contents

<b>1</b>	<b>About this manual</b>	<b>1</b>
<b>2</b>	<b>Building and Installing</b>	<b>3</b>
<b>3</b>	<b>Modeling and simulation with Adevs</b>	<b>5</b>
<b>4</b>	<b>Atomic Models</b>	<b>19</b>
<b>5</b>	<b>Network Models</b>	<b>31</b>
5.1	Parts of a Network Model . . . . .	31
5.2	Simulating a Network Model . . . . .	34
5.3	A complete example of a network model . . . . .	35
5.4	Digraph Models . . . . .	37
5.5	Cell Space Models . . . . .	44
<b>6</b>	<b>Variable Structure Models</b>	<b>49</b>
6.1	Building and Simulating Variable Structure Models . . . . .	49
6.2	A Variable Structure Example . . . . .	51
<b>7</b>	<b>Continuous Models</b>	<b>59</b>
7.1	Differential equation modeling with the <code>ode_system</code> class . . . . .	59
7.2	Modeling hybrid systems with adevs and OpenModelica . . . . .	63
<b>8</b>	<b>The Simulator Class</b>	<b>71</b>
<b>9</b>	<b>Simulation on multi-core computers</b>	<b>73</b>
9.1	Limits of the parallel simulator . . . . .	75
9.2	Modifying your models to exploit lookahead . . . . .	75
9.3	Partitioning your model . . . . .	78
9.4	Interaction between partitioning and lookahead . . . . .	79
9.5	A complete example . . . . .	80
9.6	Memory management of input and output across thread boundaries . . . . .	84
9.7	Summary of main points regarding repeatable outcomes and performance . . . . .	86
<b>10</b>	<b>Models with Many Input/Output Types</b>	<b>87</b>
<b>11</b>	<b>Random Numbers</b>	<b>91</b>
<b>12</b>	<b>Interpolation</b>	<b>93</b>



# Chapter 1

## About this manual

The purpose of this manual is to get you working with Adevs as quickly as possible. To that end, it documents the simulation engine's major features with a focus on how they are used. The table of contents summarizes those aspects of the simulator that are described here.

Among the features omitted from this manual, the Java language bindings for the Adevs simulator are both useful enough and mature enough to warrant a brief mention. Build instructions for the Java bindings are given in the “Build and Install” section. How to use these bindings will be more or less obvious once you have perused the C++ interface: the interfaces for building models and running simulations are essentially the same.

Use of the Java bindings entails several trade offs. There are three chief disadvantages. First, you pay a (usually unnoticeable) cost in execution time for the extra work that Adevs must do to manage memory associated with input and output objects and models orphaned during structure changes. Second, the facilities for combined simulation of discrete event and continuous models have not been implemented for Java. Third, this is not a ‘pure Java’ simulation engine; it uses a great deal of native code to do its work (though this is invisible to the programmer).

There are at least two very positive aspects of the Java bindings. First is that it omits the need for managing memory. The Java garbage collector (plus some extra work by the simulation engine) takes care of this for you. Second, you have access to all of the very nice features and standard libraries of the Java programming language.

Other topics not included in this manual are theory (why was the simulator built this way?) and experimental features of the simulation engine. Among the latter are support for simulation of hybrid differential-algebraic systems and conservative, parallel simulation using multi-core processors. If you are interested in any of these subjects, I offer the following (greatly abridged) list of books and articles:

1. A. M. Uhrmacher. Dynamic structures in modeling and simulation: a reflective approach, ACM Transactions on Modeling and Computer Simulation, Vol. 11, No. 2 , pp. 206-232. April 2001. The approach by Adevs to modeling and simulation of dynamic structure systems is described in this paper.
2. Bernard P. Zeigler, Tag Gon Kim and Herbert Praehofer. Theory of Modeling and Simulation, Second Edition. Academic Press. 2000. The Discrete Event System Specification (DEVS) is developed in this book from its roots in abstract systems theory.
3. James J. Nutaro. Building Software for Simulation: Theory and Algorithms, with Applications in C++. Wiley. 2010. This book presents the Discrete Event Systems Specification along side code for the (slightly abridged) Adevs simulator and examples of its use. This book describes the conservative, parallel simulator, documentation for which is not included in this manual. It also discusses the construction of new ODE and event finding modules for Adevs.

Question and comments about this software can be sent to its maintainer, Jim Nutaro, at [nutarojj@ornl.gov](mailto:nutarojj@ornl.gov).



# Chapter 2

## Building and Installing

The Adevs package is organized into the following directory structure:

```
adevs-x.y.z
+-->docs
+-->examples
+-->include
+-->src
+-->test
+-->util
```

The Adevs package consists almost entirely of template classes, and these are implemented entirely within the header files located in the *include* directory. The only exceptions to this are the random number generators, the Java language bindings, and some rarely used aspects of the parallel simulation engine. Therefore, if you do not want to use these features, it is sufficient for your program to simply include *adevs.h* and, of course, to make sure that your compiler can find the *include* directory that *adevs.h* is in.

If you do want to use features not implemented in the header files, then a static library that containing these additional features needs to be built and, if desired, you will need to build the Java libraries. To do this, start by entering the *src* directory. To build the Adevs library, which contains the random number generator and some features used by the parallel simulator, execute the command 'make' (if you are using a \*nix type system) or 'build' if you are using Windows. For Windows, the batch file will create a static link library called *adevs.lib*. For \*nix type systems, the makefile will create a static link library called *libadevs.a*.

If you are building on a Windows system, make sure that you execute the batch file from the Visual C++ command prompt. This will ensure the batch file can find the compiler, linker, etc. For \*nix type systems, make sure you have installed a recent version of the GNU C++ compiler and make. You may need to edit the makefile (i.e., the file *Makefile*) to set compiler flags, etc. though the defaults should work in most cases.

To build the Java language bindings, you need to have the Sun (now Oracle?) JDK installed or something that is compatible with it. On a Windows system, from the *src* directory enter the *adevs\_jni* directory and then execute the command 'build'. This will create three files: *adevs.jar*, *java\_adevs.dll*, and *java\_adevs.lib*. To build and run your Java programs, you will need to put *adevs.jar* into your classpath and *java\_adevs.lib* into your java.library.path (or make sure it is in your regular PATH for finding executables and dynamic link libraries).

On a \*nix type system, stay in the *src* directory and execute the command 'make java\_adevs'. This will create two files: *adevs.jar* and *libjava\_adevs.so*. As before, you need to put *adevs.jar* into your classpath and *libjava\_adevs.so* into your java.library.path or make sure it is in your regular LD\_LIBRARY\_PATH (or whatever you use) for locating dynamic link libraries.

If you want to run the test suite (which is in no way required to use the software), then first you need to build the library file and install Tcl (the test scripts need Tcl to run; if you can run 'tclsh' then you already have a working copy of Tcl). If you want to test the Java bindings, you will need to build those

too. There are three sets of tests that can be executed: one for the serial simulation engine, one for the parallel simulation engine, and one for the Java language bindings. To run the basic set of tests for the serial simulation engine, use 'make check\_cpp'. To run those tests and the ones for the serial engines use 'make check\_par'. To run the Java test cases, use 'make java\_test'. You can run all of these test by simply executing 'make'. The makefile will abort as soon as any test fails or (more likely) run to completion and thereby indicate that all of the tests were passed.

The test cases can be a bit of a bear to run on a Windows computer. If you need to edit compiler settings, executable directives, etc. to make it work, then modify the file *make.common*. For \*nix type systems using the GNU tools the test cases should work out of the box. Otherwise, edit *make.common* to fix things for your development environment.



## Chapter 3

# Modeling and simulation with Adevs

Adevs is a simulator for models described in terms of the Discrete Event System Specification (DEVS)<sup>1</sup> The key feature of models described in DEVS (and implemented with Adevs) is that their dynamic behavior is defined by events. An event is any kind of change that is significant within the context of the model being developed.

Modeling of discrete event systems is most easily introduced with an example. Suppose that we want to model the checkout line at a convenience store. There is a single clerk who serves customers in a first come-first serve fashion. The time required for the clerk to ring up each customer's bill depends on the number of items purchased. We are interested in determining the average and maximum amount of time that customers spend waiting in the checkout line.




Figure 3.1: Customers waiting in line at BusyMart.

To simulate this system, we need an object to represent each customer in the line. A **Customer** class is created for this purpose. Each customer object has three attributes. One attribute is the time needed to ring up the customer's bill. Because we want to know how long a customer has been waiting in line, we also include two attributes that record when the customer entered the queue and when the customer left the queue. The difference of these times is the amount of time the customer spent waiting in line. Here is the customer class, from which we will create customer objects as needed. The class is coded in a single header file *Customer.h*.

```
#include "adevs.h"
/// A Busy-Mart customer.
struct Customer
{
    /// Time needed for the clerk to process the customer
    double twait;
```

---

<sup>1</sup>A comprehensive introduction to the Discrete Event System Specification can be found in "Theory of Modeling and Simulation, 2nd Edition" by Bernard Zeigler *et. al.*, published by Academic Press in 2000.

```

    /// Time that the customer entered and left the queue
    double tenter, tleave;
};
/// Create an abbreviation for the Clerk's input/output type.
typedef adevs::PortValue<Customer*> IO_Type;

```

Customers are served (processed) by the clerk, which is our first example of an atomic model. The clerk has a line of people waiting at her counter. When a customer is ready to make a purchase, that customer enters the end of the line. If the clerk is not busy and the line is not empty, then the clerk rings up the bill of the customer that is first in line. That customer then leaves the line and the clerk processes the next customer or, if the line is empty, the clerk sits idly at her counter.

The DEVS model of the clerk is described in a particular way. First, we specify the type of object that the clerk consumes and produces. For this model, we use **PortValue** objects. The **PortValue** class is part of the **Digraph** class, which will be introduced later. The **PortValue** class describes a port-value pair. In this case, **Customer** objects are the value and they arrive at the clerk via an “arrive” port; the customers depart via a “depart” port. Second, we specify the state variables that describe the clerk. In this case, we need to keep track of which customers are in line. We use a **list** from the C++ Standard Template Library for this purpose.

To complete the model of the clerk, we implement four methods that model her behavior. First, let's construct the header file for the clerk. Then we can proceed to fill in the details.

```

#include "adevs.h"
#include "Customer.h"
#include <list>
/**
 * The Clerk class is derived from the adevs Atomic class.
 * The Clerk's input/output type is specified using the template
 * parameter of the base class.
 */
class Clerk: public adevs::Atomic<IO_Type>
{
public:
    /// Constructor.
    Clerk();
    /// Internal transition function.
    void delta_int();
    /// External transition function.
    void delta_ext(double e, const adevs::Bag<IO_Type>& xb);
    /// Confluent transition function.
    void delta_conf(const adevs::Bag<IO_Type>& xb);
    /// Output function.
    void output_func(adevs::Bag<IO_Type>& yb);
    /// Time advance function.
    double ta();
    /// Output value garbage collection.
    void gc_output(adevs::Bag<IO_Type>& g);
    /// Destructor.
    ~Clerk();
    /// Model input port.
    static const int arrive;
    /// Model output port.
    static const int depart;
private:

```

```

    /// The clerk's clock
    double t;
    /// List of waiting customers.
    std::list<Customer*> line;
    /// Time spent so far on the customer at the front of the line
    double t_spent;
};

```

This header file is a template for almost any atomic model that we want to create. The **Clerk** class is derived from the Adevs **Atomic** class. The **Clerk** implements six virtual methods that it derives from **Atomic**. These are the state transition functions `delta_int`, `delta_ext`, and `delta_conf`; the output function `output`; the time advance function `ta`, and the garbage collection method `gc`. The **Clerk** also has a set of static, constant port variables that correspond to the **Clerk**'s input (customer arrival) and output (customer departure) ports.

The constructor for the **Clerk** class invokes the constructor of its **Atomic** base class. The template argument of the base class defines the types of objects that the clerk uses for input and output. The **Clerk** state variables are defined as private class attributes: these are the list of customers (`line`), the clerk's clock (`t`), and the time spent so far on the first customer in line (`t_spent`).

The ports `arrive` and `depart` are assigned integer values that are unique within the scope of the **Clerk** class. Typically, the ports for a model are numbered in a way that corresponds with the order in which they are listed; for example,

```

// Assign locally unique identifiers to the ports
const int Clerk::arrive = 0;
const int Clerk::depart = 1;

```

The **Clerk** constructor places the **Clerk** into its initial state. For our experiment, this state is an empty line and the **Clerk**'s clock is initialized to zero.

```

Clerk::Clerk():
Atomic<IO_Type>(), // Initialize the parent Atomic model
t(0.0), // Set the clock to zero
t_spent(0.0) // No time spent on a customer so far
{
}

```

Because the clerk has an empty line at first, the only interesting thing that can happen is for a customer arrive. Arriving customers appear on the clerk's "arrive" input port. The arrival of a customer causes the clerk's external transition method to be activated. The arguments to the method are the time that has elapsed since the clerk last changed state and a bag of **PortValue** objects.

The external transition method updates the clerk's clock by adding to it the elapsed time. The time spent working on the current customer's order is updated by adding the elapsed time to the time spent so far. After updating these values, the input events are processed. Each **PortValue** object has two attributes. The first is the port; it contains the number of the port that the event arrived on and is equal to "arrive" in this case. The second is the **Customer** that arrived. The clerk records the time of arrival for the new customer and places him at the back of the line.

```

void Clerk::delta_ext(double e, const Bag<IO_Type>& xb)
{
    // Print a notice of the external transition
    cout << "Clerk: Computed the external transition function at t = " << t+e << endl;
    // Update the clock
    t += e;
    // Update the time spent on the customer at the front of the line

```

```

    if (!line.empty())
    {
        t_spent += e;
    }
    // Add the new customers to the back of the line.
    Bag<IO_Type>::const_iterator i = xb.begin();
    for (; i != xb.end(); i++)
    {
        // Copy the incoming Customer and place it at the back of the line.
        line.push_back(new Customer(*((*i).value)));
        // Record the time at which the customer entered the line.
        line.back()->tenter = t;
    }
    // Summarize the model state
    cout << "Clerk: There are " << line.size() << " customers waiting." << endl;
    cout << "Clerk: The next customer will leave at t = " << t+ta() << "." << endl;
}

```

The time advance function describes the amount of time that will elapse before the clerk's next internal (or self) event, barring an input that arrives in the interim. In this case, the time advance is the time remaining for the clerk to process the current customer. If there are no customers in line, then the clerk will not do anything until one arrives, and so the time advance function returns infinity (in Adevs represented by DBL\_MAX). Otherwise, the clerk's next action is when the first customer in line has been rung up, and so the time advance is the difference of the **Customer's** twait and the clerk's t\_spent.

```

double Clerk::ta()
{
    // If the list is empty, then next event is at inf
    if (line.empty()) return DBL_MAX;
    // Otherwise, return the time remaining to process the current customer
    return line.front()->twait-t_spent;
}

```

Two things happen when the clerk finishes ringing up a customer. First, the clerk sends that customer on his way. This is accomplished when the clerk's output\_func is called. When this happens, the **Clerk** places the departing customer on its "depart" output port by creating a **PortValue** object and putting it into the bag yb of output objects. The clerk's output function is shown below.

```

void Clerk::output_func(Bag<IO_Type>& yb)
{
    // Get the departing customer
    Customer* leaving = line.front();
    // Set the departure time
    leaving->tleave = t + ta();
    // Eject the customer
    IO_Type y(depart,leaving);
    yb.insert(y);
    // Print a notice of the departure
    cout << "Clerk: Computed the output function at t = " << t+ta() << endl;
    cout << "Clerk: A customer just departed!" << endl;
}

```

Second, the clerk begins to process the next customer in the line. If, indeed, there is another customer waiting in line, the clerk begins ringing that customer up just as before. Otherwise, the clerk becomes idle.

Enter checkout line	Time to process order
1	1
2	4
3	4
5	2
7	10
8	20
10	2
11	1

Table 3.1: Customer arrival times and times needed to process the customers orders.

This is accomplished when the **Clerk**'s internal transition method is called. The **Clerk**'s internal transition method changes the state of the **Clerk** by removing the departed customer from the line. The clerk's internal transition function is shown below.

```
void Clerk::delta_int()
{
    // Print a notice of the internal transition
    cout << "Clerk: Computed the internal transition function at t = " << t+ta() << endl;
    // Update the clock
    t += ta();
    // Reset the spent time
    t_spent = 0.0;
    // Remove the departing customer from the front of the line.
    line.pop_front();
    // Summarize the model state
    cout << "Clerk: There are " << line.size() << " customers waiting." << endl;
    cout << "Clerk: The next customer will leave at t = " << t+ta() << "." << endl;
}
```

At this point we have almost completely defined the behavior of the clerk; only one thing remains to be done. Suppose that a customer arrives at the clerk's line at the same time that the clerk has finished ringing up a customer. In this case we have a conflict because the internal transition function and external transition function must both be activated to handle these two events (i.e., the simultaneous arriving customer and departing customer). These types of conflicts are resolved by the confluent transition function.

For the clerk model, the confluent transition function is computed using the internal transition function first (to remove the departed customer from the list) followed by the external transition function (to add new customers to the end of the list and begin ringing up the first customer). Below is the implementation of the clerk's confluent transition function.

```
void Clerk::delta_conf(const Bag<IO_Type>& xb)
{
    delta_int();
    delta_ext(0.0,xb);
}
```

To see how this model behaves, suppose customers arrive according to the schedule shown in the table below. In this example, the first customer appears on the clerk's "arrive" port at time 1, the next customer appears on the "arrive" port at time 2, and so on. The print statements in the **Clerk**'s internal, external, and output functions let us watch the evolution of the clerk's line. Here is the output trace produced by the above sequence of inputs.

Clerk: Computed the external transition function at  $t = 1$   
 Clerk: There are 1 customers waiting.  
 Clerk: The next customer will leave at  $t = 2$ .  
 Clerk: Computed the output function at  $t = 2$   
 Clerk: A customer just departed!  
 Clerk: Computed the internal transition function at  $t = 2$   
 Clerk: There are 0 customers waiting.  
 Clerk: The next customer will leave at  $t = 1.79769e+308$ .  
 Clerk: Computed the external transition function at  $t = 2$   
 Clerk: There are 1 customers waiting.  
 Clerk: The next customer will leave at  $t = 6$ .  
 Clerk: Computed the external transition function at  $t = 3$   
 Clerk: There are 2 customers waiting.  
 Clerk: The next customer will leave at  $t = 6$ .  
 Clerk: Computed the external transition function at  $t = 5$   
 Clerk: There are 3 customers waiting.  
 Clerk: The next customer will leave at  $t = 6$ .  
 Clerk: Computed the output function at  $t = 6$   
 Clerk: A customer just departed!  
 Clerk: Computed the internal transition function at  $t = 6$   
 Clerk: There are 2 customers waiting.  
 Clerk: The next customer will leave at  $t = 10$ .  
 Clerk: Computed the external transition function at  $t = 7$   
 Clerk: There are 3 customers waiting.  
 Clerk: The next customer will leave at  $t = 10$ .  
 Clerk: Computed the external transition function at  $t = 8$   
 Clerk: There are 4 customers waiting.  
 Clerk: The next customer will leave at  $t = 10$ .  
 Clerk: Computed the output function at  $t = 10$   
 Clerk: A customer just departed!  
 Clerk: Computed the internal transition function at  $t = 10$   
 Clerk: There are 3 customers waiting.  
 Clerk: The next customer will leave at  $t = 12$ .  
 Clerk: Computed the external transition function at  $t = 10$   
 Clerk: There are 4 customers waiting.  
 Clerk: The next customer will leave at  $t = 12$ .  
 Clerk: Computed the external transition function at  $t = 11$   
 Clerk: There are 5 customers waiting.  
 Clerk: The next customer will leave at  $t = 12$ .  
 Clerk: Computed the output function at  $t = 12$   
 Clerk: A customer just departed!  
 Clerk: Computed the internal transition function at  $t = 12$   
 Clerk: There are 4 customers waiting.  
 Clerk: The next customer will leave at  $t = 22$ .  
 Clerk: Computed the output function at  $t = 22$   
 Clerk: A customer just departed!  
 Clerk: Computed the internal transition function at  $t = 22$   
 Clerk: There are 3 customers waiting.  
 Clerk: The next customer will leave at  $t = 42$ .  
 Clerk: Computed the output function at  $t = 42$   
 Clerk: A customer just departed!  
 Clerk: Computed the internal transition function at  $t = 42$

```

Clerk: There are 2 customers waiting.
Clerk: The next customer will leave at t = 44.
Clerk: Computed the output function at t = 44
Clerk: A customer just departed!
Clerk: Computed the internal transition function at t = 44
Clerk: There are 1 customers waiting.
Clerk: The next customer will leave at t = 45.
Clerk: Computed the output function at t = 45
Clerk: A customer just departed!
Clerk: Computed the internal transition function at t = 45
Clerk: There are 0 customers waiting.
Clerk: The next customer will leave at t = 1.79769e+308.

```

The basic simulation algorithm is illustrated by this example. Notice that the external transition function is always activated when an input (in this case, a customer) arrives on an input port. This is because the external transition function describes the response of the model to input events.

The internal transition function is always activated when the simulation clock has reached the model's time of next event. The internal transition function describes the autonomous behavior of the model (i.e., how the model responds to events that it has scheduled for itself). Internal transitions are scheduled with the time advance function.

A call to the internal transition function is always immediately preceded by a call to the output function. Consequently, a model can only produce output by scheduling events for itself. The value of the output is computed using the model's current state.

To complete our simulation of the convenience store, we need two other **Atomic** models. The first model creates customers for the **Clerk** to serve. The rate at which customers arrive could be modeled using a random variable or it with a table such as the one used in the example above. In either case, we hope that the model of the customer arrival process accurately reflects what happens in a typical day at the convenience store. Data used in the table for this example could come directly from observing customers at the store, or it might be produced by a statistical model in another tool (e.g., a spreadsheet program).

We will create an **Atomic** model called a **Generator** to create customers. This model is driven by a table whose format is that used in the previous example. The input file contains a line for each customer. Each such line has the customer arrival time first, followed by the customer service time. The **Generator** is an input free model since all of its events are scripted in the input file. The **Generator** has a single output port, which we will call "arrive", through which it exports arriving customers. The model state is the list of **Customers** yet to arrive at the store. Here is the header file for the **Generator**.

```

#include "adevs.h"
#include "Customer.h"
#include <list>
/**
 * This class produces Customers according to the provided schedule.
 */
class Generator: public adevs::Atomic<IO_Type>
{
public:
    /// Constructor.
    Generator(const char* data_file);
    /// Internal transition function.
    void delta_int();
    /// External transition function.
    void delta_ext(double e, const adevs::Bag<IO_Type>& xb);
    /// Confluent transition function.
    void delta_conf(const adevs::Bag<IO_Type>& xb);

```

```

    /// Output function.
    void output_func(a devs::Bag<IO_Type>& yb);
    /// Time advance function.
    double ta();
    /// Output value garbage collection.
    void gc_output(a devs::Bag<IO_Type>& g);
    /// Destructor.
    ~Generator();
    /// Model output port.
    static const int arrive;

private:
    /// List of arriving customers.
    std::list<Customer*> arrivals;
};

```

The behavior of this model is very simple. The constructor opens the file containing the customer data and uses it to create a list of **Customer** objects. The inter-arrival times of the customers are stored in their `tenter` fields. Here is the constructor that initializes the model.

```

// Assign a locally unique number to the arrival port
const int Generator::arrive = 0;

Generator::Generator(const char* sched_file):
Atomic<IO_Type>()
{
    // Open the file containing the schedule
    fstream input_strm(sched_file);
    // Store the arrivals in a list
    double next_arrival_time = 0.0;
    double last_arrival_time = 0.0;
    while (true)
    {
        Customer* customer = new Customer;
        input_strm >> next_arrival_time >> customer->twait;
        // Check for end of file
        if (input_strm.eof())
        {
            delete customer;
            break;
        }
        // The entry time holds the inter arrival times, not the
        // absolute entry time.
        customer->tenter = next_arrival_time - last_arrival_time;
        // Put the customer at the back of the line
        arrivals.push_back(customer);
        last_arrival_time = next_arrival_time;
    }
}

```

Because the generator is input free, the external transition function is empty. Similarly, the confluent transition function merely calls the internal transition function (though, in fact, it could be empty because the confluent transition will never be called).



```

void Generator::delta_ext(double e, const Bag<IO_Type>& xb)
{
    /// The generator is input free, and so it ignores external events.
}

```

```

void Generator::delta_conf(const Bag<IO_Type>& xb)
{
    /// The generator is input free, and so it ignores input.
    delta_int();
}

```

The effect of an internal event (i.e., an event scheduled for the generator by itself) is to first place the arriving **Customer** onto the **Generator**'s "arrive" output port. This is done by the output function.

```

void Generator::output_func(Bag<IO_Type>& yb)
{
    // First customer in the list is produced as output
    IO_Type output(arrive, arrivals.front());
    yb.insert(output);
}

```

After the generator has produced this output event, its internal transition function removes the newly arrived customer from the arrival list.

```

void Generator::delta_int()
{
    // Remove the first customer. Because it was used as the
    // output object, it will be deleted during the gc_output()
    // method call at the end of the simulation cycle.
    arrivals.pop_front();
}

```

Internal events are scheduled with the time advance function. The **Generator**'s time advance function returns the time remaining until the next **Customer** arrives at the store. Remember that the `tarrival` field contains **Customers**' inter-arrival times, not the absolute arrival times, and the time advance function simply returns this value.

```

double Generator::ta()
{
    // If there are not more customers, next event time is infinity
    if (arrivals.empty()) return DBL_MAX;
    // Otherwise, wait until the next arrival
    return arrivals.front()->tenter;
}

```

To conduct the simulation experiment, it is necessary to connect the **Generator**'s output to the **Clerk**'s input. By doing this, **Customer** objects output on the **Generator**'s "arrive" output port appear as input on the **Clerk**'s "arrive" input port. This input event causes the **Clerk**'s external transition function to be activated. The relationship between input and output events can be best understood by viewing the whole model as two distinct components, the **Generator** and the **Clerk**, that are connected via their input and output ports. This view of the model is depicted in Figure 3.2.

As Figure 3.2 suggests, output events produced by the **Generator** on its "arrive" port, via the output function, appear as input events on the **Clerk**'s "arrive" port when the **Clerk**'s external transition function is called. The component models and their interconnections constitute a coupled (or network) model. To create the coupled model depicted above, we create an instance of a **Digraph** model that has the **Generator** and **Clerk** as components. Shown below is the code snippet that creates this two component model.




Figure 3.2: The combined **Generator** and **Clerk** model.

```

int main(int argc, char** argv)
{
    ...
    // Create a digraph model whose components use PortValue<Customer*>
    // objects as input and output objects.
    adevs::Digraph<Customer*> store;
    // Create and add the component models
    Clerk* clrk = new Clerk();
    Generator* genr = new Generator(argv[1]);
    store.add(clrk);
    store.add(genr);
    // Couple the components
    store.couple(genr, genr->arrive, clrk, clrk->arrive);
    ...
}

```

This code snippet first creates the components models and then adds them to the **Digraph**. Next, the components are interconnected by coupling the “arrive” output port of the **Generator** to the “arrive” input port of the **Clerk**.

Having created a coupled model to represent the store, all that remains is to perform the simulation. Here is the code snippet that simulates our model.

```

...
adevs::Simulator<IO_Type> sim(&store);
while (sim.nextEventTime() < DBL_MAX)
{
    sim.execNextEvent();
}
...

```

Putting this all of this together gives the main routine for the simulation program that generated the execution traces shown in the example above.

```

#include "Clerk.h"
#include "Generator.h"
#include "Observer.h"
#include <iostream>
using namespace std;

int main(int argc, char** argv)
{
    if (argc != 3)
    {
        cout << "Need input and output files!" << endl;
        return 1;
    }
    // Create a digraph model whose components use PortValue<Customer*>
    // objects as input and output objects.
    adevs::Digraph<Customer*> store;
}

```

```

// Create and add the component models
Clerk* clrk = new Clerk();
Generator* genr = new Generator(argv[1]);
Observer* obsrv = new Observer(argv[2]);
store.add(clrk);
store.add(genr);
store.add(obsrv);
// Couple the components
store.couple(genr, genr->arrive, clrk, clrk->arrive);
store.couple(clrk, clrk->depart, obsrv, obsrv->departed);
// Create a simulator and run until its done
adevs::Simulator<IO_Type> sim(&store);
while (sim.nextEventTime() < DBL_MAX)
{
    sim.execNextEvent();
}
// Done, component models are deleted when the Digraph is
// deleted.
return 0;
}

```

We have completed our first Adevs simulation program! However, a few details have been glossed over. The first question - an essential one for a programming language without garbage collection - is what happens to objects that we create in the **Generator** and **Clerk** output functions? The answer is that each model has a garbage collection method that is called at the end of every simulation cycle (in the example above, immediately prior to the return of the method `execNextEvent()`). The argument to the garbage collection method is the bag of objects created as output in the current simulation cycle. In our store example, the **Clerk** and **Generator** models simply delete the customer pointed to by each **PortValue** object in the garbage list. The implementation of the garbage collection method is shown below. This listing is for the **Generator** model; the **Clerk**'s *gc\_output()* method is identical.

```

void Generator::gc_output(Bag<IO_Type>& g)
{
    // Delete the customer that was produced as output
    Bag<IO_Type>::iterator i;
    for (i = g.begin(); i != g.end(); i++)
    {
        delete (*i).value;
    }
}

```

A second issue that has been overlooked is how to collect the statistics that were our original objective. One approach is to modify the **Clerk** so that it writes waiting times to a file as **Customers** are processed. While this would work, it has the unfortunate effect of cluttering up the **Clerk** with code specific to our experiment.

A better approach is to have an **Observer** that is coupled to the **Clerk**'s "depart" output port. The **Observer** records the desired statistics as it receives **Customers** on its "depart" input port. The advantage of this approach is that we can create new types of clerks to perform the same experiment using, for instance, different queuing strategies, without changing the experimental setup (i.e., customer generation and data collection). Similarly, we can change the experiment (i.e., how customers are generated and what data is collected) without changing the clerk.

Below is the code for the **Observer** class. This model is driven solely by external events. The observer reacts to an external event by recording the time that the **Customer** departed the **Clerk**'s queue (i.e., the

current simulation time) and the amount of time that the **Customer** waited in line. Here is the **Observer** header file.

```
#include "adevs.h"
#include "Customer.h"
#include <fstream>
/**
 * The Observer records performance statistics for a Clerk model
 * based on its observable output.
 */
class Observer: public adevs::Atomic<IO_Type>
{
    public:
        /// Input port for receiving customers that leave the store.
        static const int departed;
        /// Constructor. Results are written to the specified file.
        Observer(const char* results_file);
        /// Internal transition function.
        void delta_int();
        /// External transition function.
        void delta_ext(double e, const adevs::Bag<IO_Type>& xb);
        /// Confluent transition function.
        void delta_conf(const adevs::Bag<IO_Type>& xb);
        /// Time advance function.
        double ta();
        /// Output function.
        void output_func(adevs::Bag<IO_Type>& yb);
        /// Output value garbage collection.
        void gc_output(adevs::Bag<IO_Type>& g);
        /// Destructor.
        ~Observer();
    private:
        /// File for storing information about departing customers.
        std::ofstream output_strm;
};
```

Below is the **Observer** source file.

```
#include "Observer.h"
using namespace std;
using namespace adevs;

// Assign a locally unique number to the input port
const int Observer::departed = 0;

Observer::Observer(const char* output_file):
Atomic<IO_Type>(),
output_strm(output_file)
{
    // Write a header describing the data fields
    output_strm << "# Col 1: Time customer enters the line" << endl;
    output_strm << "# Col 2: Time required for customer checkout" << endl;
    output_strm << "# Col 3: Time customer leaves the store" << endl;
```

```

    output_strm << "# Col 4: Time spent waiting in line" << endl;
}

double Observer::ta()
{
    // The Observer has no autonomous behavior, so its next event
    // time is always infinity.
    return DBL_MAX;
}

void Observer::delta_int()
{
    // The Observer has no autonomous behavior, so do nothing
}

void Observer::delta_ext(double e, const Bag<IO_Type>& xb)
{
    // Record the times at which the customer left the line and the
    // time spent in it.
    Bag<IO_Type>::const_iterator i;
    for (i = xb.begin(); i != xb.end(); i++)
    {
        const Customer* c = (*i).value;
        // Compute the time spent waiting in line
        double waiting_time = (c->tleave-c->tenter)-c->twait;
        // Dump stats to a file
        output_strm << c->tenter << " " << c->twait << " " << c->tleave << " " << waiting_time << endl;
    }
}

void Observer::delta_conf(const Bag<IO_Type>& xb)
{
    // The Observer has no autonomous behavior, so do nothing
}

void Observer::output_func(Bag<IO_Type>& yb)
{
    // The Observer produces no output, so do nothing
}

void Observer::gc_output(Bag<IO_Type>& g)
{
    // The Observer produces no output, so do nothing
}

Observer::~Observer()
{
    // Close the statistics file
    output_strm.close();
}

```

This model is coupled to the **Clerk**'s "depart" output port in the same manner as before. The resulting

Time that the customer leaves the store	Time spent waiting in line
2	0
6	0
10	3
12	5
22	5
42	14
44	32
45	33

Table 3.2: Customer departure times and waiting times.

coupled model is illustrated in Figure 3.3; now we have three components instead of just two.




Figure 3.3: The **Generator**, **Clerk**, and **Observer** model.

Given the customer arrival data in Table 3.1, the consequent customer departure and waiting times are shown in Table 3.2. With this output, we could use a spreadsheet or some other suitable software to find the maximum and average customer wait times.

Again notice that the customer departure times correspond exactly with the production of customer departure events by the **Clerk** model. These departing customers are delivered to the **Observer** via the **Clerk** to **Observer** coupling shown in Figure 3.3. Each entry in Table 3.2 is the result of executing the **Observer**'s external transition function. Also notice that the **Observer**'s internal and confluent transition functions will never be executed because the **Observer**'s `ta()` method always returns infinity.

This section demonstrated the most common parts of a simulation program built with Adevs. The remainder of the manual covers **Atomic** and **Network** models in greater detail, demonstrates the construction of variable structure models, and shows how continuous models can be added to your discrete event simulation.

## Chapter 4

# Atomic Models

Atomic models are the basic building blocks of a DEVS model. The behavior of an atomic model is described by its state transition functions (internal, external, and confluent), its output function, and its time advance function. Within Adevs, these aspects of an atomic model are implemented by sub-classing the **Atomic** class and implementing its pure virtual methods that correspond to the internal, external, and confluent transition functions, the output function, and the time advance function.

The state of an atomic model is realized by the attributes of the object that implements the model. The internal transition function describes the model's autonomous behavior; that is, how its state evolves in the absence of input. These types of events are called internal events because they are self-induced; i.e., internal to the model. The time advance function schedules these autonomous changes of state. The output function gives the model's output when these internal events occur.

The external transition function describes how the model changes state in response to input. The confluent transition function handles the simultaneous occurrence of an internal and external event. The types of objects that are accepted as input and produced as output are specified with a template argument to the **Atomic** base class.

The **Clerk** described in Section 3 demonstrates all of the aspects of an **Atomic** model. We'll use it to demonstrate how an **Atomic** model generates output, processes input, and schedules internal events. Here is the **Clerk**'s class definition:

```
include "adevs.h"
#include "Customer.h"
#include <list>

class Clerk: public adevs::Atomic<IO_Type>
{
    public:
        /// Constructor.
        Clerk();
        /// Internal transition function.
        void delta_int();
        /// External transition function.
        void delta_ext(double e, const adevs::Bag<IO_Type>& xb);
        /// Confluent transition function.
        void delta_conf(const adevs::Bag<IO_Type>& xb);
        /// Output function.
        void output_func(adevs::Bag<IO_Type>& yb);
        /// Time advance function.
        double ta();
        /// Output value garbage collection.
```

```

        void gc_output(a devs::Bag<IO_Type>& g);
        /// Destructor.
        ~Clerk();
        /// Model input port.
        static const int arrive;
        /// Model output port.
        static const int depart;

private:
    /// The clerk's clock
    double t;
    /// List of waiting customers.
    std::list<Customer*> line;
    /// Time spent so far on the customer at the front of the line
    double t_spent;
};

```

and here its implementation

```

#include "Clerk.h"
#include <iostream>
using namespace std;
using namespace a devs;

// Assign locally unique identifiers to the ports
const int Clerk::arrive = 0;
const int Clerk::depart = 1;

Clerk::Clerk():
    Atomic<IO_Type>(), // Initialize the parent Atomic model
    t(0.0), // Set the clock to zero
    t_spent(0.0) // No time spent on a customer so far
{
}

void Clerk::delta_ext(double e, const Bag<IO_Type>& xb)
{
    // Print a notice of the external transition
    cout << "Clerk: Computed the external transition function at t = " << t+e << endl;
    // Update the clock
    t += e;
    // Update the time spent on the customer at the front of the line
    if (!line.empty())
    {
        t_spent += e;
    }
    // Add the new customers to the back of the line.
    Bag<IO_Type>::const_iterator i = xb.begin();
    for (; i != xb.end(); i++)
    {
        // Copy the incoming Customer and place it at the back of the line.
        line.push_back(new Customer(*((*i).value)));
        // Record the time at which the customer entered the line.
    }
}

```



```

        line.back()->tenter = t;
    }
    // Summarize the model state
    cout << "Clerk: There are " << line.size() << " customers waiting." << endl;
    cout << "Clerk: The next customer will leave at t = " << t+ta() << "." << endl;
}

void Clerk::delta_int()
{
    // Print a notice of the internal transition
    cout << "Clerk: Computed the internal transition function at t = " << t+ta() << endl;
    // Update the clock
    t += ta();
    // Reset the spent time
    t_spent = 0.0;
    // Remove the departing customer from the front of the line.
    line.pop_front();
    // Summarize the model state
    cout << "Clerk: There are " << line.size() << " customers waiting." << endl;
    cout << "Clerk: The next customer will leave at t = " << t+ta() << "." << endl;
}

void Clerk::delta_conf(const Bag<IO_Type>& xb)
{
    delta_int();
    delta_ext(0.0,xb);
}

void Clerk::output_func(Bag<IO_Type>& yb)
{
    // Get the departing customer
    Customer* leaving = line.front();
    // Set the departure time
    leaving->tleave = t + ta();
    // Eject the customer
    IO_Type y(depart,leaving);
    yb.insert(y);
    // Print a notice of the departure
    cout << "Clerk: Computed the output function at t = " << t+ta() << endl;
    cout << "Clerk: A customer just departed!" << endl;
}

double Clerk::ta()
{
    // If the list is empty, then next event is at inf
    if (line.empty()) return DBL_MAX;
    // Otherwise, return the time remaining to process the current customer
    return line.front()->twait-t_spent;
}

void Clerk::gc_output(Bag<IO_Type>& g)
{

```

Enter checkout line	Time to process order
1	1
2	4
3	4
5	2
7	10
8	20
10	2
11	1

Table 4.1: Customer arrival times and times to process customers' orders.

```

// Delete the outgoing customer objects
Bag<IO_Type>::iterator i;
for (i = g.begin(); i != g.end(); i++)
{
    delete (*i).value;
}

Clerk::~Clerk()
{
    // Delete anything remaining in the customer queue
    list<Customer*>::iterator i;
    for (i = line.begin(); i != line.end(); i++)
    {
        delete *i;
    }
}

```

Consider the simulation of the convenience store described in Section 3 (i.e., with the arrivals listed in Table 3.1); this arrival data is listed again here:

Table 4.1 describes an input sequence that is fed to the **Clerk** model. The algorithm for processing this, or any other, input sequence is listed below. The **Atomic** model being simulated is called 'model',  $t$  is the current simulation time (i.e., the time of the last event - internal, external, or confluent), and  $t_{\text{input}}$  is the time of the next unprocessed event in the input sequence.

1. Set the next event time  $tN$  to the smaller of the next internal event time  $t_{\text{self}} = t + \text{model.ta}()$  and the next input event time  $t_{\text{input}}$ .
2. If  $t_{\text{self}} = tN$  and  $t_{\text{input}} > tN$  then produce an output event at time  $t_{\text{self}}$  by calling `model.output_func()` and compute the next state by calling `model.delta_int()`.
3. If  $t_{\text{self}} = t_{\text{input}} = tN$  then produce an output event at time  $t_{\text{self}}$  by calling `model.output_func()` and compute the next state by calling `model.delta_conf(x)` where  $x$  contains the input at time  $t_{\text{input}}$ .
4. If  $t_{\text{self}} > tN$  and  $t_{\text{input}} = tN$  then compute the next state by calling `model.delta_ext(t_input-t,x)` where  $x$  contains the input at time  $t_{\text{input}}$ .
5. Set  $t$  equal to  $tN$ .
6. Repeat if there are more input or internal events to process.

The simulation runs until there are no internal or external events remaining to process. The first step of the algorithm computes the time of the next event as the sooner of the next input event and the next

internal event. If the next internal event happens first, then the model produces an output and its next state is computed with the internal transition function.

If the next input event happens first, then the next state of the model is computed with the external transition function; no output is produced in this case. The elapsed time argument given to the external transition function is the amount of time that has passed since the previous event - internal, external, or confluent - at that model. If the next input and internal event happen at the same time, then the model produces an output and its next state is computed with the confluent transition function. The simulation clock is then advanced to the event time and these steps are repeated.

The output trace resulting from the input sequence in Table 4.1 is shown below. It has been broken up to show where each simulation cycle begins and ends and the type of event occurring in each cycle.

```
-External event-----
Clerk: Computed the external transition function at t = 1
Clerk: There are 1 customers waiting.
Clerk: The next customer will leave at t = 2.
-Confluent event-----
Clerk: Computed the output function at t = 2
Clerk: A customer just departed!
Clerk: Computed the internal transition function at t = 2
Clerk: There are 0 customers waiting.
Clerk: The next customer will leave at t = 1.79769e+308.
Clerk: Computed the external transition function at t = 2
Clerk: There are 1 customers waiting.
Clerk: The next customer will leave at t = 6.
-External event-----
Clerk: Computed the external transition function at t = 3
Clerk: There are 2 customers waiting.
Clerk: The next customer will leave at t = 6.
-External event-----
Clerk: Computed the external transition function at t = 5
Clerk: There are 3 customers waiting.
Clerk: The next customer will leave at t = 6.
-Internal event-----
Clerk: Computed the output function at t = 6
Clerk: A customer just departed!
Clerk: Computed the internal transition function at t = 6
Clerk: There are 2 customers waiting.
Clerk: The next customer will leave at t = 10.
-External event-----
Clerk: Computed the external transition function at t = 7
Clerk: There are 3 customers waiting.
Clerk: The next customer will leave at t = 10.
-External event-----
Clerk: Computed the external transition function at t = 8
Clerk: There are 4 customers waiting.
Clerk: The next customer will leave at t = 10.
-Confluent event-----
Clerk: Computed the output function at t = 10
Clerk: A customer just departed!
Clerk: Computed the internal transition function at t = 10
Clerk: There are 3 customers waiting.
Clerk: The next customer will leave at t = 12.
```

```

Clerk: Computed the external transition function at t = 10
Clerk: There are 4 customers waiting.
Clerk: The next customer will leave at t = 12.
-External event-----
Clerk: Computed the external transition function at t = 11
Clerk: There are 5 customers waiting.
Clerk: The next customer will leave at t = 12.
-Internal event-----
Clerk: Computed the output function at t = 12
Clerk: A customer just departed!
Clerk: Computed the internal transition function at t = 12
Clerk: There are 4 customers waiting.
Clerk: The next customer will leave at t = 22.
-Internal event-----
Clerk: Computed the output function at t = 22
Clerk: A customer just departed!
Clerk: Computed the internal transition function at t = 22
Clerk: There are 3 customers waiting.
Clerk: The next customer will leave at t = 42.
-Internal Event-----
Clerk: Computed the output function at t = 42
Clerk: A customer just departed!
Clerk: Computed the internal transition function at t = 42
Clerk: There are 2 customers waiting.
Clerk: The next customer will leave at t = 44.
-Internal event-----
Clerk: Computed the output function at t = 44
Clerk: A customer just departed!
Clerk: Computed the internal transition function at t = 44
Clerk: There are 1 customers waiting.
Clerk: The next customer will leave at t = 45.
-Internal event-----
Clerk: Computed the output function at t = 45
Clerk: A customer just departed!
Clerk: Computed the internal transition function at t = 45
Clerk: There are 0 customers waiting.
Clerk: The next customer will leave at t = 1.79769e+308.

```

Now lets create a more sophisticated clerk model; one that will interrupt the checkout of a customer with a large order to more quickly serve a customer with a small order. This clerk operates as follows. If while a customer is being served, another customer arrives whose order can be processed quickly, then the clerk stops processing the current customer and begins to process the new customer. The clerk, however, does this only occasionally. To be precise, let a small order be one requiring no more than one unit of time to process. Moreover, the clerk interrupts processing of an order at most once in every 10 units of time.

This new clerk has two state variables. The first records the time remaining before the clerk is willing to interrupt the processing of a customer. The second is the list of customers waiting to be served. Here is the header file for the new clerk model, which we will call **Clerk2**.

```

#include "adevs.h"
#include "Customer.h"
#include <list>

class Clerk2: public adevs::Atomic<IO_Type>

```

```

{
    public:
        /// Constructor.
        Clerk2();
        /// Internal transition function.
        void delta_int();
        /// External transition function.
        void delta_ext(double e, const adevs::Bag<IO_Type>& xb);
        /// Confluent transition function.
        void delta_conf(const adevs::Bag<IO_Type>& xb);
        /// Time advance function.
        double ta();
        /// Output function.
        void output_func(adevs::Bag<IO_Type>& yb);
        /// Output value garbage collection.
        void gc_output(adevs::Bag<IO_Type>& g);
        /// Destructor.
        ~Clerk2();
        /// Model input port.
        static const int arrive;
        /// Model output port.
        static const int depart;
    private:
        /// Structure for storing information about customers in the line
        struct customer_info_t
        {
            // The customer
            Customer* customer;
            // Time remaining to process the customer order
            double t_left;
        };
        /// List of waiting customers.
        std::list<customer_info_t> line;
        /// Time before we can preempt another customer
        double preempt;
        /// The clerk's clock
        double t;
        /// Threshold correspond to a 'small' order processing time
        static const double SMALL_ORDER;
        /// Minimum time between preemptions.
        static const double PREEMPT_TIME;
};

```

The **Clerk2** constructor sets the clerk's clock and interruption timers to zero.

```

Clerk2::Clerk2():
    Atomic<IO_Type>(),
    preempt(0.0),
    t(0.0)
{
}

```

The output function of the **Clerk2** model sets the exit time of the departing customer and then ejects that customer via the “depart” port.

```

void Clerk2::output_func(Bag<IO_Type>& yb)
{
    /// Set the exit time of the departing customer
    line.front().customer->tleave = t+ta();
    /// Place the customer at the front of the line onto the depart port.
    IO_Type y(depart,line.front().customer);
    yb.insert(y);
    /// Report the departure
    cout << "Clerk: A customer departed at t = " << t+ta() << endl;
}

```

The **Clerk2**'s external transition function works as follows. When a new customer arrives, the clerk first advanced its clock by the elapsed time. Next, she reduces the time remaining to process the current customer. This reduction reflects the amount of time that has already been spent on the customer's order, which is the time elapsed since the clerk's last change of state. Then the clerk decrements the time remaining before she is willing to interrupt the processing of a large order; this timer is also decremented by the elapsed time.

Now the clerk records the time at which each arriving customer enters the line; this time is the value of the clock. If any of the arriving customers has a small checkout time and the clerk is willing to interrupt the present order, then that customer with the small order goes to the front of the line. This preempts the current customer, who now has the second place in line, and causes the preempt timer to be reset. Otherwise, the new customer simply goes to the back of the line.

```

void Clerk2::delta_ext(double e, const Bag<IO_Type>& xb)
{
    /// Update the clock
    t += e;
    /// Update the time spent working on the current order
    if (!line.empty())
    {
        line.front().t_left -= e;
    }
    /// Reduce the preempt time
    preempt -= e;
    /// Place new customers into the line
    Bag<IO_Type>::const_iterator iter = xb.begin();
    for (; iter != xb.end(); iter++)
    {
        cout << "Clerk: A new customer arrived at t = " << t << endl;
        /// Create a copy of the incoming customer and set the entry time
        customer_info_t c;
        c.customer = new Customer(*((*iter).value));
        c.t_left = c.customer->twait;
        /// Record the time at which the customer enters the line
        c.customer->tenter = t;
        /// If the customer has a small order
        if (preempt <= 0.0 && c.t_left <= SMALL_ORDER)
        {
            cout << "Clerk: The new customer has preempted the current one!" << endl;
            /// We won't preempt another customer for at least this long
            preempt = PREEMPT_TIME;
            /// Put the new customer at the front of the line
            line.push_front(c);
        }
    }
}

```

```

    }
    /// otherwise just put the customer at the end of the line
    else
    {
        cout << "Clerk: The new customer is at the back of the line" << endl;
        line.push_back(c);
    }
}
}

```

The internal transition function is similar to the external transition function. It begins by decrementing the time remaining before the clerk will interrupt an order. The customer that just departed the store via the output function is then removed from the front of the line. If the line is empty, then there is nothing to do and the clerk sits idly behind her counter. If the line is not empty and the preemption time has expired then the clerk scans the line for the first customer with a small order. If such a customer can be found, that customer moves to the front of the line. Then the clerk starts ringing up the first customer in her line. Here is the internal transition function for the **Clerk2** model.

```

void Clerk2::delta_int()
{
    // Update the clerk's clock
    t += ta();
    // Update the preemption timer
    preempt -= ta();
    // Remove the departing customer from the front of the line.
    // The departing customer will be deleted later by our garbage
    // collection method.
    line.pop_front();
    // Check to see if any customers are waiting.
    if (line.empty())
    {
        cout << "Clerk: The line is empty at t = " << t << endl;
        return;
    }
    // If the preemption time has passed, then look for a small
    // order that can be promoted to the front of the line.
    list<customer_info_t>::iterator i;
    for (i = line.begin(); i != line.end() && preempt <= 0.0; i++)
    {
        if ((*i).t_left <= SMALL_ORDER)
        {
            cout << "Clerk: A queued customer has a small order at time " << t << endl;
            customer_info_t small_order = *i;
            line.erase(i);
            line.push_front(small_order);
            preempt = PREEMPT_TIME;
            break;
        }
    }
}
}

```

The time advance function returns the time remaining to process the customer at the front of the line, or infinity (i.e., DBL\_MAX) if there are no customers to process.

```
double Clerk2::ta()
{
    // If the line is empty, then there is nothing to do
    if (line.empty()) return DBL_MAX;
    // Otherwise, wait until the first customer will leave
    else return line.front().t_left;
}
```

The last function to implement is the confluent transition function. The **Clerk2** model has the same confluent transition as the **Clerk** that is described in section 3:

```
void Clerk2::delta_conf(const Bag<IO_Type>& xb)
{
    delta_int();
    delta_ext(0.0,xb);
}
```

The behavior of the **Clerk2** model is more complex than that of the **Clerk** model. To exercise the **Clerk2**, we replace the **Clerk** model in the example from section 3 with the **Clerk2** model and perform the same experiment. Here is the output trace for the **Clerk2** model in response to the input sequence shown in Table 4.1. This trace was generated by the print statements shown in the source code listings for the **Clerk2** model.

```
Clerk: A new customer arrived at t = 1
Clerk: The new customer has preempted the current one!
Clerk: A customer departed at t = 2
Clerk: The line is empty at t = 2
Clerk: A new customer arrived at t = 2
Clerk: The new customer is at the back of the line
Clerk: A new customer arrived at t = 3
Clerk: The new customer is at the back of the line
Clerk: A new customer arrived at t = 5
Clerk: The new customer is at the back of the line
Clerk: A customer departed at t = 6
Clerk: A new customer arrived at t = 7
Clerk: The new customer is at the back of the line
Clerk: A new customer arrived at t = 8
Clerk: The new customer is at the back of the line
Clerk: A customer departed at t = 10
Clerk: A new customer arrived at t = 10
Clerk: The new customer is at the back of the line
Clerk: A new customer arrived at t = 11
Clerk: The new customer has preempted the current one!
Clerk: A customer departed at t = 12
Clerk: A customer departed at t = 13
Clerk: A customer departed at t = 23
Clerk: A customer departed at t = 43
Clerk: A customer departed at t = 45
Clerk: The line is empty at t = 45
```

The evolution of the **Clerk2** line is depicted in Fig. 4.1. To time 11 the line evolves just as it did with the **Clerk** model. At time 11 the **Clerk2** changes the course of the simulation by moving a customer with a small order to the front of the line.






Figure 4.1: The evolution of the **Clerk2** line in response to the customer arrival sequence listed in Table 4.1.



## Chapter 5

# Network Models

A network model comprises atomic models and other network models that are interconnected. Network models can be components of other network models, thereby enabling the construction of very large, multi-level systems. Unlike atomic models, network models do not directly define new dynamic behavior. The dynamics of a network model are determined by the dynamics of its component parts and their interactions. Atomic models define fundamental behaviors; network models define structure.

### 5.1 Parts of a Network Model

Network models are derived from the abstract **Network** class. This class has two virtual methods: *route* and *getComponents*. The *route* method implements connections between the components of the **Network** model and between those components and the inputs and outputs of the network itself. The *getComponents* method provides the set of components that constitute the network.

The *route* method is the primary workhorse of a **Network** model. It realizes three types of connections. The first are connections between components of the network. The second are connections from the network's inputs to the inputs of its component models. The third are connections from the components' outputs to the outputs of the network.

The signature of the *route* method is

```
void route(const X& value, Devs<X>* model, Bag<Event<X> >& r)
```

for which the value argument is the object being routed, the model argument is the **Network** or **Atomic** model that originated the event (i.e., the event's source), and the r argument is a bag to be filled with the targets of the event. Each target is described by an **Event** object that carries to pieces of information: a pointer to the model that is the event's target and the object to be delivered to that target.

The simulator uses these **Event** objects in one of three ways, which depending on the relationship between the source of the event and its target. These uses are

1. If the event's source is a component of the network and the target is the network itself, then the event's value becomes an output from the network.
2. If the event's source is the network and the event's target is a component of the network, then the event's value becomes an input to that component.
3. If the event's source and target are both components of the network, then the event's value becomes an input to the target.

Any other relationship between the source and the target is illegal and will cause the simulator to raise an exception.

The simplest example of how the route method is used converts output from one **Atomic** component of a **Network** into input for another **Atomic** component of the same **Network**. Figure 5.1 illustrates this case.




Figure 5.1: Two connected **Atomic** components in a single **Network**.

The simulator begins by invoking the `output_func` method of **Atomic** model *A*. Next, the simulator iterates through the elements of *A*'s output bag and calls the **Network**'s `route` method for each one. The arguments passed to `route` at each call are

1. the output object itself, which becomes the value argument,
2. a pointer to *A*, which is the model argument, and
3. an empty **Bag** for holding **Event** objects.

Two things are done by the *route* method to cause **Atomic** model *B* to receive the output object from *A*. First, an **Event** object is created that contains the output object and a pointer to *B*. Second, this **Event** object is inserted into the **Bag** *r* of receivers of the event. If we suppose, for the sake of illustration, that input and output objects have type `int`, then the route method for this example is

```

void route(const int& value, Devs<int>* model, Bag<Event<int> >& r) {
    if (model == A) {
        Event<int> e(B,value);
        r.insert(e);
    }
}

```

where *A* and *B* are pointers to the respective components. This route method implements the network shown in Fig. 5.1.

A more complicated example is the network itself receiving input ultimately destined for one of its atomic components. This can happen, for instance, when the network is a component of another network. Suppose that input to our example network is to become input for **Atomic** model *A*. Figure 5.2 extends Fig. 5.1 to include this connection.




Figure 5.2: Two connected **Atomic** components with external input coupling to component *A*.

When an event appears at the input of the network, the simulator calls the **Network**'s route method with the following arguments:

1. the input object itself, which becomes the value argument,
2. a pointer to the **Network** that is receiving the event, and
3. an empty **Bag** for holding **Event** objects.

As before, the *route* creates an **Event** object that indicates the receiving model and the value of the event. This **Event** object is put into the **Bag** *r* of receivers. The code below implements the network shown in Fig. 5.2; note that the *this* pointer points to the **Network** itself (i.e., to the network that is receiving the input initially).

```
void route(const int& value, Devs<int>* model, Bag<Event<int> >& r) {
    if (model == A) {
        Event<int> e(B,value);
        r.insert(e);
    }
    else if (model == this) {
        Event<int> e(A,value);
        r.insert(e);
    }
}
```




Figure 5.3: A two component network model with external input, external output, and internal coupling.

For a complete example, the network is extended again, as shown in Fig. 5.2, to include two more connections: a connection from the output of model *B* to the output of the network and a feedback connection from *B* to *A*. This configuration is shown in Fig. 5.3. This new connection requires an additional case in the route method. This case checks for output from *B* and, if such an output is found, directs it to both *A* and the network itself. A **Event** object is created for each target and added to the **Bag** *r* of receivers: one of these **Events** results in an input to *A* and the other in an output from the network. Here is the implementation.

```
void route(const int& value, Devs<int>* model, Bag<Event<int> >& r) {
    if (model == A) {
        Event<int> e(B,value);
        r.insert(e);
    }
    else if (model == this) {
        Event<int> e(A,value);
        r.insert(e);
    }
    else if (model == B) {

```

```

        Event<int> e1(this,value);
        Event<int> e2(A,value);
        r.insert(e1);
        r.insert(e2);
    }
}

```

The `getComponents` method is the other virtual method that must be implemented by a **Network**. The simulator passes to this method an empty **Set** of pointers to models, and this must be filled with the network's components. The signature of the `getComponents` method is

```
void getComponents(Set<Devs<X>*>& c)
```

where `c` is the set to be filled. There isn't much else to say about this method. The code below shows how it is implemented for the two component network shown in Fig. 5.3; this code, of course, also works for the networks shown in Figs. 5.2 and 5.1.

```

void getComponents(Set<Devs<int>*>& c) {
    c.insert(A);
    c.insert(B);
}

```

There are just three other items to mention in relation to **Network** models. First, components cannot be connected to themselves. This means that direct feedback loops and connections direct through a network model are illegal. The former can always be replaced with an internal event and the latter by simply bypassing the network (which won't do anything with the input in any event). These two cases are illustrated in Fig. 5.4. Second, direct coupling can only occur between components belonging to the same network, and every component must belong to, at most, one network. Third, the *route* method is allowed to modify an output before sending it along; this can be useful in some cases.




Figure 5.4: Illegal coupling in a **Network** model.

## 5.2 Simulating a Network Model

Each iteration of the simulator has four phases: 1) advance the simulation clock to the time of the next event, 2) compute the outputs from atomic models that will change state (i.e., how will undergo an internal or confluent event) and convert these outputs into inputs for other models, 3) calculate the next state of each model with events to process, and 4) cleanup garbage left over from the output calculations. These four phase are repeated until the time of the next event is at infinity (i.e., `DBL_MAX`) or you decide to stop the simulation.

Conveniently, there are no special rules for simulating networks of network models. The simulator considers the entire collection of atomic models when determining the next event time; output from atomic

models are recursively routed to their atomic destinations; and the state transitions and garbage collection are performed over the complete set of active atomic components. Hierarchies of network models are a convenient organizing tool for the modeler, but the simulator treats (via its recursive routing of events) the multi-level network as a flat structure.

Algorithm 1 sketches the simulation procedure. Note that the procedure for simulating atomic models (see section 4) is embedded in the procedure for simulating network models. The rules for atomic models do not change in any way; each atomic model sees a sequence of input events and produces a sequence of output events just as before. The only difference here is that the input events are created by other atomic models, and so the input sequence for each atomic model is constructed as the simulation progresses.

---

**Algorithm 1** The simulation procedure for a network model.

---

```

Initialize the state of every Atomic model
Set the time of last event  $t_{l,i}$  of every Atomic model  $i$  to 0
Set the simulation time  $t$  to 0
while The smallest time of next event for the Atomic models is less than DBL_MAX do
    Set  $t$  to the smallest time of next event for the Atomic models
    Find the set of Atomic models whose next event time is equal to  $t$ . These are the imminent models.
    Get the output of each imminent model by calling its output_func
    Convert output from imminent models to input for other models using the Network models route
    methods (do this recursively if the network has more than one level)
    for each Atomic model  $i$  that is imminent or has input do
        if  $i$  is an imminent model and it does not have input then
            Compute the next model state with delta_int()
        else if  $i$  is an imminent and it has input then
            Compute the next model state with delta_conf(xb), where  $xb$  is the model input
        else if  $i$  is not an imminent model and it has input then
            Compute the next model state with delta_ext( $t - t_{l,i}$ , xb), where  $xb$  is the model input
        end if
        Set  $t_{l,i}$  to  $t$ 
    end for
end while

```

---

### 5.3 A complete example of a network model

I'll use the Adevs **SimpleDigraph** class to illustrate the process of building a network model. The **SimpleDigraph** models a network of components whose connections are represented with a directed graph. If, for example, component  $A$  is connected to component  $B$ , then all output events generated by  $A$  become input events to  $B$ .

The **SimpleDigraph** has two methods for building a network. The *add* method takes an **Atomic** or **Network** model and adds it to the set of components. The *couple* method accepts a pair of components and connects the first to the second. Below is the class definition for the model; note that it has a template parameter for setting its input and output type. The **Network**, **Devs**, **Bag**, and **Set** are in the adevs namespace, and adevs:: must precede them unless the **SimpleDigraph** is itself in the adevs namespace (which, in this case, it is).

```

template <class VALUE> class SimpleDigraph: public Network<VALUE> {
public:
    /// A component of the SimpleDigraph model
    typedef Devs<VALUE> Component;

    /// Construct a network with no components

```

```

SimpleDigraph():Network<VALUE>(){}
/// Add a model to the network.
void add(Component* model);
/// Couple the source model to the destination model
void couple(Component* src, Component* dst);
/// Assigns the model component set to c
void getComponents(Set<Component*>& c);
/// Use the coupling information to route an event
void route(const VALUE& x, Component* model, Bag<Event<VALUE> >& r);
/// The destructor destroys all of the component models
~SimpleDigraph();

private:
    // Component model set
    Set<Component*> models;
    // Coupling information
    std::map<Component*,Bag<Component*> > graph;
};

```

The **SimpleDigraph** has two member variables. The first is a set of pointers to the components of the network; these are stored in the **Set** models. The components can be **Atomic** objects, **Network** objects, or both. The **SimpleDigraph** components are the nodes of the directed graph. The second member variable is a STL map that contains the network's links, or edges; these are stored in the **map** graph.

The **SimpleGraph** has four methods beyond the required route and getComponents. One of these is the constructor, which creates an empty network. The other is the destructor, which deletes all of the network's components. The others are add and couple.

The add method does three things. First, it checks that the network is not being added to itself; this is illegal and would cause no end of trouble for the simulator. Next, it adds the new component to its set of components. Last, the **SimpleNetwork** sets the component's parent. The last step is needed so that the simulator can climb up and down the model tree. If it is omitted then the routing of events is likely fail. Here is the implementation of the *add* method.

```

template <class VALUE>
void SimpleDigraph<VALUE>::add(Component* model) {
    assert(model != this);
    models.insert(model);
    model->setParent(this);
}

```

The couple method does two things, but one of them is somewhat superfluous. First, it adds the source (src) and destination (dst) models to the set of components. We could simply have required that the user call the add method before using the couple method, but adding the components here doesn't hurt and might prevent an error. The second step is essential; the method adds the src → dst link to the graph. Notice that the *SimpleDigraph* itself is a node in the network (but it is not in the set of components!): components that are connected to the network cause network outputs. Similarly, connecting the network to a component cause input to the network to become input to that component. Here is the implementation of the couple method.

```

template <class VALUE>
void SimpleDigraph<VALUE>::couple(Component* src, Component* dst) {
    if (src != this) add(src);
    if (dst != this) add(dst);
    graph[src].insert(dst);
}

```



Of the two required methods, `route` is the more complicated. The arguments to `route` are an object that is the value to be routed, the network element (i.e., either the **SimpleDigraph** or one of its components) that created that object, and the **Bag** to be filled with **Event** objects that indicate the object's receivers. The method begins by finding the collection of components that are connected to the source of the event. Next, it iterates through this collection and for each receiver adds an **Event** to the **Bag** of receivers. When this is done the method returns. The implementation is below.

```
template <class VALUE>
void SimpleDigraph<VALUE>::route(const VALUE& x, Component* model, Bag<Event<VALUE> >& r) {
    // Find the list of target models and ports
    typename std::map<Component*, Bag<Component*> >::iterator graph_iter;
    graph_iter = graph.find(model);
    // If no target, just return
    if (graph_iter == graph.end()) return;
    // Otherwise, add the targets to the event bag
    Event<VALUE> event;
    typename Bag<Component*>::iterator node_iter;
    for (node_iter = (*graph_iter).second.begin();
        node_iter != (*graph_iter).second.end(); node_iter++) {
        event.model = *node_iter;
        event.value = x;
        r.insert(event);
    }
}
```

The second required method, `getComponents`, is trivial. If we had used some collection other than an Adevs **Set** to store the components, then the method would have needed to explicitly insert every component model into the **Set** `c`. But because the sets `models` and `c` are both **Set** objects, and the **Set** has an assignment operator, a simple call to that operator is sufficient.

```
template <class VALUE>
void SimpleDigraph<VALUE>::getComponents(Set<Component*>& c) {
    c = models;
}
```

The constructor and the destructor complete the class. The constructor implementation appears in the class definition; it only calls the superclass constructor. The destructor deletes the component models. Its implementation is shown below.

```
template <class VALUE>
SimpleDigraph<VALUE>::~SimpleDigraph() {
    typename Set<Component*>::iterator i;
    for (i = models.begin(); i != models.end(); i++) {
        delete *i;
    }
}
```

## 5.4 Digraph Models

This section introduces the **Digraph** model, which is part of the Adevs simulation library. The **Digraph** is a tool for building network models based on, or described by, a block diagram. The model of the convenience store, developed in section 3, was our first example of a **Digraph** model. The code used to construct the convenience store model (without the **Observer**) is shown below. The block diagram that corresponds to this code snippet is shown in Fig. 5.5.

```

// Create a digraph model whose components use PortValue<Customer*>
// objects as input and output objects.
adevs::Digraph<Customer*> store;
// Create and add the component models
Clerk* clrk = new Clerk();
Generator* genr = new Generator(argv[1]);
store.add(clrk);
store.add(genr);
// Couple the components
store.couple(genr, genr->arrive, clrk, clrk->arrive);

```




Figure 5.5: A **Digraph** model with two components.

The components of a **Digraph** must use **adevs::PortValue** objects for their input and output type. The **Digraph** is a template class with two template parameters. The first parameter is the type of object used for a value in the **PortValue** objects. The second parameter is the type of object used for a port in the **PortValue** objects. The port parameter is of type 'int' by default.

For the modeler, the **Digraph** has two primary methods. The add method is used to add components to the network; the argument to the add method is just the model to be included in the network. The couple method is used to connect components of the network. The first two arguments to the couple method are the source model and source port. The second two arguments are the destination model and the destination port.

The effect of coupling a source model to a destination model is that output produced by the source model on the source port appears as input to the destination model on the destination port. To illustrate this, consider the output function of the **Generator** model that is shown in Fig. 5.5.

```

void Generator::output_func(Bag<IO_Type>& yb)
{
    // First customer in the list is produced as output
    IO_Type output(arrive, arrivals.front());
    yb.insert(output);
}

```

This output function places an output value of type 'Customer\*' on the "arrive" output port of the **Generator**; recall that 'IO\_Type' is a typedef for 'PortValue<Customer\*>'. A corresponding **PortValue** object appears in the input bag of the **Clerk**. The value attribute of the **PortValue** object received by the clerk points to the **Customer** object created by the **Generator**, and the port is the **Clerk**'s "arrive" port.

The components for the network need not consist only of **Atomic** models; the **Digraph** can also have other **Network** models as its components. For instance, suppose that we want to model a convenience store that has two clerks. When customers are ready to pay their bill, they enter the shortest line. To build this model, we reuse the **Clerk**, **Generator**, and **Observer** models that were introduced in section 3, and add an additional model for the process of customers selecting a checkout line.

The header and source code for the model of the line-selection process is shown below. This model has two output ports, one for each line, and there are three input ports. One input port accepts new customers. The others are used to track the number of number of customers departing each line: a customer departing either clerk generates an event on the appropriate input port. In this way, the model is able to track the

number of customers in each line and assign new customers to the shortest one. With this information, the state transition and output functions are (I hope) self explanatory. Here is the class definition

```
#include "adevs.h"
#include "Customer.h"
#include <list>

// Number of lines to consider.
#define NUM_LINES 2

class Decision: public adevs::Atomic<IO_Type>
{
    public:
        /// Constructor.
        Decision();
        /// Internal transition function.
        void delta_int();
        /// External transition function.
        void delta_ext(double e, const adevs::Bag<IO_Type>& x);
        /// Confluent transition function.
        void delta_conf(const adevs::Bag<IO_Type>& x);
        /// Output function.
        void output_func(adevs::Bag<IO_Type>& y);
        /// Time advance function.
        double ta();
        /// Output value garbage collection.
        void gc_output(adevs::Bag<IO_Type>& g);
        /// Destructor.
        ~Decision();
        /// Input port that receives new customers
        static const int decide;
        /// Input ports that receive customers leaving the two lines
        static const int departures[NUM_LINES];
        /// Output ports that produce customers for the two lines
        static const int arrive[NUM_LINES];

    private:
        /// Lengths of the two lines
        int line_length[NUM_LINES];
        /// List of deciding customers and their decision.
        std::list<std::pair<int, Customer*> > deciding;
        /// Delete all waiting customers and clear the list.
        void clear_deciders();
        /// Returns the arrive port associated with the shortest line
        int find_shortest_line();
};
```

and here is the implementation

```
#include "Decision.h"
#include <iostream>
using namespace std;
using namespace adevs;
```

```

// Assign identifiers to ports. Assumes NUM_LINES = 2.
// The numbers are selected to allow indexing into the
// line length and port number arrays.
const int Decision::departures[NUM_LINES] = { 0, 1 };
const int Decision::arrive[NUM_LINES] = { 0, 1 };
// Inport port for arriving customer that need to make a decision
const int Decision::decide = NUM_LINES;

Decision::Decision():
Atomic<IO_Type>()
{
    // Set the initial line lengths to zero
    for (int i = 0; i < NUM_LINES; i++)
    {
        line_length[i] = 0;
    }
}

void Decision::delta_int()
{
    // Move out all of the deciders
    deciding.clear();
}

void Decision::delta_ext(double e, const Bag<IO_Type>& x)
{
    // Assign new arrivals to a line and update the line length
    Bag<IO_Type>::const_iterator iter = x.begin();
    for (; iter != x.end(); iter++)
    {
        if ((*iter).port == decide)
        {
            int line_choice = find_shortest_line();
            Customer* customer = new Customer((*iter).value);
            pair<int, Customer*> p(line_choice, customer);
            deciding.push_back(p);
            line_length[p.first]++;
        }
    }
    // Decrement the length of lines that had customers leave
    for (int i = 0; i < NUM_LINES; i++)
    {
        iter = x.begin();
        for (; iter != x.end(); iter++)
        {
            if ((*iter).port < NUM_LINES)
            {
                line_length[(*iter).port]--;
            }
        }
    }
}

```

```

}

void Decision::delta_conf(const Bag<IO_Type>& x)
{
    delta_int();
    delta_ext(0.0,x);
}

double Decision::ta()
{
    // If there are customers getting into line, then produce output
    // immediately.
    if (!deciding.empty())
    {
        return 0.0;
    }
    // Otherwise, wait for another customer
    else
    {
        return DBL_MAX;
    }
}

void Decision::output_func(Bag<IO_Type>& y)
{
    // Send all customers to their lines
    list<pair<int, Customer*> >::iterator i = deciding.begin();
    for (; i != deciding.end(); i++)
    {
        IO_Type event((*i).first, (*i).second);
        y.insert(event);
    }
}

void Decision::gc_output(Bag<IO_Type>& g)
{
    Bag<IO_Type>::iterator iter = g.begin();
    for (; iter != g.end(); iter++)
    {
        delete (*iter).value;
    }
}

Decision::~Decision()
{
    clear_deciders();
}

void Decision::clear_deciders()
{
    list<pair<int, Customer*> >::iterator i = deciding.begin();
    for (; i != deciding.end(); i++)

```

```

    {
        delete (*i).second;
    }
    deciding.clear();
}

int Decision::find_shortest_line()
{
    int shortest = 0;
    for (int i = 0; i < NUM_LINES; i++)
    {
        if (line_length[shortest] > line_length[i])
        {
            shortest = i;
        }
    }
    return shortest;
}

```

The block diagram of the store with multiple clerks is shown in Fig. 5.6. The external interface for this block diagram model is identical to that of the clerk models (i.e., it has the same inputs and outputs as the **Clerk** and **Clerk2** models), and we can therefore use the generator and observer models to conduct the same experiments as before. The external “arrive” input of the multi-clerk model is connected to the “decide” input of the **Decision** model. The “depart” output ports of each of the **Clerk** models is connected to the external “arrive” output port of the multi-clerk model. The **Decision** model has two output ports, each one producing customers for a distinct clerk. These output ports are coupled to the “arrive” port of the appropriate clerk. The **Clerk**’s “depart” output ports are then coupled to the appropriate “departures” port of the decision model.




Figure 5.6: Component models and their interconnections in the multi-clerk convenience store model.

The multi-clerk model is implemented by deriving a new class from **Digraph**. The constructor of the new class creates and adds the component models and establishes their interconnections. Here is the header file for this new multi-clerk model.

```

#include "adevs.h"
#include "Clerk.h"
#include "Decision.h"

/**
A model of a store with multiple clerks and a "shortest line"
decision process for customers.
*/

```

```

class MultiClerk: public adevs::Digraph<Customer*>
{
    public:
        // Model input port
        static const int arrive;
        // Model output port
        static const int depart;
        // Constructor.
        MultiClerk();
        // Destructor.
        ~MultiClerk();
};

```

And here is the source file

```

#include "MultiClerk.h"
using namespace std;
using namespace adevs;

// Assign identifiers to I/O ports
const int MultiClerk::arrive = 0;
const int MultiClerk::depart = 1;

MultiClerk::MultiClerk():
Digraph<Customer*>()
{
    // Create and add component models
    Decision* d = new Decision();
    add(d);
    Clerk* c[NUM_LINES];
    for (int i = 0; i < NUM_LINES; i++)
    {
        c[i] = new Clerk();
        add(c[i]);
    }
    // Create model connections
    couple(this,this->arrive,d,d->decide);
    for (int i = 0; i < NUM_LINES; i++)
    {
        couple(d,d->arrive[i],c[i],c[i]->arrive);
        couple(c[i],c[i]->depart,d,d->departures[i]);
        couple(c[i],c[i]->depart,this,this->depart);
    }
}

MultiClerk::~MultiClerk()
{
}

```

Notice that the **MultiClerk** destructor does not delete its component models. This is because the components are adopted by the base class when they are added using the **Digraph**'s add method. Consequently, the component models are deleted by the base class destructor, rather than the destructor of the derived class.

## 5.5 Cell Space Models

A cell space model is a collection of atomic and network models arranged in a regular grid and with each model connected to its neighboring models. Conway's Game of Life is a classic example of a cell space model that can be described very nicely as a discrete event system. The game is played on a flat board divided into regular cells much like a checkerboard. Each cell has a neighborhood that consists of its eight adjacent cells: above, below, left, right, and the four corners. A cell can be dead or alive, and the switch from dead to alive and vice versa occurs according to two rules:

1. (Death rule). If a cell is alive and it has less than two or more than three living neighbors then the cell dies.
2. (Rebirth rule). If a cell is dead and it has three living neighbors then the cell is reborn.

Our implementation of the Game of Life has two parts: the atomic models that implement the individual cells and the **CellSpace** that contains the cells. The **CellSpace** is a type of **Network**, and its components exchange **CellEvent** objects that have four attributes: the x, y, and z coordinates of the target cell (the cell space can have three dimensions; the game of life uses just two) and the object to deliver to that target. The **CellEvent** class is a template class whose template argument sets the type of object that the event delivers. The size of the **CellSpace** is determined when the **CellSpace** object is created, and it has methods for adding and retrieving cells by location.

The **Atomic** cells in our Game of Life have two state variables: one is the dead or alive status of the cell and the other is its count of living neighbors. Two methods are implemented to test the death and rebirth rules, and the cell sets its time advance to 1 whenever a rule is satisfied. The cell's output is its new dead or alive state. External events update the cell's count of living neighbors. In order to produce properly targeted **CellEvents**, each cell also keeps track of its own location in the cell space. In the example code, the cell space is rendered graphically using OpenGL, but I'll omit that part. Here is the abbreviated header file for our Game of Life cell.

```
/// Possible cell phases
typedef enum { Dead, Alive } Phase;
/// IO type for a cell
typedef adevs::CellEvent<Phase> CellEvent;

/// A cell in the Game of Life.
class Cell: public adevs::Atomic<CellEvent> {
public:
    /**
     * Create a cell and set the initial state.
     * The width and height fields are used to determine if a
     * cell is an edge cell. The last phase pointer is used to
     * visualize the cell space.
     */
    Cell(long int x, long int y, long int width, long int height,
         Phase phase, short int nalive, Phase* vis_phase = NULL);

    ... Required Adevs methods and destructor ...

private:
    // location of the cell in the 2D space
    long int x, y;
    // dimensions of the 2D space
    static long int w, h;
    // Current cell phase
```



```

    Phase phase;
    // number of living neighbors.
    short int nalive;
    // Output variable for visualization
    Phase* vis_phase;

    // Returns true if the cell will be born
    bool check_born_rule() const {
        return (phase == Dead && nalive == 3);
    }
    // Return true if the cell will die
    bool check_death_rule() const {
        return (phase == Alive && (nalive < 2 || nalive > 3));
    }
};

```

The template argument supplied to the base **Atomic** class is a **CellEvent** whose value attribute has the type **Phase**. The `check_born_rule` method tests the rebirth condition and the `check_death_rule` method tests the death condition. The appropriate rule, as determined by the cell's dead or alive status, is used in the time advance, output, and internal transition methods (i.e., if the cell is dead then check the rebirth rule; if alive, check the death rule). The number of living cells is updated by the cell's *delta\_ext* method when neighboring cells report a change in their status. Here are the **Cell**'s method implementations.

```

Cell::Cell(long int x, long int y, long int w, long int h,
Phase phase, short int nalive, Phase* vis_phase):
adevs::Atomic<CellEvent>(x(x),y(y),phase(phase),nalive(nalive),vis_phase(vis_phase)) {
    // Set the global cellspace dimensions
    Cell::w = w; Cell::h = h;
    // Set the initial visualization value
    if (vis_phase != NULL) *vis_phase = phase;
}

double Cell::ta() {
    // If a phase change should occur then change state
    if (check_death_rule() || check_born_rule()) return 1.0;
    // Otherwise, do nothing
    return DBL_MAX;
}

void Cell::delta_int() {
    // Change the cell state if necessary
    if (check_death_rule()) phase = Dead;
    else if (check_born_rule()) phase = Alive;
}

void Cell::delta_ext(double e, const adevs::Bag<CellEvent>& xb) {
    // Update the living neighbor count
    adevs::Bag<CellEvent>::const_iterator iter;
    for (iter = xb.begin(); iter != xb.end(); iter++) {
        if ((*iter).value == Dead) nalive--;
        else nalive++;
    }
}

```

```

void Cell::delta_conf(const adevs::Bag<CellEvent>& xb) {
    delta_int();
    delta_ext(0.0,xb);
}

void Cell::output_func(adevs::Bag<CellEvent>& yb) {
    CellEvent e;
    // Assume we are dying
    e.value = Dead;
    // Check in case this is not true
    if (check_born_rule()) e.value = Alive;
    // Set the visualization value
    if (vis_phase != NULL) *vis_phase = e.value;
    // Generate an event for each neighbor
    for (long int dx = -1; dx <= 1; dx++) {
        for (long int dy = -1; dy <= 1; dy++) {
            e.x = (x+dx)%w;
            e.y = (y+dy)%h;
            if (e.x < 0) e.x = w-1;
            if (e.y < 0) e.y = h-1;
            // Don't send to self
            if (e.x != x || e.y != y)
                yb.insert(e);
        }
    }
}

```

The *output\_func* method shows how a cell sends messages to its neighbors. The double for loop creates a **CellEvent** targeted at each adjacent cell. The location of the target cell is written to the x, y, and z attributes of the **CellEvent** object. Just like arrays, the locations can range from zero to the cell space's size minus one. The **CellSpace** routes the **CellEvents** to their targets. However, if the target of the **CellEvent** is outside of the cell space, then the **CellSpace** itself will produce the **CellEvent** as an output.

The remainder of the simulation program looks very much like the simulation programs that we've seen so far (except for some OpenGL specific code, omitted here, that is used to display the cells). A **CellSpace** object is created and we add the cells to it. Then a **Simulator** object is created and a pointer to the **CellSpace** is passed to the **Simulator**'s constructor. Last, we execute events until our stopping criteria is met. The execution part is already familiar, so let's just focus on creating the **CellSpace**. Here is the code snippet that performs the construction.

```

// Create the cellspace model
cell_space = new adevs::CellSpace<Phase>(WIDTH,HEIGHT);
for (int x = 0; x < WIDTH; x++) {
    for (int y = 0; y < HEIGHT; y++) {
        // Count the living neighbors
        short int nalive = count_living_cells(x,y);
        // The 2D phase array contains the initial Dead/Alive state of each cell
        cell_space->add(new Cell(x,y,WIDTH,HEIGHT,phase[x][y],nalive,&(phase[x][y])),x,y);
    }
}

```

Just as with the **Digraph** class, the **CellSpace** template argument determines the value type for the **CellEvents** used as input and output by the component models. The **CellSpace** constructor sets the

dimensions of the space. Every **CellSpace** is three dimensional, and the constructor accepts three arguments that set its x, y, and z dimensions; omitted arguments default to 1. The constructor signature is

```
CellSpace(long int width, long int height = 1, long int depth = 1)
```

Components are added to the cellspace with the **add** method. This method places a component at a specific x,y,z location. Its signature is

```
void add(Cell* model, long int x, long int y = 0, long int z = 0)
```

where **Cell** is a **Devs** (atomic or network) by the type definition

```
typedef Devs<CellEvent<X> > Cell;
```

Also like the **Digraph**, the **CellSpace** deletes its components when it is deleted.

The **CellSpace** has five other methods for retrieving cells and for getting the dimensionality of the cell space. These are more or less self-explanatory; the signatures are shown below.

```
const Cell* getModel(long int x, long int y = 0, long int z = 0) const;
Cell* getModel(long int x, long int y = 0, long int z = 0);
long int getWidth() const;
long int getHeight() const;
long int getDepth() const;
```

The Game of Life produces a surprising number of clearly recognizable patterns. Some of these patterns are fixed and unchanging; others oscillate, cycling through a set of patterns that always repeats itself; others seem to crawl or fly. One common pattern is the Block, which is shown in Fig. 5.7. Our discrete event implementation of the Game of Life doesn't do any work when simulating a Block. None of the cells in a Block change in any way: their phases are constant and so are their neighbor counts. The Blinker shown in

	\$	\$	
	\$	\$	

Figure 5.7: The Block.

Fig. 5.8 is more interesting. This oscillating pattern has just two stages: a vertical and a horizontal. Table 5.1 shows the input, output, and state transitions that are computed for the cell marked with a \* in Fig. 5.8. Just like the pattern it is a part of, the cells oscillates between two different states.

Time	State	Input	Output to all neighbors
0	(dead,3)	No input	No Output
1	(alive,1)	(dead,2,1,0) (dead,2,3,0)	alive
2	(dead,1)	(alive,2,1,0) (alive,2,3,0)	dead

Table 5.1: State, input, and output trajectory for the cell marked with \* in Fig. 5.8.

The confluent transition function plays a major role in the Blinker simulation. Most of the rows in Table 5.1 (all but the first row, in fact) have simultaneous input and output, which means that an internal and external event coincide. Consequently, the next state of the cell is determined by its **delta\_conf** method. It is also important that the input and output bags carry multiple values. The external transition function




Figure 5.8: The Blinker. The input, output, and state transitions for the cell marked with a \* are shown in Table 5.1. The address of each cell is shown in its upper left corner. Living cells are indicated with a \$.

(which is used in defining the confluent transition function) must be able to compute the number of living neighbors before determining its next state. If input events were provided one at a time (e.g., if the input bag were replaced by a single input event), then our discrete event Game of Life would be much more difficult to implement.

## Chapter 6

# Variable Structure Models

The composition of a variable (sometimes called dynamic) structure model changes through time. New components are added as, for example, machinery is installed in a factory, organisms reproduce, or shells are fired from a cannon. Existing components are removed as machines break, organisms dies, or shells in flight find their targets. Components are rearranged as, for example, parts move through a manufacturing process, organisms migrate, or a command and control network loses communication lines.

For modeling systems with a dynamic structure Adevs provides a simple but effective mechanism to coordinate changes in structure and changes of state. This mechanism is based on the Dynamic DEVS modeling formalism described in A.M. Uhrmacher's paper "Dynamic structures in modeling and simulation: a reflective approach", ACM Transactions on Modeling and Computer Simulation (TOMACS), Volume 11, Issue 2, pgs. 202-232, April 2001.

### 6.1 Building and Simulating Variable Structure Models

Every Adevs model, **Network** and **Atomic**, has a virtual method called *model\_transition*. This method is inherited from the **Devs** class that is at the top of the Adevs class hierarchy. The signature of the *model\_transition* method is

```
bool model_transition()
```

and its default implementation simply returns false.

At the end of every simulation cycle (that is, after computing the models' new states but prior to the garbage collection step) the simulator invokes the *model\_transition* method of every **Atomic** model that changed state in that cycle. When the *model\_transition* method is invoked the **Atomic** model can do almost anything it likes except alter the component set of a **Network** model.

If a model's *model\_transition* method returns true, then the simulator also calls the *model\_transition* method of that model's parent. The parent is, of course, a **Network** model; its *model\_transition* method may add, remove, and rearrange the network's components, but it must not delete any components! The simulator will automatically delete components removed from the model when the structure change calculations are finished.

As before, if the **Network**'s *model\_transition* method returns true then the simulator invokes the *model\_transition* method of the **Network**'s parent. Note, however, that the *model\_transition* method of any model is invoked at most once in each simulation cycle. This invocation, if it occurs, takes place after every component of the network qualifying for the evaluation of its *model\_transition* method has been activated.

After invoking every eligible model's *model\_transition* method, the simulator performs a somewhat complicated cleanup process. This process requires that simulator construct two sets. The first set contains all of the components that belonged to all of the **Network** models whose *model\_transition* method was

invoked and all of the components belonging to models that are in this set. The second set is defined in the same way, but it is computed using sets of components as they are after the *model\_transition* methods have been invoked.

The simulator deletes every model that has actually been removed; these are the models in the first set but not in the second. The simulator initializes every model that is genuinely new by computing its next event time (i.e., its creation time plus its time advance) and putting it into the event schedule; these are the models in the second set but not in the first. The simulator leaves all other models alone.

The procedure for calculating a change of structure can be summarized as follows:

1. Calculate the *model\_transition* method of every atomic model that changed state.
2. Construct the set of network models that contain an atomic model from step 1 whose *model\_transition* method returned true. These network models are sort by their depth in the tree with the bottom-most models first and top-most last. This will ensure that structure changes are calculated from the bottom up.
3. Calculate the *model\_transition* methods of the networks in order. On completing each transition, do the following:
  - (a) Remove the network from the list.
  - (b) If the network's *model\_transition* method returns true, put the parent of the network into the sorted list of networks from step 2. This ensures that a network's *model\_transition* method is invoked only after all of its eligible component's have had their *model\_transition* methods invoked.
4. When there are no more networks in the list, do the following:
  - (a) Delete the components removed from the model (i.e., the models without a parent in the tree).
  - (b) Initialize the components add to the model.

The complete procedure is illustrated in Fig. 6.1. The black models' *model\_transition* methods returned true. The set of components examined before and after the structure change are listed above the before (left) and after (right) trees. Notice that these models are in the sub-tree below the model *C*, which is the top-most model in that sub-tree that returns false from its *model\_transition* method. Also note that while the leaves of the tree may have had their model transition method's invoked, none returns true and so their parents' model transition methods are not invoked nor are the sets of components considered when determining what models have been added and removed from the model. The set of deleted components is  $\{c, D, d, e, f\} - \{e, g, d\} = \{c, D, f\}$ . The set of new components is  $\{e, g, d\} - \{c, D, d, e, f\} = \{g\}$ .

The *model\_transition* method can break the strict hierarchy and modularity that is usually observed when building **Network** models. Any **Network** model can in principle modify the component set of any other model, regardless of proximity or hierarchy. The potential for anarchy is great, and so the design of a variable structure model should be carefully considered. There are two approaches to such a design that are simple and, in many cases, entirely adequate.

The first approach is to allow only **Network** models to effect structure changes and to restrict those changes to the **Network**'s immediate sub-components. With this approach, an **Atomic** model initiates a structure change by posting a structure change request for its parent. The **Atomic** model's *model\_transition* method returns true causing its parent's *model\_transition* method to be invoked. The parent **Network** model then retrieves and acts on the request posted by its component. The **Network** repeats this process if it wants to effect structure changes involving models other than its immediate children; i.e., it posts a request for its parent and returns true from its *model\_transition* method.

The second approach allows arbitrary structure changes by forcing the model at the very top of the hierarchy to invoke its *model\_transition* method. This causes the simulator to consider every model in the aftermath of a structure change. As in the first approach, an **Atomic** model that wants to effect a structure change uses its *model\_transition* method to post a request for its parent. This request is percolated up the model hierarchy by the **Network** models whose *model\_transition* methods always return true.




Figure 6.1: Illustration of a change of structure in a variable structure model.

The first approach trades flexibility for execution time; the second approach trades execution time for flexibility. With the first approach, structure changes that involve a small number of components require a small amount of work by the simulator. The scope of change must, however, be carefully restricted. With the second approach, every structure change requires the simulator to include every part of the model in its calculations, regardless of the actual extent of the change in structure. In this case, however, the scope of a structure change may be unlimited.

## 6.2 A Variable Structure Example

The Custom Widget Company is expanding its operations. Plans are being drawn for a new factory that will make custom gizmos (and the company name will be changed to The Custom Widget and Gizmo Company). The machines for the factory are expensive to operate. To keep costs down, the factory will operate just enough machinery to fill outstanding orders for gizmos. The factory must have enough machinery to meet peak demand, but much of the machinery will be idle much of the time. The factory engineers want to answer two questions: how many machines are needed and how much will it cost to operate the them.

We are going to use a variable structure model to answer these two questions. The model has three components: a generator that creates orders for gizmos, a model of a machine, and a model of the factory that contains the machines and activates and deactivates them as required to satisfy demand. The complete model of the factory is illustrated in Fig. 6.2.

The generator creates new orders for the factory. Each order is identified with its own integer label, and the generator produces orders at the rate anticipated by the factory engineers. Demand at the factory is expected to be fairly steady with a new order arriving every 1/2 to 2 days; demand is modeled with a random variable uniformly distributed in the range [0.5,2]. Here is the code for the generator:

```
#include "adevs.h"
// The Genr models factory demand. It creates new orders every 0.5 to 2 days.
class Genr: public adevs::Atomic<int>
{
public:
    /**
     * The generator requires a seed for the random number that determines
     * the time between new orders.
     */
    Genr(unsigned long seed):
```




Figure 6.2: Block diagram of the variable structure factory model. The broken lines indicate structural elements that are subject to change.

```

        adevs::Atomic<int>(),next(1),u(seed){ set_time_to_order(); }
    // Internal transition updates the order counter and
// determines the next arrival time
    void delta_int() { next++; set_time_to_order(); }
    // Output function produces the next order
    void output_func(adevs::Bag<int>& yb) { yb.insert(next); }
    // Time advance returns the time until the next order
    double ta() { return time_to_order; }
    // Model is input free, so these methods are empty
    void delta_ext(double,const adevs::Bag<int>&){}
    void delta_conf(const adevs::Bag<int>&){}
    // No explicit memory management is needed
    void gc_output(adevs::Bag<int>&){}
private:
    // Next order ID
    int next;
    // Time until that order arrives
    double time_to_order;
    // Random variable for producing order arrival times
    adevs::rv u;
    // Method to set the order time
    void set_time_to_order() { time_to_order = u.uniform(0.5,2.0); }
};

```

The model of a machine is similar to the **Clerk** in section 3. A machine requires 3 days to make a gizmo and orders for gizmos are processed first come, first serve. The **Machine**'s *model\_transition* method is inherited from its **Atomic** base class. I'll discuss the role of the *model\_transition* method after introducing the **Factory** class; here is the code for the **Machine**.

```

#include "adevs.h"
#include <cassert>
#include <deque>
/**
 * This class models a machine as a fifo queue and server with fixed service time.
 * The model_transition method is used, in conjunction with the Factory model_transition

```



```

* method, to add and remove machines as needed to satisfy a 6 day turnaround time
* for orders.
*/
class Machine: public adevs::Atomic<int>
{
public:
    Machine():adevs::Atomic<int>(),tleft(DBL_MAX){}
    void delta_int()
    {
        q.pop_front(); // Remove the completed job
        if (q.empty()) tleft = DBL_MAX; // Is the Machine idle?
        else tleft = 3.0; // Or is it still working?
    }
    void delta_ext(double e, const adevs::Bag<int>& xb)
    {
        // Update the remaining time if the machine is working
        if (!q.empty()) tleft -= e;
        // Put new orders into the queue
        adevs::Bag<int>::const_iterator iter = xb.begin();
        for (; iter != xb.end(); iter++)
        {
            // If the machine is idle then set the service time
            if (q.empty()) tleft = 3.0;
            // Put the order into the back of the queue
            q.push_back(*iter);
        }
    }
    void delta_conf(const adevs::Bag<int>& xb)
    {
        delta_int();
        delta_ext(0.0,xb);
    }
    void output_func(adevs::Bag<int>& yb)
    {
        // Expel the completed order
        yb.insert(q.front());
    }
    double ta()
    {
        return tleft;
    }
    // The model transition function returns true if another order can not
    // be accommodated or if the machine is idle.
    bool model_transition()
    {
        // Check that the queue size is legal
        assert(q.size() <= 2);
        // Return the idle or full status
        return (q.size() == 0 || q.size() == 2);
    }
    // Get the number of orders in the queue
    unsigned int getQueueSize() const { return q.size(); }

```

```

    // No garbage collection
    void gc_output(a devs::Bag<int>&){}
private:
    // Queue for orders that are waiting to be processed
    std::deque<int> q;
    // Time remaining on the order at the front of the queue
    double tleft;
};

```

The number of **Machine** models contained in the **Factory** model at any time is determined by the current demand for gizmos. The real factory, of course, will have a fixed number of physical machines on the factory floor, but the planners do not yet know how many machines are needed. A variable structure model that creates and destroys machines as needed is a good way to accommodate this uncertainty.

The Custom Widget and Gizmo Company has built its reputation on its guaranteed service time, from order to delivery, of 15 days. This leaves only 6 days for the manufacturing process, the remaining time being consumed by order processing, delivery, etc.

A single machine can meet this schedule if it has at most one order waiting in its queue. However, it costs a dollar a day to operate a machine and so the factory engineers want to minimize the number of machines working at any time. To accomplish this goal, the factory's operating policy has two rules:

1. Assign incoming orders to the active machine that can provide the shortest turn around time and
2. keep just enough active machines to have capacity for one additional order.

The **Factory** model implements this policy in the following way. If a **Machine** becomes idle or its queue is full (i.e., the machine is working on one order and has another order waiting in its queue), then that machine's *model\_transition* method returns true. This causes the **Factory**'s *model\_transition* method to be invoked. The **Factory** first looks for and removes machines that have no work. Then it examines each remaining machine to determine if the required one unit of additional capacity is available. If the required unit of additional capacity is not available then the **Factory** creates a new machine.

This is an example of the first approach to building a variable structure model. With this design, the simulator's structure calculations are done only when the **Factory**'s *model\_transition* method is invoked, and are therefore limited to instants when **Machine** models are likely to be created or destroyed. Our design, however, is complicated somewhat by the need for **Machine** and **Factory** objects to communicate; i.e., the **Machines** must watch their own status and inform the **Factory** when there is a potential shortage of capacity.

If we had used the second approach to build our variable structure model, then the **Machines**' *model\_transition* methods could have simply returned true; no need for a status check. The **Factory** would iterate through its list of **Machines**, adding and deleting **Machines** as needed. This is more computationally expensive; the simulator looks for changes in the **Factory**'s set of components at each simulation cycle. However, the design of the model is simpler, albeit only marginally so in this instance.

The **Factory** is a **Network** model and must implement all of the **Network**'s virtual methods: *route*, *getComponents*, and *model\_transition*. The *route* method is responsible for assigning orders to the proper **Machine**. When an order arrives, it is sent to the machine that will most quickly satisfy the order. The *getComponents* method puts the current set of machines into the **Set** c of components. The *model\_transition* method examines the status of each machine, deleting idle machines and adding new machines if they are needed to maintain reserve capacity. The complete **Factory** implementation is shown below.

```

#include "adevs.h"
#include "Machine.h"
#include <list>

class Factory: public adevs::Network<int> {

```

```

public:
    Factory();
    void getComponents(a devs::Set<a devs::Devs<int>*>& c);
    void route(const int& order, a devs::Devs<int>* src,
               a devs::Bag<a devs::Event<int> >& r);
    bool model_transition();
    ~Factory();
    // Get the number of machines
    int getMachineCount();
private:
    // This is the machine set
    std::list<Machine*> machines;
    // Method for adding a machine to the factory
    void add_machine();
    // Compute time needed for a machine to finish a new job
    double compute_service_time(Machine* m);
};

#include "Factory.h"
using namespace a devs;
using namespace std;

Factory::Factory():
Network<int>() { // call the parent constructor
    add_machine(); // Add the first machine the the machine set
}

void Factory::getComponents(Set<Devs<int>*>& c) {
    // Copy the machine set to c
    list<Machine*>::iterator iter;
    for (iter = machines.begin(); iter != machines.end(); iter++)
        c.insert(*iter);
}

void Factory::route(const int& order, Devs<int>* src, Bag<Event<int> >& r) {
    // If this is a machine output, then it leaves the factory
    if (src != this) {
        r.insert(Event<int>(this,order));
        return;
    }
    // Otherwise, use the machine that can most quickly fill the order
    Machine* pick = NULL; // No machine
    double pick_time = DBL_MAX; // Infinite time for service
    list<Machine*>::iterator iter;
    for (iter = machines.begin(); iter != machines.end(); iter++) {
        // If the machine is available
        if ((*iter)->getQueueSize() <= 1) {
            double candidate_time = compute_service_time(*iter);
            // If the candidate service time is smaller than the pick service time
            if (candidate_time < pick_time) {
                pick_time = candidate_time;
                pick = *iter;
            }
        }
    }
}

```

```

        }
    }
}
// Make sure we found a machine with a small enough service time
assert(pick != NULL && pick_time <= 6.0);
// Use this machine to process the order
r.insert(Event<int>(pick,order));
}

bool Factory::model_transition() {
    // Remove idle machines
    list<Machine*>::iterator iter = machines.begin();
    while (iter != machines.end()) {
        if ((*iter)->getQueueSize() == 0) iter = machines.erase(iter);
        else iter++;
    }
    // Add the new machine if we need it
    int spare_cap = 0;
    for (iter = machines.begin(); iter != machines.end(); iter++)
        spare_cap += 2 - (*iter)->getQueueSize();
    if (spare_cap == 0) add_machine();
    return false;
}

void Factory::add_machine() {
    machines.push_back(new Machine());
    machines.back()->setParent(this);
}

double Factory::compute_service_time(Machine* m) {
    // If the machine is already working
    if (m->ta() < DBL_MAX) return 3.0+(m->getQueueSize()-1)*3.0+m->ta();
    // Otherwise it is idle
    else return 3.0;
}

int Factory::getMachineCount() {
    return machines.size();
}

Factory::~Factory() {
    // Delete all of the machines
    list<Machine*>::iterator iter;
    for (iter = machines.begin(); iter != machines.end(); iter++)
        delete *iter;
}

```

To illustrate how the *model\_transition* method works, let us manually simulate the processing of a few orders: the first order arrives at day zero, the second order at day one, and the third order at day three. At the start, on day zero, there is one idle **Machine**. When the first order arrives the **Factory**'s *route* method is invoked, and it sends the order to the idle **Machine**. The **Machine**'s *delta\_ext* method is invoked next, and it begins processing the order. Following this, the **Machine**'s *model\_transition* method is invoked, it

discovers that the **Machine** is working and has space in its queue, and returns false.

When the second order arrives on day one, the **Factory**'s *route* method is called again. There is only one **Machine** and it has space in its queue so the order is sent to that **Machine**. The **Machine**'s *delta\_ext* method is invoked next and it queues the order. The **Machine**'s *model\_transition* method is now invoked; its queue is full and so the method returns true. This causes the the **Factory**'s *model\_transition* method to be invoked; it examines the **Machine**'s status, sees that it overloaded, and creates a new **Machine**.

At this time, the working **Machine** needs two more days to finish the first order, and it will not complete its second order until a total of five days have elapsed.

There is a great deal of activity when the third order arrives on day three. First, the working **Machine**'s *output\_func* method is called, and it spits out the first completed order (i.e., the order begun on day zero). Next the **Factory**'s *route* method is called twice. In the first call, it converts the **Machine**'s output into an output from the **Factory**. In the second call, it routes the new order to the idle **Machine** (these calls to route could occur in either order).

The input and output events now routed, the state transition methods for the two **Machines** are invoked. The working **Machine**'s *delta\_int* method is called and it starts work on its queued order. The idle **Machine**'s *delta\_ext* method is called and it begins processing the new order. Last the *model\_transition* methods of both **Machines** are invoked; both **Machine**'s have room in their queue and so both methods return false.

For the sake of illustration, suppose no orders arrive in the next three days (this is impossible when orders arrive every one-half to two days, but bear with me). At day six, both machines finish their orders. The **Machines**' *output\_func* methods are invoked, producing the finished orders. These become output from the **Factory** via the **Factory**'s *route* method.

Next, the **Machines**' *delta\_int* methods are called and both **Machines** become idle. After this, the **Machines**' *model\_transition* methods are invoked and these return true because the machines are idle. This causes the **Factory**'s *model\_transition* method to be called. It examines the status of each **Machine**, sees that they are idle, and deletes both of them. Last the **Factory** computes its available capacity, which is now zero, and create a new machine. Incidentally, this returns the **Factory** to its initial configuration with one idle **Machine**.

The factory engineers have two questions: how many machines are needed and what is the factory's annual operating cost. These questions can be answered with a plot of the count of active machines versus time. The required number of machines is the maximum value of the active machine count. Each machine costs a dollar per day to operate, and so the operating cost is just the one year time integral of the active machine count. One plot of the active machine count versus time is shown in Fig. 6.3. The maximum count of active machines 4 and the annual operating cost is \$944 (this plot is from the first simulation run listed in Table 6.1).

New orders arrive at random and so the annual operating cost and maximum machine count are themselves random numbers. Consequently, data from several simulation runs are needed to make an informed decision. Somewhat arbitrarily, I have listed ten simulation runs; each uses a different sequence of random numbers and therefore produces a different outcome (i.e., another sample of the maximum active machine count and annual operating cost). The maximum active machine count and annual operating cost generated by each run is shown in Table 6.1. From this data, the factory engineers conclude that 4 machines are required and the average annual operating cost will be \$961.

Random number seed	Maximum machine count	Annual operating cost
1	4	\$944.05
234	4	\$968.58
15667	4	\$980.96
999	3	\$933.13
9090133	4	\$961.65
6113	4	\$977.33

Table 6.1: Outcomes of ten factory simulations.




Figure 6.3: Active machine count over one year.

## Chapter 7

# Continuous Models

Many engineered systems consist of computers, communications networks, and other digital (i.e., discrete event) systems whose purpose is to monitor and control a physical (electrical, mechanical, thermodynamic, etc.) process. Models of these systems have parts modeled as discrete event systems, parts modeled with continuous (differential or differential-algebraic) equations, and the interaction of these parts is crucial to understanding the system's behavior. Where the continuous models interact with the discrete event models, these interactions are necessarily discrete. For example, a digital thermometer reports temperature in discrete increments, electrical switches are either open or closed, a threshold sensor is either tripped or it is not. Discrete interactions in a combined continuous-discrete event simulation are managed just as before; the models interact by producing output events and reacting to input events.

If, on the other hand, two systems interact continuously, then those interacting parts at least are modeled with continuous equations. In this case, accurate simulation is greatly facilitated by lumping the two systems into a single assembly; in Adevs this assembly is an **Atomic** model that encapsulates the system's continuous dynamics. The essence of the approach to combined simulation in Adevs consists therefore of building atomic models that i) approximate the behavior of the continuous systems and ii) generates and consumes events at those instants when the continuous system interacts with a discrete event one.

There are three possibly outcomes of this lumping process. One possibility is that we end up with a single assembly; in this case our model is essentially continuous and we are probably better off using a simulation tool for continuous systems. At the other extreme, we find that the continuous parts of our model are very simple, yielding to analytical solutions that are easily transformed into discrete event models. Between these two extremes are models with continuous dynamics that are not simple but also do not dominate the modeling problem. The continuous system simulation part of Adevs is aimed at this type of model.

### 7.1 Differential equation modeling with the `ode_system` class

Models described by ordinary differential equations are implemented by sub-classing the **ode\_system** class. This class has two sets of methods: the first is for the model's continuous dynamics and the second is for the model's discrete event dynamics. I'll illustrate both with the simple, if somewhat contrived, example of a cherry bomb<sup>1</sup>. This bomb is dropped from a height of 1 meter and bounces until it either explodes or is doused with water. We'll assume that the cherry bomb only bounces up and down and is perfectly elastic. The cherry bomb explodes 2 seconds from the time it is lit and dropped unless doused first. Dousing the cherry bomb puts out the fuse<sup>2</sup>. Dousing is a discrete input event and the cherry bomb produces a discrete output event if it explodes.

This model has two continuous state variables: the height and velocity of the cherry bomb. Between

---

<sup>1</sup>A cherry bomb is a small red firecracker. They are dangerous, and illegal in the United States. Nonetheless, every school seems to have at least one obnoxious kid who likes to put them into toilets.

<sup>2</sup>Cherry bomb fuses are frequently water proofed.

events, these variables are governed by the pair of differential equations

$$\dot{v} = -9.8 \quad (7.1)$$

$$\dot{h} = v \quad (7.2)$$

where 9.8 meters per second per second is acceleration due to gravity,  $v$  is velocity, and  $h$  is height. In this example, it is also useful to know the current time. We keep track of this by adding one more differential equation

$$\dot{t} = 1 \quad (7.3)$$

whose solution is  $t_0 + t$  or just  $t$  if we set  $t_0 = 0$ . The ball bounces when it hits the floor, and bouncing causes the ball's velocity to change sign; specifically

$$h = 0 \ \& \ v < 0 \implies v \leftarrow -v \quad (7.4)$$

where  $\implies$  is logical implication and  $\leftarrow$  indicates an assignment.

Equations 7.1, 7.2, and 7.3 are the state variable derivatives, and these equations are implemented in the **der\_func** method of the **ode\_system** class. The signature for this method is

```
void der_func(const double* q, double* dq)
```

The  $q$  pointer is the array of state variable values; in this case  $h$ ,  $v$ , and  $t$ . The  $dq$  pointer is the array of state variable derivatives; in this case  $\dot{h}$ ,  $\dot{v}$ , and  $\dot{t}$ . When the simulator calls the **der\_func** method, it supplies  $q$ . In response, the method computes the values of  $\dot{h}$ ,  $\dot{q}$ , and  $\dot{t}$  and stores them in the  $dq$  array.

Equation 7.4 is a state event condition and it is implemented in two parts. The **state\_event\_func** method implements the ‘if’ part (left hand side) of the condition. The signature of this method is

```
void state_event_func(const double *q, double *z)
```

Again, the supplied  $q$  array contains the current state variable values. These are used to evaluate the state event condition and store the result in the  $z$  array. The simulator detects state events by looking for changes in the sign of the  $z$  array entries. Note that the event condition should be continuous in the state variables on which it depends. In the case of the cherry bomb this is simple to do: we simply use  $z = h$  if  $v < 0$  and  $z = 1$  if  $v \geq 0$ .

The ‘then’ part (right hand side) is implemented with the **internal\_event** method, which the simulator invokes when the state event condition is true. The signature of this method is

```
void internal_event(double *q, const bool *state_event)
```

where, again,  $q$  is the value of the state variables at the event. The entries of the array **state\_event** are true for each  $z$  in the state event condition array that evaluates to zero. This array therefore has one entry for each state event condition, and it has one additional entry to indicate time events, which are described below.

The cherry bomb has one discrete state variable with three possible values: the fuse is lit, the fuse is not lit, and the bomb is exploded. This variable changes in response to two events. The first event is when the bomb explodes; this is a time event that we know will occur 2 seconds from the time that the fuse is lit. The **time\_event\_func** method is used to schedule the explosion by returning the time remaining until the fuse burns out. The signature of the of this method is

```
double time_event_func(const double* q)
```

As before,  $q$  is the current value of the state variables. The **time\_event\_func** is similar to the familiar **ta** method; it is used to schedule autonomous events based on the current value of the model's state variables. When this time expires, the simulator calls the **internal\_event** method with the last flag in the state event array set to true.

The second event that can change the state of the fuse is dousing with water. This is an external event. External events, of course, are not scheduled by the model itself; they occur when and if the input event arrives. The **external\_event** method implements the response of the cherry bomb to dousing with water. Its signature is



```
void external_event(double *q, double e, const Bag<X> &xb)
```

The array `q` contains the values of the continuous state variables, `e` is the time since the last discrete event, and `xb` is the bag of input. The discrete event is an input and it appears in the input bag `xb` if the event occurs.

As before, it is possible for an external and internal event to coincide. When this happens, the simulator calls the method *confluent\_event*. Its signature is

```
void confluent_event (double *q, const bool *state_event, const Bag<X> &xb)
```

and its arguments are as described for the internal and external methods.

The cherry bomb model produces an output event when it explodes, and the *output\_func* method is used for this purpose. Its signature is

```
void output_func(const double *q, const bool *state_event, Bag<X> &yb)
```

Its `q` and `state_event` arguments are as described for the `internal_event` method, and the bag `yb` is to be filled with the model's output. As with an **Atomic** model, the `output_func` is always invoked immediately prior to the `internal_event` and `confluent_event` methods.

All that remains in the implementation is to supply a *gc\_output* for collecting garbage, a constructor, and a method for initializing the continuous state variables. The `gc_output` method works identically to that of the **Atomic** class and so needs no more discussion here. The constructor for the cherry bomb must call the constructor of its **ode\_system** base class. The signature of this method is

```
ode_system (int N_vars, int M_event_funcs)
```

where `N_vars` is the number of entries in the `q` and `dq` arrays (i.e., the number of continuous state variables) and `M_event_funcs` is the number of entries in the `z` and `state_event` arrays (plus one for the time event). For the cherry bomb, `N_vars` is three and `M_event_funcs` is one.

The constructor does not initialize the continuous state variables. Instead, the simulator calls the `init` method whose signature is

```
void init(double* q)
```

and where `q` is an array that should be filled with the initial values for, in the case of the cherry bomb,  $h$ ,  $v$ , and  $t$ . The complete implementation of the **CherryBomb** is listed below.

```
#include "adevs.h"
#include <iostream>
using namespace std;
using namespace adevs;

// Array indices for the CherryBomb state variables
#define H 0
#define V 1
#define T 2
// Discrete variable enumeration for the CherryBomb
typedef enum { FUSE_LIT, DOUSE, EXPLODE } Phase;

class CherryBomb: public ode_system<string> {
public:
    CherryBomb():ode_system<string>(
        3, // three state variables including time
        1 // 1 state event condition
    ) {
        phase = FUSE_LIT; // Light the fuse!
```

```

}
void init(double *q) {
    q[H] = 1.0; // Initial height
    q[V] = 0.0; // Initial velocity
    q[T] = 0.0; // Start time at zero
}
void der_func(const double* q, double* dq) {
    dq[V] = -9.8;
    dq[H] = q[V];
    dq[T] = 1.0;
}
void state_event_func(const double* q, double *z) {
    // Test for hitting the ground.
    if (q[V] < 0.0) z[0] = q[H];
    else z[0] = 1.0;
}
double time_event_func(const double* q) {
    if (q[T] < 2.0) return 2.0 - q[T]; // Explode at time 2
    else return DBL_MAX; // Don't do anything after that
}
void external_event(double* q, double e, const Bag<string>& xb) {
    phase = DOUSE; // Any input is a douse event
}
void internal_event(double* q, const bool* state_event) {
    if (state_event[0]) q[V] = -q[V]; // Bounce!
    if (state_event[1]) phase = EXPLODE;
}
void confluent_event(double* q, const bool* state_event,
    const Bag<string>& xb) {
    internal_event(q, state_event);
    external_event(q, 0.0, xb);
}
void output_func(const double *q, const bool* state_event,
    Bag<string>& yb) {
    if (state_event[1] && phase == FUSE_LIT)
        yb.insert("BOOM!"); // Explode!
}
void postStep(const double* q) {
    // Write the current state to std out
    cout << q[T] << " " << q[H] << " " << q[V] << " " << phase << endl;
}
// No garbage collection is needed
void gc_output(Bag<string>&){}
// Get the current value of the discrete variable
Phase getPhase() { return phase; }
private:
    Phase phase;
};

```

The **CherryBomb** itself is not derived from **Atomic** and so cannot be simulated directly. Rather, it is given to a **Hybrid** object, which is a kind of **Atomic**, that generates the trajectories for the model. This **Hybrid** object is used just like any other **Atomic** model. Input to this **Hybrid** object triggers an

input event for the **ode\_system** that it contains. Likewise, output from the **ode\_system** becomes output from the **Hybrid** object. Most importantly, the hybrid model can be part of any network of discrete event models.

A **Hybrid** object is provided with three things when it is constructed. First is the **ode\_system** itself. Second is an **ode\_solver** that produces the model's continuous trajectories. Adevs has two types of **ode\_solvers**: a **corrected\_euler** solver that uses the corrected Euler method and a **rk\_45** solver that uses a fourth order Runge-Kutta method. Third is an **event\_locator** that finds the location of state events as the simulation progresses. Adevs has only one of these: the **linear\_event\_locator**. The code below shows how these are used to create and simulate a **Hybrid** object.

```
int main() {
    // Create the model
    CherryBomb* bomb = new CherryBomb();
    // Create the ODE solver for this model. Maximum error
    // tolerance at each step is 1E-4 and the maximum
    // size of an integration step is 0.01.
    ode_solver<string>* ode_solve =
        new corrected_euler<string>(bomb,1E-4,0.01);
    // Create the event locator for this model. Maximum
    // error tolerance for the location of an event in
    // the state space is 1E-8.
    event_locator<string>* event_find =
        new linear_event_locator<string>(bomb,1E-8);
    // Create an atomic model that puts all of these
    // together to simulate the continuous system.
    Hybrid<string>* model =
        new Hybrid<string>(bomb,ode_solve,event_find);
    // Create and run a simulator for this model
    Simulator<string>* sim = new Simulator<string>(model);
    while (bomb->getPhase() == FUSE_LIT)
        sim->execNextEvent();
    delete sim; delete bomb;
    return 0;
}
```

Figure 7.1 shows the cherry bomb's trajectory from  $t = 0$  to its explosion at  $t = 2$ . This plot was produced using the simulation program listed above. There is nothing particularly surprising about it, but you can observe the discrete changes in the cherry bomb's trajectory. There are two bounce events at  $t \approx 0.45$  and  $t \approx 1.4$ . The cherry bomb explodes abruptly at the start of its third descent.

## 7.2 Modeling hybrid systems with adevs and OpenModelica

It is possible to generate models in the form described above automatically from the Modelica programming language. This has several advantages, among them are:

1. You may be able to reuse large numbers of model components from the Modelica standard library.
2. Very complicated models, including differential-algebraic systems and models that difficult to initialize, may have code for the simulation automatically generated.
3. Modification of the continuous model will be much simpler in the future because the model description is separated from the problems of its simulation (these being handled automatically by the OpenModelica compiler and runtime system).




Figure 7.1: A simulation of the cherry bomb model that terminates when the cherry bomb explodes.

To take advantage of the Modelica language, you will need to get and install from source the OpenModelica compiler, which is hosted at <http://www.openmodelica.org/>. Instructions for building the compiler, source code for the compiler, and other information on the Modelica language are available from the Open Modelica group and at the Modelica Associate website <https://www.modelica.org/>. Instructions for building the adevs runtime environment for OpenModelica may be found in the README file in the adevs/src directory of the adevs package.

The adevs runtime environment for OpenModelica is relatively new, but aims to support most of the functionality of the standard OpenModelica runtime system with the following important caveats:

1. You may only solve ordinary systems of equations in the form  $\dot{x} = f(x, t)$  or semi-explicit differential-algebraic equations in the form  $\dot{x} = f(x, y, t)$  subject to  $g(x, y, t) = 0$ .
2. Many types of discrete behaviors are not supported. There is explicit support for if and when clauses, but other discrete functions (e.g., sample statements) should not be relied upon unless they appear in examples shipped with adevs.

Assuming you have a working compiler, you may generate an adevs model with the command `'omc +s +simCodeTarget=Adevs myModel.mo'` where `myModel.mo` is the name of your Modelica model file. The adevs/examples directory gives several instances of Modelica models and instructions for compiling them. I will build one such model here to demonstrate the general approach to hybrid modeling with adevs and OpenModelica. The code and makefiles for this example can be found in the adevs/examples/modelica directory.

To begin, consider a simple circuit that consists of a voltage source connected to single resistor with its other terminal grounded (i.e., Voltage source – resistor – ground). The voltage source has a setpoint  $V_{ref}$  that is controlled discretely. The equation governing the behavior of the voltage source is

$$\dot{v} = (V_{ref} - v)^{1/3}$$

and the current through the resistor is  $i = v/R$ . Our Modelica model will have  $V_{ref}$  as an input; the source of  $V_{ref}$  will be a discrete event model implemented in adevs. With this assumption, the Modelica model 'Circuit.mo' that describes this circuit is shown below.

```
connector Pin
  flow Real i;
  Real v;
end Pin;
```

```

partial model OnePin
    Pin T;
end OnePin;

partial model TwoPin
    Pin T1;
    Pin T2;
end TwoPin;

class Resistor extends TwoPin;
    parameter Real R(start=1);
equation
    T1.v-T2.v = T1.i*R;
    T1.i + T2.i = 0;
end Resistor;

class Ground extends OnePin;
equation
    T.v = 0;
end Ground;

class VoltageSource extends OnePin;
    parameter Real Vref(start=0);
equation
    der(T.v) =
        if Vref > T.v then (abs(Vref-T.v))^(1.0/3.0)
        else -(abs(Vref-T.v))^(1.0/3.0);
initial equation
    der(T.v) = 0;
end VoltageSource;

class Circuit
    VoltageSource V;
    Resistor R;
    Ground Gnd;
equation
    connect(V.T,R.T1);
    connect(R.T2,Gnd.T);
end Circuit;

```

To compile this model, issue the command 'omc +s +simCodeTarget=Adevs Circuit.mo'. This creates a C++ class called Circuit. Two files are created that implement this class: Circuit.h and Circuit.cpp. The header file is the most informative, and an abbreviated listing is shown below.

```

#ifndef __omc_Circuit_h_
#define __omc_Circuit_h_
#include "adevs.h"

/**
 * Define the input and output type of the adevs models.
 */
#ifndef OMC_ADEVS_IO_TYPE

```

```

#define OMC_ADEVS_IO_TYPE double
#endif

/**
 * Simulation code for Circuit
 * generated by the OpenModelica Compiler.
 */
class Circuit:
    public adevs::ode_system<OMC_ADEVS_IO_TYPE>
{
    public:
        /**
         * Constructor. New state events can be added to the model by
         * passing the number of new event conditions to the constructor
         * and then extending the state_event_func method. Your state
         * events will begin at the index returned by numStateEvents().
         */
        Circuit(int extra_state_events = 0);
        /// Destructor
        ~Circuit();
        /// Index of the first extra state event
        int numStateEvents() const { return 1; }
        /**
         * This method is used to set the hysteresis value that the
         * model will use when detecting state events. The default
         * value is 1E-4. Whatever you use, it must be larger than
         * the error tolerance of the event detection algorithm
         * that you are using.
         */
        void setEventHysteresis(double eps) { epsilon = eps; }
        /**
         * These methods are generated by the OpenModelica compiler.
         */
        void init(double* q);
        void der_func(const double* q, double* dq);
        void postStep(const double* q);
        void state_event_func(const double* q, double* z);
        /**
         * These methods may be overridden by any derived class.
         */
        virtual void extra_state_event_funcs(double* z){}
        double time_event_func(const double* q) { return DBL_MAX; }
        void internal_event(double* q, const bool* state_event);
        void external_event(double* q, double e,
            const adevs::Bag<OMC_ADEVS_IO_TYPE>& xb){}
        void confluent_event(double *q, const bool* state_event,
            const adevs::Bag<OMC_ADEVS_IO_TYPE>& xb){}
        void output_func(const double *q, const bool* state_event,
            adevs::Bag<OMC_ADEVS_IO_TYPE>& yb){}
        void gc_output(adevs::Bag<OMC_ADEVS_IO_TYPE>& gb){}
        /**
         * These methods are used to access variables and

```

```

    * parameters in the modelica model by name.
    */
double get_time() const { return timeValue; }
double get_$PV$PT$Pv() const { return $PV$PT$Pv; }
double get_$P$DER$PV$PT$Pv() const { return $P$DER$PV$PT$Pv; }
double get_$PGnd$PT$Pi() const { return $PGnd$PT$Pi; }
double get_$PR$PT2$Pv() const { return $PR$PT2$Pv; }
double get_$PGnd$PT$Pv() const { return $PGnd$PT$Pv; }
double get_$PV$PVref() const { return $PV$PVref; }
double get_$PR$PR() const { return $PR$PR; }

protected:
    /**
     * Calculate the values of the state and algebraic variables.
     * State variables will be initialized to q if provided,
     * or left unchanged if not.
     */
    void calc_vars(const double* q = NULL, bool doReinit = false);
    /**
     * These methods may be used to change paramters
     * and state variables at events. Remember to call
     * calc_vars(q,true) if you change anything.
     */
    void set_$PV$PT$Pv(double* q, double val) { q[0] = $PV$PT$Pv=val; }
    void set_$PV$PVref(double val) { $PV$PVref=val; }
    void set_$PR$PR(double val) { $PR$PR=val; }
};

#endif

```

This Circuit class can be used as is or (as I will soon show) be extended to add new, discrete event dynamics to it. The major features of this class are as follows:

1. The constructor, which lets you add state event functions to the model.
2. The extra\_state\_event function that lets implement the extra state events indicated to the constructor.
3. The numStateEvents method that tells you where in the event array your extra state events begin.
4. The set and get methods, which let you read and write the continuous variables and parameters of the model.
5. The calc\_vars method, which lets you reinitialize the modelica model at discrete events.
6. The event functions (internal, external, etc.) which can be extended to add new discrete event dynamics.

For this example, the circuit class is extended to accept as input discrete changes to Vref and produce the voltage as an output when Vref is reached. These new behaviors are implemented in a derived class called CircuitExt, which adds the following:

1. A state event function that triggers on v reaching Vref.
2. An output function that generates v when the state event occurs.

3. An external event function that changes Vref to the value received.

The complete code listing for the derived class is shown below.

```
/**
 * This class extends the Modelica Circuit model to
 * add an output when v = Vref and to adjust Vref
 * on receiving discrete input.
 */
class CircuitExt:
    public Circuit
{
    public:
        CircuitExt():
            Circuit(1), // Add one state event to the base class
            atVref(false) // Flag to indicate when Vref is reached
        {
        }
        /**
         * Compute the extra state event.
         */
        void extra_state_event_funcs(double* z)
        {
            if (!atVref)
                z[0] = get_$PV$PT$Pv()-get_$PV$PVref();
            else
                z[0] = 1.0;
        }
        /**
         * Indicate that Vref has been reached when that event
         * occurs.
         */
        void internal_event(double* q, const bool* event_flags)
        {
            // Apply the internal event function of the base class
            Circuit::internal_event(q,event_flags);
            // If this is the extra state event, then set atVref to true
            if (event_flags[numStateEvents()]) atVref = true;
        }
        /**
         * Change Vref on receiving input.
         */
        void external_event(double* q, double e, const Bag<double>& xb)
        {
            // Apply the external event function of the base class
            Circuit::external_event(q,e,xb);
            // Set the reference voltage and indicate that we are no longer
            // at the reference.
            set_$PV$PVref(*(xb.begin()));
            atVref = false;
            // Reinitialize the continuous model. This is really only necessary
            // if your discrete event may result in new values for the
            // state variables (discrete or continuous) of the modelica model.

```



```

        calc_vars(q,true);
    }
    void confluent_event(double* q, const bool * event_flags,
        const Bag<double>& xb)
    {
        internal_event(q,event_flags);
        external_event(q,0.0,xb);
    }
    void output_func(const double* q, const bool* event_flags,
        Bag<double>& yb)
    {
        // If this was the reference being reached, then
        // output the current value of the voltage.
        if (event_flags[numStateEvents()])
        {
            yb.insert(get_$PV$PT$Pv());
            cerr << "Reached Vref=" << get_$PV$PT$Pv()
                << " @ t=" << get_time() << endl;
        }
    }
    void print_state()
    {
        cout <<
            get_time() << " " << // Print the time
            get_$PV$PT$Pv() << " " << // The voltage
            get_$PGnd$PT$Pi() << " " << // Current through the resistor
            endl;
    }
private:
    bool atVref;
};

```

The main function for a simulation of this model is shown below. The main simulation routine injects and event at  $t = 1$  and runs the simulation for another two seconds to  $t = 3$ . A plot of the trajectory for this model and its input and output events are shown in Fig. 7.2 below.

```

int main()
{
    // Create the circuit
    CircuitExt* model = new CircuitExt();
    // Create an atomic model to simulate it
    Hybrid<OMC_ADEVS_IO_TYPE>* hybrid_model =
        new Hybrid<OMC_ADEVS_IO_TYPE>(
            model,
            new corrected_euler<OMC_ADEVS_IO_TYPE>(model,1E-5,0.01),
            new linear_event_locator<OMC_ADEVS_IO_TYPE>(model,1E-5));
    // Create the simulator
    Simulator<OMC_ADEVS_IO_TYPE>* sim =
        new Simulator<OMC_ADEVS_IO_TYPE>(hybrid_model);
    model->print_state();
    // Simulate to t = 1
    while (sim->nextEventTime() <= 1.0)
    {

```

```

        sim->execNextEvent();
        model->print_state();
    }
    // Inject an input
    Bag<Event<double> > input_bag;
    input_bag.insert(Event<double>(hybrid_model,0.5));
    sim->computeNextState(input_bag,1.0);
    // Simulate from t=1 to t=5
    while (sim->nextEventTime() <= 3.0)
    {
        sim->execNextEvent();
        model->print_state();
    }
    // Done, cleanup
    delete sim;
    delete hybrid_model;
    return 0;
}

```




Figure 7.2: A simulation of the circuit model.

## Chapter 8

# The Simulator Class

The functionality of the **Simulator** class is broken into three pieces: determining the model's time of next event, extracting output from the model, injecting input into the model, and advancing the simulation clock. The first piece of functionality is provided by the *nextEventTime* method with which we are already familiar. I'll address the three remaining pieces in turn.

There are two essential steps for extracting output from your model. The first step is to register an **EventListener** with the simulator. This is done by i) creating a subclass of the **EventListener** and ii) given this class to the **Simulator**'s *addEventListener* method. The **EventListener**'s *outputEvent* method is then used to intercept output originating from **Atomic** and **Network** models.

The second step is to invoke the **Simulator**'s *computeNextOutput* method, which performs the output calculations and provides the results to registered **EventListeners**. The signature of *computeNextOutput* is

```
void computeNextOutput()
```

and it computes the model output at the time given by the *nextEventTime* method. The *computeNextOutput* method invokes the *output\_func* method of every imminent **Atomic** model, maps outputs to inputs by calling the *route* method of **Network** models, and calls the *outputEvent* method of every **EventListener** registered with the **Simulator**. The *computeNextOutput* method anticipates the output of your model from its current state assuming that no input events will intervene between now and the *nextEventTime*.

The *computeNextState* method is used to inject events into a model and advance the simulation clock. The method signature is

```
void computeNextState(Bag<Event<X> >& input, double t)
```

where the **Event** class is the same one that the **EventListener** accepts to its *outputEvent* method. The **Event** class has two fields: a pointer to a model of type **Devs<X>** (i.e., a **Network** or **Atomic** model) and a value of type **X**.

The *computeNextState* method applies a bag of input **Events** to the model at time *t*. If the input bag is empty and *t* is equal to the next event time, then this method has the same effect as *execNextEvent*: it calculates the output values at time *t* using the *computeNextOutput* method, computes the next state of all models undergoing internal and external events, computes structure changes, and advances the simulation clock.

If the input bag is not empty then the value of each **Event** is applied as an input to the model pointed to by that **Event**. If, in this case, *t* is equal to the next event time then the method also follows the usual steps of invoking the *computeNextOutput* method and calculating state and structure changes. However, if *t* is less than the **Simulator**'s next event time, then the procedure is nearly identical excepting that the *computeNextOutput* method is not invoked. In this case, the only input events for any model are those provided in the input bag.

The **Simulator**'s *execNextEvent* method, the workhorse of most simulation programs, actually defers its job to two other methods. The implementation takes only two lines; the **Bag** `bogus_input` is empty.

```
void execNextEvent() {  
    computeNextOutput();  
    computeNextState(bogus_input,nextEventTime());  
}
```

The **Simulator**'s *computeNextOutput*, *computeNextState*, and *execNextEvent* methods throw an exception if a model violates either of two constraints: i) the time advance is negative and ii) the coupling constraints described in section 5.1 and illustrated in Figure 5.4, are violated. The Adevs **exception** class is derived from the standard C++ **exception** class; the method *what* returns a string that describes the exception condition and the method *who* returns a pointer to the model that caused the exception.

The Adevs **exception** class is intended to assist with debugging simulations. There isn't much you can do at run-time to fix a time advance method or reorganize a model's structure (or fix the structure change logic), but the simulator tries to be friendly by identifying a problem before it becomes an obscure and difficult to find bug.

## Chapter 9

# Simulation on multi-core computers

In addition to the sequential Simulator class, Adevs has a ParSimulator class that is designed specifically to take advantage of processors that have multiple cores and shared memory machines with several processors, these possibly having several cores each. The parallel simulator is in most respects identical to the sequential simulator, and this section of the manual therefore focuses on where it is different.

The ParSimulator class is designed specifically (only, really) to support symmetric, shared memory multi-processors (SMPs). The multi-core processors that have become ubiquitous in recent years are an important instance of this class of machines. The software technology that underlies the ParSimulator is OpenMP (see <http://www.openmp.org>), which is a standardized extension for C and C++ compilers and runtime systems to support multi-threaded computing. The OpenMP standard is now supported by most (probably all) major compilers: the GNU C++ Compiler and professional editions of Microsoft Visual Studio are important examples (important to me, that is, because those are what I use for most of my simulation work).

The critical first step, therefore, to using the ParSimulator is to enable the OpenMP extensions for your compiler. For the GNU C++ compiler, simply add the flag '-fopenmp' to your linker and compiler arguments. For MS Visual Studio, this is a build option (though I forget the location of the switch, and do not at the moment have access to the professional edition). For other compilers and development environments, see your documentation. Prior to executing a simulation, the maximum number of threads that will be used by OpenMP (and, therefore, the simulator) can be set by setting the OMP\_NUM\_THREADS environment variable (this works for the GNU compilers, at least). The default in most cases is to use a number of threads equal to the number of processors or cores in your computer.

Having enabled the OpenMP options for your compiler, you are ready to start preparing your model to work with the parallel simulator. As a first step, you can do the following. This example assumes that your main simulation routine looks something like this:

```
...
Simulator<IO_Type>* sim = new Simulator<IO_Type>(my_model);
/**
 * Register listeners with the Simulator to collect statistics
 * ....
 */
while (sim->nextEventTime() < t_end)
    sim->execNextEvent();
...
```

or this

```
...
Simulator<IO_Type>* sim = new Simulator<IO_Type>(my_model);
/**
```

```

* Register listeners with the Simulator to collect statistics
* ....
*/
sim->execUntil(t_end);
...

```

which does the same thing. The reason for this assumption is described in the section on limitations. Note, however, that any Listeners you have registered with your Simulator instance will work normally (almost, I'll get to that).

Assuming you have code like the above, replace it with code like the following:

```

...
ParSimulator<IO_Type>* sim = new ParSimulator<IO_Type>(my_model);
/**
* Register listeners with the ParSimulator to collect statistics
* ....
*/
sim->execUntil(t_end);
...

```

This should work just like your previous code with the following caveats.

First, your models must not share variables; all information exchanged between models must occur via events at their input and output. Also note that the order of items in the bags of input received by your model may change from run to run, though the contents are guaranteed to be the same. Therefore, the repeatability of your simulation runs depends on your models being insensitive to the order of elements in their input bags.

Second, reports produced by your listeners may be formatted in ways you do not expect: for any individual atomic model, the listing of state transitions and output events will be in time order. Across models, however, this may not be the case. For an extreme example, suppose that you have two atomic models arranged as follows: A-;B. You may see all state transitions and output for A listed first, followed by all state transitions and outputs for B. Most likely, these will be intermingled. Note too that the output reported for network models may not be in its proper time order, though the output reported for its atomic components will be.

This latter effect is due to the fact that, while simulation of each model is done in the proper time order of its events, the ParSimulator overlaps simulation of models whenever this is possible. In the above example, it is possible to simulate model A without worrying about what B is doing because B never provides input to A. The simulator (if properly configured, as is described next) will take advantage of this to simulate A and B in parallel. Hence, we may see all of the output from A before we see anything for B. Once again, however, all callbacks to registered Listeners for any particular model will be in the proper time order.

Note too that this implies that callbacks to your listeners may occur in parallel. In the above example, it may be that the same listener receives concurrent notifications about a state change for model A and state change for model B, or output from A and state change of B, or any such combination. It is imperative therefore that the callbacks in your listeners be thread safe.

This implies also that your network models must have routing methods that are thread safe. This is the case for the network models that are included with adevs. If you have implemented your own network models, be sure that their route methods are thread safe as well.

Third, and most critical, your atomic components must implement the new lookahead method, which is inherent from the Devs base class. This method must return a positive value subject to the guarantees described in the next section on "Principles of the parallel simulator". For the purposes of getting your code to compile and run, your lookahead methods can simply return a very small value (i.e., something positive but close to zero; say 1E-8 or 1E-9). In this case, your simulator should compile and (very slowly) execute. Even if you have the patience to wait for it to complete, however, the outcome will likely be wrong. Nonetheless, such a test will let you get your build environment setup properly.

The above changes are sufficient in most cases to make your existing, sequential simulator execute with the parallel simulator. In summary, these mandatory steps are:

1. Replace your Simulator with a ParSimulator.
2. Use the ParSimulator's `execUntil` method to advance time.
3. Make sure your Listeners are thread safe.
4. Make sure your Networks' route methods are thread safe.
5. Implement the lookahead method for your Atomic models.

These steps alone are unlikely to yield an improvement in performance (or, indeed, correct results). As a general rule, correctly speeding up your simulation requires that the ParSimulator be given specific information about your model; information that only you can provide. Without this information, the synchronization overhead incurred by the parallel simulation algorithm is staggeringly huge. The majority of this document deals with the problem of creating fast and correctly executing simulations.

## 9.1 Limits of the parallel simulator

Before continuing any effort to make your simulator work with the algorithms used by the ParSimulator, you should be aware of specific capabilities of the sequential simulator that the parallel simulator does not support. These are:

1. Your compiler must support OpenMP.
2. The `execNextEvent`, `computeNextState`, and `computeNextOutput` methods are not provided. Only the `execUntil` method is provided for advancing the simulation.
3. As noted above, callbacks to a Listener for each individual atomic model will be given in the proper time order, but these may be arbitrarily interleaved with the callbacks for other atomic models.
4. Listener callbacks must be made thread safe using the OpenMP synchronization features.
5. The parallel simulator does not support models that change structure. If your models implement their `model.transition` methods, then you cannot use the parallel simulator at this time.

Beyond these purely technical limits, it should be noted that making effective use of this (or any) parallel discrete event simulator is often difficult. Applications of practical interest require the identification of lookahead for the model's atomic components (described in the next section), partitioning of the model amongst the available processors, and implementing code to enable the parallel simulator to take advantage of your model's lookahead and partitioning.

So while this guide addresses issues specific to using the adevs ParSimulator class, I strongly recommend that, if you are not already intimately familiar with parallel simulation, that you obtain a book on the topic. "Building Software for Simulation", authored by James Nutaro (the author of this manual) and published by Wiley in 2011, contains a chapter on conservative discrete event simulation with DEVS in general and adevs in particular. You may find this book to be a useful introduction, though there are other excellent texts on the subject.

## 9.2 Modifying your models to exploit lookahead

The ParSimulator takes advantage of a property intrinsic to many models: their strong causality. A model, network or atomic, is strongly causal if its output can be predicted with certainty to some future time without knowledge of its input. The length of time into the future for which this prediction can be made is called lookahead.

To illustrate, consider a very simple, atomic model that acts as follows. In the absence of input, the model neither changes state nor produces output: its time advance is infinite. Upon receiving an input, the

model retains it for exactly one unit of time and then expels it. So, for example, if the input to the system is the series of letters 'A', 'B', and 'C' at times 1, 2, and 3 respectively, the its output comprise 'A', 'B', and 'C' at times 2, 3, and 4 respectively. If the model receives an input while in the process of transcribing, then that input is discarded. So, for example, if 'A', 'B', and 'C' arrive at times 2, 2.5, and 3 then the output from the model is 'A' and 'C' at times 3 and 4.

Now observe that if I know the input to this model until some time  $t$ , then I can determine its output until time  $t+1$  and, to do this, I do not need to know its input in the interval from  $t$  to  $t+1$ . For instance, suppose I know that the input at time zero is 'A'. Clearly, the only output of the model in the interval 0 to 1 (inclusive) is the value 'A' at time 1. Moreover, suppose I know that the input at  $t=0$  is 'A' and that there is no other input until at least time 0.5. In this case, I know that the output until (but not including) time 1.5 consists only of 'A' at time 1. Any input following time 0.5 cannot occur until, at its earliest, time 1.5. This model has a lookahead of 1. Given its input to time  $t$ , its output is fixed to time  $t+1$ . Output in this interval does not depend on input in the same interval.

In network models, lookahead may accumulate. As an example, suppose that two transcribers are connected in series. In this case, the lookahead of the composite is two; i.e., the sum of the lookaheads of the components. If, however, a network model comprises two transcribers in parallel, then the lookahead of the network is only one. Generally speaking, however, larger networks tend to have larger lookaheads, and large lookahead is essential for getting good performance from the simulator.

To take advantage of lookahead in a model, the simulator must be told that it exists. This is done by overriding the lookahead method of the Devs class. All atomic and network models inherit this method from the Devs base class, and its default implementation is to return zero.

Lookahead is most useful to the simulator if it is coupled with a capability to actually calculate the model's outputs to this time. To calculate those outputs, however, requires knowledge of the intervening states. An atomic model enable the simulator to project its output into the future by implementing two method: beginLookahead and endLookahead.

The beginLookahead method is called to notify the model that further calculations of its outputs and state transitions are speculative. The default behavior of the beginLookahead method is to throw an exception, which notifies the simulator that this model does not support the projection of its output into the future. An atomic model overriding this method must be capable of restoring its state variables to their follows at the instant that beginLookahead was called. The endLookahead method is called by the simulator when it is done projecting that model's output. This method must restore the model to the same state it was in when the beginLookahead method was called.

These methods are demonstrated by the Transcribe class shown below. This atomic model implements the transcriber described above.

```
/**
 * This model copies its input to its output following a
 * delay.
 */
class Transcribe:
public adevs::Atomic<char>
{
public:
Transcribe():adevs::Atomic<char>(),ttg(DBL_MAX),to_transcribe(' '){}
void delta_int() { ttg = DBL_MAX; }
void delta_ext(double e, const adevs::Bag<char>& xb)
{
if (ttg == DBL_MAX)
{
ttg = 1.0;
// Find the largest input value
adevs::Bag<char>::const_iterator iter = xb.begin();
```



```

to_transcribe = *iter;
for (; iter != xb.end(); iter++)
{
if (to_transcribe < *iter) to_transcribe = *iter;
}
}
else ttg -= e;
}
void delta_conf(const adevs::Bag<char>& xb)
{
delta_int();
delta_ext(0.0,xb);
}
void output_func(adevs::Bag<char>& yb)
{
yb.insert(to_transcribe);
}
double ta() { return ttg; }
void gc_output(adevs::Bag<char>&){}
double lookahead() { return 1.0; }
void beginLookahead()
{
// Save the state
chkpt.ttg = ttg;
chkpt.to_transcribe = to_transcribe;
}
void endLookahead()
{
// Restore the state
ttg = chkpt.ttg;
to_transcribe = chkpt.to_transcribe;
}
char getMemory() const { return to_transcribe; }
private:
double ttg;
char to_transcribe;
struct checkpoint_t { double ttg; char to_transcribe; };
checkpoint_t chkpt;
};

```

A model comprising two transcribers connected in series could be defined as follows. Note that this network model does not have the endLookahead and beginLookahead methods. State is a property of atomic models only, and therefore only atomic models may implement these methods.

```

/**
 * This model defines a pair of transcribers connected
 * in series as shown: -> t1 -> t2 ->.
 */
class Series:
public adevs::Network<char>
{
public:
Series():

```

```

Network<char>(),
t1(),t2()
{
t1.setParent(this);
t2.setParent(this);
}
void getComponents(a devs::Set<adevs::Devs<char>*>& c)
{
c.insert(&t1);
c.insert(&t2);
}
void route(const char& x, adevs::Devs<char>* model,
adevs::Bag<adevs::Event<char> >& r)
{
adevs::Event<char> e;
e.value = x;
if (model == this) e.model = &t1;
else if (model == &t1) e.model = &t2;
else if (model == &t2) e.model = this;
}
double lookahead()
{
return t1.lookahead()+t2.lookahead();
}
private:
Transcribe t1, t2;
};

```

The above code examples illustrate all possible changes of your models - network and atomic - to facilitate parallel simulation. Of these changes, only the lookahead method of the atomic model is actually required. The ParSimulator calculates default (and very conservative) lookaheads for the Network models if these are required. Atomic models that provide the endLookahead and beginLookahead methods may improve the execution time of the simulation. But only the lookahead values of the atomic models (or their parents if the model is partitioned by hand; see the next section) are actually required for correct execution of your model.

## 9.3 Partitioning your model

Each thread in your simulator is assigned responsibility for the execution of a subset of the atomic components of your model. Models within a thread are executed sequentially. Simulation of the models within a thread proceeds just as with the sequential Simulator class. The threads execute in parallel, each stopping to synchronize with its neighbors only as necessary to exchange essential information.

In an ideal partitioning of the model, each thread is assigned roughly the same number of models, each model requires roughly the same amount of computational effort to simulate, and models assigned to separate threads exchange inputs and outputs rarely or not at all. This is the ideal that you should strive for in assigning your models to the threads of your simulation algorithm.

The actual assignment of a model to a thread is straightforward. Network and Atomic objects inherit the setProc method from their Devs base class. To assign the model to a particular thread, pass the number of that thread to the setProc method before the ParSimulator is created. Threads are numbered from 0 to the maximum number of threads (i.e., OMP\_NUM\_THREADS) minus one.

As the ParSimulator setups the simulation, it will examine the thread to which each model is assigned and take the following actions:

1. If the model's parent is assigned to thread k, then the model is also assigned to thread k, regardless of the argument passed to its setProc method.
2. If the model's parent was not assigned to specific thread, then the model will be assigned to the thread indicated by argument to its setProc method.
3. If neither the model nor its parent were assigned to a thread via the model's setProc method, then it is assigned to a thread selected at random.

This partitioning of the model tells the simulator how to distribute its computational workload. It does not tell the simulator which parts of the workload talk to which others. For example, if half of your model is assigned to thread 0 and the other half to thread 1, the simulator does not yet know if the model's in 0 send input to the models in 1, vice versa, or both. Without further information, the simulator will therefore assume that every thread contains models that must communicate with all other threads. This is the most conservative assumption, and carries with it a substantial synchronization overhead.

If you know something about how the models assigned to the thread communicate, then you can provide this information to the simulator. This is done by passing an LpGraph object to the constructor of the ParSimulator. The LpGraph is nothing more than a directed graph, and the presence of an edge from node k to node j indicates that the models for thread k send input to the models in thread j. Absence of an edge indicates no flow of information along the missing edge. An edge is added by calling addEdge(A,B) to create an edge from thread A to thread B.

The following snippet of code illustrates the partitioning procedure. This segment of code creates the block diagram model shown in Fig. 9.1. This model consists of two atomic components, A and B, and a network with two components C1 and C2. The model A is assigned to thread 0, B to thread 1, and the network C with its components C1 and C2 to thread 2. With this partition, thread 0 sends input to thread 1, thread 1 sends input to thread 2, and thread 2 to thread 0.

```
ModelA* A = new ModelA();
ModelB* B = new ModelB();
NetworkModelC* C = new NetworkModelC();
SimpleDigraph<IO_Type>* model = new SimpleDigraph<IO_Type>();
model.add(A);
model.add(C);
model.add(B);
model.couple(A,B);
model.couple(B,C);
model.couple(C,A);
A.setProc(0);
B.setProc(1);
C.setProc(2);
LpGraph lpg;
lpg.addEdge(0,1);
lpg.addEdge(1,2);
lpg.addEdge(2,3);
ParSimulator<IO_Type> sim(model,lpg);
```

## 9.4 Interaction between partitioning and lookahead

The following rules dictate which models must provide positive lookahead.

1. An atomic model whose parent is not assigned to a specific thread must provide a positive lookahead.
2. A network that is assigned to a specific thread, but whose parent is not, must provide a positive lookahead.




Figure 9.1: Partitioning a model for simulation on three processors.

3. No other model is required to provide a positive lookahead, or indeed any lookahead at all. The simulator will not use lookahead values provided by models except as indicated in cases 1 and 2 above.

## 9.5 A complete example

This example builds and simulates the network shown in Fig. 9.1 using the Transcribe model for components B, C1, and C2; a SimpleDigraph for the network C, and the following network for component A.

The network that is component A has two sub-components: a generator and a transcriber. Input to the network goes to the generator; output from the generator goes to the transcriber; and output from the transcriber becomes an output from the network. The generator operates as follows. It produces output at regular intervals of 1/2 units of time until it receives an input. At that time it stops. Note that the lookahead of the generator is zero because its pending output may be canceled at any time. The lookahead of the network, however, is one. That is, the lookahead of the network is the sum of its series components.

A Listener is constructed for this model that records the state and output trajectories of every component. Observe two facts concerning the listener. First, it has been made thread safe by the critical section placed around writes to standard output. Second, it reports state and output trajectories for each atomic model in proper time order, but interleaves the trajectories of these models with each other. This second point will be apparent when we execute the simulation.

The complete implementation of the model, listener, and simulator are shown below.

```

#include "Transcribe.h"
#include "Genr.h"
#include "adevs.h"
#include <iostream>
using namespace adevs;
using namespace std;

/**
 * Extend the SimpleDigraph class to allow its lookahead
 * to be set manually.
 */
class SimpleDigraphWithLookahead:
public SimpleDigraph<char>

```

```

{
public:
SimpleDigraphWithLookahead():
SimpleDigraph<char>(),
look_ahead(0.0)
{
}
void setLookahead(double look_ahead)
{
this->look_ahead = look_ahead;
}
double lookahead() { return look_ahead; }
private:
double look_ahead;
};

/**
 * Listener to record the output and state trajectories of the
 * component models.
 */
Genr* A_g;
Transcribe *A_t, *B, *C1, *C2;
SimpleDigraphWithLookahead *A, *C;

class Listener:
public EventListener<char>
{
public:
Listener(){}
void outputEvent(Event<char> y, double t)
{
string which = which_model(y.model);
#pragma omp critical
cout << which << " @ t = " << t << ", y(t)= " << y.value << endl;
}
void stateChange(Atomic<char>* model, double t)
{
if (model == A_g)
#pragma omp critical
cout << which_model(A_g) << " @ t = " << t << ", running= "
<< A_g->isRunning() << ", next output= " <<
A_g->getNextOutput() << endl;
else if (model == A_t)
#pragma omp critical
cout << which_model(A_t) << " @ t = " << t << ", memory= "
<< A_t->getMemory() << ", ta()= " <<
A_t->ta() << endl;
else if (model == C1)
#pragma omp critical
cout << which_model(C1) << " @ t = " << t << ", memory= "
<< C1->getMemory() << ", ta()= " <<
C1->ta() << endl;
}
}

```

```

else if (model == C2)
#pragma omp critical
cout << which_model(C2) << " @ t = " << t << ", memory= "
<< C2->getMemory() << ", ta()= " <<
C2->ta() << endl;
else if (model == B)
#pragma omp critical
cout << which_model(B) << " @ t = " << t << ", memory= "
<< B->getMemory() << ", ta()= " <<
B->ta() << endl;
else assert(false);
}
private:
string which_model(Devs<char>* model)
{
if (model == A_g) return "A.A_g";
if (model == A_t) return "A.A_t";
if (model == A) return "A";
if (model == B) return "B";
if (model == C1) return "C.C1";
if (model == C2) return "C.C2";
if (model == C) return "C";
assert(false);
return "";
}
};

int main(int argc, char** argv)
{
// Component A
A_g = new Genr();
A_t = new Transcribe();
A = new SimpleDigraphWithLookahead();
A->setLookahead(A_t->lookahead()+A_g->lookahead());
A->add(A_g);
A->add(A_t);
A->couple(A,A_g); // A -> A_g
A->couple(A_g,A_t); // A_g -> A_t
A->couple(A_t,A); // A_t -> A
A->setProc(0); // Assign to thread zero
// Component B
B = new Transcribe();
B->setProc(1); // Assign to thread one
// Component C
C1 = new Transcribe();
C2 = new Transcribe();
C = new SimpleDigraphWithLookahead();
C->setLookahead(C1->lookahead()+C2->lookahead());
C->add(C1);
C->add(C2);
C->couple(C,C1); // C -> C1
C->couple(C2,C); // C2 -> C

```

```

C->couple(C1,C2); // C1 -> C2
C->couple(C2,C1); // C2 -> C1
C->setProc(2); // Assign to thread two
// Create the overarching model
SimpleDigraph<char>* model = new SimpleDigraph<char>();
model->add(A);
model->add(B);
model->add(C);
model->couple(A,B);
model->couple(B,C);
model->couple(C,A);
// Create the corresponding LPGraph
LpGraph lpg;
lpg.addEdge(0,1);
lpg.addEdge(1,2);
lpg.addEdge(2,0);
// Create the simulator
ParSimulator<char>* sim = new ParSimulator<char>(model,lpg);
// Register the listener
Listener* listener = new Listener();
sim->addEventListener(listener);
// Run the simulation until t=10
sim->execUntil(10.0);
// Cleanup and exit
delete sim;
delete listener;
delete model;
return 0;
}

```

A subset of the output produced by the simulator is shown below. The intermingling of reported events in time is immediately apparent.

```

...
C.C2 @ t = 7.5, y(t)= G
C @ t = 7.5, y(t)= G
C.C1 @ t = 7.5, y(t)= I
C.C2 @ t = 7.5, memory= I, ta()= 1
C.C1 @ t = 7.5, memory= G, ta()= 1
C.C1 @ t = 8.5, y(t)= G
C.C2 @ t = 8.5, y(t)= I
C @ t = 8.5, y(t)= I
C.C1 @ t = 8.5, memory= I, ta()= 1
C.C2 @ t = 8.5, memory= G, ta()= 1
A.A_g @ t = 7.5, running= 0, next output= I
A.A_g @ t = 8.5, running= 0, next output= I
A.A_g @ t = 9.5, running= 0, next output= I
C.C1 @ t = 9.5, y(t)= I
C.C2 @ t = 9.5, y(t)= G
C @ t = 9.5, y(t)= G
C.C1 @ t = 9.5, memory= G, ta()= 1
C.C2 @ t = 9.5, memory= I, ta()= 1
...

```

However, a search for just the events for model C1 gives the expected result: all of its events are listed in their proper time order. Specifically, the command 'grep C2 output', where 'output' is the result of the simulation, yields the following:

```
C.C2 @ t = 3.5, memory= A, ta()= 1
C.C2 @ t = 4.5, y(t)= A
C.C2 @ t = 4.5, memory= C, ta()= 1
C.C2 @ t = 5.5, y(t)= C
C.C2 @ t = 5.5, memory= E, ta()= 1
C.C2 @ t = 6.5, y(t)= E
C.C2 @ t = 6.5, memory= G, ta()= 1
C.C2 @ t = 7.5, y(t)= G
C.C2 @ t = 7.5, memory= I, ta()= 1
C.C2 @ t = 8.5, y(t)= I
C.C2 @ t = 8.5, memory= G, ta()= 1
C.C2 @ t = 9.5, y(t)= G
C.C2 @ t = 9.5, memory= I, ta()= 1
```

So too for the output of model 'A.g', which is shown below:

```
A.A_g @ t = 0.5, y(t)= A
A.A_g @ t = 0.5, running= 1, next output= B
A.A_g @ t = 1, y(t)= B
A.A_g @ t = 1, running= 1, next output= C
A.A_g @ t = 1.5, y(t)= C
A.A_g @ t = 1.5, running= 1, next output= D
A.A_g @ t = 2, y(t)= D
A.A_g @ t = 2, running= 1, next output= E
A.A_g @ t = 2.5, y(t)= E
A.A_g @ t = 2.5, running= 1, next output= F
A.A_g @ t = 3, y(t)= F
A.A_g @ t = 3, running= 1, next output= G
A.A_g @ t = 3.5, y(t)= G
A.A_g @ t = 3.5, running= 1, next output= H
A.A_g @ t = 4, y(t)= H
A.A_g @ t = 4, running= 1, next output= I
A.A_g @ t = 4.5, y(t)= I
A.A_g @ t = 4.5, running= 0, next output= I
A.A_g @ t = 5.5, running= 0, next output= I
A.A_g @ t = 6.5, running= 0, next output= I
A.A_g @ t = 7.5, running= 0, next output= I
A.A_g @ t = 8.5, running= 0, next output= I
A.A_g @ t = 9.5, running= 0, next output= I
```

So the individual traces for these components appear in the proper order in the output, but they are intermingled in an arbitrary way.

## 9.6 Memory management of input and output across thread boundaries

Because the atomic moels in a simulator are executed at different rates, it often happens that an output object produced by a model in one thread will be used as some later time by models in other threads. To



manage the memory associated with these objects, it is necessary for the simulator to be able to determine when any such object can be safely deleted. This is done most easily when every thread has its own copy of the object, and the MessageManager interface is used by the simulator for this purpose.

If your input and output types are primitive objects (ints, chars, etc.) or simple structures, then the default approach to memory management is sufficient. Indeed, this memory manager should be sufficient for any types of objects for which the compiler's default copy constructor and assignment operator create deep copies. If you are passing pointers to complex objects or objects that use their own internal scheme for managing memory (e.g., that use copy on write semantics, reference counting, etc.), then you will need to build a custom memory manager. The MessageManager is used for this purpose, and a custom MessageManager is provided to your ParSimulator as the final argument to its constructor.

The MessageManager has two pure virtual methods that must be overridden by its specialization. The first is the clone methods, which has the signature

```
X clone(X& value)
```

where X is the type of object that your simulator uses for input and output. This method must make a deep copy of the object value and return that copy to the caller via the return value. The second method is destroy, and it has the signature

```
void destroy(X& value)
```

where X is as before. This method must free the memory associated with the value.

The implementation of the default memory manager is as follows:

```
template <typename X> class NullMessageManager:
    public MessageManager<X>
{
    public:
        /// Uses the objects default copy constructor
        X clone(X& value) { return value; }
        /// Takes no action on the value
        void destroy(X& value){}
};
```

To illustrate the construction of a new MessageManager, the implementation below is for a model that uses C style strings (i.e., null terminate arrays of characters) for its input and output. The clone method allocated memory for a string and the copies to it the contents of the value string. The destroy method frees the memory allocated for the string.

```
class CStringMessageManager:
    public adevs::MessageManager<char*>
{
    public:
        char* clone(char* & value)
        {
            char* new_string = new char[strlen(value)];
            strcpy(new_string,value);
            return new_string;
        }
        void destroy(char* & value)
        {
            delete [] value;
        }
};
```

If this message manager were to be supplied to the simulator in the example of the previous section, then the ParSimulator would be constructed as follows:

```
. . .  
CStringMessageManager* msg_mgr = new CStringMessageManager();  
ParSimulator<char*>* sim = new ParSimulator<char*>(model,lpg,msg_mgr);  
. . .
```

The simulator would then use the supplied message manager to manage input and output objects that must exist simultaneously in the simulators many threads.

## 9.7 Summary of main points regarding repeatable outcomes and performance

If your model and simulator have been setup properly then the outcomes produced by the parallel and sequential simulators will be identical. To this end, keep the following rules in mind:

1. Models must not shared variables.
2. The state transition functions of your models must not depend on the order of items in their input bag.
3. Listeners must be thread safe.
4. Listeners must not expect events to be reported in a global time order. Only the events associated with individual atomic models will be reported in their proper order; all other events will be interleaved in time.

In general you should not expect to speedup relatively small simulations by use of the parallel simulator. Rather, its purpose is chiefly to enable your model to grow in its size, complexity, or both without a corresponding increase in execution time. With this in mind, to achieve good execution times with the parallel simulation requires the following:

1. The size of the model must be sufficient to keep all of your processors busy all of the time with useful work. Moreover, the amount of useful work to be done by each processor must substantially exceed the overhead of parallel simulation algorithm.
2. Your model must have parallelism that the parallel algorithm can exploit. In practice, this means that your model must be partitioned to both maximize the lookahead of the network assigned to each processor and to minimize the communication between processors.

## Chapter 10

# Models with Many Input/Output Types

It would be surprising if every component in a large model (or even a small one) had the same input and output requirements. Some models can be satisfactorily constructed with a single type of input/output object and, if this is the case, it will simplify the design of your simulator. If not, you'll need to address this problem when you design your simulation program.

One solution to this problem is to establish a base class for all input and output types and derive specific types from the common base. The simulator and all of its components exchange pointers to the base class and down cast specific objects as needed. The C++ `dynamic_cast` operator is particularly useful for this purpose. Although it is not without its problems, I have used this solution in many designs and it works well.

It is not always possible for every component in a model to share a common base class for its input and output type. This can happen if different sub-model have very different input and output needs or when models from earlier projects are reused. For example, to use a **CellSpace** model as a component of a **Digraph** model requires some means of converting the **CellSpace**'s **CellEvent** objects into the **PortValue** objects required by the **Digraph**. Happily, there is a simple solution to this problem that makes clever use of the **Simulator** and **EventListener** classes to wrap a model with one input and output type inside of an atomic model with a different input and output type.

The Adevs **ModelWrapper** class is an **Atomic** model that encapsulates another model. The encapsulated model can be a **Network** or **Atomic** model. The **ModelWrapper** uses input/output objects of type `ExternalType`, but the encapsulated class uses input/output objects of type `InternalType`. Two abstract methods are provided for converting objects with one type into objects with the other type; these methods are

```
void translateInput(const Bag<ExternalType>& external_input, Bag<Event<InternalType> >& internal_input)
void translateOutput(const Bag<Event<InternalType> >& internal_output, Bag<ExternalType>& external_output)
```

Clean up of converted objects are managed with the *gc\_output* method, which is inherited from the **ModelWrapper**'s **Atomic** base class, and a new *gc\_input* method for cleaning up objects created by the *translateInput* method; its signature is

```
void gc_input(Bag<Event<InternalType> >& g)
```

The model to encapsulate is passed to the **ModelWrapper** constructor. The **ModelWrapper** creates a **Simulator** for the model that is used to control its evolution; the **ModelWrapper** is a simulator inside of a model inside of a simulator! The **ModelWrapper** keeps track of the wrapped model's last event time, and it uses this information and the **Simulator**'s *nextEventTime* method to compute its time advance. Internal, external, and confluent events cause the **WrappedModel** to invoke its **Simulator**'s

*computeNextState* method and thereby advance the state of the wrapped model. Internal events are simplest; the *computeNextState* method is invoked with the wrapped model's next event time and an empty input bag.

The *delta\_conf* and *delta\_ext*, however, must convert the incoming input events, which have the type *ExternalType*, into input events for the wrapped model, which have the type *InternalEvent*. This is accomplished with the *translateInput* method. The first argument to the method is the input bag passed to the **ModelWrapper**'s *delta\_ext* or *delta\_conf* method. The second argument is an empty bag that the method implementation must fill. When the *translateInput* method returns this bag will be passed to the *computeNextState* method of the **ModelWrapper**'s simulator. Notice that the *internal\_input* argument is a **Bag** filled with **Event** objects; if the wrapped model is a **Network** then the translated events can be targeted at any of the **Network**'s components. The **ModelWrapper** invokes the *gc\_input* method when it is done with the events in the *internal\_input* bag. This gives you the opportunity to delete objects that you created when *translateInput* was called.

A similar process occurs when the **ModelWrapper**'s *output\_func* is invoked, but in this case it is necessary to convert output objects from the wrapped model, which have type *InternalType*, to output objects from the **ModelWrapper**, which have type *ExternalType*. This is accomplished by invoking the *translateOutput* method. The method's first argument is the bag of output events produced collectively by all of the wrapped model's components; notice that the contents of the *internal\_output* bag are **Event** objects. The *model* field points to the component of the wrapped model (or the wrapped model itself) that produced the event and the *value* field contains an output object produced by that model. These **Events** must be converted to objects of type *ExternalType* and stored in the *external\_output* bag. The *external\_output* bag is, in fact, the bag passed to the wrapper's *output\_func*, and so its contents become the output objects produced by the wrapper. The *gc\_output* method is used in the usual way to clean up any objects created by this process.

The **Wrapper** class shown below illustrates how to use the Adevs **WrapperModel** class. The **Wrapper** is derived from the **WrapperModel** and implements its four abstract methods: *translateInput*, *translateOutput*, *gc\_input*, and *gc\_output*. This class wraps an **Atomic** model that uses *int\** objects as its input/output. The **Wrapper** uses C strings as its input and output type. The translation methods merely convert integers to strings and vice versa. The **Wrapper** can be used just like any **Atomic** model; it can be a component in a network model or simulated by itself. The behavior of the **Wrapper** is identical to the model it wraps; the only change is in the interface.

```
// This class converts between char* and int* event types.
class Wrapper: public adevs::ModelWrapper<char*,int*> {
public:
    Wrapper(adevs::Atomic<int*>* model):
        // Pass the model to the base class constructor
        adevs::ModelWrapper<char*,int*>(model){}
    void translateInput(const adevs::Bag<char*>& external,
        adevs::Bag<adevs::Event<int*> >& internal) {
        // Iterate through the incoming events
        adevs::Bag<char*>::const_iterator iter;
        for (iter = external.begin(); iter != external.end(); iter++) {
            // Convert each one into an int* and send it to the
            // wrapped model
            adevs::Event<int*> event;
            // Set the event value
            event.value = new int(atoi(*iter));
            // Set the event target
            event.model = getWrappedModel();
            // Put it into the bag of translated objects
            internal.insert(event);
        }
    }
};
```

```

    }
}
void translateOutput(const adevs::Bag<adevs::Event<int*> >& internal,
    adevs::Bag<char*>& external) {
    // Iterate through the incoming events
    adevs::Bag<adevs::Event<int*> >::const_iterator iter;
    for (iter = internal.begin(); iter != internal.end(); iter++) {
        // Convert the incoming event value to a string
        char* str = new char[100];
        sprintf(str,"%d",*((*iter).value));
        // Put it into the bag of translated objects
        external.insert(str);
    }
}
void gc_output(adevs::Bag<char*>& g) {
    // Delete strings allocated in the translateOutput method
    adevs::Bag<char*>::iterator iter;
    for (iter = g.begin(); iter != g.end(); iter++)
        delete [] *iter;
}
void gc_input(adevs::Bag<adevs::Event<int*> >& g) {
    // Delete integers allocated in the translateInput method
    adevs::Bag<adevs::Event<int*> >::iterator iter;;
    for (iter = g.begin(); iter != g.end(); iter++)
        delete (*iter).value;
}
};

```



## Chapter 11

# Random Numbers

Adevs has two classes that work together to generate many types of random numbers. These two classes are the **random\_seq** class and the **rv** class. The **random\_seq** class provides uniformly distributed random numbers to the **rv** class, and the **rv** transforms this uniform stream of random numbers into a variety of random number distributions.

The **random\_seq** class is an interface for a random number generator. Its derived classes produce uniformly distributed pseudo-random numbers. The underlying random number stream is accessed with two methods. The *next\_long* returns a raw random number as an unsigned long. The *next\_dbl* refines the *next\_long* method by reducing the raw random number to a double precision number in the interval  $[0, 1]$ . The random number sequence is initialized with the *set\_seed* method, and the entire random number generator can be copied with the *copy* method. To summarize, the **random\_seq** class has four abstract methods

```
void set_seed(unsigned long seed)
double next_dbl()
random_seq* copy() const
unsigned long next_long()
```

that must be implemented by any derived class.

Adevs comes with two implementations of the **random\_seq** class: the **crand** class and the **mtrand** class. The **crand** class uses the `rand` function from the standard C library to implement the required methods. Its implementation is trivial; I've listed it below as an example of how to implement the **random\_seq** interface.

```
class crand: public random_seq {
public:
    /// Create a generator with the default seed
    crand(){}
    /// Create a generator with the given seed
    crand(unsigned long seed) { srand (seed); }
    /// Set the seed for the random number generator
    void set_seed(unsigned long seed) { srand (seed); }
    /// Get the next double uniformly distributed in  $[0, 1]$ 
    double next_dbl() { return (double)rand()/(double)RAND_MAX; }
    /// Copy the random number generator
    unsigned long next_long() { return (unsigned long)rand(); }

    random_seq* copy() const { return new crand (); }
    /// Destructor
    ~crand(){}
};
```

The **mtrand** class implements the Mersenne Twister random number generator<sup>1</sup>. The code is based on their open source implementation of the Mersenne Twister. Aside from its potential advantages as a random number generator, the **mtrand** class differs from the **crand** class by its ability to make deep copies; every instance of the **mtrand** class has its own random number stream.

The **rv** class uses the uniform random numbers provided by a **random\_seq** object to produce several different random number distributions: triangular, uniform, normal, exponential, lognormal, Poisson, Weibull, binomial, and many others. Every instance of the **rv** class is created with a **random\_seq**; the default is an **mtrand** object, but any type of **random\_seq** object can be passed to the **rv** constructor. The different random distributions are sampled by calling the appropriate method: ***triangular*** for a triangular distribution, ***exponential*** for an exponential distribution, ***poisson*** for a Poisson distribution, etc. Adevs is open source software; if a new distribution is needed then you can add a method that implements it to the **rv** class (and, I hope, contribute the expansion to the Adevs project).

---

<sup>1</sup>M. Matsumoto and T. Nishimura, “Mersenne Twister: A 623-Dimensionally Equidistributed Uniform Pseudo-Random Number Generator”, ACM Transactions on Modeling and Computer Simulation, Vol. 8, No. 1, January 1998, pgs. 3-30.



## Chapter 12

# Interpolation

Data is often made available as a set of tabulated points. Hourly temperature data, millisecond samples of a voltage signal in an electric circuit, and a minute by minute record of a radar track are examples of continuous signals recorded at discrete points in time. But if you need temperate data every half hour, the circuit voltage a fractions of a millisecond, or track information at quarter minutes then a method for approximating values between tabulated points will be very useful. The **InterPoly** class exists for this purpose.

The **InterPoly** class approximates smooth continuous signal by fitting a polynomial to the available data points. The approximating polynomial is guaranteed to pass through every available data point, and in many cases it provides a reasonable approximation to the original signal between the available data points. The most familiar example of an interpolating polynomial is a line that connects two data points  $(t_1, x_1)$  and  $(t_2, x_2)$ ; the connecting line is

$$p(t) = \frac{t - t_2}{t_1 - t_2}x_1 + \frac{t - t_1}{t_2 - t_1}x_2$$

and it is easy to check that  $p(t_1) = x_1$  and  $p(t_2) = x_2$ . If more data points are available then quadratic, cubic, quartic, and even higher degree polynomials can be used to obtain (in principle) better approximations.

An interpolating polynomial can be constructed with the **InterPoly** class in three ways. The first way is to provide the sample data to the **InterPoly** constructor

```
InterPoly(const double* u, const double* t, unsigned int n)
```

where  $u$  is an array of data values,  $t$  is an array of associated time points, and  $n$  is the number of data points (i.e., the size of the  $u$  and  $t$  array). The constructor will build an  $n - 1$  degree polynomial that fits the supplied data. The second way is supply just the data values, a time step, the first time value, and number of data points to the constructor

```
InterPoly(const double* u, double dt, unsigned int n, double t0 = 0.0)
```

where  $u$  is an array of data values,  $dt$  is the time spacing of the data points,  $n$  is the number of data points, and  $t_0$  is the time instant of the first data point (i.e., the data point  $i$  is at time  $t_0 + i \cdot dt$ ). Both constructors make copies of the supplied arrays, and changes to the array values will not be reflected by the **InterPoly** object. The third way is to assign new data point values to an existing polynomial by calling the **InterPoly** method

```
void setData(const double* u, const double* t = NULL)
```

where  $u$  is the new set of data values and  $t$  is (optionally) the new set of time points. This method requires that the number of data points in  $u$  (and  $t$  if used) be equal to the number of points supplied to the **InterPoly** constructor.

There are three methods for computing interpolated values: the *interpolate* method, the overloaded *operator()*, and the *derivative* method. The method signatures are listed below:

```
double interpolate(double t) const
double operator()(double t) const
double derivative(double t) const
```

The *interpolate* method and *operator()* give the value of the interpolating polynomial at the time point  $t$ . The *derivative* method gives the value of the first time derivative of the interpolating polynomial as an approximation of the first time derivative of the original function. For example, if the data describes the position of an object through time then the *derivative* method gives an approximation of the object's velocity.

To demonstrate the **InterPoly** class and give you a sense of what the interpolating polynomials look like, I've listed below a program that computes  $\sin(t)$ , its time derivative  $\cos(t)$ , and interpolated approximations of both. Interpolating polynomials built with 4, 5, and 6 data points in the interval  $[0, 2\pi]$  are illustrated in Figs. 12.1 and 12.2. The quality of the approximation generally improves as more data points are added, but the function and interpolating polynomial diverge significantly outside of the interval spanned by the data points. Be careful if you extrapolate!

```
#include "adevs.h"
#include <cmath>
#include <iostream>
using namespace std;

int main(int argc, char** argv) {
    // Get the number of data points to use and allocate
    // memory for the data points arrays
    int N = atoi(argv[1]);
    double* t = new double[N];
    double* u = new double[N];
    // Compute data points using the sin function
    for (int i = 0; i < N; i++) {
        t[i] = i*(2.0*3.14159/(N-1));
        u[i] = sin(t[i]);
    }
    // Create the interpolating polynomial
    adevs::InterPoly p(u,t,N);
    // The data arrays can be deleted now
    delete [] t; delete [] u;
    // Compute several points with sin, its derivative, and the polynomial
    // inside and a little beyond the interval spanned by the data
    for (double t = 0; t < 2.0*3.14159+0.5; t += 0.01)
        cout << t
              << " " << sin(t) << " " << p(t)
              << " " << cos(t) << " " << p.derivative(t)
              << endl;
    // Done
    return 0;
}
```




Figure 12.1: The function  $\sin(t)$  and some interpolating polynomials with data spanning the interval  $[0, 2\pi]$ .




Figure 12.2: The time derivative of  $\sin(t)$  and the time derivative of some interpolating polynomials with data spanning the interval  $[0, 2\pi]$ .