Employees.xls

	A	B	C	D	E
1	Name	Annual Salary	Monthly Salary	Location	Date Hired
2	James Brackman	42,400	3,533	New York	2/1/93
3	Michael Orenthal	28,900	2,408	Arizona	4/5/94
4	Francis Jenkins	67,800	5,650	New York	10/12/93
5	Peter Yolanda	19,850	1,654	Minnesota	1/4/95
6	Walter Franklin	45,000	3,750	Arizona	2/28/90
7	Louise Victor	52,000	4,333	New York	5/2/94
8	Sally Rice	48,500	4,042	New York	11/21/92
9	Charles K. Barkley	24,500	2,042	Minnesota	6/4/90
10	Melinda Hintquest	56,400	4,700	Arizona	6/1/87
11	Linda Harper	75,000	6,250	Minnesota	8/7/91
12					
13					
14					

Sheet1

Figure 23-1: An example of a list.

People often refer to the columns in a list as *fields* and to the rows as *records*. Using this terminology, the list shown in the figure has five fields (Name, Annual Salary, Monthly Salary, Location, and Date Hired) and ten records.

The size of the lists that you develop in Excel is limited by the size of a single worksheet. In other words, a list can have no more than 256 fields and can consist of no more than 65,535 records (one row contains the field names). A list of this size would require a great deal of memory and even then may not be possible. At the other extreme, a list can consist of a single cell—not very useful, but it's still considered a list.

In versions of Excel prior to Excel 97, a list was limited to 16,383 records.

What Can You Do with a List?

Excel provides several tools to help you manage and manipulate lists. Consequently, people use lists for a wide variety of purposes. For some users, a list is simply a method to keep track of information (for example, customer lists); others use lists to store data that ultimately will appear in a report. Common list operations include:

- ✦ Entering data into the list
- ✦ Filtering the list to display only the rows that meet certain criteria
- ✦ Sorting the list
- ✦ Inserting formulas to calculate subtotals
- ✦ Creating formulas to calculate results on the list filtered by certain criteria
- ✦ Creating a summary table of the data in the list (this is done using a pivot table; see Chapter 25).

With the exception of the last item, these operations are covered in this chapter.

Designing a List

Although Excel is quite accommodating when it comes to the information that is stored in a list, planning the organization of your list information will pay off. The following are some guidelines to keep in mind when creating lists:

- ✦ Insert descriptive labels (one for each column) in the first row of the list, called the *header row*. If you use lengthy labels, consider using the Wrap Text format so that you don't have to widen the columns.

See Chapter 11 for information on the Wrap Text format.

- ✦ Make sure each column contains the same type of information. For example, don't mix dates and text in a single column.
- ✦ You can use formulas that perform calculations on other fields in the same record. If you use formulas that refer to cells outside the list, make these absolute references; otherwise, you get unexpected results when you sort the list.
- ✦ Don't leave any empty rows within the list. For list operations, Excel determines the list boundaries automatically, and an empty row signals the end of the list.
- ✦ For best results, try to keep the list on a worksheet by itself. If you must place other information on the same worksheet as the list, place the information above or below the list. In other words, don't use the cells to the left or the right of a list.
- ✦ Select Window ➪ Freeze Panes to make sure that you can see the headings when you scroll the list.
- ✦ You can preformat entire columns to ensure that the data has the same format. For example, if a column contains dates, format the entire column with the desired date format.

Many people find working in spreadsheets most appealing because changing the layout is relatively easy. Lists behave no differently than any other kind of data in Excel; changing a list's layout is also easy. For example, you may create a list and then decide that it needs another column (field). No problem. Just insert a new column, give it a field name, and Excel expands your list. If you've ever used a database management program, you can appreciate the simplicity of this layout change.

Entering Data into a List

You can enter data into a list in three ways:

- ✦ Manually, using all standard data entry techniques
- ✦ By importing it or copying it from another file
- ✦ By using a dialog box

There's really nothing special about entering data into a list. You just navigate through the worksheet and enter the data into the appropriate cells.

Excel has two features that assist with repetitive data entry:

- ✦ **AutoComplete.** When you begin to type in a cell, Excel scans up and down the column for entries that match what you're typing. If it finds a match, Excel fills in the rest of the text automatically. Press Enter to make the entry. You can turn this feature on or off in the Edit tab of the Options dialog box.
- ✦ **Pick Lists.** You can right-click on a cell and select Pick from list from the shortcut menu. Excel displays a list box that shows all entries in the column (see Figure 23-2). Click on the one that you want to enter into the cell (no typing is required).

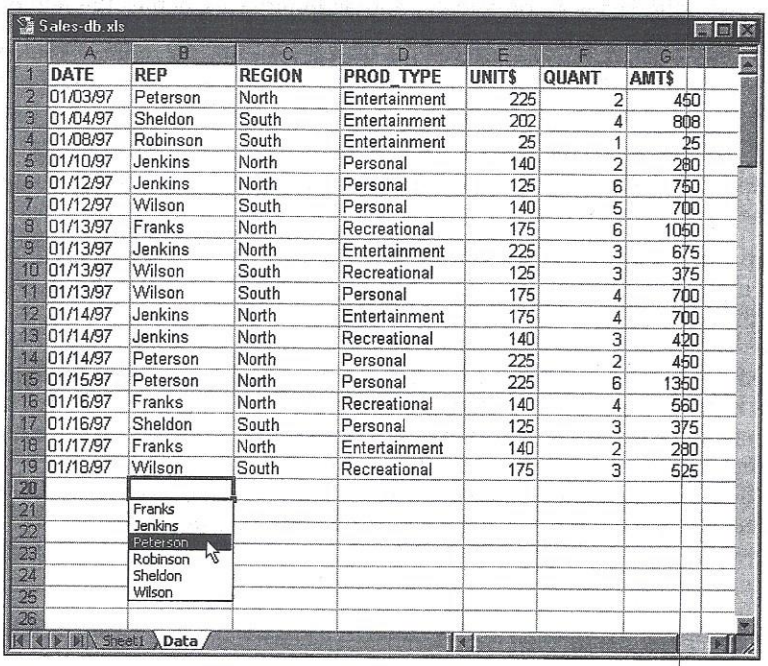

Figure 23-2: Choosing the Pick from list command on the shortcut menu gives you a list of all items in the current column.

If you prefer to use a dialog box for your data entry, Excel accommodates you. To display a data entry dialog box, move the cell pointer anywhere within the list and choose Data ⇨ Form. Excel determines the boundaries of your list and displays a dialog box showing each field in the list. Figure 23-3 depicts an example of such a dialog box. Fields that have a formula don't have an edit box.

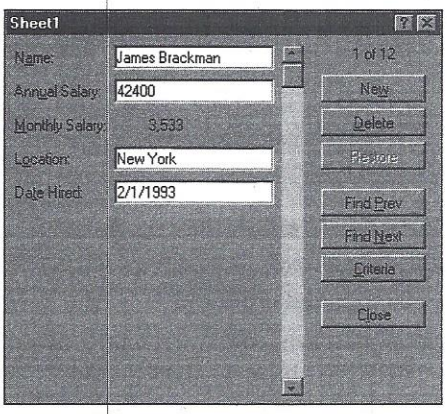

Figure 23-3: The Data ⇨ Form command gives you a handy data entry dialog box.

If the number of fields exceeds the limit of your display, the dialog box contains two columns of field names. If your list consists of more than 32 fields, however, the Data ⇨ Form command doesn't work. You must forego this method of data entry and enter the information directly into the cells.

Entering Data with the Data Form Dialog Box

When the data form dialog box appears, Excel displays the first record in the list. Notice the indicator in the upper-right corner of the dialog box that tells you the number of the selected record and the total number of records in the list.

To enter a new record, click on the New button to clear the fields. Then you can enter the new information into the appropriate fields. Use Tab or Shift+Tab to move among the fields. When you click on New (or Close), Excel appends the data that you entered to the bottom of the list. You also can press Enter, which is equivalent to clicking on the New button. If the list contains any formulas, Excel enters them for you automatically into the new record.

If you named the range of your list Database, Excel automatically extends the range definition to include the new row(s) that you add to the list using the data form dialog box. Note that this works only if you name the list Database; any other name doesn't work.

Other Uses for the Data Form Dialog Box

You can use the data form dialog box for more than just data entry. You can edit existing data in the list, view data one record at a time, delete records, and display records that meet certain criteria.

The dialog box contains a number of additional buttons, which are described as follows:

- ✦ **Delete:** Deletes the displayed record.
- ✦ **Restore:** Restores any information that you edited. You must click on this button before you click on the New button.
- ✦ **Find Prev:** Displays the previous record in the list. If you entered a criterion, this button displays the previous record that matches the criterion.
- ✦ **Find Next:** Displays the next record in the list. If you entered a criterion, this button displays the next record that matches the criterion.
- ✦ **Criteria:** Clears the fields and lets you enter a criterion upon which to search for records. For example, to locate records that have a salary greater than $50,000, enter **>50000** into the Salary field. Then you can use the Find Next and Find Prev buttons to display the qualifying records.
- ✦ **Close:** Closes the dialog box (and enters the data that you were entering, if any).

Using Microsoft Access Forms for Data Entry

If you have Microsoft Access installed on your system, you can use its form creation tools to develop a data entry form for an Excel worksheet. This feature uses the Access Links add-in, which must be loaded. When the add-in is loaded, you have a new command: Data ⇨ Access Form.

Choosing this command starts Access (if it's not already running) and begins its Form Wizard. Use the Form Wizard to create the data entry form. You can then use this form to add data to your Excel worksheet. Access's Form Wizard places a button on your worksheet that contains the text View Access Form. Click on this button to use the form. Figure 23-4 shows an Access form being used to enter data into an Excel worksheet.

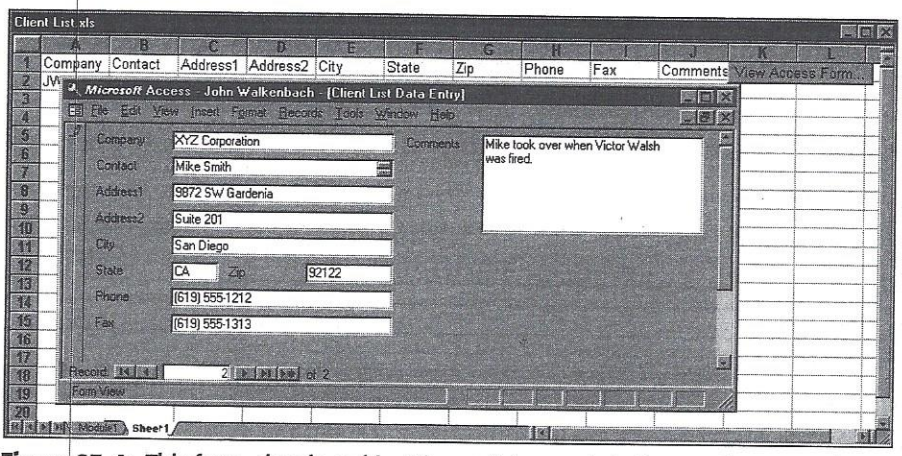

Figure 23-4: This form, developed in Microsoft Access, is being used to enter data into an Excel worksheet.

Filtering a List

Filtering a list is the process of hiding all rows in the list except those that meet some criteria that you specify. For example, if you have a list of customers, you can filter the list to show only those who live in New Jersey. Filtering is a common (and very useful) technique. Excel provides two ways to filter a list:

- ✦ AutoFilter, for simple filtering criteria
- ✦ Advance Filter, for more complex filtering

AutoFiltering

To use Excel's AutoFilter feature to filter a list, place the cell pointer anywhere within the list and then choose Data ➪ Filter ➪ AutoFilter. Excel analyzes your list and adds drop-down arrows to the field names in the header row, as shown in Figure 23-5.

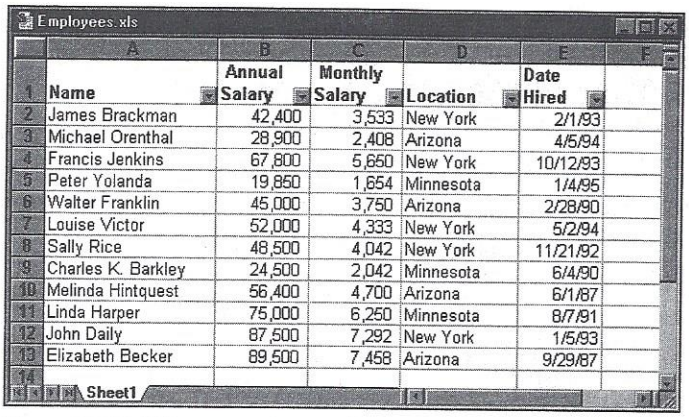

Employees.xls

	A	B	C	D	E	F
1	Name	Annual Salary	Monthly Salary	Location	Date Hired	
2	James Brackman	42,400	3,533	New York	2/1/93	
3	Michael Orenthal	28,900	2,408	Arizona	4/5/94	
4	Francis Jenkins	67,800	5,650	New York	10/12/93	
5	Peter Yolanda	19,850	1,654	Minnesota	1/4/95	
6	Walter Franklin	45,000	3,750	Arizona	2/28/90	
7	Louise Victor	52,000	4,333	New York	5/2/94	
8	Sally Rice	48,500	4,042	New York	11/21/92	
9	Charles K. Barkley	24,500	2,042	Minnesota	6/4/90	
10	Melinda Hintquest	56,400	4,700	Arizona	6/1/87	
11	Linda Harper	75,000	6,250	Minnesota	8/7/91	
12	John Daily	87,500	7,292	New York	1/5/93	
13	Elizabeth Becker	89,500	7,458	Arizona	9/29/87	
14						

Sheet1

Figure 23-5: When you choose the Data ⇨ Filter ⇨ AutoFilter command, Excel adds drop-down arrows to the field names in the header row.

When you click on the arrow in one of these drop-down lists, the list expands to show the unique items in that column. Select an item, and Excel hides all rows except those that include the selected item. In other words, Excel filters the list by the item that you selected.

After you filter the list, the status bar displays a message that tells you how many rows qualified. In addition, the drop-down arrow changes color to remind you that you filtered the list by a value in that column.

AutoFiltering has a limit. Only the first 999 unique items in the column appear in the drop-down list. If your list exceeds this limit, you can use advanced filtering, which is described later.

Besides showing every item in the column, the drop-down list includes five other items:

- ✦ **All:** Displays all items in the column. Use this to remove filtering for a column.
- ✦ **Top 10:** Filters to display the "top 10" items in the list; this is discussed later.
- ✦ **Custom:** Lets you filter the list by multiple items; this is discussed later.
- ✦ **Blanks:** Filters the list by showing rows that contain blanks in this column.
- ✦ **NonBlanks:** Filters the list by showing rows that contain non-blanks in this column.

To display the entire list again, click on the arrow and choose All — the first item in the drop-down list. Or, you can select Data ⇨ Filter ⇨ Show All.

To move out of Autofilter mode and remove the drop-down arrows from the field names, choose Data ⇨ Filter ⇨ AutoFilter again to remove the check mark from the AutoFilter menu item and restore the list to its normal state.

If you have any formulas that refer to data in a filtered list, be aware that the formulas don't adjust to use only the visible cells. For example, if a cell contains a formula that sums values in column C, the formula continues to show the sum for *all* the values in column C—not just those in the visible rows. To solve this problem, use database functions, which I describe later in this chapter.

Multicolumn AutoFiltering

Sometimes you may need to filter a list by values in more than one column. Figure 23-6 shows a list comprised of several fields.

sales.xls

	A	B	C	D	E
1	Month	State	Product	Price	From Ad
2	Jan	CA	Printer	208	Yes
3	Jan	CA	Printer	203	No
4	Jan	IL	Printer	468	No
5	Jan	IL	Printer	226	No
6	Jan	NY	Printer	484	Yes
7	Jan	NY	Printer	373	Yes
8	Jan	CA	Modem	249	Yes
9	Jan	CA	Modem	329	No
10	Jan	IL	Modem	760	Yes
11	Jan	IL	Modem	959	No
12	Jan	NY	Modem	419	No
13	Jan	NY	Modem	555	No
14	Jan	CA	HardDrive	287	Yes
15	Jan	CA	HardDrive	758	No
16	Jan	IL	HardDrive	651	Yes
17	Jan	IL	HardDrive	233	No
18	Jan	NY	HardDrive	332	Yes
19	Jan	NY	HardDrive	852	Yes
20	Jan	CA	Mouse	748	No
21	Jan	CA	Mouse	811	No

Product Sales

Figure 23-6: The list before filtering by multiple columns.

Assume that you want to see the records that show modems sold in February. In other words, you want to filter out all records except those in which the Month field is *Feb* and the Product field is *Modem*.

First, get into Autofilter mode. Then click on the drop-down arrow in the Month field and select *Feb* to filter the list to show only records with *Feb* in the Month field. Then click on the drop-down arrow in the Product field and select *Modem*, filtering the filtered list to show only records that contain *Modem* in the Product column—resulting in a list filtered by values in two columns. Figure 23-7 shows the result.

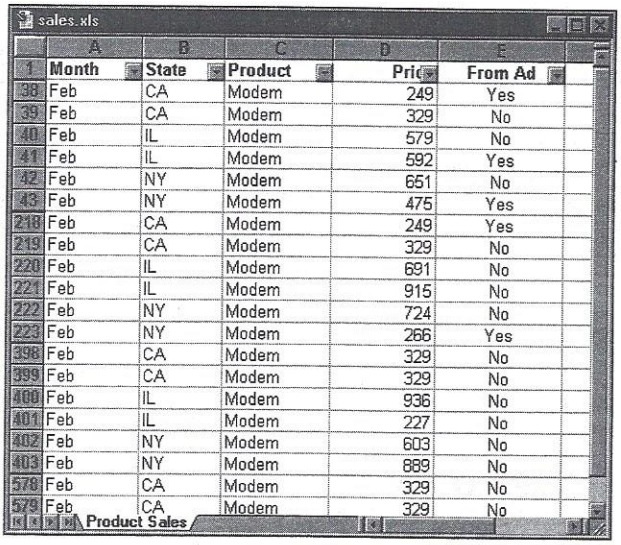

Figure 23-7: The same list filtered by values in two columns.

You can filter a list by any number of columns. Excel applies a different color to the drop-down arrows in the columns that have a filter applied.

Custom AutoFiltering

Usually, AutoFiltering involves selecting a single value for one or more columns. If you choose the Custom option in a drop-down list, you gain a bit more flexibility in filtering the list; Excel displays a dialog box like the one shown in Figure 23-8.

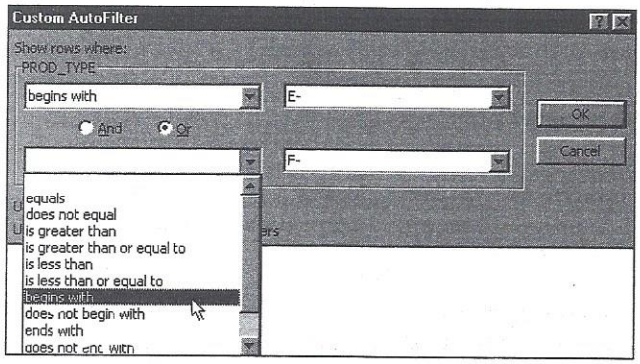

Figure 23-8: The Custom AutoFilter dialog box gives you more filtering options.

The Custom AutoFilter dialog box lets you filter in several ways:

- **Values above or below a specified value.** For example, sales amounts greater than 10,000.
- **Values within a range.** For example, sales amounts greater than 10,000 AND sales amounts less than 50,000.
- **Two discrete values.** For example, state equal to *New York* OR state equal to *New Jersey*.
- **Approximate matches.** You can use the * and ? wildcards to filter in a number of other ways. For example, to display only those customers whose last name begins with *B*, use **B***.

Custom AutoFiltering can be useful, but it definitely has limitations. For example, if you want to filter the list to show only three values in a field (such as New York or New Jersey or Connecticut), you can't do it by AutoFiltering. Such filtering tasks require the advanced filtering feature, which I discuss later in this chapter.

Top 10 AutoFiltering

Sometimes you may want to use a filter on numerical fields to show only the highest or lowest values in the list. For example, if you have a list of employees, you may want to identify the 12 employees with the longest tenure. You could use the custom AutoFilter option, but then you must supply a cutoff date (which you may not know). The solution is to use Top 10 AutoFiltering.

Top 10 AutoFiltering is a generic term; it doesn't limit you to the top *10* items. In fact, it doesn't even limit you to the *top* items. When you choose the Top 10 option from a drop-down list, you see dialog box that is shown in Figure 23-9.

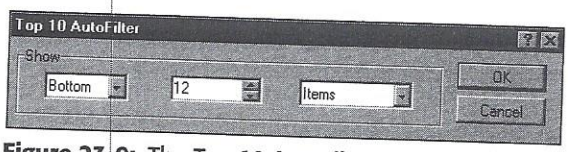

Figure 23-9: The Top 10 AutoFilter gives you more AutoFilter options.

You can choose either Top or Bottom and specify any number. Suppose, for example, that you want to see the 12 employees with the longest tenure. Choose Bottom and 12 to filter the list and show the 12 rows with the smallest values in the Date Hired field. You also can choose Percent or Value in this dialog box. For example, you can filter the list to show the Bottom 5 percent of the records.

Charting filtered list data

You can create some interesting multipurpose charts that use data in a filtered list. The technique is useful because only the visible data appears in the chart. When you change the AutoFilter criteria, the chart updates itself to show only the visible cells.

For this technique to work, select the chart and make sure that the Plot Visible Cells Only option is enabled on the Chart tab of the Options dialog box.

Figure 23-10 shows an example of a chart created with an unfiltered list. It shows sales data for three months for each of four sales regions.

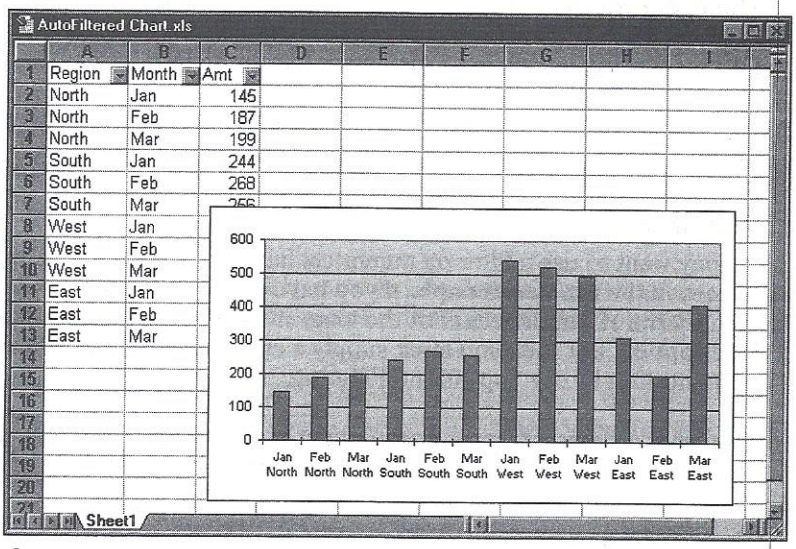

Figure 23-10: This chart was created from an unfiltered list.

Figure 23-11 shows the same chart, but the list was filtered to show only the North sales region. You can apply other filters, and the chart updates automatically. This technique lets a single chart show several different views of the data.

Advanced Filtering

In many cases, AutoFiltering does the job. But if you run up against its limitations, you need to use advanced filtering. Advanced filtering is much more flexible than AutoFiltering, but it takes a bit of up-front work to use it. Advanced filtering provides you with the following capabilities:

- ✦ You can specify more complex filtering criteria.
- ✦ You can specify computed filtering criteria.
- ✦ You can extract a copy of the rows that meet the criteria to another location.

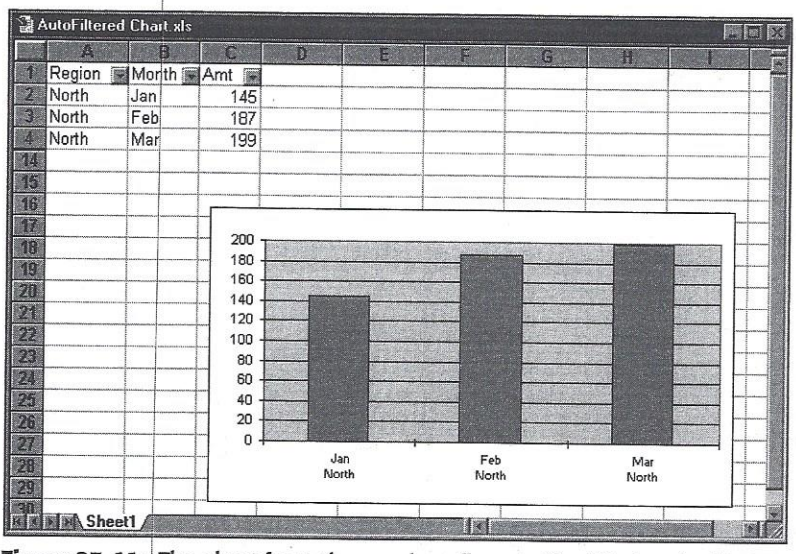

Figure 23-11: The chart from the previous figure, after filtering the list.

Setting up a criteria range

Before you can use the advanced filtering feature, you must set up a *criteria range,* a designated range on a worksheet that conforms to certain requirements. The criteria range holds the information that Excel uses to filter the list. It must conform to the following specifications:

- ✦ It must consist of at least two rows, and the first row must contain some or all field names from the list.
- ✦ The other rows of the criteria range must consist of your filtering criteria.

Although you can put the criteria range anywhere in the worksheet, it's a good idea not to put it in rows where you placed the list. Because Excel hides some of these rows when filtering the list, you may find that your criteria range is no longer visible after filtering. Therefore, you should generally place the criteria range above or below the list.

Figure 23-12 shows a criteria range, located in A1:D2, above the list that it uses. Only some field names appear in the criteria range. You don't need to include, in the criteria range, field names for fields that you don't use in the selection criteria.

Sales-db.xls

	A	B	C	D	E	F	G
1	Month	SalesRep	Type	TotalSale			
2	January		New				
3							
4							
5							
6	Month	SalesRep	Type	UnitCost	Quantity	TotalSale	
7	March	Wilson	New	175	5	875	
8	March	Wilson	New	140	3	420	
9	February	Franks	Existing	225	1	225	
10	March	Wilson	New	125	5	625	
11	January	Peterson	Existing	225	2	450	
12	March	Sheldon	New	140	2	280	
13	February	Peterson	Existing	225	6	1350	
14	March	Jenkins	Existing	140	2	280	
15	February	Sheldon	New	225	4	900	
16	January	Wilson	New	140	4	560	
17	January	Wilson	New	125	3	375	
18	January	Sheldon	New	225	6	1350	
19	February	Sheldon	New	175	5	875	
20	January	Robinson	New	140	3	420	
21	February	Sheldon	New	125	2	250	
22	March	Sheldon	New	140	6	840	
23	March	Jenkins	Existing	225	3	675	
24	January	Robinson	New	225	2	450	
25	March	Sheldon	New	225	6	1350	
26	February	Wilson	New	140	3	420	

Sheet1 Data

Figure 23-12: A criteria range for a list.

In this example, the criteria range has only one row of criteria. The fields in each row of the criteria range (except for the header row) are joined with an AND operator. Therefore, the filtered list shows rows in which the Month column equals *January* AND the Type column equals *New*. In other words, the list displays only sales to new customers made in January.

To perform the filtering, choose Data ➪ Filter ➪ Advanced filter. Excel displays the dialog box that is shown in Figure 23-13. Specify the list range and the criteria range, and make sure that you select the option labeled Filter the List in-place. Click on OK, and Excel filters the list by the criteria that you specified.

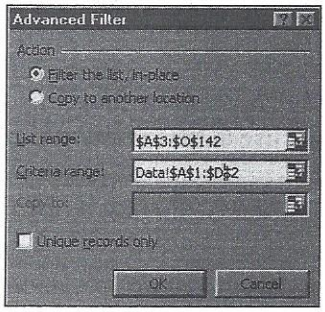

Figure 23-13: The Advanced Filter dialog box.

Multiple criteria

If you use more than one row in the criteria range, the criteria in each row are joined with an OR operator. A criteria range can have any number of rows, each of which is joined to the others with an OR operator. Figure 23-14 shows a criteria range (A1:D3) with two rows of criteria.

Sales-db.xls

	A	B	C	D	E	F	G
1	**Month**	**SalesRep**	**Type**	**TotalSale**			
2	January		New				
3	February			>1000			
4							
5							
6	**Month**	**SalesRep**	**Type**	**UnitCost**	**Quantity**	**TotalSale**	
7	March	Wilson	New	175	5	875	
8	March	Wilson	New	140	3	420	
9	February	Franks	Existing	225	1	225	
10	March	Wilson	New	125	5	625	
11	January	Peterson	Existing	225	2	450	
12	March	Sheldon	New	140	2	280	
13	February	Peterson	Existing	225	6	1350	
14	March	Jenkins	Existing	140	2	280	
15	February	Sheldon	New	225	4	900	
16	January	Wilson	New	140	4	560	
17	January	Wilson	New	125	3	375	
18	January	Sheldon	New	225	6	1350	
19	February	Sheldon	New	175	5	875	
20	January	Robinson	New	140	3	420	
21	February	Sheldon	New	125	2	250	
22	March	Sheldon	New	140	6	840	
23	March	Jenkins	Existing	225	3	675	
24	January	Robinson	New	225	2	450	
25	March	Sheldon	New	225	6	1350	
26	February	Wilson	New	140	3	420	

Sheet1 Data

Figure 23-14: This criteria range has two sets of criteria.

In this example, the filtered list shows rows in either of the following:

- ✦ *The Month field is January AND the Type field is New.*
- ✦ The Month field is *February* AND the Total Sale field is greater than 1000.

You cannot filter this way with AutoFiltering.

Types of criteria

The entries that you make in the criteria range can be either of the following:

- ✦ Text or value criteria. The filtering involves comparisons to a value or string, using operators such as equal (=), greater than (>), not equal to (<>), and so on.
- ✦ **Computed criteria.** The filtering involves a computation of some sort.

Text or value criteria

Table 23-1 lists the comparison operators that you can use with text or value criteria.

Table 23-1
Comparison Operators

Operator	*Comparison Type*
=	Equal to
>	Greater than
>=	Greater than or equal to
<	Less than
<=	Less than or equal to
<>	Not equal to

Table 23-2 shows examples of criteria that use strings.

Table 23-2
Examples of String Criteria

Criteria	*Effect*
>K	Text that begins with *L* through *Z*
<>C	All text, except text that begins with C
="January"	Text that matches January
Sm*	Text that begins with Sm
s*s	Text that begins with s and ends with s
s?s	Three-letter text that begins with *s* and ends with *s*

The text comparisons are not case sensitive. For example, si* matches *Simpson* as well as *sick*.

Computed criteria

Using computed criteria can make filtering even more powerful. Computed criteria filter the list based one or more calculations. Figure 23-15 shows a simple list that consists of project numbers, start dates, end dates, and resources. Above the list, in range A1:A2, is the criteria range. Notice, however, that this criteria range does

not use a field header from the list—it uses a new field header. A computed criteria essentially computes a new field for the list. Therefore, you must supply new field names in the first row of the criteria range.

Project Tracking.xls

	A	B	C	D	E
1	ProjLength				
2	TRUE				
3					
4	Project Number	Start Date	End Date	Resources	
5	AS-109	03/05/97	04/09/97	3,395	
6	AS-110	03/12/97	03/17/97	485	
7	AS-111	04/01/97	04/10/97	873	
8	AS-112	04/01/97	05/03/97	3,104	
9	AS-113	04/12/97	05/01/97	1,843	
10	AS-114	04/21/97	06/05/97	4,365	
11	AS-115	05/03/97	05/15/97	1,164	
12	AS-116	05/21/97	06/09/97	1,843	
13	AS-117	06/02/97	08/01/97	5,820	

Project List

Figure 23-15: This list is to be filtered using computed criteria.

Cell A2 contains the following formula:

```
=C5-B5+1>=30
```

This formula returns a logical value of either *True* or *False*. The result of the formula refers to cells in the first row of data in the list; it does *not* refer to the header row. When you filter the list by this criterion, the list shows only rows in which the project length (End Date–Start Date+1) is greater than or equal to 30 days. In other words, Excel bases the comparison on a computation.

You could accomplish the same effect, without using a computed criterion, by adding a new column to the list that contains a formula to calculate the project length. Using a computed criterion, however, eliminates the need to add a new column.

To filter the list to show only the projects that use above average resources, you could use the following computed criteria formula:

```
=D5>AVERAGE(D:D)
```

This filters the list to show only the rows in which the value of the Resources field is greater than the average of the Resources field.

Keep in mind the following items when using computed criteria:

- ✦ Don't use a field name in the criteria range that appears in the list. Create a new field name or just leave the cell blank.

- ✦ You can use any number of computed criteria and mix and match them with noncomputed criteria.
- ✦ Don't pay attention to the values returned by formulas in the criteria range. These refer to the first row of the list.
- ✦ If your computed formula refers to a value outside the list, use an absolute reference rather than a relative reference. For example, use C1 rather than C1.
- ✦ Create your computed criteria formulas using the first row of data in the list (not the field names). Make these references relative, not absolute. For example, use C5 rather than C5.

Other advanced filtering operations

The Advanced Filter dialog box gives you two other options:

- ✦ Copy to Another Location
- ✦ Unique Records Only

Both of these advanced filtering options are discussed below.

Copying qualifying rows

If you choose the Copy to Another Location option in the Advanced Filter dialog box, Excel copies the qualifying rows to another location in the worksheet or a different worksheet. You specify the location for the copied rows in the Copy to edit box. Note that the list itself is not filtered when you use this option.

Displaying only unique rows

Choosing the option labeled Unique records only hides all duplicate rows that meet the criteria that you specify. If you don't specify a criteria range, this option hides all duplicate rows in the list.

Using Database Functions with Lists

It's important to understand that Excel's worksheet functions don't ignore hidden cells. Therefore, if you have a SUM formula that calculates the total of the values in a column of a list, the formula returns the same value when you filter the list.

To create formulas that return results based on filtering criteria, you need to use Excel's database worksheet functions. For example, you can create a formula that calculates the sum of values in a list that meets certain criteria. Set up a criteria range as described previously. Then enter a formula such as the following:

```
=DSUM(ListRange,FieldName,Criteria)
```

In this case, ListRange refers to the list, FieldName refers to the field name cell of the column that you are summing, and Criteria refers to the criteria range.

Excel's database functions are listed in Table 23-3.

Table 23-3
Excel's Database Worksheet Functions

Function	*Description*
DAVERAGE	Returns the average of selected database entries
DCOUNT	Counts the cells containing numbers from a specified database and criteria
DCOUNTA	Counts nonblank cells from a specified database and criteria
DGET	Extracts from a database a single record that matches the specified criteria
DMAX	Returns the maximum value from selected database entries
DMIN	Returns the minimum value from selected database entries
DPRODUCT	Multiplies the values in a particular field of records that match the criteria in a database
DSTDEV	Estimates the standard deviation based on a sample of selected database entries
DSTDEVP	Calculates the standard deviation based on the entire population of selected database entries
DSUM	Adds the numbers in the field column of records in the database that match the criteria
DVAR	Estimates variance based on a sample from selected database entries
DVARP	Calculates variance based on the entire population of selected database entries

Refer to Chapter 10 for general information about using worksheet functions.

Sorting a List

In some cases, the order of the rows in your list doesn't matter. But in other cases, you want the rows to appear in a specific order. For example, in a price list, you may want the rows to appear in alphabetical order by product name. This makes the products easier to locate in the list. Or, if you have a list of accounts receivable information, you may want to sort the list so that the higher amounts appear at the top of the list (in descending order).

Rearranging the order of the rows in a list is called *sorting*. Excel is quite flexible when it comes to sorting lists, and you can often accomplish this task with the click of a mouse button.

Simple Sorting

To quickly sort a list in ascending order, move the cell pointer into the column that you want to sort. Then click on the Sort Ascending button the Standard toolbar. The Sort Descending button works the same way, but it sorts the list in descending order. In both cases, Excel determines the extent of your list and sorts all the rows in the list.

When you sort a filtered list, Excel sorts only the visible rows. When you remove the filtering from the list, the list is no longer sorted.

Be careful if you sort a list that contains formulas. If the formulas refer to cells in the list that are in the same row, you don't have any problems. But if the formulas refer to cells in other rows in the list or to cells outside the list, the formulas will not be correct after you sort the list. If formulas in your list refer to cells outside the list, make sure that the formulas use an absolute cell reference.

More Complex Sorting

Sometimes, you may want to sort by two or more columns. This is relevant to break ties. A tie occurs when rows with duplicate data remain unsorted. Figure 23-16 shows an example of an unsorted list. If you sort this list by Month, Excel places the rows for each month together. But you may also want to show the Sales Reps in ascending order within each month. In this case, you would need to sort by two columns (Month and Sales Rep). Figure 23-17 shows the list after sorting by these two columns.

You can use the Sort Ascending and Sort Descending buttons to do this — but you need to do two sorts. First, sort by the Sales Reps column, and then sort by the Month column. As I explain in the next section, Excel provides a way to accomplish multicolumn sorting with a single command.

Sales-db.xls

	A	B	C	D	E	F
7	Month	Sales Rep	Type	Unit Cost	Quantity	Total Sale
8	May	Sheldon	Existing	125	1	125
9	January	Sheldon	Existing	175	1	175
10	January	Sheldon	New	140	6	840
11	January	Jenkins	New	225	1	225
12	February	Robinson	New	225	1	225
13	March	Wilson	Existing	125	4	500
14	April	Robinson	Existing	125	2	250
15	February	Sheldon	Existing	175	1	175
16	March	Robinson	Existing	125	1	125
17	May	Jenkins	New	225	3	675
18	April	Jenkins	New	225	2	450
19	February	Wilson	Existing	125	5	625
20	February	Jenkins	New	225	2	450
21	January	Franks	New	225	4	900
22	May	Wilson	New	225	1	225
23	January	Sheldon	New	225	1	225
24	March	Jenkins	New	225	2	450
25	March	Jenkins	Existing	125	5	625
26	April	Peterson	New	140	2	280
27	February	Franks	Existing	175	2	350
28	May	Robinson	New	140	3	420
29	April	Peterson	Existing	175	6	1050
30	February	Robinson	New	225	3	675

SALES-DB

Figure 23-16: This list is unsorted.

Sales-db.xls

	A	B	C	D	E	F
7	Month	Sales Rep	Type	Unit Cost	Quantity	Total Sale
8	January	Franks	New	225	4	900
9	January	Franks	Existing	175	1	175
10	January	Franks	Existing	175	5	875
11	January	Franks	New	225	1	225
12	January	Franks	Existing	175	1	175
13	January	Franks	Existing	125	3	375
14	January	Jenkins	New	225	1	225
15	January	Jenkins	Existing	125	1	125
16	January	Jenkins	New	140	3	420
17	January	Jenkins	Existing	175	2	350
18	January	Jenkins	New	140	1	140
19	January	Peterson	Existing	125	1	125
20	January	Peterson	Existing	125	3	375
21	January	Peterson	New	140	1	140
22	January	Peterson	New	225	1	225
23	January	Robinson	New	140	2	280
24	January	Robinson	Existing	125	5	625
25	January	Robinson	Existing	175	4	700
26	January	Sheldon	Existing	175	1	175
27	January	Sheldon	New	140	6	840
28	January	Sheldon	New	225	1	225
29	January	Sheldon	Existing	125	2	250
30	January	Sheldon	Existing	175	5	875

SALES-DB

Figure 23-17: The list after sorting on two fields.

Excel's Sorting Rules

Because cells can contain different types of information, you may be curious about how Excel sorts this information. For an ascending sort, the information appears in the following order:

1. Values: Excel sorts numbers from smallest negative to largest positive, and treats dates and times as values. In all cases, Excel sorts using the actual values in cells (not their formatted appearance).
2. Text: In alphabetical order, as follows: 0 1 2 3 4 5 6 7 8 9 (space) ! " # $ % & ' ( ) * + , - . / : ; < = > ? @ [\ ] ^ _ ` { | } ~ 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.

   By default, sorting is not case sensitive. You can change this behavior, however, in the Sort Options dialog box (described in this chapter).
3. Logical values: False comes before True.
4. Error values: Error values (such as #VALUE! and #NA) appear in their original order; Excel does not sort them by error type.
5. Blank cells: Blanks cells always appear last.

Sorting in descending order reverses this sequence—except that blank cells *still* appear last.

The Sort dialog box

If you want to sort by more than one field, choose Data ⇨ Sort. Excel displays the dialog box that is shown in Figure 23-18. Simply select the first sort field from the drop-down list labeled Sort By, and specify Ascending or Descending order. Then, do the same for the second sort field. If you want to sort by a third field, specify the field in the third section. If the Header Row option is set, the first row (field names) is not affected by the sort. Click on OK, and the list's rows rearrange in a flash.

If the sorting didn't occur as you expected, select Edit ⇨ Undo (or press Ctrl+Z) to undo the sorting.

What if you need to sort your list by more than three fields? It can be done, but it takes an additional step. For example, assume that you want to sort your list by five fields: Field1, Field2, Field3, Field4, and Field5. Start by sorting by Field3, Field4, and Field5. Then re-sort the list by Field1 and Field2. In other words, sort the three "least important" fields first; they remain in sequence when you do the second sort.

Figure 23-18: The Sort dialog box lets you sort by up to three columns.

Tip Often, you want to keep the records in their original order but perform a temporary sort just to see how it looks. The solution is to add an additional column to the list with sequential numbers in it (don't use formulas to generate these numbers, but you can use the Fill command). Then, after you sort, you can return to the original order by resorting on the field that contains the sequential numbers. You can also use Excel's undo feature to return the list to its original order. If you use an additional column, you can perform other operations while the list is temporarily sorted (and these operations won't be undone when you undo the sort operation).

Sort options

When you click on the Options button in the Sort dialog box, Excel displays the Sort Options dialog, shown in Figure 23-19.

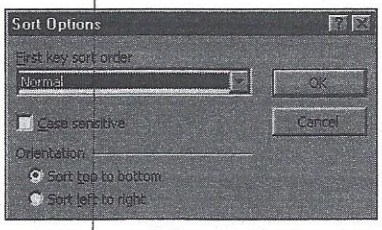

Figure 23-19: The Sort Options dialog gives you some additional sorting options.

These options are described as follows:

- **First key sort order:** Lets you specify a custom sort order for the sort (see the next section).

How Excel Identifies a Header Row

When you use the Data⇨Sort command, there's no need to select the list before you choose the command. That's because Excel examines the active cell position and then establishes the list's boundaries for you. In addition, Excel attempts to determine whether the list contains a header row. If the list has a header row, Excel excludes this row from the sort.

How does this happen? I'm not sure exactly, but the following seems to be Excel's "thought" process:

1. Select the current region. (You can do this manually: press F5, click on the Special button, select the Current Region option, and click on OK.)
2. Examine the first row of the selection.
3. Determine whether the first row contains any blanks. If so, this list has no header row.
4. Determine whether the first row contains text. If so, check the other cells. If they also contain text, this list has no header row.
5. Determine whether the first row contains uppercase text while the list itself contains lowercase or proper case text. If so, this list has a header row.
6. Determine whether the cells in the first row are formatted differently from the other cells in the list. If so, this list has a header row.

Knowing this information can help you eliminate incorrect sorting. For example, if you want to sort a range that doesn't have a header row, you need to make sure that Excel doesn't sort the data as if it had a header row. For best results, use the Sort Ascending and Sort Descending toolbar buttons only when the data that you're sorting has headers. If your data contains no headers, select Data⇨Sort and make sure that the No Header Row option is selected.

- ✦ **Case sensitive:** Makes the sorting case sensitive so that uppercase letters appear before lowercase letters in an ascending sort. Normally, sorting ignores the case of letters.
- ✦ **Orientation:** Enables you to sort by columns rather than by rows (the default).

Using a Custom Sort Order

Excel typically sorts either numerically or alphabetically, depending on the data being sorted. In some cases, however, you may want to sort your data in other ways. For example, if your data consists of month names, you usually want it to appear in month order rather than alphabetically. You can use the Sort Options dialog box to perform such a sort. Select the appropriate list from the drop-down list labeled First key sort order. Excel, by default, has four "custom lists," and you can define your own. Excel's custom lists are as follows:

- **Abbreviated days:** Sun, Mon, Tue, Wed, Thu, Fri, Sat
- **Days:** Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday
- **Abbreviated months:** Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, Dec
- **Months:** January, February, March, April, May, June, July, August, September, October, November, December

Note that the abbreviated days and months do not have periods after them. If you use periods for these abbreviations, Excel doesn't recognize them (and doesn't sort them correctly).

You may want to create a custom list. For example, your company may have several stores, and you want the stores to be listed in a particular order (not alphabetically). If you create a custom list, sorting puts the items in the order that you specify in the list. You must use the Data ➪ Sort command to sort by a custom list (click on the Options button to specify the custom list).

To create a custom list, use the Custom Lists tab of the Options dialog box, as shown in Figure 23-20. Select the NEW LIST option, and make your entries (in order) in the List Entries box. Or, you can import your custom list from a range of cells by selecting the range and then clicking the Import button.

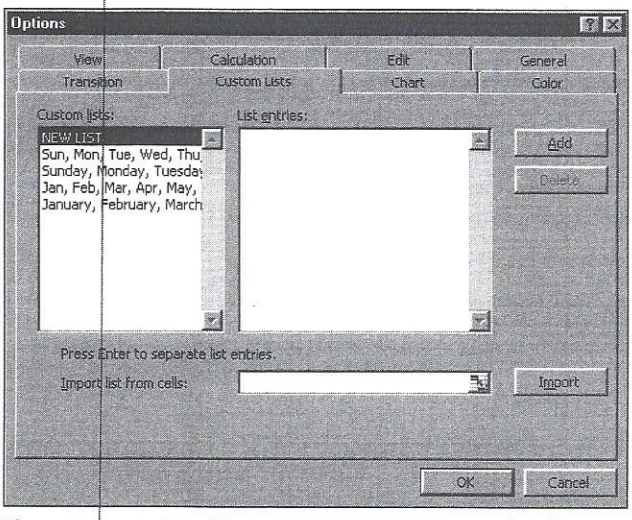

Figure 23-20: Excel lets you create custom sorting lists.

Custom lists also work with the AutoFill handle in cells. If you enter the first item of a custom list and then drag the cell's AutoFill handle, Excel fills in the remaining list items automatically.

Sorting Nonlists

You can, of course, sort any range in a worksheet — it doesn't have to be a list. You need to be aware of a few things, however. The Sort Ascending and Sort Descending toolbar buttons may assume (erroneously) that the top row is a header row and not include these cells in the sort (see the sidebar, "How Excel identifies a header row," earlier in this chapter).

Therefore, to avoid potential errors when sorting non-lists, don't use these toolbar buttons. Rather, select the entire range, and select Data ⇨ Sort (making sure that you choose the No Header Row option).

Creating Subtotals

The final topic of this chapter is automatic subtotals — a handy feature that can save you a great deal of time. To use this feature, your list must be sorted, because the subtotals are inserted whenever the value in a specified field changes. Figure 23-21 shows an example of a list, sorted by the Month field, which is appropriate for subtotals.

Sales-db.xls

	A	B	C	D	E	F	G
1	Month	Sales Rep	Type	Unit Cost	Quantity	Total Sale	
2	January	Franks	New	225	4	900	
3	January	Franks	Existing	175	5	875	
4	January	Franks	New	225	1	225	
5	January	Franks	Existing	175	1	175	
6	January	Jenkins	New	225	1	225	
7	January	Jenkins	Existing	125	1	125	
8	February	Franks	New	225	4	900	
9	February	Jenkins	New	225	2	450	
10	February	Jenkins	New	225	3	675	
11	February	Jenkins	New	225	3	675	
12	February	Jenkins	New	225	3	675	
13	February	Jenkins	Existing	175	1	175	
14	February	Peterson	New	225	1	225	
15	February	Peterson	New	225	2	450	
16	March	Peterson	Existing	125	2	250	
17	March	Peterson	New	225	2	450	
18	March	Robinson	Existing	125	1	125	
19	March	Robinson	Existing	125	5	625	
20	March	Robinson	New	225	4	900	
21	April	Franks	New	175	4	700	
22	April	Franks	New	175	3	525	
23	April	Jenkins	New	225	2	450	
24	April	Jenkins	New	140	3	420	

SALES-DB

Figure 23-21: This list is a good candidate for subtotals, which are inserted at each change of the month.

To insert subtotal formulas into a list automatically, move the cell pointer anywhere in the list and choose Data ⇨ Subtotals. You see the dialog box shown in Figure 23-22.

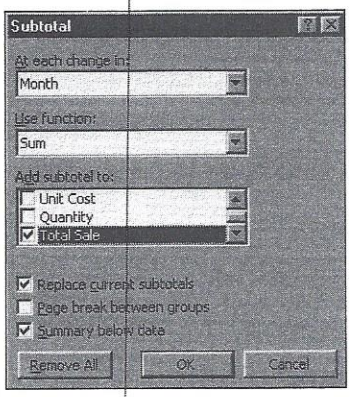

Figure 23-22: The Subtotal dialog box automatically inserts subtotal formulas into a sorted list.

This dialog box offers the following choices:

- **At Each Change in:** This drop-down list displays all fields in your list. You must have sorted the list by the field that you choose.
- **Use Function: Choose from 11 functions:** You should normally use Sum (the default).
- **Add Subtotal to:** This list box shows all the fields in your list. Place a check mark next to the field or fields that you want to subtotal.
- **Replace Current Subtotals:** If this box is checked, Excel removes any existing subtotal formulas and replaces them with the new subtotals.
- **Page Break Between Groups:** If this box is checked, Excel inserts a manual page break after each subtotal.
- **Summary Below Data:** If this box is checked, Excel places the subtotals below the data (the default). Otherwise, the subtotal formulas appear above the totals.
- **Remove All:** This button removes all subtotal formulas in the list.

When you click on OK, Excel analyzes the list and inserts formulas as specified—and creates an outline for you. The formulas all use the SUBTOTAL worksheet function.

When you add subtotals to a filtered list, the subtotals may no longer be accurate when the filter is removed.

Figure 23-23 shows a worksheet after adding subtotals.

Sales-db.xls

	A	B	C	D	E	F
1	Month	Sales Rep	Type	Unit Cost	Quantity	Total Sale
2	January	Franks	New	225	4	900
3	January	Franks	Existing	175	5	875
4	January	Franks	New	225	1	225
5	January	Franks	Existing	175	1	175
6	January	Jenkins	New	225	1	225
7	January	Jenkins	Existing	125	1	125
8	**January Total**					2525
9	February	Franks	New	225	4	900
10	February	Jenkins	New	225	2	450
11	February	Jenkins	New	225	3	675
12	February	Jenkins	New	225	3	675
13	February	Jenkins	New	225	3	675
14	February	Jenkins	Existing	175	1	175
15	February	Peterson	New	225	1	225
16	February	Peterson	New	225	2	450
17	**February Total**					4225
18	March	Peterson	Existing	125	2	250
19	March	Peterson	New	225	2	450
20	March	Robinson	Existing	125	1	125
21	March	Robinson	Existing	125	5	625
22	March	Robinson	New	225	4	900
23	**March Total**					2350
24	April	Franks	New	175	4	700

SALES-DB

Figure 23-23: Excel added the subtotal formulas automatically—and even created an outline.

Summary

In this chapter, I discuss lists. A list is simply a database table that is stored on a worksheet. The first row of the list (the header row) contains field names, and subsequent rows contain data (records). I offer some pointers on data entry and discuss two ways to filter a list to show only rows that meet certain criteria. AutoFiltering is adequate for many tasks, but if your filtering needs are more complex, you need to use advanced filtering. I end the chapter with a discussion of sorting and Excel's automatic subtotal feature.

✦ ✦ ✦

CHAPTER 27

Analyzing Data Using Goal Seeking and Solver

In This Chapter

The preceding chapter discusses *what-if analysis* — the process of changing input cells to observe the results on other dependent cells. This chapter looks at that process from the opposite perspective — finding the value of one or more input cells that produces a desired result in a formula cell.

What-If Analysis — In Reverse

Consider the following what-if question: "What is the total profit if sales increase by 20 percent?" If you set up your worksheet properly, you can change the value in one cell to see what happens to the profit cell. Goal seeking takes the opposite approach. If you know what a formula result *should* be, Excel can tell you the values that you need to enter in one or more input cells to produce that result. In other words, you can ask a question such as, "How much do sales need to increase to produce a profit of $1.2 million?" Excel provides two tools that are relevant:

- **Goal seeking:** Determines the value that you need to enter in a single input cell to produce a result that you want in a dependent (formula) cell.
- **Solver:** Determines the values that you need to enter in multiple input cells to produce a result that you want. Moreover, because you can specify certain constraints to the problem, you gain significant problem-solving ability.

Single-Cell Goal Seeking

Single-cell goal seeking (also known as *backsolving*) is a rather simple concept. Excel determines what value in an input cell produces a desired result in a formula cell. Walk through the following example to understand how single-cell goal seeking works.

A Goal-Seeking Example

Figure 27-1 shows the mortgage loan worksheet that was used in the preceding chapter. This worksheet has four input cells and four formula cells. Originally, this worksheet was used for a what-if analysis example. In this section, the opposite approach is taken—rather than supply different input cell values to look at the calculated formulas, this example lets Excel determine one of the input values.

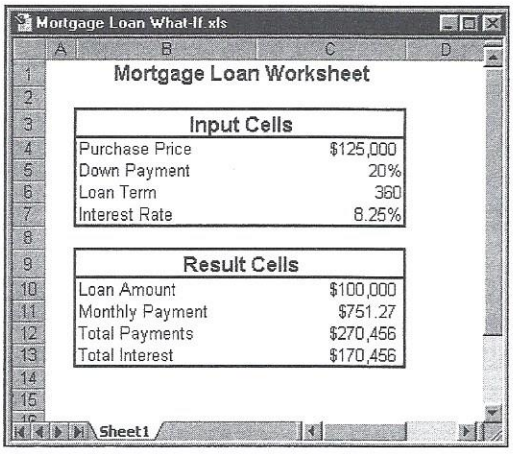

Figure 27-1: This worksheet is a good demonstration of goal seeking.

Assume that you're in the market for a new home and you know that you can afford $1,200 per month in mortgage payments. You also know that a lender can issue a fixed-rate mortgage loan for 8.25 percent, based on an 80 percent loan-to-value (that is, a 20-percent down payment). The question is, "What is the maximum purchase price I can handle?" In other words, what value in cell C4 causes the formula in cell C11 to result in $1,200? You could plug values into cell C4 until C11 displays $1,200; however, Excel can determine the answer much more efficiently.

To answer the question posed in the preceding paragraph, select Tools ⇨ Goal Seek. Excel displays the dialog box shown in Figure 27-2. Completing this dialog box is similar to forming a sentence. You want to set cell C11 to 1200 by changing cell C4. Enter this information in the dialog box either by typing the cell references or by pointing with the mouse. Click OK to begin the goal-seeking process.

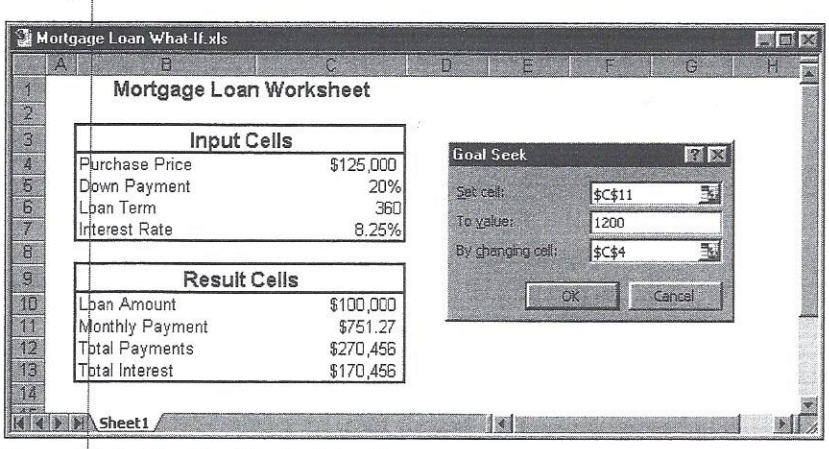

Figure 27-2: The Goal Seek dialog box.

In about a second, Excel announces that it has found the solution and displays the Goal Seek Status box, which shows the target value and the value that Excel calculated. In this case, Excel found an exact value. The worksheet now displays the found value in cell C4 ($199,663). As a result of this value, the monthly payment amount is $1,200. At this point, you have two options:

- ✦ Click OK to replace the original value with the found value.
- ✦ Click Cancel to restore your worksheet to the form that it had before you chose Tools ⇨ Goal Seek.

More About Goal Seeking

Excel can't always find a value that produces the result for which you're looking—sometimes, a solution simply doesn't exist. In such a case, the Goal Seek Status box informs you of that fact (see Figure 27-3).

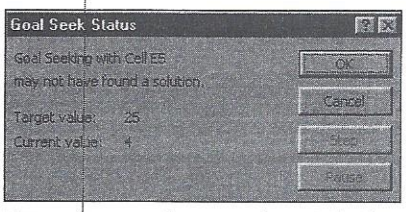

Figure 27-3: When Excel can't find a solution to your goal-seeking problem, it tells you so.

Other times, however, Excel may report that it can't find a solution, but you're pretty sure that one exists. If that's the case, you can try the following options:

- ✦ Change the current value of the By changing cell box in the Goal Seek dialog box to a value that is closer to the solution, and then reissue the command.
- ✦ Adjust the Maximum iterations setting in the Calculation tab of the Options dialog box. Increasing the number of iterations makes Excel try more possible solutions.
- ✦ Double-check your logic and make sure that the formula cell does, indeed, depend on the specified changing cell.

Like all computer programs, Excel has limited precision. To demonstrate this limitation, enter **=A1^2** into cell A2. Then, select Tools ⇨ Goal Seek to find the value in cell A1 (which is empty) that makes the formula return 16. Excel comes up with a value of 4.00002269 (you may need to widen the column to see the complete value), which is close to the square root of 16, but certainly not exact. You can adjust the precision in the Calculation tab of the Options dialog box (make the Maximum change value smaller).

In some cases, multiple values of the input cell produce the same desired result. For example, the formula =A1^2 returns 16 if cell A1 contains either –4 or +4. If you use goal seeking when two solutions are possible, Excel gives you the solution that has the same sign as the current value in the cell.

Perhaps the main limitation of the Tools ⇨ Goal Seek command is its inability to find the value for more than one input cell. For example, it can't tell you what purchase price *and* what down-payment percent will result in a particular monthly payment. If you want to change more than one variable at a time, use Solver (discussed later in this chapter).

Graphical Goal Seeking

Excel provides another way to perform goal seeking — by manipulating a graph. Figure 27-4 shows a worksheet that projects sales for a startup company. The CFO knows from experience that companies in this industry can grow exponentially according to a formula such as this one:

$$y*(b_x)$$

Table 27-1 lists and describes the variables.

Table 27-1
Variables Used in the Sales Growth Formula

Variable	*Description*
y	A constant equal to the first year's sales
b	A growth coefficient
x	A variable relating to time

The company managers know that sales during the first year are going to be $250,000, and they want to increase the company's sales to $10 million by the year 2005. The financial modelers want to know the exact growth coefficient that meets this goal. The worksheet that is shown in Figure 27-4 uses formulas to forecast the annual sales, based on the growth coefficient in cell B1. The worksheet has an embedded chart that plots the annual sales.

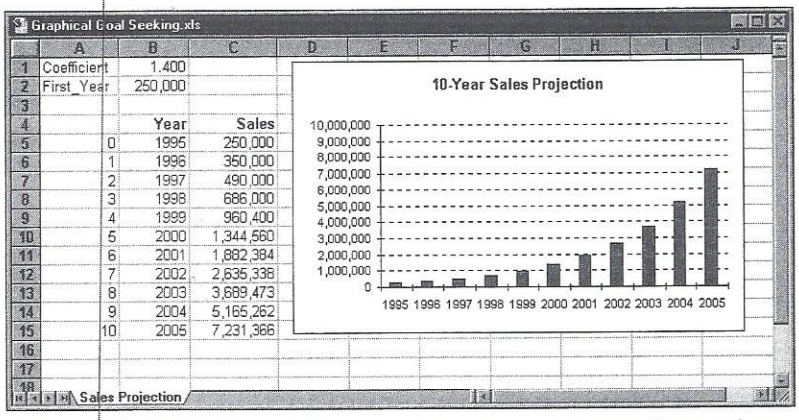

Figure 27-4: This sales projection predicts exponential growth, based on the growth coefficient in cell B1.

The initial guess for the growth coefficient is 1.40. As you can see, this number is too low—it results in sales of only $7.231 million for the year 2005. Although you can select Tools ⇨ Goal Seek to arrive at the exact coefficient, you have another way to do it.

Click the chart so that you can edit it and then select the chart series. Now, click the last data column to select only that column in the series. Point to the top of the column, and the mouse pointer changes shape. Drag the column upward and watch the value change in the small box displayed next to the mouse pointer. When the value is exactly $10 million, release the mouse button.

Excel responds with the Goal Seek dialog box, with two fields completed, as shown in Figure 27-5. Excel just needs to know which cell to use for the input cell. Specify cell B1 or enter **Coefficient** in the By changing cell edit box. Excel calculates the value of Coefficient that is necessary to produce the result that you pointed out on the chart. If you want to keep that number (which, by the way, is 1.44612554959182), click OK. Excel replaces the current value of Coefficient with the new value, and the chart is updated automatically. You can probably appreciate the fact that it would take quite a while to arrive at this number by plugging in successive approximations.

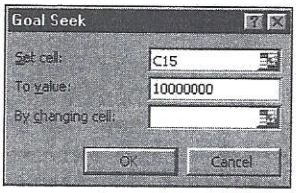

Figure 27-5: The Goal Seek dialog box appears when you directly manipulate a point on a chart that contains a formula.

You don't want to use this graphical method all the time, however, because the normal Tools ⇨ Goal Seek command is more efficient. But, it does demonstrate another way to approach problems that is helpful for those who are more visually oriented.

As you may expect, goal seeking can get much more impressive when it's used with complex worksheets that have many dependent cells. In any event, it sure beats trial and error.

Introducing Solver

Excel's goal-seeking feature is a useful tool, but it clearly has limitations. It can solve for only one adjustable cell, for example, and it returns only a single solution. Excel's powerful Solver tool extends this concept by enabling you to do the following:

- ✦ Specify multiple adjustable cells.
- ✦ Specify constraints on the values that the adjustable cells can have.
- ✦ Generate a solution that maximizes or minimizes a particular worksheet cell.
- ✦ Generate multiple solutions to a problem.

Although goal seeking is a relatively simple operation, using Solver can be much more complicated. In fact, Solver is probably one of the most difficult (and potentially frustrating) features in Excel. I'm the first to admit that Solver isn't for

everyone. In fact, most Excel users have no use for this feature. However, many users find that having this much power is worth spending the extra time to learn about it.

Appropriate Problems for Solver

Problems that are appropriate for Solver fall into a relatively narrow range. They typically involve situations that meet the following criteria:

- ✦ A target cell depends on other cells and formulas. Typically, you want to maximize or minimize this target cell or set it equal to some value.
- ✦ The target cell depends on a group of cells (called *changing cells*) that Solver can adjust to affect the target cell.
- ✦ The solution must adhere to certain limitations, or *constraints*.

After you set up your worksheet appropriately, you can use Solver to adjust the changing cells and produce the result that you want in your target cell — and, simultaneously meet all the constraints that you have defined.

You can find all the Solver examples in this chapter on this book's CD-ROM.

A Simple Solver Example

I start with a simple example to introduce Solver and then present some increasingly complex examples to demonstrate what it can do.

Figure 27-6 shows a worksheet that is set up to calculate the profit for three products. Column B shows the number of units of each product, column C shows the profit per unit for each product, and column C contains formulas that calculate the profit for each product by multiplying the units by the profit per unit.

Solver Production Model.xls

	A	B	C	D	E
1					
2		Units	Profit/Unit	Profit	
3	Product A	100	$13	$1,300	
4	Product B	100	$18	$1,800	
5	Product C	100	$22	$2,200	
6	Total	300		$5,300	
7					
8					
9					
10					
11					

Sheet1

Figure 27-6: Use Solver to determine the number of units to maximize the total profit.

It doesn't take an MBA degree to realize that the greatest profit per unit comes from Product C. Therefore, the logical solution is to produce only Product C. If things were really this simple, you wouldn't need tools such as Solver. As in most situations, this company has some constraints to which it must adhere:

- ✦ The combined production capacity is 300 total units per day.
- ✦ The company needs 50 units of Product A to fill an existing order.
- ✦ The company needs 40 units of Product B to fill an anticipated order.
- ✦ Because the market for Product C is relatively limited, the company doesn't want to produce more than 40 units of this product.

These four constraints make the problem more realistic and challenging. In fact, it's a perfect problem for Solver.

The basic procedure for using Solver is as follows:

1. Set up the worksheet with values and formulas. Make sure that you format cells logically; for example, if you cannot produce portions of your products, format those cells to contain numbers with no decimal values.
2. Bring up the Solver dialog box.
3. Specify the target cell.
4. Specify the changing cells.
5. Specify the constraints.
6. Change the Solver options, if necessary.
7. Let Solver solve the problem.

To start Solver, select Tools ⇨ Solver. Excel displays its Solver Parameters dialog box, shown in Figure 27-7.

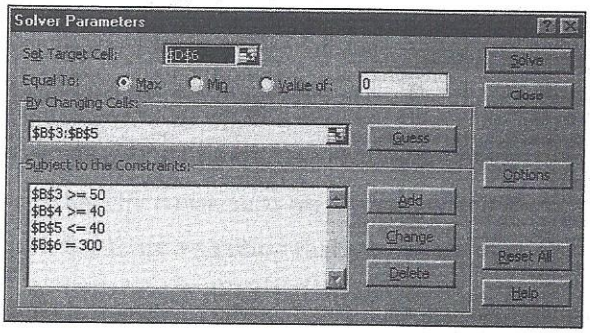

Figure 27-7: The Solver Parameters dialog box.

No Tools ⇨ Solver Command?

Solver is an add-in, so it's available only when the add-in is installed. If the Tools menu doesn't show a Solver command, you need to install the add-in before you can use it.

Select Tools ⇨ Add-Ins. Excel displays its Add-Ins dialog box. Scroll down the list of add-ins and place a check mark next to the item named Solver Add-In. Click OK, and Excel installs the add-in and makes the Tools ⇨ Solver command available. If Solver isn't available on your computer, you'll be asked if you want to install it.

In this example, the target cell is D6—the cell that calculates the total profit for three products. Enter (or point to) cell D6 in the Set Target Cell field of the Solver Parameters dialog box. Because the objective is to maximize this cell, click the Max option. Next, specify the changing cells, which are in the range B3:B5, in the By Changing Cells box.

The next step is to specify the constraints on the problem. The constraints are added one at a time and appear in the box labeled Subject to the Constraints. To add a constraint, click the Add button. Excel displays the Add Constraint dialog box, shown in Figure 27-8. This dialog box has three parts: a cell reference, an operator, and a value. To set the first constraint—that the total production capacity is 300 units—enter B6 as the Cell Reference, choose equal (=) from the drop-down list of operators, and enter 300 as the Constraint value. Click Add to add the remaining constraints. Table 27-2 summarizes the constraints for this problem.

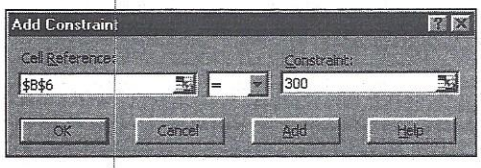

Figure 27-8: The Add Constraint dialog box.

Table 27-2
Constraints Summary

Constraint	*Expressed As*
Capacity is 300 units	B6=300
At least 50 units of Product A	B3>=50
At least 40 units of Product B	B4>=40
No more than 40 units of Product C	B5<=40

After you enter the last constraint, click OK to return to the Solver Parameters dialog box—which now lists the four constraints.

At this point, Solver knows everything about the problem. Click the Solver button to start the solution process. You can watch the progress onscreen, and Excel soon announces that it has found a solution. The Solver Results dialog box is shown in Figure 27-9.

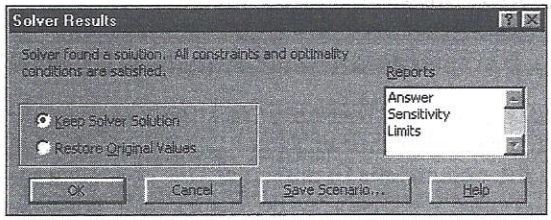

Figure 27-9: Solver displays this dialog box when it finds a solution to the problem.

At this point, you have the following options:

- ✦ Replace the original changing cell values with the values that Solver found
- ✦ Restore the original changing cell values
- ✦ Create any or all three reports that describe what Solver did (press Shift to select multiple reports from this list)
- ✦ Click the Save Scenario button to save the solution as a scenario, so that the Scenario Manager can use it (see Chapter 26)

If you specify any report options, Excel creates each report on a new worksheet, with an appropriate name. Figure 27-10 shows an Answer Report. In the Constraints section of the report, all the constraints except one are *binding*, which means that the constraint was satisfied at its limit, with no more room to change.

This simple example illustrates how Solver works. The fact is, you could probably solve this particular problem manually just as quickly. That, of course, isn't always the case.

Solver Production Model.xls

Microsoft Excel 9.0 Answer Report
Worksheet: [Solver Production Model.xls]Sheet1
Report Created: 9/4/1998 11:52:06 AM

Target Cell (Max)

Cell	Name	Original Value	Final Value
D6	Profit	$ 5,300	$ 5,310

Adjustable Cells

Cell	Name	Original Value	Final Value
B3	Product A Units	100	50
B4	Product B Units	100	210
B5	Product C Units	100	40

Constraints

Cell	Name	Cell Value	Formula	Status	Slack
B6	Units	300	B6=300	Binding	0
B5	Product C Units	40	B5<=40	Binding	0
B3	Product A Units	50	B3>=50	Binding	0
B4	Product B Units	210	B4>=40	Not Binding	170

Answer Report 1 | Sheet1

Figure 27-10: One of three reports that Solver can produce.

More About Solver

Before presenting complex examples, this section discusses the Solver Options dialog box—one of the more feature-packed dialog boxes in Excel. From this dialog box, you control many aspects of the solution process, as well as load and save model specifications in a worksheet range.

Having Solver report to you that it can't find a solution isn't unusual—even when you know that one should exist. Often, you can change one or more of the Solver options and try again. When you choose the Options button in the Solver Parameters dialog box, Excel displays the Solver Options dialog box shown in Figure 27-11.

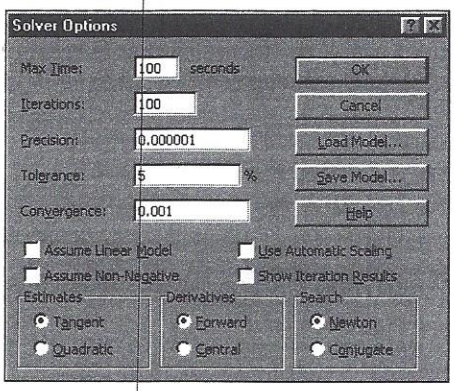

Figure 27-11: You can control many aspects of how Solver solves a problem.

This list describes Solver's options:

- **Max Time:** Specify the maximum amount of time (in seconds) that you want Solver to spend on a problem. If Solver reports that it exceeded the time limit, you can increase the amount of time that it spends searching for a solution.
- **Iterations:** Enter the maximum number of trial solutions that you want Solver to perform.
- **Precision:** Specify how close the Cell Reference and Constraint formulas must be to satisfy a constraint. Excel may solve the problem more quickly if you specify less precision.
- **Tolerance:** Designate the maximum percentage of error allowed for integer solutions (relevant only if an integer constraint is used).
- **Assume Linear Model:** Choose this option to speed the solution process, but you can use it only if all the relationships in the model are linear. You can't use this option if the adjustable cells are multiplied or divided, or if the problem uses exponents.
- **Use Automatic Scaling:** Use when the problem involves large differences in magnitude—when you attempt to maximize a percentage, for example, by varying cells that are very large.
- **Show Iteration Results:** Instruct Solver to pause and display the results after each iteration, by checking this box.
- **Estimates, Derivatives, and Search group boxes:** Use these options to control some technical aspects of the solution. In most cases, you don't need to change these settings.
- **Load Model:** Click this button to make Excel display the Load Model dialog box, in which you specify a range containing the model that you want to load.
- **Save Model:** Click this button to make Excel display the Save Model dialog box, in which you specify a range where Excel should save the model parameters.

Usually, you want to save a model only when you're using more than one set of Solver parameters with your worksheet, because Excel saves the first Solver model automatically with your worksheet (using hidden names). If you save additional models, Excel stores the information in the form of the formulas that correspond to the specification that you make (the last cell in the saved range is an array formula that holds the options settings).

Solver Examples

The remainder of this chapter consists of examples of using Solver for various types of problems.

Minimizing Shipping Costs

This example involves finding alternative options for shipping materials while keeping total shipping costs at a minimum (see Figure 27-12). A company has warehouses in Los Angeles, St. Louis, and Boston. Retail outlets throughout the United States place orders, which the company then ships from one of the warehouses. Ideally, the company wants to meet the product needs of all six retail outlets from available inventory in the warehouses — and keep total shipping charges as low as possible.

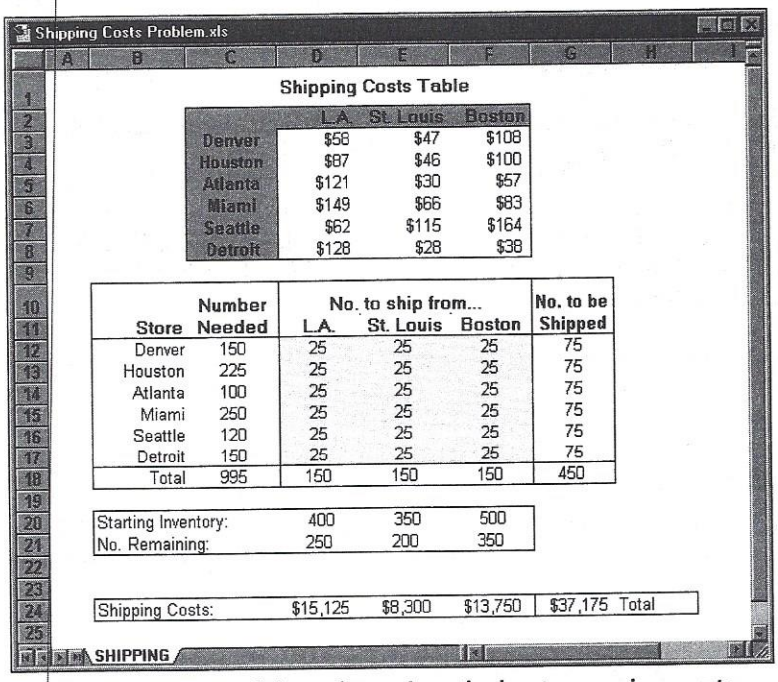

	L.A.	St. Louis	Boston
Denver	$58	$47	$108
Houston	$87	$46	$100
Atlanta	$121	$30	$57
Miami	$149	$66	$83
Seattle	$62	$115	$164
Detroit	$128	$28	$38

Store	Number Needed	No. to ship from... L.A.	St. Louis	Boston	No. to be Shipped
Denver	150	25	25	25	75
Houston	225	25	25	25	75
Atlanta	100	25	25	25	75
Miami	250	25	25	25	75
Seattle	120	25	25	25	75
Detroit	150	25	25	25	75
Total	995	150	150	150	450

Starting Inventory:	400	350	500
No. Remaining:	250	200	350

Shipping Costs:	$15,125	$8,300	$13,750	$37,175 Total

Figure 27-12: This worksheet determines the least expensive way to ship products from warehouses to retail outlets.

This workbook is rather complicated, so each part is explained individually:

- **Shipping Costs Table:** This table, at the top of the worksheet, contains per-unit shipping costs from each warehouse to each retail outlet. The cost to ship a unit from Los Angeles to Denver, for example, is $58.
- **Product needs of each retail store:** This information appears in C12:C17. For example, Denver needs 150 units, Houston needs 225, and so on. C18 holds the total needed.