

## Chapter 1

# Introduction to MATLAB

This book is an introduction to two subjects: MATLAB and numerical computing. This first chapter introduces MATLAB by presenting several programs that investigate elementary, but interesting, mathematical problems. If you already have some experience programming in another language, we hope that you can see how MATLAB works by simply studying these programs.

If you want a more comprehensive introduction, an on-line manual from The MathWorks is available. Select **Help** in the toolbar atop the MATLAB command window, then select **MATLAB Help** and **Getting Started**. A PDF version is available under **Printable versions**. The document is also available from The MathWorks Web site [10]. Many other manuals produced by The MathWorks are available on line and from the Web site.

A list of over 600 MATLAB-based books by other authors and publishers, in several languages, is available at [11]. Three introductions to MATLAB are of particular interest here: a relatively short primer by Sigmon and Davis [8], a medium-sized, mathematically oriented text by Higham and Higham [3], and a large, comprehensive manual by Hanselman and Littlefield [2].

You should have a copy of MATLAB close at hand so you can run our sample programs as you read about them. All of the programs used in this book have been collected in a directory (or folder) named

`NCM`

(The directory name is the initials of the book title.) You can either start MATLAB in this directory or use

`pathtool`

to add the directory to the MATLAB path.

---

February 15, 2008

## 1.1 The Golden Ratio

What is the world's most interesting number? Perhaps you like  $\pi$ , or  $e$ , or 17. Some people might vote for  $\phi$ , the *golden ratio*, computed here by our first MATLAB statement.

```
phi = (1 + sqrt(5))/2
```

This produces

```
phi =
    1.6180
```


Let's see more digits.

```
format long
phi

phi =
1.61803398874989
```

This didn't recompute  $\phi$ , it just displayed 15 significant digits instead of 5.

The golden ratio shows up in many places in mathematics; we'll see several in this book. The golden ratio gets its name from the golden rectangle, shown in Figure 1.1. The golden rectangle has the property that removing a square leaves a smaller rectangle with the same shape.



**Figure 1.1.** *The golden rectangle.*

Equating the aspect ratios of the rectangles gives a defining equation for  $\phi$ :

$$\frac{1}{\phi} = \frac{\phi - 1}{1}.$$

This equation says that you can compute the reciprocal of  $\phi$  by simply subtracting one. How many numbers have that property?

Multiplying the aspect ratio equation by  $\phi$  produces the polynomial equation

$$\phi^2 - \phi - 1 = 0.$$

The roots of this equation are given by the quadratic formula:

$$\phi = \frac{1 \pm \sqrt{5}}{2}.$$

The positive root is the golden ratio.

If you have forgotten the quadratic formula, you can ask MATLAB to find the roots of the polynomial. MATLAB represents a polynomial by the vector of its coefficients, in descending order. So the vector

```
p = [1 -1 -1]
```

represents the polynomial

$$p(x) = x^2 - x - 1.$$

The roots are computed by the `roots` function.

```
r = roots(p)
```

produces

```
r =
-0.61803398874989
 1.61803398874989
```

These two numbers are the only numbers whose reciprocal can be computed by subtracting one.

You can use the Symbolic Toolbox, which connects MATLAB to a computer algebra system, to solve the aspect ratio equation without converting it to a polynomial. The equation is represented by a character string. The `solve` function finds two solutions.

```
r = solve('1/x = x-1')
```

produces

```
r =
[ 1/2*5^(1/2)+1/2]
[ 1/2-1/2*5^(1/2)]
```

The `pretty` function displays the results in a way that resembles typeset mathematics.

```
pretty(r)
```

produces

```
[      1/2      ]
[1/2 5      + 1/2]
[      ]
[      1/2      ]
[1/2 - 1/2 5      ]
```

The variable `r` is a vector with two components, the symbolic forms of the two solutions. You can pick off the first component with

```
phi = r(1)
```

which produces

```
phi =
1/2*5^(1/2)+1/2
```

This expression can be converted to a numerical value in two different ways. It can be evaluated to any number of digits using variable-precision arithmetic with the `vpa` function.

```
vpa(phi,50)
```

produces 50 digits.

```
1.6180339887498948482045868343656381177203091798058
```

It can also be converted to double-precision floating point, which is the principal way that MATLAB represents numbers, with the `double` function.

```
phi = double(phi)
```

produces

```
phi =
1.61803398874989
```

The aspect ratio equation is simple enough to have closed-form symbolic solutions. More complicated equations have to be solved approximately. In MATLAB an *anonymous function* is a convenient way to define an object that can be used as an argument to other functions. The statement

```
f = @(x) 1./x-(x-1)
```

defines  $f(x) = 1/x - (x - 1)$  and produces

```
f =
@(x) 1./x-(x-1)
```


The graph of  $f(x)$  over the interval  $0 \leq x \leq 4$  shown in Figure 1.2 is obtained with

```
ezplot(f,0,4)
```

The name `ezplot` stands for “easy plot,” although some of the English-speaking world would pronounce it “e-zed plot.” Even though  $f(x)$  becomes infinite as  $x \rightarrow 0$ , `ezplot` automatically picks a reasonable vertical scale.

The statement

```
phi = fzero(f,1)
```



**Figure 1.2.**  $f(\phi) = 0$ .

looks for a zero of  $f(x)$  near  $x = 1$ . It produces an approximation to  $\phi$  that is accurate to almost full precision. The result can be inserted in Figure 1.2 with

```
hold on
plot(phi,0,'o')
```

The following MATLAB program produces the picture of the golden rectangle shown in Figure 1.1. The program is contained in an M-file named `goldrect.m`, so issuing the command

```
goldrect
runs the script and creates the picture.

% GOLDRECT Plot the golden rectangle

phi = (1+sqrt(5))/2;
x = [0 phi phi 0 0];
y = [0 0 1 1 0];
u = [1 1];
v = [0 1];
plot(x,y,'b',u,v,'b--')
text(phi/2,1.05,'\phi')
text((1+phi)/2,-.05,'\phi - 1')
text(-.05,.5,'1')
text(.5,-.05,'1')
axis equal
axis off
set(gcf,'color','white')
```

The vectors  $\mathbf{x}$  and  $\mathbf{y}$  each contain five elements. Connecting consecutive  $(x_k, y_k)$  pairs with straight lines produces the outside rectangle. The vectors  $\mathbf{u}$  and  $\mathbf{v}$  each contain two elements. The line connecting  $(u_1, v_1)$  with  $(u_2, v_2)$  separates the rectangle into the square and the smaller rectangle. The `plot` command draws these lines—the  $x - y$  lines in solid blue and the  $u - v$  line in dashed blue. The next four statements place text at various points; the string `'\phi'` denotes the Greek letter. The two `axis` statements cause the scaling in the  $x$  and  $y$  directions to be equal and then turn off the display of the axes. The last statement sets the background color of `gcf`, which stands for *get current figure*, to white.

A *continued fraction* is an infinite expression of the form

$$a_0 + \frac{1}{a_1 + \frac{1}{a_2 + \frac{1}{a_3 + \dots}}}$$

If all the  $a_k$ 's are equal to 1, the continued fraction is another representation of the golden ratio:

$$\phi = 1 + \frac{1}{1 + \frac{1}{1 + \frac{1}{\dots}}}$$

The following MATLAB function generates and evaluates truncated continued fraction approximations to  $\phi$ . The code is stored in an M-file named `goldfract.m`.

```
function goldfract(n)
%GOLDFRACT Golden ratio continued fraction.
% GOLDFRACT(n) displays n terms.

p = '1';
for k = 1:n
    p = ['1+1/(' p ')'];
end
p

p = 1;
q = 1;
for k = 1:n
    s = p;
    p = p + q;
    q = s;
end
p = sprintf('%d/%d',p,q)

format long
p = eval(p)

format short
err = (1+sqrt(5))/2 - p
```

The statement

```

goldfract(6)

produces

p =
1+1/(1+1/(1+1/(1+1/(1+1/(1))))))

p =
21/13

p =
1.61538461538462

err =
0.0026

```

The three `p`'s are all different representations of the same approximation to  $\phi$ .

The first `p` is the continued fraction truncated to six terms. There are six right parentheses. This `p` is a string generated by starting with a single '1' (that's `goldfract(0)`) and repeatedly inserting the string '1+1/' in front and the string ')' in back. No matter how long this string becomes, it is a valid MATLAB expression.

The second `p` is an "ordinary" fraction with a single integer numerator and denominator obtained by collapsing the first `p`. The basis for the reformulation is

$$1 + \frac{1}{\frac{p}{q}} = \frac{p+q}{p}.$$

So the iteration starts with

$$\frac{1}{1}$$

and repeatedly replaces the fraction

$$\frac{p}{q}$$

with

$$\frac{p+q}{p}.$$

The statement

```
p = sprintf('%d/%d',p,q)
```

prints the final fraction by formatting `p` and `q` as decimal integers and placing a '/' between them.

The third `p` is the same number as the first two `p`'s, but is represented as a conventional decimal expansion, obtained by having the MATLAB `eval` function actually do the division expressed in the second `p`.

The final quantity `err` is the difference between `p` and  $\phi$ . With only 6 terms, the approximation is accurate to less than 3 digits. How many terms does it take to get 10 digits of accuracy?

As the number of terms  $n$  increases, the truncated continued fraction generated by `goldfract(n)` theoretically approaches  $\phi$ . But limitations on the size of the integers in the numerator and denominator, as well as roundoff error in the actual floating-point division, eventually intervene. Exercise 1.3 asks you to investigate the limiting accuracy of `goldfract(n)`.

## 1.2 Fibonacci Numbers

Leonardo Pisano Fibonacci was born around 1170 and died around 1250 in Pisa in what is now Italy. He traveled extensively in Europe and Northern Africa. He wrote several mathematical texts that, among other things, introduced Europe to the Hindu-Arabic notation for numbers. Even though his books had to be transcribed by hand, they were widely circulated. In his best known book, *Liber Abaci*, published in 1202, he posed the following problem:

*A man put a pair of rabbits in a place surrounded on all sides by a wall. How many pairs of rabbits can be produced from that pair in a year if it is supposed that every month each pair begets a new pair which from the second month on becomes productive?*

Today the solution to this problem is known as the *Fibonacci sequence*, or *Fibonacci numbers*. There is a small mathematical industry based on Fibonacci numbers. A search of the Internet for “Fibonacci” will find dozens of Web sites and hundreds of pages of material. There is even a Fibonacci Association that publishes a scholarly journal, the *Fibonacci Quarterly*.

If Fibonacci had not specified a month for the newborn pair to mature, he would not have a sequence named after him. The number of pairs would simply double each month. After  $n$  months there would be  $2^n$  pairs of rabbits. That’s a lot of rabbits, but not distinctive mathematics.

Let  $f_n$  denote the number of pairs of rabbits after  $n$  months. The key fact is that the number of rabbits at the end of a month is the number at the beginning of the month plus the number of births produced by the mature pairs:

$$f_n = f_{n-1} + f_{n-2}.$$

The initial conditions are that in the first month there is one pair of rabbits and in the second there are two pairs:

$$f_1 = 1, \quad f_2 = 2.$$

The following MATLAB function, stored in the M-file `fibonacci.m`, produces a vector containing the first  $n$  Fibonacci numbers.

```
function f = fibonacci(n)
% FIBONACCI Fibonacci sequence
% f = FIBONACCI(n) generates the first n Fibonacci numbers.
f = zeros(n,1);
```



```
f(1) = 1;
f(2) = 2;
for k = 3:n
    f(k) = f(k-1) + f(k-2);
end
```

With these initial conditions, the answer to Fibonacci's original question about the size of the rabbit population after one year is given by

```
fibonacci(12)
```

This produces

```
1
2
3
5
8
13
21
34
55
89
144
233
```

The answer is 233 pairs of rabbits. (It would be 4096 pairs if the number doubled every month for 12 months.)

Let's look carefully at `fibonacci.m`. It's a good example of how to create a MATLAB function. The first line is

```
function f = fibonacci(n)
```

The first word on the first line says this is a `function` M-file, not a script. The remainder of the first line says this particular function produces one output result, `f`, and takes one input argument, `n`. The name of the function specified on the first line is not actually used, because MATLAB looks for the name of the M-file, but it is common practice to have the two match. The next two lines are comments that provide the text displayed when you ask for `help`.

```
help fibonacci
```

produces

```
FIBONACCI Fibonacci sequence
f = FIBONACCI(n) generates the first n Fibonacci numbers.
```

The name of the function is in uppercase because historically MATLAB was case insensitive and ran on terminals with only a single font. The use of capital letters may be confusing to some first-time MATLAB users, but the convention persists. It

is important to repeat the input and output arguments in these comments because the first line is not displayed when you ask for `help` on the function.

The next line

```
f = zeros(n,1);
```

creates an  $n$ -by-1 matrix containing all zeros and assigns it to `f`. In MATLAB, a matrix with only one column is a column vector and a matrix with only one row is a row vector.

The next two lines,

```
f(1) = 1;
f(2) = 2;
```

provide the initial conditions.

The last three lines are the `for` statement that does all the work.

```
for k = 3:n
    f(k) = f(k-1) + f(k-2);
end
```

We like to use three spaces to indent the body of `for` and `if` statements, but other people prefer two or four spaces, or a tab. You can also put the entire construction on one line if you provide a comma after the first clause.

This particular function looks a lot like functions in other programming languages. It produces a vector, but it does not use any of the MATLAB vector or matrix operations. We will see some of these operations soon.

Here is another Fibonacci function, `fibnum.m`. Its output is simply the  $n$ th Fibonacci number.

```
function f = fibnum(n)
% FIBNUM Fibonacci number.
% FIBNUM(n) generates the nth Fibonacci number.
if n <= 1
    f = 1;
else
    f = fibnum(n-1) + fibnum(n-2);
end
```

The statement

```
fibnum(12)
```

produces

```
ans =
    233
```

The `fibnum` function is *recursive*. In fact, the term *recursive* is used in both a mathematical and a computer science sense. The relationship  $f_n = f_{n-1} + f_{n-2}$  is known as a *recursion relation* and a function that calls itself is a *recursive function*.

A recursive program is elegant, but expensive. You can measure execution time with `tic` and `toc`. Try

```
tic, fibnum(24), toc
```

Do *not* try

```
tic, fibnum(50), toc
```

Now compare the results produced by `goldfract(6)` and `fibonacci(7)`. The first contains the fraction 21/13 while the second ends with 13 and 21. This is not just a coincidence. The continued fraction is collapsed by repeating the statement

```
p = p + q;
```

while the Fibonacci numbers are generated by

```
f(k) = f(k-1) + f(k-2);
```

In fact, if we let  $\phi_n$  denote the golden ratio continued fraction truncated at  $n$  terms, then

$$\frac{f_{n+1}}{f_n} = \phi_n.$$

In the infinite limit, the ratio of successive Fibonacci numbers approaches the golden ratio:

$$\lim_{n \rightarrow \infty} \frac{f_{n+1}}{f_n} = \phi.$$

To see this, compute 40 Fibonacci numbers.

```
n = 40;
f = fibonacci(n);
```

Then compute their ratios.

```
f(2:n) ./ f(1:n-1)
```

This takes the vector containing `f(2)` through `f(n)` and divides it, element by element, by the vector containing `f(1)` through `f(n-1)`. The output begins with

```
2.000000000000000
1.500000000000000
1.666666666666667
1.600000000000000
1.625000000000000
1.61538461538462
1.61904761904762
1.61764705882353
1.61818181818182
```

and ends with

```
1.61803398874990
1.61803398874989
1.61803398874990
1.61803398874989
1.61803398874989
```

Do you see why we chose `n = 40`? Use the up arrow key on your keyboard to bring back the previous expression. Change it to

```
f(2:n)./f(1:n-1) - phi
```

and then press the Enter key. What is the value of the last element?

The population of Fibonacci's rabbit pen doesn't double every month; it is multiplied by the golden ratio every month.

It is possible to find a closed-form solution to the Fibonacci number recurrence relation. The key is to look for solutions of the form

$$f_n = c\rho^n$$

for some constants  $c$  and  $\rho$ . The recurrence relation

$$f_n = f_{n-1} + f_{n-2}$$

becomes

$$\rho^2 = \rho + 1.$$

We've seen this equation before. There are two possible values of  $\rho$ , namely  $\phi$  and  $1 - \phi$ . The general solution to the recurrence is

$$f_n = c_1\phi^n + c_2(1 - \phi)^n.$$

The constants  $c_1$  and  $c_2$  are determined by initial conditions, which are now conveniently written

$$\begin{aligned} f_0 &= c_1 + c_2 = 1, \\ f_1 &= c_1\phi + c_2(1 - \phi) = 1. \end{aligned}$$

Exercise 1.4 asks you to use the MATLAB backslash operator to solve this 2-by-2 system of simultaneous linear equations, but it is actually easier to solve the system by hand:

$$\begin{aligned} c_1 &= \frac{\phi}{2\phi - 1}, \\ c_2 &= -\frac{(1 - \phi)}{2\phi - 1}. \end{aligned}$$

Inserting these in the general solution gives

$$f_n = \frac{1}{2\phi - 1}(\phi^{n+1} - (1 - \phi)^{n+1}).$$

This is an amazing equation. The right-hand side involves powers and quotients of irrational numbers, but the result is a sequence of integers. You can check this with MATLAB, displaying the results in scientific notation.

```
format long e
n = (1:40)';
f = (phi.^(n+1) - (1-phi).^(n+1))/(2*phi-1)
```

The `.` operator is an element-by-element power operator. It is not necessary to use `./` for the final division because  $(2*\text{phi}-1)$  is a scalar quantity. The computed result starts with

```
f =
 1.0000000000000000e+000
 2.0000000000000000e+000
 3.0000000000000000e+000
 5.0000000000000001e+000
 8.0000000000000002e+000
 1.3000000000000000e+001
 2.1000000000000000e+001
 3.4000000000000001e+001
```

and ends with

```
5.7028870000000007e+006
9.2274650000000011e+006
1.4930352000000002e+007
2.4157817000000003e+007
3.9088169000000005e+007
6.3245986000000007e+007
1.0233415500000001e+008
1.6558014100000002e+008
```

Roundoff error prevents the results from being exact integers, but

```
f = round(f)
```

finishes the job.

## 1.3 Fractal Fern


The M-files `fern.m` and `finitefern.m` produce the “Fractal Fern” described by Michael Barnsley in *Fractals Everywhere* [1]. They generate and plot a potentially infinite sequence of random, but carefully choreographed, points in the plane. The command

```
fern
```

runs forever, producing an increasingly dense plot. The command

```
finitefern(n)
```

generates `n` points and a plot like Figure 1.3. The command



**Figure 1.3.** *Fractal fern.*

```
finitefern(n,'s')
```

shows the generation of the points one at a time. The command

```
F = finitefern(n);
```

generates, but does not plot,  $n$  points and returns an array of zeros and ones for use with sparse matrix and image-processing functions.

The NCM collection also includes `fern.png`, a 768-by-1024 color image with half a million points that you can view with a browser or a paint program. You can also view the file with

```
F = imread('fern.png');  
image(F)
```

If you like the image, you might even choose to make it your computer desktop background. However, you should really run `fern` on your own computer to see the dynamics of the emerging fern in high resolution.

The fern is generated by repeated transformations of a point in the plane. Let  $x$  be a vector with two components,  $x_1$  and  $x_2$ , representing the point. There are

four different transformations, all of them of the form

$$x \rightarrow Ax + b,$$

with different matrices  $A$  and vectors  $b$ . These are known as *affine transformations*. The most frequently used transformation has

$$A = \begin{pmatrix} 0.85 & 0.04 \\ -0.04 & 0.85 \end{pmatrix}, \quad b = \begin{pmatrix} 0 \\ 1.6 \end{pmatrix}.$$

This transformation shortens and rotates  $x$  a little bit, then adds 1.6 to its second component. Repeated application of this transformation moves the point up and to the right, heading toward the upper tip of the fern. Every once in a while, one of the other three transformations is picked at random. These transformations move the point into the lower subfern on the right, the lower subfern on the left, or the stem.

Here is the complete fractal fern program.

```
function fern
%FERN MATLAB implementation of the Fractal Fern
%Michael Barnsley, Fractals Everywhere, Academic Press, 1993
%This version runs forever, or until stop is toggled.
%See also: FINITEFERN.

shg
clf reset
set(gcf,'color','white','menubar','none',...
    'numbertitle','off','name','Fractal Fern')
x = [.5; .5];
h = plot(x(1),x(2),'.');
darkgreen = [0 2/3 0];
set(h,'markersize',1,'color',darkgreen,'erasemode','none');
axis([-3 3 0 10])
axis off
stop = uicontrol('style','toggle','string','stop',...
    'background','white');
drawnow

p = [.85 .92 .99 1.00];
A1 = [.85 .04; -.04 .85]; b1 = [0; 1.6];
A2 = [.20 -.26; .23 .22]; b2 = [0; 1.6];
A3 = [-.15 .28; .26 .24]; b3 = [0; .44];
A4 = [ 0  0;  0 .16];

cnt = 1;
tic
while ~get(stop,'value')
```

```

r = rand;
if r < p(1)
    x = A1*x + b1;
elseif r < p(2)
    x = A2*x + b2;
elseif r < p(3)
    x = A3*x + b3;
else
    x = A4*x;
end
set(h,'xdata',x(1),'ydata',x(2));
cnt = cnt + 1;
drawnow
end
t = toc;
s = sprintf('%8.0f points in %6.3f seconds',cnt,t);
text(-1.5,-0.5,s,'fontweight','bold');
set(stop,'style','pushbutton','string','close', ...
    'callback','close(gcf)')

```

Let's examine this program a few statements at a time.

`shg`

stands for “show graph window.” It brings an existing graphics window forward, or creates a new one if necessary.

`clf reset`

resets most of the figure properties to their default values.

```

set(gcf,'color','white','menubar','none', ...
    'numbertitle','off','name','Fractal Fern')

```

changes the background color of the figure window from the default gray to white and provides a customized title for the window.

```

x = [.5; .5];

```

provides the initial coordinates of the point.

```

h = plot(x(1),x(2),'.');

```

plots a single dot in the plane and saves a *handle*, `h`, so that we can later modify the properties of the plot.

```

darkgreen = [0 2/3 0];

```

defines a color where the red and blue components are zero and the green component is two-thirds of its full intensity.



```
set(h,'markersize',1,'color','darkgreen','erasemode','none');
```

makes the dot referenced by `h` smaller, changes its color, and specifies that the image of the dot on the screen should not be erased when its coordinates are changed. A record of these old points is kept by the computer's graphics hardware (until the figure is reset), but MATLAB itself does not remember them.

```
axis([-3 3 0 10])
axis off
```

specifies that the plot should cover the region

$$-3 \leq x_1 \leq 3, \quad 0 \leq x_2 \leq 10,$$

but that the axes should not be drawn.

```
stop = uicontrol('style','toggle','string','stop', ...
    'background','white');
```

creates a toggle user interface control, labeled `'stop'` and colored white, in the default position near the lower left corner of the figure. The handle for the control is saved in the variable `stop`.

```
drawnow
```

causes the initial figure, including the initial point, to actually be plotted on the computer screen.

The statement

```
p = [ .85 .92 .99 1.00];
```

sets up a vector of probabilities. The statements

```
A1 = [ .85 .04; -.04 .85]; b1 = [0; 1.6];
A2 = [ .20 -.26; .23 .22]; b2 = [0; 1.6];
A3 = [-.15 .28; .26 .24]; b3 = [0; .44];
A4 = [ 0 0; 0 .16];
```

define the four affine transformations. The statement

```
cnt = 1;
```

initializes a counter that keeps track of the number of points plotted. The statement

```
tic
```

initializes a stopwatch timer. The statement

```
while ~get(stop,'value')
```

begins a `while` loop that runs as long as the `'value'` property of the `stop` toggle is equal to 0. Clicking the `stop` toggle changes the value from 0 to 1 and terminates the loop.

```
r = rand;
```

generates a *pseudorandom* value between 0 and 1. The compound `if` statement

```
if r < p(1)
    x = A1*x + b1;
elseif r < p(2)
    x = A2*x + b2;
elseif r < p(3)
    x = A3*x + b3;
else
    x = A4*x;
end
```

picks one of the four affine transformations. Because `p(1)` is 0.85, the first transformation is chosen 85% of the time. The other three transformations are chosen relatively infrequently.

```
set(h, 'xdata', x(1), 'ydata', x(2));
```

changes the coordinates of the point `h` to the new  $(x_1, x_2)$  and plots this new point. But `get(h, 'erasemode')` is `'none'`, so the old point also remains on the screen.

```
cnt = cnt + 1;
```

counts one more point.

```
drawnow
```

tells MATLAB to take the time to redraw the figure, showing the new point along with all the old ones. Without this command, nothing would be plotted until `stop` is toggled.

```
end
```

matches the `while` at the beginning of the loop. Finally,

```
t = toc;
```

reads the timer.

```
s = sprintf('%8.0f points in %6.3f seconds', cnt, t);
text(-1.5, -0.5, s, 'fontweight', 'bold');
```

displays the elapsed time since `tic` was called and the final count of the number of points plotted. Finally,

```
set(stop, 'style', 'pushbutton', 'string', 'close', ...
    'callback', 'close(gcf)')
```

changes the control to a push button that closes the window.

## 1.4 Magic Squares

MATLAB stands for *Matrix Laboratory*. Over the years, MATLAB has evolved into a general-purpose technical computing environment, but operations involving vectors, matrices, and linear algebra continue to be its most distinguishing feature.

Magic squares provide an interesting set of sample matrices. The command `help magic` tells us the following:

```
MAGIC(N) is an N-by-N matrix constructed from the integers
1 through N^2 with equal row, column, and diagonal sums.
Produces valid magic squares for all N > 0 except N = 2.
```

Magic squares were known in China over 2,000 years before the birth of Christ. The 3-by-3 magic square is known as *Lo Shu*. Legend has it that Lo Shu was discovered on the shell of a turtle that crawled out of the Lo River in the 23rd century B.C. Lo Shu provides a mathematical basis for *feng shui*, the ancient Chinese philosophy of balance and harmony. MATLAB can generate Lo Shu with

```
A = magic(3)
```

which produces

```
A =
     8     1     6
     3     5     7
     4     9     2
```

The command

```
sum(A)
```

sums the elements in each column to produce

```
    15    15    15
```

The command

```
sum(A')'
```

transposes the matrix, sums the columns of the transpose, and then transposes the results to produce the row sums

```
    15
    15
    15
```

The command

```
sum(diag(A))
```

sums the main diagonal of A, which runs from upper left to lower right, to produce

```
    15
```

The opposite diagonal, which runs from upper right to lower left, is less important in linear algebra, so finding its sum is a little trickier. One way to do it makes use of the function that “flips” a matrix “upside-down.”

```
sum(diag(flipud(A)))
```

produces

```
15
```

This verifies that **A** has equal row, column, and diagonal sums.

Why is the magic sum equal to 15? The command

```
sum(1:9)
```

tells us that the sum of the integers from 1 to 9 is 45. If these integers are allocated to 3 columns with equal sums, that sum must be

```
sum(1:9)/3
```

which is 15.

There are eight possible ways to place a transparency on an overhead projector. Similarly, there are eight magic squares of order three that are rotations and reflections of **A**. The statements

```
for k = 0:3
    rot90(A,k)
    rot90(A',k)
end
```

display all eight of them.

8	1	6	8	3	4
3	5	7	1	5	9
4	9	2	6	7	2
6	7	2	4	9	2
1	5	9	3	5	7
8	3	4	8	1	6
2	9	4	2	7	6
7	5	3	9	5	1
6	1	8	4	3	8
4	3	8	6	1	8
9	5	1	7	5	3
2	7	6	2	9	4

These are all the magic squares of order three.

Now for some linear algebra. The determinant of our magic square,

```
det(A)
```

is

```
-360
```

The inverse,

```
X = inv(A)
```

is

```
X =
    0.1472   -0.1444    0.0639
   -0.0611    0.0222    0.1056
   -0.0194    0.1889   -0.1028
```

The inverse looks better if it is displayed with a rational format.

```
format rat
X
```

shows that the elements of  $X$  are fractions with  $\det(A)$  in the denominator.

```
X =
    53/360   -13/90     23/360
   -11/180     1/45     19/180
    -7/360    17/90    -37/360
```

The statement

```
format short
```

restores the output format to its default.

Three other important quantities in computational linear algebra are *matrix norms*, *eigenvalues*, and *singular values*. The statements

```
r = norm(A)
e = eig(A)
s = svd(A)
```

produce

```
r =
    15

e =
    15.0000
     4.8990
    -4.8990

s =
    15.0000
     6.9282
     3.4641
```

The magic sum occurs in all three because the vector of all ones is an eigenvector and is also a left and right singular vector.

So far, all the computations in this section have been done using floating-point arithmetic. This is the arithmetic used for almost all scientific and engineering computation, especially for large matrices. But for a 3-by-3 matrix, it is easy to repeat the computations using symbolic arithmetic and the Symbolic Toolbox. The statement

```
A = sym(A)
```

changes the internal representation of **A** to a symbolic form that is displayed as

```
A =
 [ 8, 1, 6]
 [ 3, 5, 7]
 [ 4, 9, 2]
```

Now commands like

```
sum(A), sum(A')', det(A), inv(A), eig(A), svd(A)
```

produce symbolic results. In particular, the eigenvalue problem for this matrix can be solved exactly, and

```
e =
 [          15]
 [ 2*6^(1/2)]
 [-2*6^(1/2)]
```

A 4-by-4 magic square is one of several mathematical objects on display in *Melancolia*, a Renaissance etching by Albrecht Dürer. An electronic copy of the etching is available in a MATLAB data file.

```
load durer
whos
```

produces

```
 X          648x509          2638656 double array
caption      2x28              112 char array
map          128x3            3072 double array
```

The elements of the matrix **X** are indices into the gray-scale color map named **map**. The image is displayed with

```
image(X)
colormap(map)
axis image
```

Click the magnifying glass with a “+” in the toolbar and use the mouse to zoom in on the magic square in the upper right-hand corner. The scanning resolution becomes evident as you zoom in. The commands

```
load detail
image(X)
colormap(map)
axis image
```

display a higher resolution scan of the area around the magic square.

The command

```
A = magic(4)
```

produces a 4-by-4 magic square.

```
A =
    16     2     3    13
     5    11    10     8
     9     7     6    12
     4    14    15     1
```

The commands

```
sum(A), sum(A'), sum(diag(A)), sum(diag(flipud(A)))
```

yield enough 34's to verify that **A** is indeed a magic square.

The 4-by-4 magic square generated by MATLAB is not the same as Dürer's magic square. We need to interchange the second and third columns.

```
A = A(:, [1 3 2 4])
```

changes **A** to

```
A =
    16     3     2    13
     5    10    11     8
     9     6     7    12
     4    15    14     1
```

Interchanging columns does not change the column sums or the row sums. It usually changes the diagonal sums, but in this case both diagonal sums are still 34. So now our magic square matches the one in Dürer's etching. Dürer probably chose this particular 4-by-4 square because the date he did the work, 1514, occurs in the middle of the bottom row.

We have seen two different 4-by-4 magic squares. It turns out that there are 880 different magic squares of order 4 and 275305224 different magic squares of order 5. Determining the number of different magic squares of order 6 or larger is an unsolved mathematical problem.

The determinant of our 4-by-4 magic square,  $\det(\mathbf{A})$ , is 0. If we try to compute its inverse

```
inv(A)
```

we get

```
Warning: Matrix is close to singular or badly scaled.
Results may be inaccurate.
```

So some magic squares represent singular matrices. Which ones? The *rank* of a square matrix is the number of linearly independent rows or columns. An  $n$ -by- $n$  matrix is singular if and only if its rank is less than  $n$ .

The statements

```
for n = 1:24, r(n) = rank(magic(n)); end
[(1:24)' r']
```

produce a table of order versus rank.

1	1
2	2
3	3
4	3
5	5
6	5
7	7
8	3
9	9
10	7
11	11
12	3
13	13
14	9
15	15
16	3
17	17
18	11
19	19
20	3
21	21
22	13
23	23
24	3

Look carefully at this table. Ignore  $n = 2$  because `magic(2)` is not really a magic square. What patterns do you see? A bar graph makes the patterns easier to see.

```
bar(r)
title('Rank of magic squares')
```

produces Figure 1.4.

The rank considerations show that there are three different kinds of magic squares:






Figure 1.4. Rank of magic squares.

- *Odd order*:  $n$  is odd.
- *Singly even order*:  $n$  is a multiple of 2, but not 4.
- *Doubly even order*:  $n$  is a multiple of 4.

Odd-ordered magic squares,  $n = 3, 5, 7, \dots$ , have full rank  $n$ . They are nonsingular and have inverses. Doubly even magic squares,  $n = 4, 8, 12, \dots$ , have rank three no matter how large  $n$  is. They might be called *very singular*. Singly even magic squares,  $n = 6, 10, 14, \dots$ , have rank  $n/2 + 2$ . They are also singular, but have fewer row and column dependencies than the doubly even squares.

If you have MATLAB Version 6 or later, you can look at the M-file that generates magic squares with

```
edit magic.m
```

or

```
type magic.m
```

You will see the three different cases in the code.

The different kinds of magic squares also produce different three-dimensional surface plots. Try the following for various values of  $n$ .

```
surf(magic(n))
axis off
set(gcf, 'doublebuffer', 'on')
cameratoolbar
```

Double buffering prevents flicker when you use the various camera tools to move the viewpoint.

The following code produces Figure 1.5.

```
for n = 8:11
    subplot(2,2,n-7)
    surf(magic(n))
    title(num2str(n))
    axis off
    view(30,45)
    axis tight
end
```

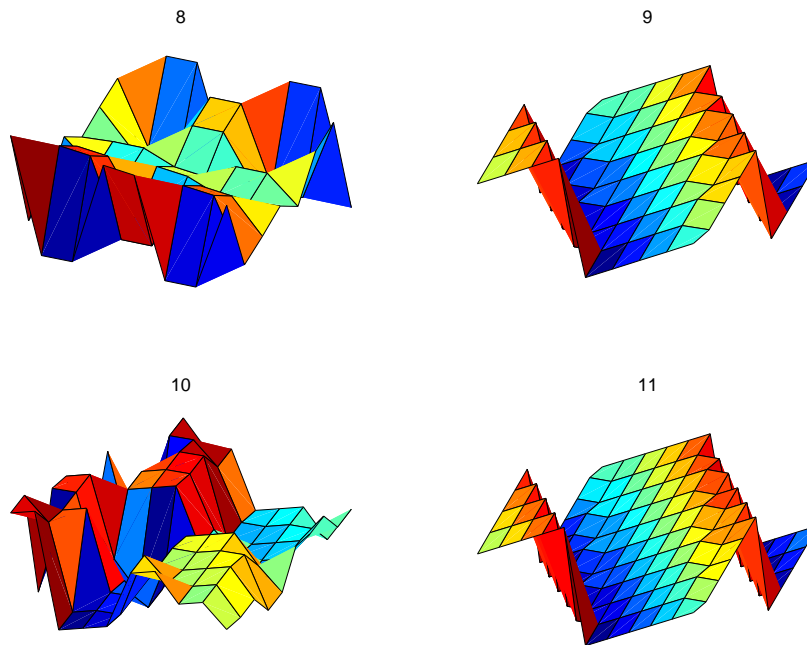


Figure 1.5. *Surface plots of magic squares.*

## 1.5 Cryptography

This section uses a cryptography example to show how MATLAB deals with text and character strings. The cryptographic technique, which is known as a *Hill cipher*, involves arithmetic in a *finite field*.

Almost all modern computers use the ASCII character set to store basic text. ASCII stands for *American Standard Code for Information Interchange*. The character set uses 7 of the 8 bits in a byte to encode 128 characters. The first 32 characters are nonprinting control characters, such as tab, backspace, and end-of-line. The 128th character is another nonprinting character that corresponds to the Delete key on your keyboard. In between these control characters are 95 printable

characters, including a space, 10 digits, 26 lowercase letters, 26 uppercase letters, and 32 punctuation marks.

MATLAB can easily display all the printable characters in the order determined by their ASCII encoding. Start with

```
x = reshape(32:127,32,3)'
```

This produces a 3-by-32 matrix.

```
x =
    32    33    34    ...    61    62    63
    64    65    66    ...    93    94    95
    96    97    98    ...   125   126   127
```

The `char` function converts numbers to characters. The statement

```
c = char(x)
```

produces

```
c =
! "# $ % & ' ( ) * + , - . / 0 1 2 3 4 5 6 7 8 9 : ; < = > ?
@ A B C D E F G H I J K L M N O P Q R S T U V W X Y Z [ \ ] ^ _
` a b c d e f g h i j k l m n o p q r s t u v w x y z { | } ~
```

We have cheated a little bit because the last element of `x` is 127, which corresponds to the nonprinting delete character, and we have not shown the last character in `c`. You can try this on your computer and see what is actually displayed.

The first character in `c` is blank, indicating that

```
char(32)
```

is the same as

```
, ,
```

The last printable character in `c` is the *tilde*, indicating that

```
char(126)
```

is the same as

```
, ~ ,
```

The characters representing digits are in the first line of `c`. In fact,

```
d = char(48:57)
```

displays a 10-character string

```
d =
0123456789
```

This string can be converted to the corresponding numerical values with `double` or `real`. The statement

```
double(d) - '0'
```

produces

```
0    1    2    3    4    5    6    7    8    9
```

Comparing the second and third lines of `c`, we see that the ASCII encoding of the lowercase letters is obtained by adding 32 to the ASCII encoding of the uppercase letters. Understanding this encoding allows us to use vector and matrix operations in MATLAB to manipulate text.

The ASCII standard is often extended to make use of all eight bits in a byte, but the characters that are displayed depend on the computer and operating system you are using, the font you have chosen, and even the country you live in. Try

```
char(reshape(160:255,32,3)')
```

and see what happens on your machine.

Our encryption technique involves *modular arithmetic*. All the quantities involved are integers and the result of any arithmetic operation is reduced by taking the remainder or *modulus* with respect to a prime number  $p$ . The functions `rem(x,y)` and `mod(x,y)` both compute the remainder if  $x$  is divided by  $y$ . They produce the same result if  $x$  and  $y$  have the same sign; the result also has that sign. But if  $x$  and  $y$  have opposite signs, then `rem(x,y)` has the same sign as  $x$ , while `mod(x,y)` has the same sign as  $y$ . Here is a table:

```
x = [37 -37 37 -37]';
y = [10 10 -10 -10]';
r = [ x  y  rem(x,y) mod(x,y)]
```

produces

```
37    10     7     7
-37   10    -7     3
37   -10     7    -3
-37  -10    -7    -7
```

We have chosen to encrypt text that uses the entire ASCII character set, not just the letters. There are 95 such characters. The next larger prime number is  $p = 97$ , so we represent the  $p$  characters by the integers  $0:p-1$  and do arithmetic mod  $p$ .

The characters are encoded two at a time. Each pair of characters is represented by a 2-vector,  $x$ . For example, suppose the text contains the pair of letters 'TV'. The ASCII values for this pair of letters are 84 and 86. Subtracting 32 to make the representation start at 0 produces the column vector

$$x = \begin{pmatrix} 52 \\ 54 \end{pmatrix}.$$

The encryption is done with a 2-by-2 matrix-vector multiplication over the integers mod  $p$ . The symbol  $\equiv$  is used to indicate that two integers have the same remainder, modulo the specified prime:

$$y \equiv Ax, \text{ mod } p,$$

where  $A$  is the matrix

$$A = \begin{pmatrix} 71 & 2 \\ 2 & 26 \end{pmatrix}.$$

For our example, the product  $Ax$  is

$$Ax = \begin{pmatrix} 3800 \\ 1508 \end{pmatrix}.$$

If this is reduced mod  $p$ , the result is

$$y = \begin{pmatrix} 17 \\ 53 \end{pmatrix}.$$

Converting this back to characters by adding 32 produces '1U'.

Now comes the interesting part. Over the integers modulo  $p$ , the matrix  $A$  is its own inverse. If

$$y \equiv Ax, \text{ mod } p,$$

then

$$x \equiv Ay, \text{ mod } p.$$

In other words, in arithmetic mod  $p$ ,  $A^2$  is the identity matrix. You can check this with MATLAB.

```
p = 97;
A = [71 2; 2 26]
I = mod(A^2,p)
```

produces

```
A =
    71     2
     2    26

I =
     1     0
     0     1
```

This means that the encryption process is its own inverse. The same function can be used to both encrypt and decrypt a message.

The M-file `crypto.m` begins with a preamble.

```
function y = crypto(x)
% CRYPTO Cryptography example.
% y = crypto(x) converts an ASCII text string into another
% coded string. The function is its own inverse, so
% crypto(crypto(x)) gives x back.
% See also: ENCRYPT.
```

A comment precedes the statement that assigns the prime  $p$ .

```
% Use a two-character Hill cipher with arithmetic
% modulo 97, a prime.
p = 97;
```

Choose two characters above ASCII 128 to expand the size of the character set from 95 to 97.

```
c1 = char(169);
c2 = char(174);
x(x==c1) = 127;
x(x==c2) = 128;
```

The conversion from characters to numerical values is done by

```
x = mod(real(x-32),p);
```

Prepare for the matrix-vector product by forming a matrix with two rows and lots of columns.

```
n = 2*floor(length(x)/2);
X = reshape(x(1:n),2,n/2);
```

All this preparation has been so that we can do the actual finite field arithmetic quickly and easily.

```
% Encode with matrix multiplication modulo p.
A = [71 2; 2 26];
Y = mod(A*X,p);
```

Reshape into a single row.

```
y = reshape(Y,1,n);
```

If `length(x)` is odd, encode the last character

```
if length(x) > n
    y(n+1) = mod((p-1)*x(n+1),p);
end
```

Finally, convert the numbers back to characters.

```
y = char(y+32);
y(y==127) = c1;
y(y==128) = c2;
```

Let's follow the computation of  $y = \text{crypto}(\text{'Hello world'})$ . We begin with a character string.

```
x = 'Hello world'
```

This is converted to an integer vector.

```
x =
  40  69  76  76  79  0  87  79  82  76  68
```

`length(x)` is odd, so the reshaping temporarily ignores the last element

```
X =
  40  76  79  87  82
  69  76  0  79  76
```

A conventional matrix-vector multiplication  $A*X$  produces an intermediate matrix.

```
2978  5548  5609  6335  5974
1874  2128  158  2228  2140
```

Then the `mod(.,p)` operation produces

```
Y =
  68  19  80  30  57
  31  91  61  94  6
```

This is rearranged to a row vector.

```
y =
  68  31  19  91  80  61  30  94  57  6
```

Now the last element of  $x$  is encoded by itself and attached to the end of  $y$ .

```
y =
  68  31  19  91  80  61  30  94  57  6  29
```

Finally,  $y$  is converted back to a character string to produce the encrypted result.

```
y = 'd?3{p]>~Y&='
```

If we now compute `crypto(y)`, we get back our original `'Hello world'`.

## 1.6 The $3n+1$ Sequence

This section describes a famous unsolved problem in number theory. Start with any positive integer  $n$ . Repeat the following steps:

- If  $n = 1$ , stop.
- If  $n$  is even, replace it with  $n/2$ .
- If  $n$  is odd, replace it with  $3n + 1$ .

For example, starting with  $n = 7$  produces

7, 22, 11, 34, 17, 52, 26, 13, 40, 20, 10, 5, 16, 8, 4, 2, 1.

The sequence terminates after 17 steps. Note that whenever  $n$  reaches a power of 2, the sequence terminates in  $\log_2 n$  more steps.

The unanswered question is, does the process *always* terminate? Or is there some starting value that causes the process to go on forever, either because the numbers get larger and larger, or because some periodic cycle is generated?

This problem is known as the  $3n + 1$  problem. It has been studied by many eminent mathematicians, including Collatz, Ulam, and Kakatani, and is discussed in a survey paper by Jeffrey Lagarias [5].

The following MATLAB code fragment generates the sequence starting with any specified  $n$ .

```
y = n;
while n > 1
    if rem(n,2)==0
        n = n/2;
    else
        n = 3*n+1;
    end
    y = [y n];
end
```

We don't know ahead of time how long the resulting vector  $y$  is going to be. But the statement

```
y = [y n];
```

automatically increases `length(y)` each time it is executed.

In principle, the unsolved mathematical problem is, Can this code fragment run forever? In actual fact, floating-point roundoff error causes the calculation to misbehave whenever  $3n + 1$  becomes greater than  $2^{53}$ , but it is still interesting to investigate modest values of  $n$ .

Let's embed our code fragment in a GUI. The complete function is in the M-file `threenplus1.m`. For example, the statement

```
threenplus1(7)
```

produces Figure 1.6.

The M-file begins with a preamble containing the function header and the help information.

```
function threenplus1(n)
% 'Three n plus 1'.
% Study the 3n+1 sequence.
% threenplus1(n) plots the sequence starting with n.
% threenplus1 with no arguments starts with n = 1.
% uicontrols decrement or increment the starting n.
% Is it possible for this to run forever?
```






Figure 1.6. threenplus1.

The next section of code brings the current graphics window forward and resets it. Two push buttons, which are the default `uicontrols`, are positioned near the bottom center of the figure at pixel coordinates `[260,5]` and `[300,5]`. Their size is 25 by 22 pixels and they are labeled with `'<'` and `'>'`. If either button is subsequently pushed, the `'callback'` string is executed, calling the function recursively with a corresponding `'-1'` or `'+1'` string argument. The `'tag'` property of the current figure, `gcf`, is set to a characteristic string that prevents this section of code from being reexecuted on subsequent calls.

```

if ~isequal(get(gcf,'tag'),'3n+1')
    shg
    clf reset
    uicontrol( ...
        'position',[260 5 25 22], ...
        'string','<', ...
        'callback','threenplus1(''-1'')');
    uicontrol( ...
        'position',[300 5 25 22], ...
        'string','>', ...
        'callback','threenplus1(''+1'')');
    set(gcf,'tag','3n+1');
end

```

The next section of code sets `n`. If `nargin`, the number of input arguments, is 0, then `n` is set to 1. If the input argument is either of the strings from the push button callbacks, then `n` is retrieved from the `'userdata'` field of the `figure` and decremented or incremented. If the input argument is not a string, then it is the desired `n`. In all situations, `n` is saved in `'userdata'` for use on subsequent calls.

```

if nargin == 0
    n = 1;
elseif isequal(n,'-1')
    n = get(gcf,'userdata') - 1;
elseif isequal(n,'+1')
    n = get(gcf,'userdata') + 1;
end
if n < 1, n = 1; end
set(gcf,'userdata',n)

```

We've seen the next section of code before; it does the actual computation.

```

y = n;
while n > 1
    if rem(n,2)==0
        n = n/2;
    else
        n = 3*n+1;
    end
    y = [y n];
end

```

The final section of code plots the generated sequence with dots connected by straight lines, using a logarithmic vertical scale and customized tick labels.

```

semilogy(y, '-.')
axis tight
ymax = max(y);
ytick = [2.^(0:ceil(log2(ymax))-1) ymax];
if length(ytick) > 8, ytick(end-1) = []; end
set(gca,'ytick',ytick)
title(['n = ' num2str(y(1))]);

```

## 1.7 Floating-Point Arithmetic

Some people believe that

- numerical analysis is the study of floating-point arithmetic;
- floating-point arithmetic is unpredictable and hard to understand.

We intend to convince you that both of these assertions are false. Very little of this book is actually about floating-point arithmetic. But when the subject does arise, we hope you will find floating-point arithmetic is not only computationally powerful, but also mathematically elegant.

If you look carefully at the definitions of fundamental arithmetic operations like addition and multiplication, you soon encounter the mathematical abstraction known as real numbers. But actual computation with real numbers is not very

practical because it involves limits and infinities. Instead, MATLAB and most other technical computing environments use floating-point arithmetic, which involves a finite set of numbers with finite precision. This leads to the phenomena of *roundoff*, *underflow*, and *overflow*. Most of the time, it is possible to use MATLAB effectively without worrying about these details, but, every once in a while, it pays to know something about the properties and limitations of floating-point numbers.

Twenty years ago, the situation was far more complicated than it is today. Each computer had its own floating-point number system. Some were binary; some were decimal. There was even a Russian computer that used trinary arithmetic. Among the binary computers, some used 2 as the base; others used 8 or 16. And everybody had a different precision. In 1985, the IEEE Standards Board and the American National Standards Institute adopted the ANSI/IEEE Standard 754-1985 for Binary Floating-Point Arithmetic. This was the culmination of almost a decade of work by a 92-person working group of mathematicians, computer scientists, and engineers from universities, computer manufacturers, and microprocessor companies.

All computers designed since 1985 use IEEE floating-point arithmetic. This doesn't mean that they all get exactly the same results, because there is some flexibility within the standard. But it does mean that we now have a machine-independent model of how floating-point arithmetic behaves.

MATLAB has traditionally used the IEEE double-precision format. There is a single-precision format that saves space, but that isn't much faster on modern machines. MATLAB 7 will have support for single-precision arithmetic, but we will deal exclusively with double precision in this book. There is also an extended-precision format, which is optional and therefore is one of the reasons for lack of uniformity among different machines.

Most nonzero floating-point numbers are normalized. This means they can be expressed as

$$x = \pm(1 + f) \cdot 2^e.$$

The quantity  $f$  is the fraction or mantissa and  $e$  is the exponent. The fraction satisfies

$$0 \leq f < 1$$

and must be representable in binary using at most 52 bits. In other words,  $2^{52}f$  is an integer in the interval

$$0 \leq 2^{52}f < 2^{52}.$$

The exponent  $e$  is an integer in the interval

$$-1022 \leq e \leq 1023.$$

The finiteness of  $f$  is a limitation on *precision*. The finiteness of  $e$  is a limitation on *range*. Any numbers that don't meet these limitations must be approximated by ones that do.

Double-precision floating-point numbers are stored in a 64-bit word, with 52 bits for  $f$ , 11 bits for  $e$ , and 1 bit for the sign of the number. The sign of  $e$  is accommodated by storing  $e + 1023$ , which is between 1 and  $2^{11} - 2$ . The 2 extreme

values for the exponent field, 0 and  $2^{11}-1$ , are reserved for exceptional floating-point numbers that we will describe later.

The entire fractional part of a floating-point number is not  $f$ , but  $1+f$ , which has 53 bits. However, the leading 1 doesn't need to be stored. In effect, the IEEE format packs 65 bits of information into a 64-bit word.

The program `floatgui` shows the distribution of the positive numbers in a model floating-point system with variable parameters. The parameter  $t$  specifies the number of bits used to store  $f$ . In other words,  $2^t f$  is an integer. The parameters  $e_{min}$  and  $e_{max}$  specify the range of the exponent, so  $e_{min} \leq e \leq e_{max}$ . Initially, `floatgui` sets  $t = 3$ ,  $e_{min} = -4$ , and  $e_{max} = 3$  and produces the distribution shown in Figure 1.7.




Figure 1.7. `floatgui`.

Within each binary interval  $2^e \leq x \leq 2^{e+1}$ , the numbers are equally spaced with an increment of  $2^{e-t}$ . If  $e = 0$  and  $t = 3$ , for example, the spacing of the numbers between 1 and 2 is  $1/8$ . As  $e$  increases, the spacing increases.

It is also instructive to display the floating-point numbers with a logarithmic scale. Figure 1.8 shows `floatgui` with `logscale` checked and  $t = 5$ ,  $e_{min} = -4$ , and  $e_{max} = 3$ . With this logarithmic scale, it is more apparent that the distribution in each binary interval is the same.

A very important quantity associated with floating-point arithmetic is highlighted in red by `floatgui`. MATLAB calls this quantity `eps`, which is short for *machine epsilon*.

*eps* is the distance from 1 to the next larger floating-point number.

For the `floatgui` model floating-point system, `eps = 2-t`.




Figure 1.8. `floatgui(logscale)`.

Before the IEEE standard, different machines had different values of `eps`. Now, for IEEE double-precision,

$$\text{eps} = 2^{-52}.$$

The approximate decimal value of `eps` is  $2.2204 \cdot 10^{-16}$ . Either `eps/2` or `eps` can be called the roundoff level. The maximum relative error incurred when the result of an arithmetic operation is rounded to the nearest floating-point number is `eps/2`. The maximum relative spacing between numbers is `eps`. In either case, you can say that the roundoff level is about 16 decimal digits.

A frequent instance of roundoff occurs with the simple MATLAB statement

```
t = 0.1
```

The mathematical value  $t$  stored in `t` is not exactly 0.1 because expressing the decimal fraction  $1/10$  in binary requires an infinite series. In fact,

$$\frac{1}{10} = \frac{1}{2^4} + \frac{1}{2^5} + \frac{0}{2^6} + \frac{0}{2^7} + \frac{1}{2^8} + \frac{1}{2^9} + \frac{0}{2^{10}} + \frac{0}{2^{11}} + \frac{1}{2^{12}} + \dots$$

After the first term, the sequence of coefficients 1, 0, 0, 1 is repeated infinitely often. Grouping the resulting terms together four at a time expresses  $1/10$  in a base 16, or *hexadecimal*, series.

$$\frac{1}{10} = 2^{-4} \cdot \left( 1 + \frac{9}{16} + \frac{9}{16^2} + \frac{9}{16^3} + \frac{9}{16^4} + \dots \right)$$

Floating-point numbers on either side of  $1/10$  are obtained by terminating the fractional part of this series after 52 binary terms, or 13 hexadecimal terms, and rounding the last term up or down. Thus

$$t_1 < 1/10 < t_2,$$

where

$$t_1 = 2^{-4} \cdot \left( 1 + \frac{9}{16} + \frac{9}{16^2} + \frac{9}{16^3} + \dots + \frac{9}{16^{12}} + \frac{9}{16^{13}} \right),$$

$$t_2 = 2^{-4} \cdot \left( 1 + \frac{9}{16} + \frac{9}{16^2} + \frac{9}{16^3} + \dots + \frac{9}{16^{12}} + \frac{10}{16^{13}} \right).$$

It turns out that  $1/10$  is closer to  $t_2$  than to  $t_1$ , so  $t$  is equal to  $t_2$ . In other words,

$$t = (1 + f) \cdot 2^e,$$

where

$$f = \frac{9}{16} + \frac{9}{16^2} + \frac{9}{16^3} + \dots + \frac{9}{16^{12}} + \frac{10}{16^{13}},$$

$$e = -4.$$

The MATLAB command

```
format hex
```

causes `t` to be displayed as

```
3fb999999999999a
```

The characters **a** through **f** represent the hexadecimal “digits” 10 through 15. The first three characters, **3fb**, give the hexadecimal representation of decimal 1019, which is the value of the biased exponent  $e+1023$  if  $e$  is  $-4$ . The other 13 characters are the hexadecimal representation of the fraction  $f$ .

In summary, the value stored in **t** is very close to, but not exactly equal to, 0.1. The distinction is occasionally important. For example, the quantity

```
0.3/0.1
```

is not exactly equal to 3 because the actual numerator is a little less than 0.3 and the actual denominator is a little greater than 0.1.

Ten steps of length **t** are not precisely the same as one step of length 1. MATLAB is careful to arrange that the last element of the vector

```
0:0.1:1
```

is exactly equal to 1, but if you form this vector yourself by repeated additions of 0.1, you will miss hitting the final 1 exactly.

What does the floating-point approximation to the golden ratio look like?

```
format hex
phi = (1 + sqrt(5))/2
```

produces

```
phi =
3ff9e3779b97f4a8
```

The first hex digit, 3, is 0011 in binary. The first bit is the sign of the floating-point number; 0 is positive, 1 is negative. So **phi** is positive. The remaining bits of the first three hex digits contain  $e + 1023$ . In this example, **3ff** in base 16 is  $3 \cdot 16^2 + 15 \cdot 16 + 15 = 1023$  in decimal. So

$$e = 0.$$

In fact, any floating-point number between 1.0 and 2.0 has  $e = 0$ , so its **hex** output begins with **3ff**. The other 13 hex digits contain  $f$ . In this example,

$$f = \frac{9}{16} + \frac{14}{16^2} + \frac{3}{16^3} + \cdots + \frac{10}{16^{12}} + \frac{8}{16^{13}}.$$

With these values of  $f$  and  $e$ ,

$$(1 + f)2^e \approx \phi.$$

Another example is provided by the following code segment.

```
format long
a = 4/3
b = a - 1
c = 3*b
e = 1 - c
```

With exact computation, `e` would be 0. But with floating-point, the output produced is

```
a =
  1.333333333333333
b =
  0.333333333333333
c =
  1.000000000000000
e =
  2.220446049250313e-016
```

It turns out that the only roundoff occurs in the division in the first statement. The quotient cannot be exactly  $4/3$ , except on that Russian trinary computer. Consequently the value stored in `a` is close to, but not exactly equal to,  $4/3$ . The subtraction `b = a - 1` produces a `b` whose last bit is 0. This means that the multiplication `3*b` can be done without any roundoff. The value stored in `c` is not exactly equal to 1, and so the value stored in `e` is not 0. Before the IEEE standard, this code was used as a quick way to estimate the roundoff level on various computers.

The roundoff level `eps` is sometimes called “floating-point zero,” but that’s a misnomer. There are many floating-point numbers much smaller than `eps`. The smallest positive normalized floating-point number has  $f = 0$  and  $e = -1022$ . The largest floating-point number has  $f$  a little less than 1 and  $e = 1023$ . MATLAB calls these numbers `realmin` and `realmax`. Together with `eps`, they characterize the standard system.

	Binary	Decimal
<code>eps</code>	$2^{-52}$	2.2204e-16
<code>realmin</code>	$2^{-1022}$	2.2251e-308
<code>realmax</code>	$(2-\text{eps}) \cdot 2^{1023}$	1.7977e+308

If any computation tries to produce a value larger than `realmax`, it is said to *overflow*. The result is an exceptional floating-point value called *infinity* or `Inf`. It is represented by taking  $e = 1024$  and  $f = 0$  and satisfies relations like  $1/\text{Inf} = 0$  and  $\text{Inf} + \text{Inf} = \text{Inf}$ .

If any computation tries to produce a value that is undefined even in the real number system, the result is an exceptional value known as Not-a-Number, or `NaN`. Examples include `0/0` and `Inf-Inf`. `NaN` is represented by taking  $e = 1024$  and  $f$  nonzero.

If any computation tries to produce a value smaller than `realmin`, it is said to *underflow*. This involves one of the optional, and controversial, aspects of the IEEE standard. Many, but not all, machines allow exceptional denormal or subnormal floating-point numbers in the interval between `realmin` and `eps*realmin`. The smallest positive subnormal number is about  $0.494\text{e-}323$ . Any results smaller than this are set to 0. On machines without subnormals, any results less than `realmin` are set to 0. The subnormal numbers fill in the gap you can see in the `floatgui` model system between 0 and the smallest positive number. They do provide an

elegant way to handle underflow, but their practical importance for MATLAB-style computation is very rare. Denormal numbers are represented by taking  $e = -1023$ , so the biased exponent  $e + 1023$  is 0.

MATLAB uses the floating-point system to handle integers. Mathematically, the numbers 3 and 3.0 are the same, but many programming languages would use different representations for the two. MATLAB does not distinguish between them. We sometimes use the term *flint* to describe a floating-point number whose value is an integer. Floating-point operations on flints do not introduce any roundoff error, as long as the results are not too large. Addition, subtraction, and multiplication of flints produce the exact flint result if it is not larger than  $2^{53}$ . Division and square root involving flints also produce a flint if the result is an integer. For example, `sqrt(363/3)` produces 11, with no roundoff.

Two MATLAB functions that take apart and put together floating-point numbers are `log2` and `pow2`.

```
help log2
help pow2
```

produces

```
[F,E] = LOG2(X) for a real array X, returns an array F
of real numbers, usually in the range 0.5 <= abs(F) < 1,
and an array E of integers, so that X = F .* 2.^E.
Any zeros in X produce F = 0 and E = 0.
```

```
X = POW2(F,E) for a real array F and an integer array E
computes X = F .* (2.^E). The result is computed quickly
by simply adding E to the floating-point exponent of F.
```

The quantities  $F$  and  $E$  used by `log2` and `pow2` predate the IEEE floating-point standard and so are slightly different from the  $f$  and  $e$  we are using in this section. In fact,  $f = 2 * F - 1$  and  $e = E - 1$ .

```
[F,E] = log2(phi)
```

produces

```
F =
    0.80901699437495
E =
     1
```

Then

```
phi = pow2(F,E)
```

gives back

```
phi =
    1.61803398874989
```



As an example of how roundoff error affects matrix computations, consider the 2-by-2 set of linear equations

$$\begin{aligned} 17x_1 + 5x_2 &= 22, \\ 1.7x_1 + 0.5x_2 &= 2.2. \end{aligned}$$

The obvious solution is  $x_1 = 1, x_2 = 1$ . But the MATLAB statements

```
A = [17 5; 1.7 0.5]
b = [22; 2.2]
x = A\b
```

produce

```
x =
-1.0588
 8.0000
```

Where did this come from? Well, the equations are singular, but consistent. The second equation is just 0.1 times the first. The computed  $\mathbf{x}$  is one of infinitely many possible solutions. But the floating-point representation of the matrix  $A$  is not exactly singular because  $A(2,1)$  is not exactly  $17/10$ .

The solution process subtracts a multiple of the first equation from the second. The multiplier is  $\mu = 1.7/17$ , which turns out to be the floating-point number obtained by truncating, rather than rounding, the binary expansion of  $1/10$ . The matrix  $A$  and the right-hand side  $\mathbf{b}$  are modified by

```
A(2,:) = A(2,:) - mu*A(1,:)
b(2) = b(2) - mu*b(1)
```

With exact computation, both  $A(2,2)$  and  $b(2)$  would become zero, but with floating-point arithmetic, they both become nonzero multiples of `eps`.

```
A(2,2) = (1/4)*eps
        = 5.5511e-17
b(2) = 2*eps
      = 4.4408e-16
```

MATLAB notices the tiny value of the new  $A(2,2)$  and displays a message warning that the matrix is close to singular. It then computes the solution of the modified second equation by dividing one roundoff error by another.

```
x(2) = b(2)/A(2,2)
      = 8
```

This value is substituted back into the first equation to give

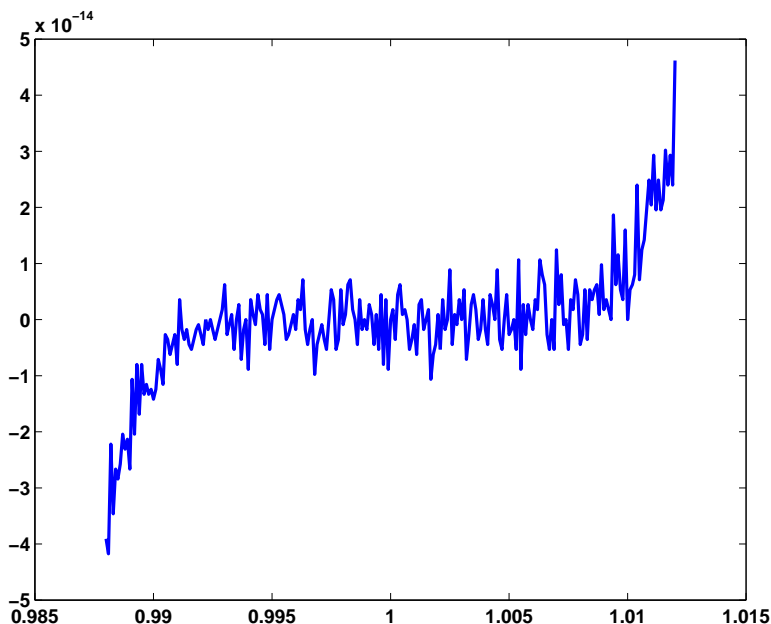
```
x(1) = (22 - 5*x(2))/17
      = -1.0588
```

The details of the roundoff error lead MATLAB to pick out one particular solution from among the infinitely many possible solutions to the singular system.

Our final example plots a seventh-degree polynomial.

```
x = 0.988:.0001:1.012;
y = x.^7-7*x.^6+21*x.^5-35*x.^4+35*x.^3-21*x.^2+7*x-1;
plot(x,y)
```

The resulting plot in Figure 1.9 doesn't look anything like a polynomial. It isn't smooth. You are seeing roundoff error in action. The  $y$ -axis scale factor is tiny,  $10^{-14}$ . The tiny values of  $y$  are being computed by taking sums and differences of



**Figure 1.9.** *Is this a polynomial?*

numbers as large as  $35 \cdot 1.012^4$ . There is severe subtractive cancellation. The example was contrived by using the Symbolic Toolbox to expand  $(x - 1)^7$  and carefully choosing the range for the  $x$ -axis to be near  $x = 1$ . If the values of  $y$  are computed instead by

```
y = (x-1).^7;
```

then a smooth (but very flat) plot results.

## 1.8 Further Reading

Additional information about floating-point arithmetic and roundoff error can be found in Higham [4] and Overton [6].

---

## Exercises

- 1.1. Which of these familiar rectangles is closest to a golden rectangle? Use MATLAB to do the calculations with an element-by-element vector division, `w./h`.
  - 3-by-5 inch index card,
  - 8.5-by-11 inch U.S. letter paper,
  - 8.5-by-14 inch U.S. legal paper,
  - 9-by-12 foot rug,
  - 9:16 “letterbox” TV picture,
  - 768-by-1024 pixel computer monitor.
- 1.2. ISO standard A4 paper is commonly used throughout most of the world, except in the United States and Canada. Its dimensions are 210 by 297 mm. This is not a golden rectangle, but the aspect ratio is close to another familiar irrational mathematical quantity. What is that quantity? If you fold a piece of A4 paper in half, what is the aspect ratio of each of the halves? Modify the M-file `goldrect.m` to illustrate this property.
- 1.3. How many terms in the truncated continued fraction does it take to approximate  $\phi$  with an error less than  $10^{-10}$ ? As the number of terms increases beyond this roundoff, error eventually intervenes. What is the best accuracy you can hope to achieve with double-precision floating-point arithmetic and how many terms does it take?
- 1.4. Use the MATLAB *backslash* operator to solve the 2-by-2 system of simultaneous linear equations

$$\begin{aligned}c_1 + c_2 &= 1, \\c_1\phi + c_2(1 - \phi) &= 1\end{aligned}$$

for  $c_1$  and  $c_2$ . You can find out about the backslash operator by taking a peek at the next chapter of this book, or with the commands

```
help \
help slash
```

- 1.5. The statement

```
semilogy(fibonacci(18), '-o')
```

makes a logarithmic plot of Fibonacci numbers versus their index. The graph is close to a straight line. What is the slope of this line?

- 1.6. How does the execution time of `fibnum(n)` depend on the execution time for `fibnum(n-1)` and `fibnum(n-2)`? Use this relationship to obtain an approximate formula for the execution time of `fibnum(n)` as a function of  $n$ . Estimate how long it would take your computer to compute `fibnum(50)`. Warning: You probably do not want to actually run `fibnum(50)`.

- 1.7. What is the index of the largest Fibonacci number that can be represented *exactly* as a MATLAB double-precision quantity without roundoff error? What is the index of the largest Fibonacci number that can be represented *approximately* as a MATLAB double-precision quantity without overflowing?

- 1.8. Enter the statements

```
A = [1 1; 1 0]
X = [1 0; 0 1]
```

Then enter the statement

```
X = A*X
```

Now repeatedly press the up arrow key, followed by the Enter key. What happens? Do you recognize the matrix elements being generated? How many times would you have to repeat this iteration before X overflows?

- 1.9. Change the fern color scheme to use pink on a black background. Don't forget the stop button.
- 1.10. (a) What happens if you resize the figure window while the fern is being generated? Why?  
 (b) The M-file `finitefern.m` can be used to produce printed output of the fern. Explain why printing is possible with `finitefern.m` but not with `fern.m`.
- 1.11. Flip the fern by interchanging its  $x$ - and  $y$ -coordinates.
- 1.12. What happens to the fern if you change the only nonzero element in the matrix `A4`?
- 1.13. What are the coordinates of the lower end of the fern's stem?
- 1.14. The coordinates of the point at the upper tip end of the fern can be computed by solving a certain 2-by-2 system of simultaneous linear equations. What is that system and what are the coordinates of the tip?
- 1.15. The fern algorithm involves repeated random choices from four different formulas for advancing the point. If the  $k$ th formula is used repeatedly by itself, without random choices, it defines a deterministic trajectory in the  $(x, y)$  plane. Modify `finitefern.m` so that plots of each of these four trajectories are superimposed on the plot of the fern. Start each trajectory at the point  $(-1, 5)$ . Plot o's connected with straight lines for the steps along each trajectory. Take as many steps as are needed to show each trajectory's limit point. You can superimpose several plots with

```
plot(...)
hold on
plot(...)
plot(...)
hold off
```

- 1.16. Use the following code to make your own Portable Network Graphics file from the fern. Then compare your image with one obtained from `ncm/fern.png`.

```

bg = [0 0 85];          % Dark blue background
fg = [255 255 255];    % White dots
sz = get(0,'screensize');
rand('state',0)
X = finitefern(500000,sz(4),sz(3));
d = fg - bg;
R = uint8(bg(1) + d(1)*X);
G = uint8(bg(2) + d(2)*X);
B = uint8(bg(3) + d(3)*X);
F = cat(3,R,G,B);
imwrite(F,'myfern.png','png','bitdepth',8)

```

- 1.17. Modify `fern.m` or `finitefern.m` so that it produces *Sierpinski's triangle*. Start at

$$x = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

At each iterative step, the current point  $x$  is replaced with  $Ax + b$ , where the matrix  $A$  is always

$$A = \begin{pmatrix} 1/2 & 0 \\ 0 & 1/2 \end{pmatrix}$$

and the vector  $b$  is chosen at random with equal probability from among the three vectors

$$b = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \quad b = \begin{pmatrix} 1/2 \\ 0 \end{pmatrix}, \quad \text{and} \quad b = \begin{pmatrix} 1/4 \\ \sqrt{3}/4 \end{pmatrix}.$$

- 1.18. `greetings(phi)` generates a seasonal holiday fractal that depends upon the parameter `phi`. The default value of `phi` is the golden ratio. What happens for other values of `phi`? Try both simple fractions and floating-point approximations to irrational values.
- 1.19.  $A = \text{magic}(4)$  is singular. Its columns are linearly dependent. What do `null(A)`, `null(A,'r')`, `null(sym(A))`, and `rref(A)` tell you about that dependence?
- 1.20. Let  $A = \text{magic}(n)$  for  $n = 3, 4$ , or  $5$ . What does

```
p = randperm(n); q = randperm(n); A = A(p,q);
```

do to

```

sum(A)
sum(A')
sum(diag(A))
sum(diag(flipud(A)))
rank(A)

```

- 1.21. The character `char(7)` is a control character. What does it do?
- 1.22. What does `char([169 174])` display on your computer?

- 1.23. What fundamental physical law is hidden in this string?

```
s = '/b_t3{$H~M06JTQI>v~#3Giew*1(p,nF'
```

- 1.24. Find the two files `encrypt.m` and `gettysburg.txt`. Use `encrypt` to encrypt `gettysburg.txt`. Then decrypt the result. Use `encrypt` to encrypt itself.
- 1.25. With the NCM directory on your path, you can read the text of Lincoln's Gettysburg Address with

```
fp = fopen('gettysburg.txt');
G = char(fread(fp))'
fclose(fp);
```

- (a) How many characters are in the text?
- (b) Use the `unique` function to find the unique characters in the text.
- (c) How many blanks are in the text? What punctuation characters, and how many of each, are there?
- (d) Remove the blanks and the punctuation and convert the text to all upper- or lowercase. Use the `histc` function to count the number of letters. What is the most frequent letter? What letters are missing?
- (e) Use the `bar` function as described in `help histc` to plot a histogram of the letter frequencies.
- (f) Use `get(gca,'xtick')` and `get(gca,'xticklabel')` to see how the  $x$ -axis of the histogram is labeled. Then use

```
set(gca,'xtick',...,'xticklabel',...)
```

to relabel the  $x$ -axis with the letters in the text.

- 1.26. If `x` is the character string consisting of just two blanks,

```
x = ' '
```

then `crypto(x)` is actually equal to `x`. Why does this happen? Are there any other two-character strings that `crypto` does not change?

- 1.27. Find another 2-by-2 integer matrix `A` for which

```
mod(A*A,97)
```

is the identity matrix. Replace the matrix in `crypto.m` with your matrix and verify that the function still works correctly.

- 1.28. The function `crypto` works with 97 characters instead of 95. It can produce output, and correctly handle input, that contains two characters with ASCII values greater than 127. What are these characters? Why are they necessary? What happens to other characters with ASCII values greater than 127?
- 1.29. Create a new `crypto` function that works with just 29 characters: the 26 lowercase letters, plus blank, period, and comma. You will need to find a 2-by-2 integer matrix `A` for which `mod(A*A,29)` is the identity matrix.
- 1.30. The graph of the  $3n + 1$  sequence has a particular characteristic shape if the starting  $n$  is 5, 10, 20, 40,  $\dots$ , that is,  $n$  is five times a power of 2. What is this shape and why does it happen?

- 1.31. The graphs of the  $3n + 1$  sequences starting at  $n = 108, 109,$  and  $110$  are very similar to each other. Why?
- 1.32. Let  $L(n)$  be the number of terms in the  $3n + 1$  sequence that starts with  $n$ . Write a MATLAB function that computes  $L(n)$  *without* using any vectors or unpredictable amounts of storage. Plot  $L(n)$  for  $1 \leq n \leq 1000$ . What is the maximum value of  $L(n)$  for  $n$  in this range, and for what value of  $n$  does it occur? Use `threenplus1` to plot the sequence that starts with this particular value of  $n$ .
- 1.33. Modify `floatgui.m` by changing its last line from a comment to an executable statement and changing the question mark to a simple expression that counts the number of floating-point numbers in the model system.
- 1.34. Explain the output produced by

```
t = 0.1
n = 1:10
e = n/10 - n*t
```

- 1.35. What does each of these programs do? How many lines of output does each program produce? What are the last two values of  $x$  printed?
- ```
x = 1; while 1+x > 1, x = x/2, pause(.02), end

x = 1; while x+x > x, x = 2*x, pause(.02), end

x = 1; while x+x > x, x = x/2, pause(.02), end
```
- 1.36. Which familiar real numbers are approximated by floating-point numbers that display the following values with `format hex`?

```
4059000000000000
3f847ae147ae147b
3fe921fb54442d18
```

- 1.37. Let  $\mathcal{F}$  be the set of all IEEE double-precision floating-point numbers, except NaNs and Infs, which have biased exponent `7ff` (hex), and denormals, which have biased exponent `000` (hex).
- How many elements are there in  $\mathcal{F}$ ?
  - What fraction of the elements of  $\mathcal{F}$  are in the interval  $1 \leq x < 2$ ?
  - What fraction of the elements of  $\mathcal{F}$  are in the interval  $1/64 \leq x < 1/32$ ?
  - Determine by random sampling approximately what fraction of the elements  $x$  of  $\mathcal{F}$  satisfy the MATLAB logical relation

```
x*(1/x) == 1
```

- 1.38. The classic quadratic formula says that the two roots of the quadratic equation

$$ax^2 + bx + c = 0$$

are

$$x_1, x_2 = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}.$$

Use this formula in MATLAB to compute both roots for

$$a = 1, \quad b = -100000000, \quad c = 1.$$

Compare your computed results with

```
roots([a b c])
```

What happens if you try to compute the roots by hand or with a hand calculator?

You should find that the classic formula is good for computing one root, but not the other. So use it to compute one root accurately and then use the fact that

$$x_1 x_2 = \frac{c}{a}$$

to compute the other.

1.39. The power series for  $\sin x$  is

$$\sin x = x - \frac{x^3}{3!} + \frac{x^5}{5!} - \frac{x^7}{7!} + \cdots$$

This MATLAB function uses the series to compute  $\sin x$ .

```
function s = powersin(x)
% POWERSIN. Power series for sin(x).
% POWERSIN(x) tries to compute sin(x)
% from a power series
s = 0;
t = x;
n = 1;
while s+t ~= s;
    s = s + t;
    t = -x.^2/((n+1)*(n+2)).*t;
    n = n + 2;
end
```

What causes the `while` loop to terminate?

Answer the following questions for  $x = \pi/2$ ,  $11\pi/2$ ,  $21\pi/2$ , and  $31\pi/2$ :

How accurate is the computed result?

How many terms are required?

What is the largest term in the series?

What do you conclude about the use of floating-point arithmetic and power series to evaluate functions?

1.40. *Steganography* is the technique of hiding messages or other images in the low-order bits of the data for an image. The MATLAB `image` function has a hidden image that contains other hidden images. To see the top-level image, just execute the single command



```
image
```


Then, to improve its appearance,

```
colormap(gray(32))
truesize
axis ij
axis image
axis off
```

But that's just the beginning. The NCM program `stegano` helps you continue the investigation.

- (a) How many images are hidden in the `cdata` for the default image?
- (b) What does this have to do with the structure of floating-point numbers?

- 1.41. *Prime spirals.* A Ulam prime spiral is a plot of the location of the prime numbers using a numbering scheme that spirals outward from the center of a grid. Our NCM file `primespiral(n,c)` generates an  $n$ -by- $n$  prime spiral starting with the number  $c$  in the center. The default is  $c = 1$ . Figure 1.10 is `primespiral(7)` and Figure 1.11 is `primespiral(250)`.



**Figure 1.10.** `primespiral(7)`.

The concentration of primes on some diagonal segments is remarkable, and not completely understood. The value of the element at position  $(i, j)$  is a piecewise quadratic function of  $i$  and  $j$ , so each diagonal segment represents a mini-theorem about the distribution of primes. The phenomenon was discovered by Stanislaw Ulam in 1963 and appeared on the cover of *Scientific American* in 1964. There are a number of interesting Web pages devoted to prime spirals. Start with [7] and [9].

- (a) The MATLAB `demos` directory contains an M-file `spiral.m`. The integers from 1 to  $n^2$  are arranged in a spiral pattern, starting in the center of the matrix. The code in `demos/spiral.m` is not very elegant. Here is a better version.




Figure 1.11. `primespiral(250)`.

```
function S = spiral(n)
%SPIRAL SPIRAL(n) is an n-by-n matrix with elements
% 1:n^2 arranged in a rectangular spiral pattern.
S = [];
for m = 1:n
    S = rot90(S,2);
    S(m,m) = 0;
    p = ???
    v = (m-1:-1:0);
    S(:,m) = p-v';
    S(m,:) = p+v;
end
if mod(n,2)==1
    S = rot90(S,2);
end
```

What value should be assigned to `p` each time through the loop so that this function generates the same matrices as `spiral.m` in the `demos` directory?

- (b) Why do half of the diagonals of `spiral(n)` contain no primes?
- (c) Let `S = spiral(2*n)` and let `r1` and `r2` be rows that go nearly halfway across the middle of the matrix:

```
r1 = S(n+1,1:n-2)
r2 = S(n-1,n+2:end)
```

Why do these rows contain no primes?

(d) There is something particularly remarkable about

```
primespiral(17,17)
primespiral(41,41)
```


What is it?

(e) Find values of  $n$  and  $c$ , both less than 50, and not equal to 17 or 41, so that

```
[S,P] = primespiral(n,c)
```

contains a diagonal segment with 8 or more primes.

- 1.42. *Triangular numbers* are integers of the form  $n(n+1)/2$ . The term comes from the fact that a triangular grid with  $n$  points on a side has a total of  $n(n+1)/2$  points. Write a function `trinums(m)` that generates all the triangular numbers less than or equal to  $m$ . Modify `primespiral` to use your `trinums` and become `trinumspriral`.
- 1.43. Here is a puzzle that does not have much to do with this chapter, but you might find it interesting nevertheless. What familiar property of the integers is represented by the following plot?



- 1.44. In the Gregorian calendar, a year  $y$  is a *leap year* if and only if

$$(\text{mod}(y,4) == 0) \ \& \ (\text{mod}(y,100) \neq 0) \ | \ (\text{mod}(y,400) == 0)$$

Thus 2000 was a leap year, but 2100 will not be a leap year. This rule implies that the Gregorian calendar repeats itself every 400 years. In that 400-year period, there are 97 leap years, 4800 months, 20871 weeks, and 146097 days. The MATLAB functions `datenum`, `datevec`, `datestr`, and `weekday` use these facts to facilitate computations involving calendar dates. For example, either of the statements

```
[d,w] = weekday('Aug. 17, 2003')
```

and

```
[d,w] = weekday(datenum([2003 8 17]))
```

tells me that my birthday was on a Sunday in 2003.

Use MATLAB to answer the following questions.

(a) On which day of the week were you born?

- (b) In a 400-year Gregorian calendar cycle, which weekday is the most likely for your birthday?
- (c) What is the probability that the 13th of any month falls on a Friday? The answer is close to, but not exactly equal to,  $1/7$ .
- 1.45. *Biorhythms* were very popular in the 1960s. You can still find many Web sites today that offer to prepare personalized biorhythms, or that sell software to compute them. Biorhythms are based on the notion that three sinusoidal cycles influence our lives. The physical cycle has a period of 23 days, the emotional cycle has a period of 28 days, and the intellectual cycle has a period of 33 days. For any individual, the cycles are initialized at birth. Figure 1.12 is my biorhythm, which begins on August 17, 1939, plotted for an eight-week period centered around the date this is being written, October 19, 2003. It shows that my intellectual power reached a peak yesterday, that my physical strength and emotional wellbeing will reach their peaks within 6 h of each other on the same day next week, and that all three cycles will be at their low point within a few days of each other early in November.




Figure 1.12. *My biorhythm.*

The date and graphics functions in MATLAB make the computation and display of biorhythms particularly convenient. Dates are represented by their *date number*, which is the number of days since the zeroth day of a theoretical calendar year zero. The function `datenum` returns the date number for any given date and time. For example, `datenum('Oct. 19, 2003')` is 731873. The expression `fix(now)` returns the date number of the current date. The following code segment is part of a program that plots a biorhythm for an eight-week period centered on the current date.

```
t0 = datenum(mybirthday);
t1 = fix(now);
t = (t1-28):1:(t1+28);
y = 100*[sin(2*pi*(t-t0)/23)
        sin(2*pi*(t-t0)/28)
        sin(2*pi*(t-t0)/33)];
plot(t,y)
```

- 
- (a) Complete this program, using your own birthday, and the `line`, `datetick`, `title`, `datestr`, and `legend` functions. Your program should produce something like Figure 1.12.
- (b) All three cycles started at zero when you were born. How long does it take until all three simultaneously return to that initial condition? How old were you, or will you be, on that date? Plot your biorhythm near that date. You should find the `lcm` function helpful.
- (c) Is it possible for all three cycles to reach their maximum or minimum at exactly the same time?



# Bibliography

- [1] M. BARNESLEY, *Fractals Everywhere*, Academic Press, Boston, 1993.
- [2] D. C. HANSELMAN AND B. LITTLEFIELD, *Mastering MATLAB 6, A Comprehensive Tutorial and Reference*, Prentice-Hall, Upper Saddle River, NJ, 2000.
- [3] D. J. HIGHAM AND N. J. HIGHAM, *MATLAB Guide*, SIAM, Philadelphia, 2000.
- [4] N. J. HIGHAM, *Accuracy and Stability of Numerical Algorithms*, SIAM, Philadelphia, 2002.
- [5] J. LAGARIAS, *The  $3x + 1$  problem and its generalizations*, American Mathematical Monthly, 92 (1985), pp. 3–23.  
<http://www.cecm.sfu.ca/organics/papers/lagarias>
- [6] M. OVERTON, *Numerical Computing with IEEE Floating Point Arithmetic*, SIAM, Philadelphia, 2001.
- [7] I. PETERSON, *Prime Spirals*, Science News Online, 161 (2002).  
<http://www.sciencenews.org/20020504/mathtrek.asp>
- [8] K. SIGMON AND T. A. DAVIS, *MATLAB Primer*, Sixth Edition, Chapman and Hall/CRC, Boca Raton, FL, 2002.
- [9] E. WEISSTEIN, *World of Mathematics, Prime Spiral*,  
<http://mathworld.wolfram.com/PrimeSpiral.html>
- [10] THE MATHWORKS, INC., *Getting Started with MATLAB*.  
[http://www.mathworks.com/access/helpdesk/help/techdoc/learn\\_matlab/learn\\_matlab.shtml](http://www.mathworks.com/access/helpdesk/help/techdoc/learn_matlab/learn_matlab.shtml)
- [11] THE MATHWORKS, INC., List of MATLAB-based books.  
<http://www.mathworks.com/support/books/index.jsp>