

## Chapter 4

# Simulation Programming with VBASim in MATLAB

This chapter shows how simulations of some of the examples in Chap. 3 can be programmed in VBASim in MATLAB. The goals of the chapter are to introduce VBASim, and to hint at the experiment design and analysis issues that will be covered in later chapters.

### 4.1 VBASim Overview

VBASim in MATLAB is a collection of MATLAB Scripts, Functions and Class M-files that aid in developing discrete-event simulations. They are entirely open source and can be modified to suit the user. The random-number and random-variate generation routines are MATLAB translations of the corresponding routines in simlib (Law 2007) which is written in C. VBASim is designed to be easy to understand and use, but not necessarily efficient.

Here is a brief description of the Scripts and Functions in VBASim:

```
function VBASimInit:  Initializes VBASim for use, typically called before
                      the start of each replication.
function Schedule:    Schedules future events.
function SchedulePlus: Schedules future events and allows an object to
                      be stored with the event.
function Report:      Writes a result to a specific row and column of an Excel
                      worksheet.
function ClearStats:  Clears certain statistics being recorded by VBASim.
function InitializeRNSeed: Initializes the random-number generator;
                      typically called only once in a simulation.
function Expon:       Generates exponentially distributed random variates.
function Uniform:     Generates uniformly distributed random variates.
function RandomInteger: Generates a random integer.
function Erlang:       Generates Erlang distributed random variates.
```

```
function Normal:   Generates normally distributed random variates.
function Lognormal: Generates lognormally distributed random variates.
function Triangular: Generates triangularly distributed random variates.
```

The random-variate generation functions take two types of arguments: parameters, and a random-number stream; the random-number stream is always the last argument. For instance

```
X = Uniform(10, 45, 2)
```

generates random-variates that are uniformly distributed between 10 and 45, using stream 2.

As you already know, the pseudorandom numbers we use to generate random variates in simulations are essentially a long list. Random-number streams are just different starting places in this list, spaced far apart. The generator in VBASim (which is a translation of the generator in Law 2007) has 100 streams. Calling function `InitializeRNSeed` sets the initial position of each stream. Then each subsequent call to a variate-generation routine using stream # advances stream # to the next pseudorandom number. The need for streams is discussed in Chap. 7, but it is important to know that any stream is as good as any other stream in a well-tested generator.

Here is a brief description of the Class M-files in VBASim:

```
Entity:   Object used to model transient items that flow through the system.
FIFOQueue: Object to hold entities in first-in-first-out order.
Resource: Object used to model a scarce quantity.
EventNotice: Object used to represent a future event.
EventCalendar: Data structure holding event notices in chronological order.
CTStat:   Object used to record continuous-time statistics.
DTStat:   Object used to record discrete-time statistics.
```

## 4.2 Simulating the $M(t)/M/\infty$ Queue

Here we consider the parking lot example of Sect. 3.1, a queueing system with time-varying car arrival rate, exponentially distributed parking time and an infinite number of parking spaces. The simulation program consists of some global declarations (Fig. 4.1), a main program and some event routines (Fig. 4.2), an initialization sub (Fig. 4.3), a function to generate car arrivals (Fig. 4.4), and the support functionality provided by VBASim. Two important aspects of VBASim are illustrated by this model: event scheduling and collecting time-averaged statistics.

The key state variable that we need to track is the number of cars in the lot. There are two essential events that affect the state, the arrival of a car and the departure of a car, and a third event we will use to stop the simulation at a specific point in time. The tricky part is that there can be an unlimited number of pending “departure” events; in fact, there are as many pending departure events as there are cars in the lot.

Therefore, having a unique variable to represent the scheduled time of each pending event, as was used for the TTF example in Chap. 2, will not work.

A class variable MyGV in MyGVAR class is defined to store event related global variables in its properties.

```
classdef MyGVAR < handle
    properties
        % Set up event handling
        Clock      %simulation global clock
        Calendar    %event calendar
        % Set up Collections to be reinitialized between replications
        TheCTStats  % continuous-time statistics
        TheDTStats  % discrete-time statistics
        TheQueues   % queues
        TheResources % resources
    end
    methods
        % Initialize MyGVAR object variables
        function obj = MyGVAR(varargin)
            switch nargin
                case 0
                    % Set default values
                    obj.Clock = 0;
                    obj.Calendar = EventCalendar;
                    obj.TheCTStats = [];
                    obj.TheDTStats = [];
                    obj.TheQueues = [];
                    obj.TheResources = [];
                otherwise
                    error('Cannot set property values')
            end
        end
    end
end
```

To handle event scheduling, VBASim has an event calendar data structure called MyGV.Calendar, a function called Schedule for putting events on the MyGV.Calendar, and a method called Remove for extracting the next event from the calendar. An event in VBASim is a MATLAB object of class EventNotice which has (at least) two properties: EventType and EventTime. The statement

```
Schedule(Name, Increment, MyGV)
```

creates an EventNotice, assigns the character string Name to the EventType property, assigns the value MyGV.Clock + Increment to the EventTime property, and schedules the EventNotice on the MyGV.Calendar in the correct chronological order. The method Calendar.Remove extracts the chronologically next EventNotice from the MyGV.Calendar, making its EventType and EventTime properties available to advance time and execute the proper event.

```

% Example illustrating use of VBASim for simulation of M(t)/M/infinity Queue
% parking lot example.
% In this version parking time averages 2 hours;
% the arrival rate varies around 100 per hour;
% the lot starts empty, and we look at a 24-hour period.

% See VBASim for generic declarations
% See Class Definitions for the supporting VBASim classes

%Initialize global variables
global MyGV
MyGV = MyGVVar;

global QueueLength      %Use to keep statistics on N
global N                %Number in queue
global MaxQueue         %Largest observed value of N

```

**Fig. 4.1** Declarations for the parking lot simulation.

The simulation main program `MtMInf.m` in Fig. 4.2 illustrates how the event-related features of VBA<sub>Sim</sub> work. The following four statements, or ones very similar, will appear in all simulations using VBA<sub>Sim</sub>:

```

NextEvent = MyGV.Calendar.Remove
MyGV.Clock = NextEvent.EventTime
switch NextEvent.EventType

```

Since our simulation will repeatedly remove the next `EventNotice` from `MyGV.Calendar`, we need an object of this class to which to assign it. The statement `NextEvent = MyGV.Calendar.Remove` illustrates how the `Remove` method extracts the next event, after which we advance the simulation `MyGV.Clock` to time `NextEvent.EventTime` and execute the event indicated by `NextEvent.EventType`.

VBA<sub>Sim</sub>'s `Schedule` function places events on the `MyGV.Calendar`; for instance, the statement

```
Schedule('EndSimulation', 24, MyGV)
```

creates an `EventNotice` of type `EndSimulation` to occur 24 hours from the time currently on `MyGV.Clock` (which is 0 at the beginning of the simulation). Notice that VBA<sub>Sim</sub> requires the user to decide on the base time unit in the simulation and to use it consistently throughout. In this simulation the base time unit is hours.

The key state variable in this simulation is `N`, the current number of cars in the parking lot, and we want statistics on it. VBA<sub>Sim</sub> contains a `CTStat` class that can be used to record time-averaged statistics. Here is how it works: First, we declare a new `CTStat` using the statment

```
QueueLength = CTStat
```

## MtmInf.m

```

% Initialize
QueueLength = CTStat;
MyInit(MyGV); % special one-time initializations

for Reps = 1:1000
    % Initializations for each replication
    N = 0; %Number in queue
    MaxQueue = 0; %Largest observed value of N
    VBASimInit(MyGV); %initialize VBASim for each replication

    Schedule('Arrival', NSPP(1, MyGV), MyGV);
    Schedule('EndSimulation', 24, MyGV);

    endsimulation = false;
    while ~endsimulation
        NextEvent = MyGV.Calendar.Remove;
        MyGV.Clock = NextEvent.EventTime;
        switch NextEvent.EventType
            case 'Arrival'
                Arrival(MyGV);
            case 'Departure'
                Departure;
        end
        endsimulation = strcmp(NextEvent.EventType, 'EndSimulation');
    end

    % Write output report for each replication

    Report(QueueLength.Mean, 'MtMInf.dat', '%f');
    Report(' ', 'MtMInf.dat');
    Report(MaxQueue, 'MtMInf.dat', '%f');
    Report('\n', 'MtMInf.dat');
end

```

## Arrival.m

```

function Arrival(MyGV)

global QueueLength %Use to keep statistics on N
global N %Number in queue
global MaxQueue %Largest observed value of N
global MeanParkingTime %Parameters we may want to change

% Arrival event

% Schedule next arrival
Schedule('Arrival', NSPP(1, MyGV), MyGV);

% Update number in queue and max
N = N + 1;
QueueLength.Record(N);
if N > MaxQueue
    MaxQueue = N;
end

% Schedule departure
Schedule('Departure', Expon(MeanParkingTime, 2), MyGV);

end

```

Departure.m

```
function Departure()

global QueueLength      %Use to keep statistics on N
global N                %Number in queue
% End of service event

% Update number in queue
N = N - 1;
QueueLength.Record(N);

end
```

**Fig. 4.2** Main program and event functions for the parking lot simulation.

as shown in the global declarations (Fig. 4.1). Second, the `CTStat` can (and usually should) be added to a special collection called `MyGV.TheCTStats`, as shown in Fig. 4.3. VBA<sub>Sim</sub> reinitializes any `CTStat` in `MyGV.TheCTStats` collection whenever `VBASimInit` is executed, which is typically at the beginning of each replication. Next, whenever the value of the variable of interest changes, the `Record` method of the `CTStat` is employed to record the change (which means it is called just after the change occurs); in this simulation the statement is `QueueLength.Record(N)`, as shown in Fig. 4.2. Finally, the time-averaged value of the `CTStat` can be computed using the `Mean` method of the `CTStat`, as in `QueueLength.Mean`.

Notice that VBA<sub>Sim</sub> also has a `Report` sub that writes a value or character string to a text file. The syntax is

```
Report(string, filename) Report(value, filename, format)
```

See Figs. 4.2 and 4.3.

Recall that the arrival rate (in cars per hour) to the parking lot was modeled by the function  $\lambda(t) = 1000 + 100 \sin(\pi t/12)$ . To make the simulation execute more quickly for the purpose of this introduction, we change that rate to  $\lambda(t) = 100 + 10 \sin(\pi t/12)$ , so that the arrival rate varies between 90 and 110 cars per hour, depending on the hour of the day  $t$ .

The function `NSPP` shown in Fig. 4.4 generates interarrival times (time gaps between arrivals) from a nonstationary Poisson process with this arrival rate. The formal definition of a nonstationary Poisson process is a topic of Chap. 6. However, we provide an intuitive justification for the working of function `NSPP` here:

A stationary Poisson process has times between arrivals that are exponentially distributed with a fixed rate  $\lambda$  (or equivalently a constant mean time between arrivals  $1/\lambda$ ). The inverse cdf method for generating exponentially distributed random variates was described in Chap. 2.2.1. The maximum arrival rate for  $\lambda(t)$  is 110 cars per hour, so function `NSPP` generates possible arrivals using a stationary arrival process with rate  $\lambda = 110$ . To achieve the time-varying arrival rate, it

```

function MyInit (MyGV)

global QueueLength      %Use to keep statistics on N
global MeanParkingTime  %Parameters we may want to change

% Initialize the simulation
InitializeRNSeed;
MeanParkingTime = 2;
MyGV.TheCTStats = [MyGV.TheCTStats, QueueLength];

% Write headings for the output reports

Report('Average Number In Queue', 'MtMInf.dat');
Report(',', 'MtMInf.dat');
Report('Maximum Number in Queue', 'MtMInf.dat');
Report('\n', 'MtMInf.dat');
end

```

**Fig. 4.3** Initializing the parking lot simulation.

```

function result = NSPP(Stream, MyGV)
% This function implements thinning to generate interarrival times
% from the nonstationary Poisson arrival process representing
% car arrivals. Time units are minutes.

PossibleArrival = MyGV.Clock + Expon(1 / 110, Stream);

while ~(Uniform(0, 1, Stream) < (100 + 10 * sin(pi * ...
    PossibleArrival / 12)) / 110)

PossibleArrival = PossibleArrival + Expon(1 / 110, Stream);
end

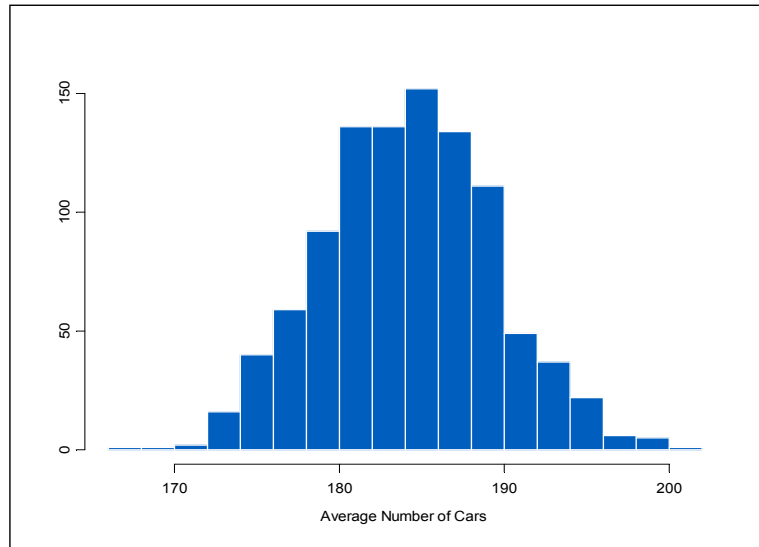
result = PossibleArrival - MyGV.Clock;
end

```

**Fig. 4.4** Function to generate interarrival times to the parking lot.

only accepts a *possible* arrival at time  $t$  as an *actual* arrival with probability  $\lambda(t)/\lambda$ . Thus, a possible arrival that is to occur at a time when  $\lambda(t) = 110$  will always be an actual arrival, while a possible arrival that is to occur when  $\lambda(t) = 90$  will only have a  $90/110$  probability of being an actual arrival. That this method, which is called “thinning,” generates a nonstationary Poisson process with the desired rate is discussed in Chap. 6.

Fig. 4.5 shows a histogram of the 1000 daily averages of the number of cars in the parking lot obtained by running the simulation for 1000 replications; the overall average of these averages is  $184.2 \pm 0.3$  cars, where the “ $\pm 0.3$ ” part comes from a



**Fig. 4.5** Histogram of the daily average number of cars.


95% confidence interval on the mean (confidence intervals are a subject of Chap. 7). Thus, the simulation provides a pretty precise estimate of the time-average mean number of cars that would be found in the (conceptually) infinite-size garage during a day.

The histogram shows that the average number of cars in the garage can vary substantially from day to day, so we certainly would not want to build a garage with a capacity of, say, 185 cars. Further, averaging over the day masks the largest number of cars in the garage during the day, and that number is more useful for selecting a finite capacity for the garage.

VBA<sub>Sim</sub> provides no special support for the maximum statistic, but since we have access to everything in a VBA<sub>Sim</sub> simulation we can easily record whatever we want. Here we define a variable `MaxQueue` which is initialized to 0 at the beginning of each replication, and is increased to the current value of `N` whenever `N` exceeds the previous value of `MaxQueue`. See in particular `Arrival.m` in Fig. 4.2.

Suppose that we want the parking lot to be of adequate size 99% of the time. Since we record the maximum size on 1000 replications, we could use the 990th value (sorted from smallest to largest) as the size of the garage, which turned out to be 263 cars in this simulation. Figure 4.6 shows the empirical cumulative distribution (ecdf) of the 1000 maximums recorded by the simulation. The ecdf treats each observed value as equally likely (and therefore as having probability 1/1000), and plots the sorted maximum values on the horizontal axis and the cumulative probability of each observation on the vertical axis. The plot shows how the cumulative probability of 0.99 maps to the value of 263 cars, which might be a very reasonable capacity for the garage. Putting a confidence interval on this value is quite different





**Fig. 4.6** Empirical cdf of the daily maximum number of cars in the parking lot.

than putting one on the mean, and will be discussed in Chap. 7. Without a confidence interval (or some measure of error) we cannot be sure if 1000 replications is really enough to estimate the 99th percentile of the maximum.

#### 4.2.1 Issues and Extensions

1. The  $M(t)/M/\infty$  simulation presented here simulates 24 hours of parking lot operation, and treats each 24-hour period as an independent replication starting with an empty garage. This only makes sense if the garage is emptied each day, for instance if the mall closes at night. Is the assumed arrival rate  $\lambda(t)$  appropriate for a mall that closes at night?
2. Suppose that the parking lot serves a facility that is actually in operation 24 hours a day, seven days per week (that is, all the time). How should the simulation be initialized, and how long should the run length be in this case?
3. How could the simulation be initialized so that there are 100 cars in the parking lot at time 0?
4. When this example was introduced in Sect. 3.1, it was suggested that we size the garage based on the (Poisson) distribution of the number of cars in the garage at the point in time when the mean number in the garage was maximum. Is that what we did, empirically, here? If not, how is the quantity we estimated by simulation related to the suggestion in Sect. 3.1 (for instance, will the simulation tend to suggest a bigger or smaller garage than the analytical solution in Sect. 3.1)?

5. One reason that this simulation executes quite slowly when  $\lambda(t) = 1000 + 100\sin(\pi t/12)$  is that the thinning method we used is very inefficient (lots of possible arrivals are rejected). Speculate about ways to make it faster.
6. For stochastic processes experts: Another reason that the simulation is slow when  $\lambda(t) = 1000 + 100\sin(\pi t/12)$  is that there can be 1000 or more pending departure events on `MyGV.Calendar` at any time, which means that scheduling a new event in chronological order involves a slow search. However, it is possible to exploit the memoryless property of the exponential distribution of parking time to create an equivalent simulation that has only two pending events (the next car arrival and next car departure) at any point in time. Describe how to do this.

### 4.3 Simulating the $M/G/1$ Queue

Here we consider the hospital example of Sect. 3.2, a queueing system with Poisson arrival process, some (as yet unspecified) service-time distribution, and a single server (either a receptionist or an electronic kiosk); in other words, an  $M/G/1$  queue. Patient waiting time is the key system performance measure, and the long-run average waiting time in particular.

Recall that Lindley's Equation (3.3) provides a shortcut way to simulate successive customer waiting times:

$$Y_0 = 0 \quad X_0 = 0$$

$$Y_i = \max\{0, Y_{i-1} + X_{i-1} - A_i\}, \quad i = 1, 2, \dots$$

where  $Y_i$  is the  $i$ th customer's waiting time,  $X_i$  is that customer's service time, and  $A_i$  is the interarrival time between customers  $i-1$  and  $i$ . Lindley's equation avoids the need for an event-based simulation, but is limited in what it produces (how would you track the time-average number of customers in the queue?). In this section we will describe both recursion-based and event-based simulations of this queue, starting with the recursion.

#### 4.3.1 Lindley Simulation of the $M/G/1$ Queue

To be specific, suppose that the mean time between arrivals is 1 minute, with the distribution being exponential, and the mean time to use the kiosk is 0.8 minutes (48 seconds), with the distribution being an Erlang-3. An Erlang- $p$  is the sum of  $p$  i.i.d. exponentially distributed random variables, so an Erlang-3 with mean 0.8 is the sum of 3 exponentially distributed random variables each with mean  $0.8/3$ .

In Sect. 4.3 we noted that the waiting-time random variables  $Y_1, Y_2, \dots$  converge in distribution to a random-variable  $Y$ , and it is  $\mu = E(Y)$  that we will use to summarize the performance of the queueing system. We also noted that  $\bar{Y}(m) = m^{-1} \sum_{i=1}^m Y_i$

```

Report('Average Wait', 'Lindley.dat');
Report('\n', 'Lindley.dat');
InitializeRNSeed;
m = 55000;
d = 5000;

for Rep = 1:1
    Y = 0;
    SumY = 0;
    tic
    for i = 1:d
        a = Expon(1, 1);
        X = Erlang(3, 0.8, 2);
        Y = max(0, Y + X - a);
    end
    for i = d + 1:m
        a = Expon(1, 1);
        X = Erlang(3, 0.8, 2);
        Y = max(0, Y + X - a);
        SumY = SumY + Y;
    end
    Report(SumY / double(m - d), 'Lindley.dat', '%f\n');
end

```

**Fig. 4.7** Simulation of the  $M/G/1$  queue using Lindley's equation.

converges with probability 1 to  $\mu$  as the number of customers simulated  $m$  goes to infinity.

All of this suggests that we make a very long simulation run (large  $m$ ) and estimate  $\mu$  by the average of the observed waiting times  $Y_1, Y_2, \dots, Y_m$ . But this is not what we will do, and here is why: Any  $m$  we pick is not  $\infty$ , so the waiting times early in the run—which will tend to be smaller than  $\mu$  because the queue starts empty—will likely pull  $\bar{Y}(m)$  down. To reduce this effect, we will let the simulation generate waiting times for a while (say  $d$  of them) before starting to actually include them in our average. We will still make  $m$  large, but our average will only include the last  $m - d$  waiting times. That is, we will use as our estimator the truncated average

$$\bar{Y}(m, d) = \frac{1}{m - d} \sum_{i=d+1}^m Y_i. \quad (4.1)$$

In addition, we will not make a single run of  $m$  customers, but instead will make  $n$  replications. This yields  $n$  i.i.d. averages  $\bar{Y}_1(m, d), \bar{Y}_2(m, d), \dots, \bar{Y}_n(m, d)$  to which we can apply standard statistical analysis. This avoids the need to directly estimate the asymptotic variance  $\gamma^2$ , a topic we defer to later chapters.

Figure 4.7 shows a VBA<sup>Sim</sup> simulation of the  $M/G/1$  queue using Lindley's equation. In this simulation  $m = 55,000$  customers, we discard the first  $d = 5000$  of

**Table 4.1** Ten replications of the  $M/G/1$  queue using Lindley's equation.

replication	$Y(55000, 5000)$
1	2.191902442
2	2.291913404
3	2.147858324
4	2.114346960
5	2.031447995
6	2.110924602
7	2.132711743
8	2.180662859
9	2.139610760
10	2.146212039
average	2.148759113
std dev	0.066748617

them, and make  $n = 10$  replications. The ten replication averages are individually written to a text file called "Lindley.dat" and are displayed in Table 4.1.

Notice that the average waiting time is a bit over 2 minutes, and that MATLAB, like all programming languages, displays a very large number of output digits. How many are really meaningful? A confidence interval is one way to provide an answer.

Since the across-replication averages are i.i.d., and since each across-replication average is itself the within-replication average of a large number of individual waiting times (50,000 to be exact), the assumption of independent, normally distributed output data is reasonable. This justifies a  $t$ -distribution confidence interval on  $\mu$ . The key ingredient is  $t_{1-\alpha/2, n-1}$ , the  $1 - \alpha/2$  quantile of the  $t$  distribution with  $n - 1$  degrees of freedom. If we want a 95% confidence interval, then  $1 - \alpha/2 = 0.975$ , and our degrees of freedom are  $10 - 1 = 9$ . Since  $t_{0.975, 9} = 2.26$ , we get  $2.148759113 \pm (2.26)(0.066748617)/\sqrt{10}$  or  $2.148759113 \pm 0.047703552$ . This implies that we can claim with high confidence that  $\mu$  is around 2.1 minutes, or we could give a little more complete information as  $2.14 \pm 0.05$  minutes. Any additional digits are statistically meaningless.

Is an average of 2 minutes too long to wait? To actually answer that question would require some estimate of the corresponding wait to see the receptionist, either from observational data or a simulation of the current system. Statistical comparison of alternatives is a topic of Chap. 8.

### 4.3.2 Event-based Simulation of the $M/G/1$ Queue

The simulation program consists of some global declarations (Fig. 4.8), a main program (Fig. 4.9), some event functions (Fig. 4.10), an initialization function (Fig. 4.11), and the support functionality provided by VBA<sub>Sim</sub>. Four VBA<sub>Sim</sub> class objects and one function are illustrated by this model: `Entity`, `FIFOQueue`, `Resource`, `DTStat` and `ClearStats`. At a high level, here is what they do:

```

% Example illustrating use of VBASim for
% simulation of M/G/1 Queue.

% See VBASim for generic declarations
% See Class Definitions for the supporting VBASim classes

% Parameters we may want to change
global MeanTBA           % mean time between arrivals
global RunLength         % run length
global WarmUp            % "warm-up" time

% Global objects needed for simulation
% These will usually be queues and statistics

global Queue             %customer queue
global Wait              %discrete-time statistics on customer waiting
global Server            %server resource

%Initialize global variables
global MyGV
MyGV = MyGVar;

```

**Fig. 4.8** Declarations for the hospital simulation.

- Entity objects are used to model transient items, such as transactions or customers that pass through the system. They can have attributes (MATLAB class properties) that they carry with them; by default they have an attribute called `CreateTime` which is set to the value of `MyGV.Clock` when an Entity object is created. In this simulation the Entity objects represent patients or visitors.
- `FIFOQueue` is a MATLAB class collection, much like `MyGV.Calendar`, that is used for holding Entity objects in first-in-first-out order; it also records time-average number in the queue statistics. In this simulation the queue represents the patients using or waiting for the kiosk.
- `Resource` objects represent scarce quantities like workers, machines and computers that are needed (typically) to serve or process an Entity in some way; they also record the average number of busy resource units. In this simulation the `Resource` object represents the kiosk.
- `DTStat` is an object for recording discrete-time statistics; it is the companion to `CTStat` for continuous-time statistics. In this simulation a `DTStat` is used to record total waiting-time statistics, where we will define “total waiting time” to be the time from patient arrival until they are finished with the kiosk. Notice that this is (intentionally) different from the definition in Sect. 4.3.1, where the wait only included the time to reach the front of the line.
- `ClearStats` is a function that reinitializes all statistical variables found in two collections, `MyGV.TheCTStats` and `TheDTStats`. `FIFOQueue` and

```

MyInit(MyGV); % special one-time initializations

for Reps = 1:10

    VBA SimInit(MyGV); %initialize VBA Sim for each replication

    Schedule('Arrival', Expon(MeanTBA, 1), MyGV);
    Schedule('EndSimulation', RunLength, MyGV);
    Schedule('ClearIt', WarmUp, MyGV);

    endsimulation = false;
    while ~endsimulation
        NextEvent = MyGV.Calendar.Remove;
        MyGV.Clock = NextEvent.EventTime;
        switch NextEvent.EventType
            case 'Arrival'
                Arrival(MyGV);
            case 'EndOfService'
                EndOfService(MyGV);
            case 'ClearIt'
                ClearStats(MyGV);
            end
        endsimulation = strcmp(NextEvent.EventType, 'EndSimulation');
    end

    % Write output report for each replication

    Report(Wait.Mean, 'MG1.dat', '%f,');
    Report(Queue.Mean, 'MG1.dat', '%f,');
    Report(Queue.NumQueue, 'MG1.dat', '%f,');
    Report(Server.Mean, 'MG1.dat', '%f\n');

end

```

**Fig. 4.9** Main program for the hospital simulation.

Resource objects each create a CTStat which is automatically added to MyGV.TheCTStats. When the programmer creates a custom CTStat or DTStat then they must add it to the appropriate collection.

Figure 4.8 shows the declarations of the FIFOQueue, DTStat and Resource objects, specifically:

```

global Queue           %customer queue
global Wait            %discrete-time statistics on customer waiting
global Server          %server resource

```

These are declared globally because there is only one of each and they will be referenced from many places in the simulation code.

This contrasts with the `Entity` objects that will be created and discarded as needed to represent patients coming and going. For instance, consider these four statements in `Arrival`:

```
Customer = Entity;
Queue.Add(Customer);
clearvars Customer;
```

We can access its attributes using the `.` extension, as in `Customer.CreateTime`. The `Add` method of the `FIFOQueue` object `Queue` places `Entity` objects into the queue in order of arrival. Finally, the `clearvars Customer` statement frees the name `Customer` to be used again for a new `Entity`. This statement is not absolutely necessary unless we execute the statement `Customer = Entity` a second time within function `Arrival` because the name `Customer` will automatically be released when we exit the function. However, we think it is good programming practice to release the object name when it is no longer needed so there is no confusion. Notice that the `Entity` object is not lost because it is stored in `Queue`.

Moving into the function `EndOfService` event routine, the following three statements remove the first customer from the queue, use its `CreateTime` attribute to compute the total waiting time, record this value using the `DTStat` object `Wait`, and finally release the name `DepartingCustomer`.

```
DepartingCustomer = Queue.Remove;
Wait.Record(MyGV.Clock - DepartingCustomer.CreateTime);
clearvars DepartingCustomer;
```

Before a `Resource` object can be used, its capacity (number of identical units) must be set. In function `MyInit` this is accomplished by using the object's `SetUnit` method, `Server.SetUnits (1)`. If, for instance, there were 3 identical kiosks then this statement would be `Server.SetUnits (3)`. To make one (or more) units of a `Resource` busy, the `Seize` method is employed, while idling a `Resource` is accomplished via the `Free` method, as shown in the event functions.

In this simulation we are interested in long-run performance, so the length of a replication is determined by whatever we decide is long enough (which is not an easy decision, actually). When we used Lindley's equation it was natural to specify the replication length in terms of the number of customers simulated. However, in more complex simulations with many different outputs, it is far more common to specify the replication length by a stopping time  $T$  chosen so that it will be long enough for all of the outputs. Similarly, if we plan to discard data, then it is easier to specify a time  $T_d$  at which all statistics will be cleared.

To control the replication length by time, we schedule two events: an "EndSimulation" event scheduled to occur at time  $T = 55,000$  minutes, and a statistics clearing event "ClearIt" that calls the sub `ClearStats` at time  $T_d = 5000$  minutes. *Because the arrival rate of patients and visitors is 1 per minute, these times will ap-*

## Arrival.m

```

function Arrival(MyGV)

% Parameters we may want to change
global MeanTBA      % mean time between arrivals
global MeanST       % mean service time
global Phases       % number of phases in service distribution

% Global objects needed for simulation
% These will usually be queues and statistics

global Queue        %customer queue
global Server       %server resource

% Arrival event
% Schedule next arrival
    Schedule('Arrival', Expon(MeanTBA, 1), MyGV);

% Process the newly arriving customer
    Customer = Entity;
    Queue.Add(Customer);
    clearvars Customer;
% If server is not busy, start service by seizing the server
    if Server.Busy == 0
        Server.Seize(1);
        Schedule('EndOfService', Erlang(Phases, MeanST, 2), MyGV);
    end

end

```

## EndofService.m

```

function EndOfService(MyGV)

% Parameters we may want to change
global MeanST       % mean service time
global Phases       % number of phases in service distribution

% Global objects needed for simulation
% These will usually be queues and statistics
global Queue        %customer queue
global Wait         %discrete-time statistics on customer waiting
global Server       %server resource
% End of service event

% Remove departing customer from queue and record wait time
    DepartingCustomer = Queue.Remove;
    Wait.Record(MyGV.Clock - DepartingCustomer.CreateTime);
    clearvars DepartingCustomer;    % be sure to free up memory

% Check to see if there is another customer; if yes start service
% otherwise free the server

    if Queue.NumQueue > 0
        Schedule('EndOfService', Erlang(Phases, MeanST, 2), MyGV);
    else
        Server.Free(1);
    end

end

```

**Fig. 4.10** Event functions for the hospital simulation.



```

function MyInit(MyGV)
% Parameters we may want to change
global MeanTBA      % mean time between arrivals
global MeanST       % mean service time
global Phases       % number of phases in service distribution
global RunLength    % run length
global WarmUp       % "warm-up" time

% Global objects needed for simulation
% These will usually be queues and statistics

global Queue        %customer queue
global Wait         %discrete-time statistics on customer waiting
global Server       %server resource

% Initialize the simulation
    InitializeRNSeed;
    Server = Resource;
    Queue = FIFOQueue;
    Wait = DTStat;
    Server.SetUnits(1); % set the number of servers to 1
    MeanTBA = 1;
    MeanST = 0.8;
    Phases = 3;
    RunLength = 55000;
    WarmUp = 5000;

% Add queues, resources and statistics that need to be
% initialized between replications to the global collections

    MyGV.TheDTStats = [MyGV.TheDTStats Wait];
    MyGV.TheQueues = [MyGV.TheQueues Queue];
    MyGV.TheResources = [MyGV.TheResources Server];

% Write headings for the output reports

    Report('Average Wait,', 'MG1.dat');
    Report('Average Number in Queue,', 'MG1.dat');
    Report('Number Remaining in Queue,', 'MG1.dat');
    Report('Server Utilization\n', 'MG1.dat');

end

```

**Fig. 4.11** Initializing the hospital simulation.

**Table 4.2** Ten replications of the  $M/G/1$  queue using the event-based simulation.

Rep	Total Wait	Queue	Remaining	Utilization
1	2.996682542	3.01742345	1	0.806654823
2	3.103155842	3.127773149	7	0.807276539
3	2.951068607	2.948352414	2	0.799341091
4	2.848547497	2.846673097	0	0.794701908
5	2.908913572	2.900437432	2	0.798615617
6	2.895622648	2.896900635	3	0.801983001
7	2.909777649	2.901219722	0	0.798658678
8	2.914666297	2.908612119	4	0.795440658
9	2.922193762	2.923535588	0	0.799157017
10	2.873148311	2.862172885	0	0.799143690
average	2.932377673	2.933310049	1.9	0.800097302
stdev	0.07227642	0.082882288	2.282785822	0.004156967
95% ci	0.051654132	0.059233879	1.631449390	0.002970879

proximately correspond to the run lengths of 55,000 and 5000 patients; however the actual number of patients will be random and vary from replication to replication.

Let  $\bar{Y}(T, T_d)$  be a replication average of all waiting times recorded between times  $T_d$  and  $T$ . Clearly  $\bar{Y}(m, d)$  based on count, and  $\bar{Y}(T, T_d)$  based on time, will have different statistical properties, but it is intuitively clear that both will be good estimators of  $\mu$  if their arguments are fixed (not a function of the data) and large enough. Notice also that a run time and deletion time are ideal for continuous-time statistics like the time-average number in queue.

For this event-based simulation it is easy to record and report a number of statistics. `FIFOQueue`, `Resource` and `DTStat` objects all have `Mean` methods for reporting average values, which for this simulation are `Queue.Mean`, `Server.Mean` and `Wait.Mean`, respectively. In addition, the `FIFOQueue` objects have a `NumQueue` method to deliver the current number of entities in the queue; here we use it to report the number of patients in the queue at time 55,000 when the replication ends. Table 4.2 shows the results from 10 replications, along with the overall averages and 95% confidence interval halfwidths.

Again there are meaningless digits, but the confidence intervals can be used to prune them. For instance, for the mean total wait we could report  $2.93 \pm 0.05$  minutes. How does this statistic relate to the  $2.14 \pm 0.05$  minutes reported for the simulation via Lindley's equation? Mean total time in the kiosk system (which is what the event-based simulation estimates) consists of mean time waiting to be served (which is what the Lindley simulation estimates) plus the mean service time (which we know to be 0.8 minutes). So it is not surprising that these two estimates differ by about 0.8 minutes.

### 4.3.3 Issues and Extensions

1. In what situations does it make more sense to compare the simulated kiosk system to simulated data from the current receptionist system rather than real data from the current receptionist system?
2. It is clear that if all we are interested in is mean waiting time, defined either as time until service begins or total time including service, the Lindley approach is superior (since it is clearly faster, and we can always add in the mean service time to the Lindley estimate). However, if we are interested in the distribution of total waiting time, then adding in the mean service time does not work. How can the Lindley recursion be modified to simulate total waiting times?
3. How can the event-based simulation be modified so that it also records waiting time until service begins?
4. How can the event-based simulation be modified to clear statistics after exactly 5000 patients, and to stop at exactly 55,000 patients?
5. The experiment design method illustrated in the event-based simulation is often called the “replication-deletion” method. If we only had time to generate 500,000 waiting times, what issues should be considered in deciding the values of  $n$  (replications),  $m$  (run length) and  $d$  (deletion amount)? Notice that we must have  $nm = 500,000$ , and only  $n(m - d)$  observations will be used for estimating  $\mu$ .
6. An argument against summarizing system performance by long-run measures is that no system stays unchanged forever (55,000 patients is approximately 38 24-hour days during which time there could be staff changes, construction or emergencies), so a measure like  $\mu$  is not a reflection of reality. The counter argument is that it is difficult, if not impossible, to model all of the detailed changes that occur over any time horizon (even the time-dependent arrival process in the  $M(t)/M/\infty$  simulation is difficult to estimate in practice), so long-run performance at least provides an understandable summary measure (“If our process stayed the same, then over the long run...”). Also, it is often mathematically easier to obtain long-run measures than it is to estimate them by simulation (since simulations have to stop). Considering these issues, what sort of analysis makes the most sense for the hospital problem?

## 4.4 Simulating the Stochastic Activity Network

Here we consider the construction example of Sect. 3.4 which is represented as a stochastic activity network (SAN). Recall that the time to complete the project,  $Y$ , can be represented as

$$Y = \max\{X_1 + X_4, X_1 + X_3 + X_5, X_2 + X_5\}$$

where  $X_i$  is the duration of the  $i$ th activity. This simple representation requires that we enumerate all paths through the SAN, so that the project completion time is the longest of these paths. Path enumeration itself can be time consuming, and this approach does not easily generalize to projects that have resources shared between activities, for instance. Therefore, we also present a discrete-event representation which is more complicated, but also more general.

#### 4.4.1 Maximum Path Simulation of the SAN

Figure 4.12 shows a VBASim implementation of the algorithm displayed in Sect. 3.4 and repeated here:

1. set  $s = 0$
2. repeat  $n$  times:
  - a. generate  $X_1, X_2, \dots, X_5$
  - b. set  $Y = \max\{X_1 + X_4, X_1 + X_3 + X_5, X_2 + X_5\}$
  - c. if  $Y > t_p$  then set  $s = s + 1$
3. estimate  $\theta$  by  $\hat{\theta} = s/n$

Since  $\Pr\{Y \leq t_p\}$  is known for this example (see Eq. (3.12)), the true  $\theta = \Pr\{Y > t_p\} = 0.165329707$  when  $t_p = 5$  is also computed by the program so that we can compare it to the simulation estimate. Of course, in a practical problem we would not know the answer, and we would be wasting our time simulating it if we did. Notice that all of the digits in this probability are correct—assuming that the numerical functions in MATLAB did their job—although certainly not practically useful.

The simulation estimate turns out to be  $\hat{\theta} = 0.163$ . A nice feature of a probability estimate that is based on i.i.d. outputs is that an estimate of its standard error is easily computed:

$$\widehat{\text{se}} = \sqrt{\frac{\hat{\theta}(1 - \hat{\theta})}{n}}.$$

Thus,  $\widehat{\text{se}}$  is approximately 0.011, and the simulation has done its job since the true value  $\theta$  is well within  $\pm 1.96\widehat{\text{se}}$  of  $\hat{\theta}$ . This is a reminder that simulations do not deliver *the answer*, like Eq. (3.12), but do provide the capability to estimate the simulation error, and to reduce that error to an acceptable level by increasing the simulation effort (number of replications).

```

% Simple simulation of the SAN using the max function
X = zeros(1,5);

Report('Pr{Y > tp}', 'SANMax.dat');
Report('True Theta\n', 'SANMax.dat');
InitializeRNSeed;
N = 1000;
c = 0;
tp = 5;

for Rep = 1:N
    for i = 1:5
        X(i) = Expon(1, 7);
    end
    Y = max([X(1) + X(4), X(1) + X(3) + X(5), X(2) + X(5)]);
    if Y > tp
        c = c + 1;
    end
end

Report(c / N, 'SANMax.dat', '%f\n');

Theta = 1 - ((tp ^ 2 / 2 - 3 * tp - 3) * exp(-2 * tp) ...
    + (-tp ^ 2 / 2 - 3 * tp + 3) * exp(-tp) + 1 - exp(-3 * tp));
Report(Theta, 'SANMax.dat', '%f\n');

```

**Fig. 4.12** Simulation of the SAN as the maximum path through the network.

#### 4.4.2 Discrete-event Simulation of the SAN

This section uses more advanced MATLAB constructs than any other part of the book and may be skipped without loss of continuity.

As noted in Sect. 3.4, we can think of the completion of a project activity as an event, and when all of the inbound activities  $\mathcal{I}(j)$  to a milestone  $j$  are completed then the outbound activities  $i \in \mathcal{O}(j)$  can be scheduled, where the destination milestone of activity  $i$  is  $\mathcal{D}(i)$ . Thus, the following generic milestone event is the only one needed:

**event milestone** (activity  $\ell$  inbound to node  $j$ )

$\mathcal{I}(j) = \mathcal{I}(j) - \ell$

if  $\mathcal{I}(j) = \emptyset$  then

    for each activity  $i \in \mathcal{O}(j)$

        schedule milestone(activity  $i$  inbound to node  $\mathcal{D}(i)$  to occur  $X_i$  time units

    later)

end if

Of course, this approach shifts the effort from enumerating all of the paths through the SAN to creating the sets  $\mathcal{I}$ ,  $\mathcal{O}$ ,  $\mathcal{D}$ , but these sets have to be either explicitly or implicitly defined to define the project itself. The key lesson from this example, which applies to many simulations, is that it is possible to program a single event routine to handle many simulation events that are conceptually distinct, and this is done by passing event-specific information to the event routine. In this case we need to pass the inbound activity and the target node, and since this information is needed when the event is executed, not when it is scheduled, we need to store it with the event notice. To do this, we will create a new class and make use of a feature of the EventNotice class.

We define a new class `Activity` which has two properties: `WhichActivity` and `WhichNode`:

```
% Object to model an activity-destination node pair
classdef Activity
    properties
        WhichActivity %Integer
        WhichNode %Integer
    end
end
```

This object represents an activity, which has a number  $1, 2, \dots, 5$ , and also a destination node, which we will number with  $j = 1$  for  $a$ ,  $j = 2$  for  $b$ ,  $j = 3$  for  $c$  and  $j = 4$  for  $d$ .

When we schedule an activity to be completed at some time in the future, we will associate an `Activity` with the `EventNotice` using its `WhichObject` property; the `EventNotice` class is reproduced below:

```
% This is a generic EventNotice object with EventTime, EventType
% and WhichObject attributes
classdef EventNotice < handle
    properties
        EventTime %Double
        EventType %String
        WhichObject %Object
        % Add additional problem specific attributes here
    end
end
```

Perhaps the most complex feature of this simulation is how nodes are represented as a two-dimensional array of MATLAB Cell group. A Cell group is a generic MATLAB data structure, very much like a one-dimensional array, except that Cell groups may contain objects, not just numbers. In this case, each element of the array `Nodes{i, j}` is a list of activities, if  $i = 1$  it is inbound activities, and if  $i = 2$  it is outbound activities, for node  $j$ . Thus, `Nodes{1, j}` plays the role of the set  $\mathcal{I}(j)$ , while `Nodes{2, j}` corresponds to the set  $\mathcal{O}(j)$ . Another Cell group called `Destination` associates each activity with its destination and plays the role of

## SAN.m

```

% SAN Simulation using a discrete-event approach

% Nodes is a two-dimensional array of Collections, where each
% Collection is a list of inbound or outbound activities to that node
% For Nodes(i, j) = inbound i=1 or outbound i=2
% node j = 1 for a, j = 2 for b, j = 3 for c, j = 4 for d

global a b c d InTo OutOf
a = 1; b = 2; c = 3; d = 4; InTo = 1; OutOf = 2;

%Initialize global variables
global MyGV
MyGV = MyGVar;

% Initialize
MyInit;

for Rep = 1:1000
    VBASimInit(MyGV);
    SANInit; % initializes the activities in the SAN
    % causes outbound activities of node a to be scheduled
    Milestone(0, a, MyGV);
    endsimulation = false;
    while ~endsimulation
        NextEvent = MyGV.Calendar.Remove;
        MyGV.Clock = NextEvent.EventTime;
        ThisActivity = NextEvent.WhichObject;
        Milestone(ThisActivity.WhichActivity, ThisActivity.WhichNode,...
            MyGV)
        % stop when event calendar is empty
        endsimulation = (MyGV.Calendar.N == 0);
    end
    Report(MyGV.Clock, 'SAN.dat', '%f\n');
end

MyInit.m

function MyInit()
    InitializeRNSeed;
    Report('Completion Time\n', 'SAN.dat');
end

```

**Fig. 4.13** Main program for the discrete-event SAN simulation.

```

function SANInit()

global Nodes Destination
Nodes=cell(2,4);
global a b c d InTo OutOf

% destinations
% Destination(i) corresponds to the destination of activity i
Destination = [b c c d d];

% node a
Inbound = [];
Outbound = [1 2];
Nodes{InTo, a} = Inbound;
Nodes{OutOf, a} = Outbound;
clearvars Inbound Outbound;

% node b
Inbound = 1;
Outbound = [3 4];
Nodes{InTo, b} = Inbound;
Nodes{OutOf, b} = Outbound;
clearvars Inbound Outbound;

% node c
Inbound = [2 3];
Outbound = 5;
Nodes{InTo, c} = Inbound;
Nodes{OutOf, c} = Outbound;
clearvars Inbound Outbound;

% node d
Inbound = [4 5];
Outbound = [];
Nodes{InTo, d} = Inbound;
Nodes{OutOf, d} = Outbound;
clearvars Inbound Outbound;

end

```

**Fig. 4.14** Network description for the discrete-event SAN simulation.



```

function Milestone(ActIn, Node, MyGV)

global Nodes Destination
global InTo OutOf

Inbound = Nodes{InTo, Node};
Outbound = Nodes{OutOf, Node};
m = length(Inbound);

for Incoming = 1:m
    if Inbound(Incoming) == ActIn
        Inbound(Incoming) = [];
        break;
    end
end
Nodes{InTo, Node} = Inbound;

if isempty(Inbound)
    m = length(Outbound);
    for ActOut = 1:m
        ThisActivity = Activity;
        ThisActivity.WhichActivity = Outbound(1);
        ThisActivity.WhichNode = Destination(Outbound(1));
        SchedulePlus('Milestone', Expon(1, 1), ThisActivity, MyGV)
        clearvars ThisActivity;
        Outbound(1) = [];
    end
end
end


```

**Fig. 4.15** Milestone event for the discrete-event SAN simulation.

the set  $\mathcal{D}$ . The key initializations that define the SAN take place in function `SANInit`, shown in Fig. 4.14.

Notice (see Fig. 4.13) that the simulation ends when there are no additional activities remaining to be completed. This can be checked via `MyGV.Calendar.N`, a method that provides the number of events currently on the event calendar.

A difference between this implementation of the SAN simulation and the one in Sect. 4.4.1 is that here we write out the actual time the project completes on each replication. By doing so, we can estimate  $\Pr\{Y > t_p\}$  for any value of  $t_p$  by sorting the data and counting how many out of 1000 replications were greater than  $t_p$ . Figure 4.16 shows the empirical cdf of the 1000 project completion times, which is the simulation estimate of Eq. (3.12).



**Fig. 4.16** Empirical cdf of the project completion times.

#### 4.4.3 Issues and Extensions

1. In real projects there are not only activities, but also limited and often shared resources that are needed to complete the activities. Further, there may be specific resource allocation rules when multiple activities contend for the same resource. How might this be modeled in VBA Sim?
2. Time to complete the project is an important overall measure, but at the planning stage it may be more important to discover which activities or resources are the most critical to on-time completion of the project. What additional output measures might be useful for deciding which activities are “critical?”

### 4.5 Simulating the Asian Option

Here we consider estimating the value of an Asian option

$$v = E \left[ e^{-rT} (\bar{X}(T) - K)^+ \right]$$

as described in Sect. 3.5, where the maturity is  $T = 1$  year, the risk-free interest rate is  $r = 0.05$  and the strike price is  $K = \$55$ . The underlying asset has an initial value of  $X(0) = \$50$  and the volatility is  $\sigma^2 = (0.3)^2$ . Recall that the key quantity is

$$\bar{X}(T) = \frac{1}{T} \int_0^T X(t) dt$$

the time average of a continuous-time, continuous-state geometric Brownian motion process which we cannot truly simulate on a digital computer. Thus, we approximate it by dividing the interval  $[0, T]$  into  $m$  steps of size  $\Delta t = T/m$  and using the discrete approximation

$$\widehat{\bar{X}}(T) = \frac{1}{m} \sum_{i=1}^m X(i\Delta t).$$

This makes simulation possible, since

$$X(t_{i+1}) = X(t_i) \exp \left\{ \left( r - \frac{1}{2} \sigma^2 \right) (t_{i+1} - t_i) + \sigma \sqrt{t_{i+1} - t_i} Z_{i+1} \right\}$$

for any increasing sequence of times  $\{t_0, t_1, \dots, t_m\}$ , where  $Z_1, Z_2, \dots, Z_m$  are i.i.d.  $N(0, 1)$ .

Figure 4.17 is VBASim code that uses  $m = 32$  steps in the approximation, and makes 10,000 replications to estimate  $v$ . Discrete-event structure would slow execution without any obvious benefit, so a simple loop is used to advance time. The value of the option from each replication is written to a text file for post-simulation analysis.

The estimated value of  $v$  is \$2.20 with a relative error of just over 2% (recall that the relative error is the standard error divided by the mean). As the histogram in Fig. 4.18 shows, the option is frequently worthless (approximately 68% of the time), but the average payoff, conditional on the payoff being positive, is approximately \$6.95.

## 4.6 Case Study: Service Center Simulation

This section presents a simulation case based on a project provided by a former student. While still relatively simple, it is more complex than the previous stylized examples, and the answer is not known without simulating. The purpose of this section is to illustrate how one might attack simulation modeling and programming for a realistic problem.

### *Example 4.1 (Fax Center Staffing).*

A service center receives faxed orders throughout the day, with the rate of arrival varying hour by hour. The arrivals are modeled by a nonstationary Poisson process with the rates shown in Table 4.3.

A team of Entry Agents select faxes on a first-come-first-served basis from the fax queue. Their time to process a fax is modeled as normally distributed with mean 2.5 minutes and standard deviation 1 minute. There are two possible outcomes after the Entry Agent finishes processing a fax: either it was a simple fax and the work on it is complete, or it was not simple and it needs to go to a Specialist for further processing. Over the course of a day, approximately 20% of the faxes require a Spe-

```

% Simulation of an Asian Option

% Set option parameters
Replications = 10000;
Maturity = 1;
Steps = 32;
Sigma = 0.3;
InterestRate = 0.05;
InitialValue = 50;
StrikePrice = 55;
Interval = Maturity / Steps;
Sigma2 = Sigma * Sigma / 2;

Report('Option Value\n', 'Option.dat');
InitializeRNSeed;
for i = 1:Replications
    Sum = 0;
    X = InitialValue;
    for j = 1:Steps
        Z = Normal(0, 1, 12);
        X = X * exp((InterestRate - Sigma2) * Interval ...
            + Sigma * sqrt(Interval) * Z);
        Sum = Sum + X;
    end
    Value = exp(-InterestRate * Maturity) * max(Sum / Steps...
        - StrikePrice, 0);
    Report(Value, 'Option.dat', '%f\n');
end

```


**Fig. 4.17** MATLAB Simulation of the Asian option problem.

**Table 4.3** Arrival rate of faxes by hour.

Time	Rate (faxes/minute)
8 AM–9 AM	4.37
9 AM–10 AM	6.24
10 AM–11 AM	5.29
11 AM–12 PM	2.97
12 PM–1 PM	2.03
1 PM–2 PM	2.79
2 PM–3 PM	2.36
3 PM–4 PM	1.04

cialist. The time for a Specialist to process a fax is modeled as normally distributed with mean 4.0 minutes and standard deviation 1 minute.

Minimizing the number of staff minimizes cost, but certain service-level requirements must be achieved. In particular, 96% of all simple faxes should be completed



**Fig. 4.18** Histogram of the realized value of the Asian option from 10,000 replications.

within 10 minutes of their arrival, while 80% of faxes requiring a Specialist should also be completed (by both the Entry Agent and the Specialist) within 10 minutes of their arrival.

The service center is open from 8 AM to 4 PM daily, and it is possible to change the staffing level at 12 PM. Thus, a staffing policy consists of four numbers: the number of Entry Agents and Specialists before noon, and the number of Entry Agents and Specialists after noon. Any fax that starts its processing before noon completes processing by that same agent before the agent goes off duty; and faxes in the queues at the end of the day are processed before the agents leave work and therefore are not carried over to the next day.

*The first step in building any simulation model is deciding what question or questions that the model should answer.* Knowing the questions helps identify the system performance measures that the simulation needs to estimate, which in turn drives the scope and level of detail in the simulation model.

The grand question for the service center is, what is the minimum number of Entry Agents and Specialists needed for both time periods to meet the service-level requirements? Therefore, the simulation must at least provide an estimate of the percentage of faxes of each type entered within 10 minutes, given a specific staff assignment.

Even when there seems to be a clear overall objective (minimize the staff required to achieve the service-level requirement), we often want to consider trade offs around that objective. For instance, if meeting the requirement requires a staff that is so large that they are frequently underutilized, or if employing the minimal staff means that the Entry Agents or Specialists frequently have to work well past

the end of the day, then we might be willing to alter the service requirement a bit. Statistics on the number and the time spent by faxes in queue, and when the last fax of each day is actually completed, provide this information. Including additional measures of system performance, beyond the most critical ones, makes the simulation more useful.

Many discrete-event, stochastic simulations involve entities that dynamically flow through some sort of queueing network where they compete for resources. In such simulations, identifying the entities and resources is a good place to start the model. For this service center the faxes are clearly the dynamic entities, while the Entry Agents and Specialists are resources. The fax machines themselves might also be considered a resource, especially if they are heavily utilized or if outgoing as well as incoming faxes use the same machines. It turns out that for this service center there is a bank of fax machines dedicated to incoming faxes, so it is reasonable to treat the arrival of faxes as an unconstrained external arrival process. This fact was not stated in the original description of the problem; follow-up questions are often needed to fully understand the system of interest.

Whenever there are scarce resources, queues may form. Queues are often first-in-first-out, with one queue for each resource, as they are in this service center. However, queues may have priorities, and multiple queues may be served by the same resource, or a single queue may feed multiple resources. Queueing behavior is often a critical part of the model.

When the simulation involves entities flowing through a network of queues, then there can be two types of arrivals: arrivals from outside of the network and arrivals internal to the network. Outside arrivals are like those we have already seen in the  $M(t)/M/\infty$  and  $M/G/1$  examples. Internal arrivals are departures from one queue that become arrivals to others. How these are modeled depends largely on whether the departure from one queue is an immediate arrival to the next—in which case the departure and arrival events are effectively the same thing—or whether there is some sort of transportation delay—in which case the arrival to the next queue should be scheduled as a distinct event. For the service center the arrival of faxes to the Entry Agents is an outside arrival process, while the 20% of faxes that require a Specialist are internal arrivals from the Entry Agents to the Specialists.

Critical to experiment design is defining what constitutes a replication. Replications should be independent and identically distributed. Since the service center does not carry faxes over from one day to the next, a “day” defines a replication. If faxes did carry over, but all faxes are cleared weekly, then a replication might be defined by a work week. However, if there is always significant carry over from one day to the next, then a replication might have to be defined arbitrarily.

The work day at the service center is eight hours; however the staff does not leave until all faxes that arrive before 4 PM are processed. If we defined a replication to be exactly eight hours then we could be fooled by a staffing policy that allows a large queue of faxes to build up toward the end of the day, since the entry of those faxes would not be included in our statistics. To model a replication that ends when there is no additional work remaining, we will cut off the fax arrivals at 4 PM and then

```

% FaxCenter Simulation
% Parameters we may want to change

global NumAgents          % number of regular agents
global NumSpecialists     % number of special agents
global NumAgentsPM        % number of regular agents after noon
global NumSpecialistsPM   % number of special agents after noon
% Global objects needed for simulation
global RegularQ           % queue for all faxes
global SpecialQ           % queue for special faxes
global RegularWait        % discrete-time statistics on fax waiting
global SpecialWait        % discrete-time statistics on special fax waiting
global Regular10          % discrete-time statistics on < 10 minutes threshold
global Special10          % discrete-time statistics on < 10 minutes threshold
global Agents             % entry agents resource
global Specialists        % specialists resource
global ARate              % arrival rates
global MaxRate            % maximum arrival rate
global Period             % period for which arrival rate stays constant
global NPeriods           % number of periods in a "day"

```

**Fig. 4.19** Global declarations for service center simulation.

end the simulation when the event calendar is empty. This works because idle Entry Agents and Specialists will always take a fax from their queue if one is available.

Rather than walk through the VBASim code line by line, we will point out some highlights to facilitate the reader's understanding of the code.

Figure 4.19 shows the global declarations for the service center simulation. Of particular note are the two `DTStat` statements defining `Regular10` and `Special10`. These will be used to obtain the fraction of regular and special faxes that are processed within the 10-minute requirement by recording a 1 for any fax that meets the requirement, and a 0 otherwise. The mean of these values is the desired fraction.

The main program for the simulation is in Fig. 4.20. Of particular importance is the condition that ends the main simulation loop:

```
endofsimulation = (Calendar.N == 0)
```

The `N` method of `MyGV.Calendar` returns the current number of pending events. When the event calendar is empty, then there are no additional faxes to process, and no pending arrival of a fax. This condition will only hold after 4 PM and once all remaining faxes have been entered.

Also notice the event `ChangeStaff`, which is scheduled to occur at noon (240 minutes). Here we use the `SetUnits` method of the `Resource` to change the staffing levels. The reader should look at the VBASim `Resource` class object and convince themselves that even if we reduce the number of staff at noon the faxes in process will not be affected.

The staffing policy to be simulated can be changed in `FaxCenter.m`. We can simulate different staffing policies, and view the results, without actually touching the MATLAB code.

The function `NSPP_Fax` generates the interarrival times for faxes with the desired time-varying rate; this function will be described in Chap. 6.



```

% FaxCenter Simulation
% Parameters we may want to change
global NumAgents          % number of regular agents
global NumSpecialists     % number of special agents
global NumAgentsPM        % number of regular agents after noon
global NumSpecialistsPM   % number of special agents after noon
% Global objects needed for simulation
global RegularQ           % queue for all faxes
global SpecialQ           % queue for special faxes
global RegularWait        % discrete-time statistics on fax waiting
global SpecialWait        % discrete-time statistics on special fax waiting
global Regularl0          % discrete-time statistics on < 10 minutes threshold
global Speciall0          % discrete-time statistics on < 10 minutes threshold
global Agents             % entry agents resource
global Specialists        % specialists resource
global ARate              % arrival rates
global MaxRate            % maximum arrival rate
global Period             % period for which arrival rate stays constant
global NPeriods           % number of periods in a "day"
%Initialize global variables
global MyGV;
MyGV = MyGVVar; RegularQ = FIFOQueue; SpecialQ = FIFOQueue;
RegularWait = DTStat; SpecialWait = DTStat; Regularl0 = DTStat;
Speciall0 = DTStat; Agents = Resource; Specialists = Resource;
% Read in staffing policy
NumAgents = 15; NumAgentsPM = 9; NumSpecialists = 6; NumSpecialistsPM = 3;
MyInit(MyGV);
for Reps = 1:10
    VBASimInit(MyGV);
    Agents.SetUnits(NumAgents);
    Specialists.SetUnits(NumSpecialists);
    Schedule('Arrival', NSPP_Fax(ARate, MaxRate, NPeriods, ...
        Period, 1, MyGV), MyGV);
    Schedule('ChangeStaff', 4 * 60, MyGV);
    endsimulation = false;
    while ~endsimulation
        NextEvent = MyGV.Calendar.Remove;
        MyGV.Clock = NextEvent.EventTime;
        switch NextEvent.EventType
            case 'Arrival'
                Arrival(MyGV);
            case 'EndOfEntry'
                EndOfEntry(NextEvent.WhichObject, MyGV);
            case 'EndOfEntrySpecial'
                EndOfEntrySpecial(NextEvent.WhichObject, MyGV);
            case 'ChangeStaff'
                Agents.SetUnits(NumAgentsPM);
                Specialists.SetUnits(NumSpecialistsPM);
        end
        % stop when event calendar empty
        endsimulation = (MyGV.Calendar.N == 0);
    end
    % Write output report for each replication
    Report(RegularWait.Mean, 'Fax.dat', '%f,');
    Report(RegularQ.Mean, 'Fax.dat', '%f,');
    Report(Agents.Mean, 'Fax.dat', '%f,');
    Report(SpecialWait.Mean, 'Fax.dat', '%f,');
    Report(SpecialQ.Mean, 'Fax.dat', '%f,');
    Report(Specialists.Mean, 'Fax.dat', '%f,');
    Report(Regularl0.Mean, 'Fax.dat', '%f,');
    Report(Speciall0.Mean, 'Fax.dat', '%f,');
    Report(MyGV.Clock, 'Fax.dat', '%f\n');
end

```

Fig. 4.20 Main program for service center simulation.

Figure 4.21 contains the arrival and end-of-entry events for faxes at the Entry Agents. The next arrival is scheduled only if `MyGV.Clock < RunLength`; in this way we cut off fax arrivals after 4 PM. The `EndOfEntry` event passes 20% of the faxes directly and immediately to the Specialists by executing `SpecialArrival (DepartingFax, MyGV)` and passing the pointer to the `DepartingFax` entity. Equivalently, we could have scheduled a `SpecialArrival` event to occur zero time units into the future (or nonzero time units if it takes time to transport the fax).

The `Record` method of the `DTStat Regular10` is used to collect a 0 or 1 depending on whether the total wait was less than 10 minutes.

The arrival and end-of-entry events for the Specialists, shown in Fig. 4.22, work similarly to those of the Entry Agents.

Initializations that occur once are shown in Fig. 4.23.

Ten replications of this simulation with a staffing policy of 15 Entry Agents in the morning and 9 in the afternoon, and 6 Specialists in the morning and 3 in the afternoon, gives  $0.98 \pm 0.02$  for the fraction of regular faxes entered in 10 minutes or less, and  $0.81 \pm 0.06$  for the special faxes. The “ $\pm$ ” are 95% confidence intervals. This policy appears to be close to the requirements, although if we absolutely insist on 80% for the special faxes then additional replications are needed to narrow the confidence interval.

## Arrival.m

```

function Arrival(MyGV)
global MeanRegular      % mean entry time regular faxes
global VarRegular       % variance entry time regular faxes
global RunLength        % length of the working day
% Global objects needed for simulation
global RegularQ         % queue for all faxes
global Agents           % entry agents resource
global ARate            % arrival rates
global MaxRate          % maximum arrival rate
global Period           % period for which arrival rate stays constant
global NPeriods         % number of periods in a "day"
% Schedule next fax arrival if < 4 PM
if MyGV.Clock < RunLength
    Schedule('Arrival', NSPP_Fax(ARate, MaxRate, NPeriods,...
        Period, 1, MyGV), MyGV);
else
    return;
end
% Process the newly arriving Fax
Fax = Entity;
if Agents.Busy < Agents.NumberOfUnits
    Agents.Seize(1);
    SchedulePlus('EndOfEntry', Normal(MeanRegular, VarRegular, 2),...
        Fax, MyGV);
else
    RegularQ.Add(Fax);
end
clearvars Fax;
end

```

## EndOfEntry.m

```

function EndOfEntry(DepartingFax, MyGV)
global MeanRegular      % mean entry time regular faxes
global VarRegular       % variance entry time regular faxes
% Global objects needed for simulation
global RegularQ         % queue for all faxes
global RegularWait      % discrete-time statistics on fax waiting
global Regular10        % discrete-time statistics on < 10 minutes threshold
global Agents           % entry agents resource
% Record wait time if regular; move on if special
if Uniform(0, 1, 3) < 0.2
    SpecialArrival(DepartingFax, MyGV);
else
    Wait = MyGV.Clock - DepartingFax.CreateTime;
    RegularWait.Record(Wait);
    if Wait < 10
        Regular10.Record(1);
    else
        Regular10.Record(0);
    end
end
clearvars DepartingFax;
% Check to see if there is another Fax; if yes start entry otherwise free the agent
if (RegularQ.NumQueue > 0) && (Agents.NumberOfUnits ≥
Agents.Busy)
    DepartingFax = RegularQ.Remove;
    SchedulePlus('EndOfEntry', Normal(MeanRegular, VarRegular, 2),...
        DepartingFax, MyGV);
    clearvars DepartingFax;
else
    Agents.Free(1);
end
end

```

Fig. 4.21 Events for Entry Agents.

## SpecialArrival.m

```

function SpecialArrival(SpecialFax, MyGV)
% Parameters we may want to change

global MeanSpecial      % mean entry time special faxes
global VarSpecial       % variance entry time special faxes

% Global objects needed for simulation

global SpecialQ         % queue for special faxes
global Specialists      % specialists resource

% If special agent available, start entry by seizing the special agent

    if Specialists.Busy < Specialists.NumberOfUnits
        Specialists.Seize(1);
        SchedulePlus('EndOfEntrySpecial', Normal(MeanSpecial, ...
            VarSpecial, 4), SpecialFax, MyGV);
    else
        SpecialQ.Add(SpecialFax);
    end
    clearvars SpecialFax;

end

```

## EndOfEntrySpecial.m

```

function EndOfEntrySpecial(DepartingFax, MyGV)
% Parameters we may want to change
global MeanSpecial      % mean entry time special faxes
global VarSpecial       % variance entry time special faxes
% Global objects needed for simulation
global SpecialQ         % queue for special faxes
global SpecialWait      % discrete-time statistics on special fax waiting
global Special10        % discrete-time statistics on < 10 minutes threshold
global Specialists      % specialists resource
% Record wait time and indicator if < 10 minutes
    Wait = MyGV.Clock - DepartingFax.CreateTime;
    SpecialWait.Record(Wait);
    if Wait < 10
        Special10.Record(1);
    else
        Special10.Record(0);
    end
    clearvars DepartingFax;
% Check to see if there is another Fax; if yes start entry
% otherwise free the specialist

    if (SpecialQ.NumQueue > 0) && (Specialists.NumberOfUnits ≥
Specialists.Busy)
        DepartingFax = SpecialQ.Remove;
        SchedulePlus('EndOfEntrySpecial', Normal(MeanSpecial, VarSpecial, 4), ...
            DepartingFax, MyGV);
        clearvars DepartingFax;
    else
        Specialists.Free(1);
    end
end

```

**Fig. 4.22** Events for Specialists.

### 4.6.1 Issues and Extensions

1. There are many similarities between the programming for this simulation and the event-based simulation of the  $M/G/1$  queue. However, there is an important difference that is due to having multiple agents. For the  $M/G/1$  queue, a single `FIFOQueue` object held both the customer in service (who is at the front of the queue) and the customers waiting for service. This approach does not work for the Fax Center because when there are multiple agents the faxes need not complete entry in the same order in which they arrive. To accommodate this, the `FIFOQueue` holds only those faxes waiting for entry, and the `Entity` representing a fax that is being entered is stored with the `EventNotice` for the end-of-entry event. This is accomplished by the statement

```
SchedulePlus('EndOfEntry', _
    Normal(MeanRegular, VarRegular, 2), DepartingFax, MyGV)
```

`SchedulePlus` allows an object (`DepartingFax` in this case) to be assigned to the `.WhichObject` property of the `EventNotice`. The `Entity` can then be passed to the event using the statement

```
EndOfEntry(NextEvent.WhichObject)
```

2. The fax entry times were modeled as being normally distributed. However, the normal distribution admits negative values, which certainly does not make sense. What should be done about this? Consider mapping negative values to 0, or generating a new value whenever a negative value occurs. Which is more likely to be realistic and why?

## Exercises

1. For the hospital problem, simulate the current system in which the receptionist's service time is well modeled as having an Erlang-4 distribution with mean 0.6 minutes. Compare the waiting time to the proposed electronic kiosk alternative.
2. Simulate an  $M(t)/G/\infty$  queue where  $G$  corresponds to an Erlang distribution with fixed mean but try different numbers of phases. That is, keep the mean service time fixed but change the variability. Is the expected number in queue sensitive to the variance in the service time?
3. Modify the SAN simulation to allow each activity to have a different mean time to complete (currently they all have mean time 1). Use a `Collection` to hold these mean times.
4. Try the following numbers of steps for approximating the value of the Asian option to see how sensitive the value is to the step size:  $m = 8, 16, 32, 64, 128$ .

```

function MyInit(MyGV)
% Parameters we may want to change
global MeanRegular      % mean entry time regular faxes
global VarRegular       % variance entry time regular faxes
global MeanSpecial      % mean entry time special faxes
global VarSpecial       % variance entry time special faxes
global RunLength        % length of the working day
% Global objects needed for simulation
global RegularQ         % queue for all faxes
global SpecialQ         % queue for special faxes
global RegularWait      % discrete-time statistics on fax waiting
global SpecialWait      % discrete-time statistics on special fax waiting
global Regular10        % discrete-time statistics on < 10 minutes threshold
global Special10        % discrete-time statistics on < 10 minutes threshold
global Agents           % entry agents resource
global Specialists      % specialists resource
global ARate            % arrival rates
global MaxRate          % maximum arrival rate
global Period           % period for which arrival rate stays constant
global NPeriods         % number of periods in a "day"
% Initialize the simulation
    InitializeRNSeed;
    MeanRegular = 2.5; VarRegular = 1; MeanSpecial = 4;
    VarSpecial = 1; RunLength = 480;
% Add queues, resources and statistics that need to be
% initialized between replications to the global collections
    MyGV.TheDTStats = [MyGV.TheDTStats RegularWait];
    MyGV.TheDTStats = [MyGV.TheDTStats SpecialWait];
    MyGV.TheDTStats = [MyGV.TheDTStats Regular10];
    MyGV.TheDTStats = [MyGV.TheDTStats Special10];
    MyGV.TheQueues = [MyGV.TheQueues RegularQ];
    MyGV.TheQueues = [MyGV.TheQueues SpecialQ];
    MyGV.TheResources = [MyGV.TheResources Agents];
    MyGV.TheResources = [MyGV.TheResources Specialists];
% Write headings for the output reports
    Report('Ave Reg Wait', 'Fax.dat', '%f,');
    Report('Ave Num Reg Q', 'Fax.dat', '%f, ');
    Report('Agents Busy', 'Fax.dat', '%f,');
    Report('Ave Spec Wait', 'Fax.dat', '%f,');
    Report('Ave Num Spec Q', 'Fax.dat', '%f,');
    Report('Specialists Busy', 'Fax.dat', '%f,');
    Report('Reg < 10', 'Fax.dat', '%f,');
    Report('Spec < 10', 'Fax.dat', '%f,');
    Report('End Time', 'Fax.dat', '%f,');
% Arrival process data
    NPeriods = 8; Period = 60; MaxRate = 6.24;
    ARate(1) = 4.37; ARate(2) = 6.24; ARate(3) = 5.29; ARate(4) = 2.97;
    ARate(5) = 2.03; ARate(6) = 2.79; ARate(7) = 2.36; ARate(8) = 1.04;
end

```

**Fig. 4.23** Initializations for service center simulation.

5. In the simulation of the Asian option, the sample mean of 10,000 replications was 2.198270479, and the standard deviation was 4.770393202. Approximately how many replications would it take to decrease the relative error to less than 1%?
6. For the service center, increase the number of replications until you can be confident that that suggested policy does or does not achieve the 80% entry in less than 10 minutes requirement for special faxes.
7. For the service center, find the minimum staffing policy (in terms of total number of staff) that achieves the service-level requirement. Examine the other statistics generated by the simulation to make sure you are satisfied with this policy.
8. For the service center, suppose that Specialists earn twice as much as Entry Agents. Find the minimum cost staffing policy that achieves the service-level requirement. Examine the other statistics generated by the simulation to make sure you are satisfied with this policy.
9. For the service center, suppose that the staffing level can change hourly, but once an Agent or Specialist comes on duty they must work for four hours. Find the minimum staffing policy (in terms of total number of staff) that achieves the service-level requirement.
10. For the service center, pick a staffing policy that fails to achieve the service level requirements by 20% or more. Rerun the simulation with a replication being defined as exactly 8 hours, but do not carry waiting faxes over to the next day. How much do the statistics differ using the two different ways to end a replication?
11. The function NSPP\_Fax is listed below. This function implements the thinning method described in Sect. 4.2 for a nonstationary Poisson process with piecewise-constant rate function. Study it and describe how it works.

```
function y=NSPP_Fax(ARate, MaxRate, NPeriods,Period, Stream, MyGV)
% This function generates interarrival times from a NSPP with piecewise constant
% arrival rate over a fixed time of Period*NPeriod time units

% ARate = array of arrival rates over a common length Period
% MaxRate = maximum value of ARate
% Period = time units between (possible) changes in arrival rate
% NPeriods = number of time periods in ARate

PossibleArrival = MyGV.Clock + Expon(1 / MaxRate, Stream);
i = min(NPeriods, ceil(PossibleArrival / Period));

while Uniform(0, 1, Stream) ≥ ARate(i) / MaxRate
    PossibleArrival = PossibleArrival + Expon(1 / MaxRate, Stream);
    i = min(NPeriods, ceil(PossibleArrival / Period));
end

y = PossibleArrival - MyGV.Clock;

end
```

12. Beginning with the event-based  $M/G/1$  simulation, implement the changes necessary to make it an  $M/G/s$  simulation (a single queue with any number of

servers). Keeping  $\lambda = 1$  and  $\tau/s = 0.8$ , simulate  $s = 1, 2, 3$  servers and compare the results. What you are doing is comparing queues with the same service capacity, but with 1 fast server as compared to two or more slower servers. State clearly what you observe.

13. Modify the VBA<sub>Sim</sub> event-based simulation of the  $M/G/1$  queue to simulate an  $M/G/1/c$  retrial queue. This means that customers who arrive to find  $c$  customers in the system (including the customer in service) leave immediately, but arrive again after an exponentially distributed amount of time with mean `MeanTR`. Hint: The existence of retrial customers should not affect the arrival process for first-time arrivals.
14. This problem assumes a more advanced background in stochastic processes. In the simulation of the  $M(t)/M/\infty$  queue there could be a very large number of events on the event calendar: one “Arrival” and one “Departure” for *each* car currently in the garage. However, properties of the exponential distribution can reduce this to no more than two events. Let  $\beta = 1/\tau$  be the departure rate for a car (recall that  $\tau$  is the mean parking time). If at any time we observe that there are  $N$  car in the garage (no matter how long they have been there), then the time until the first of these cars departs is exponentially distributed with mean  $1/(N\beta)$ . Use this insight to build an  $M(t)/M/\infty$  simulation with at most two pending events, next arrival and next departure. Hint: Whenever an arrival occurs the distribution of the time until the next departure changes, so the scheduled next departure time must again be generated.
15. The phone desk for a small office is staffed from 8 AM to 4 PM by a single operator. Calls arrive according to a Poisson process with rate 6 per hour, and the time to serve a call is uniformly distributed between 5 and 12 minutes. Callers who find the operator busy are placed on hold, if there is space available, otherwise they receive a busy signal and the call is considered “lost.” In addition, 10% of callers who do not immediately get the operator decide to hang up rather than go on hold; they are not considered lost, since it was their choice. Because the hold queue occupies resources, the company would like to know the smallest capacity (number of callers) for the hold queue that keeps the daily fraction of lost calls under 5%. In addition, they would like to know the long-run utilization of the operator to make sure he or she will not be too busy. Use VBA<sub>Sim</sub> to simulate this system and find the required capacity for the hold queue. Model the callers as class `Entity`, the hold queue as class `FIFOQueue` and the operator as class `Resource`. Use the VBA<sub>Sim</sub> functions `Expon` and `Uniform` for random-variate generation. Use class `DTStat` to estimate the fraction of calls lost (record a 0 for calls not lost, a 1 for those that are lost so that the sample mean is the fraction lost). Use the statistics collected by class `Resource` to estimate the utilization.
16. Software Made Personal (SMP) customizes software products in two areas: financial tracking and contact management. They currently have a customer support call center that handles technical questions for owners of their software from the hours of 8 AM to 4 PM Eastern Time.



When a customer calls they first listen to a recording that asks them to select among the product lines; historically 59% are financial products and 41% contact management products. The number of customers who can be connected (talking to an agent or on hold) at any one time is essentially unlimited. Each product line has its own agents. If an appropriate agent is available then the call is immediately routed to the agent; if an appropriate agent is not available, then the caller is placed in a hold queue (and listens to a combination of music and ads). SMP has observed that hang ups very rarely happen.

SMP is hoping to reduce the total number of agents they need by cross-training agents so that they can answer calls for any product line. Since the agents will not be experts across all products, this is expected to increase the time to process a call by about 5%. The question that SMP has asked you to answer is how many cross-trained agents are needed to provide service at the same level as the current system.

Incoming calls can be modeled as a Poisson arrival process with a rate of 60 per hour. The mean time required for an agent to answer a question is 5 minutes, with the actual time being Erlang-2 for financial calls, and Erlang-3 for contact management calls. The current assignment of agents is 4 for financial and 3 for contact management. Simulate the system to find out how many agents are needed to deliver the same level of service in the cross-trained system as in the current system.