

Borland® Pascal with Objects

Version 7.0

Language Guide

Copyright © 1983, 1992 by Borland International. All rights reserved. All Borland products are trademarks or registered trademarks of Borland International, Inc. Windows, as used in this manual, shall refer to Microsoft's implementation of a windows system. Other brand and product names are trademarks or registered trademarks of their respective holders.

C O N T E N T S

Introduction	1	Real types	28
What's in this manual	2	Software floating point	28
Part 1 The Borland Pascal language		80x87 floating point	29
Chapter 1 What is a Borland Pascal program?	5	String types	29
A Borland Pascal program	5	Structured types	30
Procedures and functions	6	Array types	30
Statements	8	Record types	32
Expressions	9	Object types	34
Tokens	10	Components and scope	36
Types, variables, constants, and typed constants	10	Methods	37
Putting it all together	11	Virtual methods	37
Units	13	Dynamic methods	38
Syntax diagrams	14	Instantiating objects	38
Chapter 2 Tokens	15	Method activations	40
Special symbols	15	Qualified-method activations	41
Reserved words and standard directives	16	Set types	42
Identifiers	17	File types	42
Numbers	18	Pointer types	43
Labels	19	Type Pointer	43
Character strings	19	Type PChar	44
Comments	20	Procedural types	44
Program lines	20	Procedural values	45
Chapter 3 Constants	21	Type compatibility	46
Chapter 4 Types	23	Identical and compatible types	47
Simple types	23	Type identity	47
Ordinal types	24	Type compatibility	48
Integer types	25	Assignment compatibility	48
Boolean types	25	The type declaration part	49
Char type	26	Chapter 5 Variables and typed constants	51
Enumerated types	26	Variable declarations	51
Subrange types	27	The data segment	52
		The stack segment	52
		Absolute variables	53
		Variable references	54

Qualifiers	54	Chapter 7 Statements	81
Arrays, strings, and indexes	55	Simple statements	81
Records and field designators	56	Assignment statements	82
Object component designators	56	Object-type assignments	82
Pointers and dynamic variables	56	Procedure statements	82
Variable typecasts	57	Goto statements	83
Typed constants	59	Structured statements	83
Simple-type constants	59	Compound statements	84
String-type constants	60	Conditional statements	84
Structured-type constants	60	If statements	84
Array-type constants	60	Case statements	85
Record-type constants	62	Repetitive statements	86
Object-type constants	63	Repeat statements	87
Set-type constants	63	While statements	87
Pointer-type constants	63	For statements	88
Procedural-type constants	64	With statements	90
Chapter 6 Expressions	65	Chapter 8 Blocks, locality, and scope	93
Expression syntax	66	Blocks	93
Operators	68	Rules of scope	95
Arithmetic operators	68	Block scope	95
Logical operators	69	Record scope	95
Boolean operators	70	Object scope	96
String operator	71	Unit scope	96
Character-pointer operators	71	Chapter 9 Procedures and functions	97
Set operators	72	Procedure declarations	97
Relational operators	73	Near and far declarations	98
Comparing simple types	73	Export declarations	99
Comparing strings	74	Interrupt declarations	100
Comparing packed strings	74	Forward declarations	100
Comparing pointers	74	External declarations	101
Comparing character pointers	74	Assembler declarations	102
Comparing sets	74	Inline declarations	102
Testing set membership	75	Function declarations	103
The @ operator	75	Method declarations	104
@ with a variable	75	Constructors and destructors	106
@ with a procedure, function, or method	75	Constructor error recovery	108
Function calls	76	Parameters	109
Set constructors	76	Value parameters	110
Value typecasts	77	Constant parameters	110
Procedural types in expressions	78	Variable parameters	111
		Untyped parameters	111

Open parameters	113	String procedures and functions ..	149
Open-string parameters	113	Dynamic-allocation procedures and	
Open-array parameters	114	functions	149
Dynamic object-type variables	116	Pointer and address functions	150
Chapter 10 Programs and units	119	Miscellaneous standard procedures	
Program syntax	119	and functions	150
The program heading	119	Predeclared variables	151
The uses clause	119	Chapter 14 Input and output	155
Unit syntax	120	File input and output	156
The unit heading	120	Text files	157
The interface part	121	Untyped files	159
The implementation part	121	The FileMode variable	159
The initialization part	122	Devices in Borland Pascal	160
Indirect unit references	122	DOS devices	160
Circular unit references	123	The CON device	160
Sharing other declarations	125	The LPT1, LPT2, and LPT3 devices .	161
Chapter 11 Dynamic-link libraries	127	The COM1 and COM2 devices	161
What is a DLL?	127	The NUL device	161
Using DLLs	128	Text-file devices	161
Import units	129	Input and output with the Crt unit	162
Static and dynamic imports	131	Using the Crt unit	162
Writing DLLs	132	CRT windows	163
The export procedure directive	133	Special characters	163
The exports clause	133	Line input	163
Library initialization code	135	Crt procedures and functions	164
Library programming notes	136	Crt unit constants and variables	165
Global variables in a DLL	137	Input and output with the WinCrt unit .	166
Global memory and files in a DLL ...	137	Using the WinCrt unit	166
DLLs and the System unit	137	Special characters	168
Run-time errors in DLLs	138	Line input	168
DLLs and stack segments	138	WinCrt procedures and functions ...	168
Writing shared DLLs	139	WinCrt unit variables	169
Part 2 The run-time libraries		Printing from a Windows program	170
Chapter 12 Overview of the		Changing titles	170
run-time libraries	143	Changing fonts	170
Chapter 13 Standard procedures		Stopping a print job	171
and functions	147	Special characters	171
Flow-control procedures	148	WinPrn procedures and functions ...	172
Transfer functions	148	Text-file device drivers	172
Arithmetic functions	148	The Open function	173
Ordinal procedures and functions .	149	The InOut function	174
		The Flush function	174
		The Close function	174

Chapter 15 Using the 80x87	175	Using segment registers as temporary variables	199
The 80x87 data types	177	Accessing memory beyond a segment's limit	199
Extended range arithmetic	178	Writing to a code segment	199
Comparing reals	179	Dereferencing nil pointers	199
The 80x87 evaluation stack	180	Code and data segments	199
Writing reals with the 80x87	181	Heap management	200
Units using the 80x87	181	Predefined selectors	200
Detecting the 80x87 in a DOS program	182	The SelectorInc variable	201
Detecting the 80x87 in a Windows program	183	The WinAPI unit	203
Emulation in assembly language	183	Memory management	204
Chapter 16 Interfacing with DOS	185	Module management	207
Dos unit procedures and functions	186	Resource management	208
Dos unit constants, types, and variables	188	Selector management	209
Constants	188	Other API routines	209
Types	189	Accessing the DPMServer directly	210
Variables	189	Compiling a protected-mode application	210
WinDos unit procedures and functions	189	Running a DOS protected-mode application	211
WinDos unit constants, types, and variables	191	Controlling the amount of memory RTM uses	212
Constants	191	Chapter 18 Using null-terminated strings	215
Types	192	What is a null-terminated string?	215
Variables	192	Strings unit functions	215
Chapter 17 Programming in DOS protected mode	193	Using null-terminated strings	217
What is protected mode?	193	Character pointers and string literals	217
Borland protected-mode DOS extensions	196	Character pointers and character arrays	219
The DPMServer	196	Character pointer indexing	219
The run-time manager	196	Null-terminated strings and standard procedures	221
Developing a DOS protected-mode application	197	An example using string-handling functions	221
Safe programming in protected mode	197	Chapter 19 Using the Borland Graphics Interface	223
Loading invalid values into segment registers	198	Drivers	223
The Ptr function and the Mem arrays	198	IBM 8514 support	224
Absolute variables	198	Coordinate system	225
Using segment arithmetic	198	Current pointer	226

Text	226	MOVEABLE or FIXED	266
Figures and styles	227	PRELOAD or DEMANDLOAD .	266
Viewports and bit images	227	DISCARDABLE or	
Paging and colors	228	PERMANENT	267
Error handling	228	The data and stack segments	267
Getting started	229	Changing attributes	268
Heap management routines	231	The DOS protected-mode heap	
Graph procedures and functions	233	manager	268
Graph unit constants, types, and		The HeapError variable	269
variables	236	Memory issues for Windows programs .	270
Constants	236	Code segments	270
Types	237	Segment attributes	270
Variables	238	MOVEABLE or FIXED	270
Chapter 20 Using overlays	239	PRELOAD or DEMANDLOAD .	270
The overlay manager	240	DISCARDABLE or	
Overlay-buffer management	241	PERMANENT	271
Overlay procedures and functions	244	Changing attributes	271
Variables and constants	244	The automatic data segment	271
Result codes	245	The heap manager	273
Designing overlaid programs	245	The HeapError variable	275
Overlay code generation	245	Internal data formats	276
The far call requirement	246	Integer types	276
Initializing the overlay manager	246	Char types	276
Initialization sections	249	Boolean types	276
What not to overlay	250	Enumerated types	276
Debugging overlays	251	Floating-point types	277
External routines in overlays	251	The Real type	277
Installing an overlay read function	252	The Single type	277
Overlays in .EXE files	254	The Double type	278
		The Extended type	278
		The Comp type	278
		Pointer types	279
		String types	279
		Set types	279
		Array types	280
		Record types	280
		Object types	280
		Virtual method tables	281
		Dynamic method tables	283
		File types	286
		Procedural types	288
		Direct memory access	288
		Direct port access	288
Part 3 Inside Borland Pascal			
Chapter 21 Memory issues	257		
Memory issues for DOS real-mode			
programs	257		
The DOS heap manager	259		
Disposal methods	260		
The free list	263		
The HeapError variable	265		
Memory issues for DOS protected-mode			
programs	266		
Code segments	266		
Segment attributes	266		

Chapter 22 Control issues	291
Calling conventions	291
Variable parameters	292
Value parameters	292
Open parameters	293
Function results	293
NEAR and FAR calls	294
Nested procedures and functions	294
Method calling conventions	295
Virtual method calls	296
Dynamic method calls	297
Constructors and destructors	298
Entry and exit code	298
Register-saving conventions	301
Exit procedures	301
Interrupt handling	303
Writing interrupt procedures	303
Chapter 23 Optimizing your code	305
Constant folding	305
Constant merging	306
Short-circuit evaluation	306
Constant parameters	306
Redundant pointer-load elimination	307
Constant set inlining	307
Small sets	308
Order of evaluation	308
Range checking	309
Shift instead of multiply or divide	309
Automatic word alignment	309
Eliminating dead code	310
Smart linking	310

Part 4 Using Borland Pascal with assembly language

Chapter 24 The built-in assembler	315
The asm statement	316
Register use	316
Assembler statement syntax	316
Labels	317
Instruction opcodes	317
RET instruction sizing	318
Automatic jump sizing	318
Assembler directives	319
Operands	321
Expressions	322
Differences between Pascal and	
Assembler expressions	322
Expression elements	323
Constants	323
Numeric constants	323
String constants	324
Registers	325
Symbols	325
Expression classes	329
Expression types	330
Expression operators	332
Assembler procedures and functions	334
Chapter 25 Linking assembler code	339
Turbo Assembler and Borland Pascal	340
Examples of assembly language	
routines	341
Assembly language methods	343
Inline machine code	344
Inline statements	344
Inline directives	345
Index	347

T A B L E S

2.1: Borland Pascal reserved words	16	14.7: Special characters in the WinCrt window	168
2.2: Borland Pascal directives	17	14.8: WinCrt procedures and functions . .	168
4.1: Predefined integer types	25	14.9: WinCrt variables	169
4.2: Real data types	28	14.10: Special characters in the WinPrn unit	171
6.1: Precedence of operators	65	14.11: WinPrn procedures and functions .	172
6.2: Binary arithmetic operations	68	15.1: Test8087 variable values	183
6.3: Unary arithmetic operations	69	16.1: Dos date and time procedures	186
6.4: Logical operations	70	16.2: Dos interrupt support procedures .	186
6.5: Boolean operations	70	16.3: Dos disk status functions	187
6.6: String operation	71	16.4: Dos file-handling procedures and functions	187
6.7: Permitted PChar constructs	72	16.5: Dos environment-handling functions	187
6.8: Set operations	72	16.6: Dos process-handling procedures . .	187
6.9: Relational operations	73	16.7: Dos miscellaneous procedures and functions	188
13.1: Flow-control procedures	148	16.8: Dos constants	188
13.2: Transfer functions	148	16.9: Dos types	189
13.3: Arithmetic functions	148	16.10: WinDos date and time procedures	189
13.4: Ordinal procedures and functions .	149	16.11: WinDos interrupt support procedures	190
13.5: String procedures and functions . .	149	16.12: WinDos disk status functions	190
13.6: Dynamic-allocation procedures and functions	150	16.13: File-handling procedures and functions	190
13.7: Pointer and address functions	150	16.14: WinDos directory-handling procedures and functions	191
13.8: Miscellaneous standard procedures and functions	151	16.15: WinDos environment-handling functions	191
13.9: Variables in the DOS real-mode System unit	151	16.16: WinDos miscellaneous procedures and functions	191
13.10: Variables in the Windows System unit	152	16.17: WinDos constants	192
13.11: Variables in the DOS protected-mode System unit	153	16.18: WinDos types	192
14.1: Input and output procedures and functions	155	17.1: Predefined selectors	200
14.2: Crt unit special characters	163		
14.3: Crt unit editing keys	164		
14.4: Crt unit procedures and functions .	164		
14.5: Crt unit constants	165		
14.6: Crt unit variables	165		

17.2: Memory management API routines	204	19.4: Graph unit constant groups	236
17.3: Module management API routines	207	19.5: Graph unit types	237
17.4: Resource management API functions	208	20.1: Overlay unit procedures and functions	244
17.5: Selector management API functions	209	20.2: Overlay unit variables	244
17.6: Other API routines	209	24.1: Built-in assembler reserved words	321
17.7: Environment variable options to control RTM's memory allocation	213	24.2: String examples and their values	325
18.1: Strings unit functions	216	24.3: CPU registers	325
19.1: Library and Graph unit file names	223	24.4: Values, classes, and types of symbols	327
19.2: BGI drivers	224	24.5: Predefined type symbols	332
19.3: Graph unit procedures and functions	233	24.6: Summary of built-in assembler expression operators	332
		24.7: Definitions of built-in assembler expression operators	333

F I G U R E S

1.1: Procedure or function diagram	6	21.2: Disposal method using Mark and Release	260
1.2: Simple Pascal program diagram	7	21.3: Heap layout with Release(P) executed	261
1.3: Statement diagram	10	21.4: Creating a "hole" in the heap	262
1.4: An expanded Pascal program diagram	12	21.5: Enlarging the free block	262
17.1: Real-mode physical address generation	194	21.6: Releasing the free block	263
17.2: Protected-mode physical address generation	195	21.7: Automatic data segment	272
19.1: Screen with xy-coordinates	225	21.8: Layouts of instances of TLocation, TPoint, and TCircle	281
20.1: Loading and disposing of overlays .	242	21.9: TPoint and TCircle's VMT layouts .	283
21.1: Memory map for a DOS real-mode program	258	21.10: TBase's VMT and DMT layouts . .	285
		21.11: TDerived's VMT and DMT layouts	286

To get an overview of the entire Borland Pascal with Objects documentation set, read the introduction in the User's Guide and find out how to use the Borland Pascal manuals most effectively.

This manual is about the Pascal language as it is used in Borland Pascal for Objects. It

- Presents the formal definition of the Borland Pascal language
- Explains how to use and write dynamic-link libraries
- Introduces the run-time library
- Explains how to write programs for the DOS protected-mode platform
- Describes what goes on inside Borland Pascal in regards to memory, data formats, calling conventions, input and output, and automatic optimizations
- Explains how to use Borland Pascal with assembly language

You'll find this manual most useful if you're an experienced Pascal programmer.

Read the *User's Guide* if

- You want to know how to install Borland Pascal
- You've used Turbo Pascal or Turbo Pascal for Windows before and you want to know what's new in this release
- You're not familiar with Borland's integrated development environments (the IDEs)
- You want an introduction to object-oriented programming
- You want an introduction to programming for Windows
- You're new to programming Windows in Pascal
- You want an introduction to ObjectWindows

Read the *Programmer's Reference* to look up reference material on

- The run-time libraries
- Compiler directives
- Error messages
- The editor

What's in this manual

This book is split into four parts: language grammar, libraries, advanced programming issues, and using assembly language with Borland Pascal.

Part I, "The Borland Pascal language," defines the Borland Pascal language. First you're introduced to the overall structure of a Borland Pascal program; then you examine each element of a program in detail.

Part II, "The run-time libraries," contains information about all the standard units that make up the run-time libraries and how to use them. It also tells you how to write DOS protected-mode programs.

Part III, "Inside Borland Pascal," presents technical information for advanced users about

- How Borland Pascal uses memory
- How Borland Pascal implements program control
- The details of input and output
- Optimizing your code

Part IV, "Using Borland Pascal with assembly language," explains how to use the built-in assembler and how to link your Borland Pascal programs with Turbo Assembler code.

P

A

R

T

1

The Borland Pascal language

What is a Borland Pascal program?

The next several chapters present the formal definition of the Borland Pascal language. Each chapter discusses an element of Borland Pascal. Together, these elements make up a Borland Pascal program.

It's difficult to gain an understanding of the whole by examining only the parts, however. This chapter presents an overview of a Borland Pascal program and omits the details. It gives you a brief description of each of the elements of a program and then shows you how they all fit together. You can then refer to Chapters 2 through 11 to find the details of the language.

A Borland Pascal program

In its simplest form, a Borland Pascal program is made up of a *program heading*, which names the program, and the *main program block*, which accomplishes the purpose of the program. Within the main program block is a section of code that occurs between two key words: **begin** and **end**. Here is a very simple program that illustrates these concepts:

```
program Welcome;  
begin  
    WriteLn('Welcome to Borland Pascal');  
end.
```

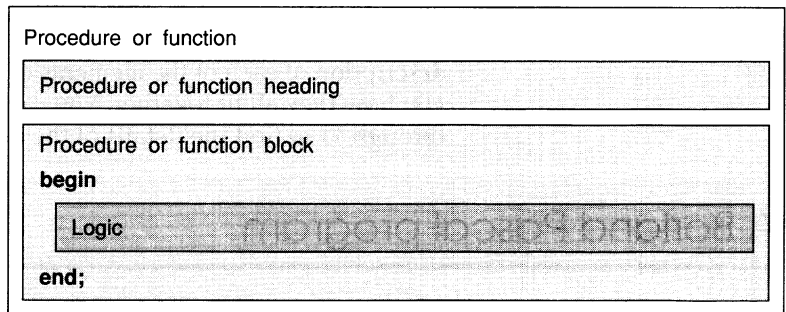
The first line is the program heading, which names the program. The remainder of the program is the code that starts with **begin** and stops with **end**. Although this particular code section contains only one line, it could contain many. In any Borland Pascal program, all the action occurs between **begin** and **end**.

Procedures and functions

The code between the last **begin** and **end** in a program drives the logic of the program. In a very simple program, this section of code might be all you need. In larger, more complex programs, putting all your code here can make your program harder to read and understand—and more difficult to develop.

Procedures and *functions* let you divide the logic of a program into smaller, more manageable chunks, and are similar to subroutines in some other languages. All the action in a procedure or function occurs in the code between its **begin** and **end** just like in the main program block. Each of these segments of code performs a small, discrete task.

Figure 1.1
Procedure or function
diagram



If you find your program does the same thing many times, you might want to put the logic into a procedure or function. You write the code in a procedure or function once and your program can use it as often as necessary.

Here is an example of a function. This *GetNumber* function gets a number from the user:

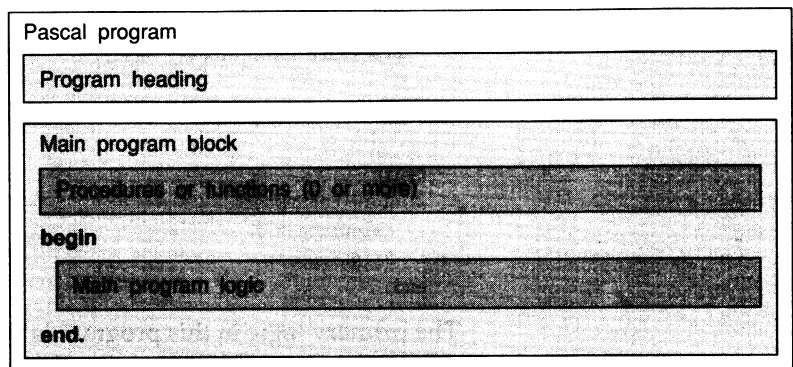
```

function GetNumber: Real;
var
    Response: Real;
begin
    Write('Enter a number: ');
    Readln(Response);
    GetNumber := Response;
end;

```

A procedure or function must appear before the main code section in the main program block. The main code section can then use the procedure or function.

Figure 1.2
Simple Pascal program
diagram



The following example is an outline of a program that uses the *GetNumber* function. The programmer has divided the logic of this program into three tasks:

1. Get a number from the user.
2. Perform the necessary calculations with the user-supplied number.
3. Print a report.

The main logic of the program is found between the last **begin** and **end**.

```

program Report;

var
    A: Real;

{more declarations}
:
function GetNumber: Real;
var
    Response: Real;
begin
    Write('Enter a number: ');
    Readln(Response);
    GetNumber := Response;
end;

procedure Calculate(X: Real);
:
procedure PrintReport;
:
begin
    A := GetNumber;
    Calculate(A);
    PrintReport;
end.

```

The primary logic in this program is very simple to understand. All the details are hidden within the bodies of the procedures and functions. Using procedures and functions encourages you to think about your program in a logical, modular way.

Statements

The code section between **begin** and **end** contains statements that describe the actions the program can take and is called the *statement part*. These are examples of statements:

```

A := B + C;           {Assign a value}

Calculate(Length, Height); {Activate a procedure}

if X < 2 then       {Conditional statement}
    Answer := X * Y;

begin               {Compound statement}
    X := 3;
    Y := 4;
    Z := 5;
end;

```



```

while not EOF(InFile) do           {Repetitive statement}
begin
  Readln(InFile, Line);
  Process(Line);
end;

```

Simple statements can either assign a value, activate a procedure or function, or transfer the running of the program to another statement in the code. The first two examples shown in the examples are simple statements.

Structured statements can be compound statements that contain multiple statements, conditional and repetitive statements that control the flow of logic within a program, and **with** statements that simplify access to data in a record.

You might compare a Pascal statement to a sentence in a human language such as English, Danish, or Greek. Simple Pascal statements and simple human sentences hold one complete thought. Structured Pascal statements and complex sentences contain more complicated logic.

Expressions

Just as a sentence is made up of phrases, so is a Pascal statement made up of expressions. The phrases of a sentence are made up of words, and the expressions of a statement are composed of elements called factors and operators. Expressions usually compare things or perform arithmetic, logical, or Boolean operations.

Just as phrases in a human language can be made up of smaller phrases, so can expressions in Pascal be made up of simpler expressions. You can read about all the combinations of factors and operators that make up expressions in Chapter 6. They can be quite complex. For now, it might help to see some examples of expressions:

```

X + Y
Done <> Error
I <= Length
-X

```


Tokens

Tokens are the smallest meaningful elements in a Pascal program. They make up the factors and operators of expressions. Tokens are special symbols, reserved words, identifiers, labels, numbers, and string constants; they are akin to the words and punctuation of a written human language. These are examples of Pascal tokens:

<code>function</code>	{reserved word}
<code>(</code>	{special symbol}
<code>:=</code>	{special symbol}
<code>Calculate</code>	{identifier for a procedure}
<code>9</code>	{number}

Here is an illustration of a statement. You can see that statements are made up of expressions, which are made up of tokens.

Figure 1.3
Statement diagram



Types, variables, constants, and typed constants

A *variable* can hold a value that can change. Every variable must have a *type*. A variable's type specifies the set of values the variable can have.

For example, this next program declares that variables *X* and *Y* are of type *Integer*; therefore, the only values *X* and *Y* can contain are integers, which are whole numbers. Borland Pascal displays an error message if your program tries to assign any other type of value to these variables.

```

program Example;

const
  A = 12;           {Constant A never changes in value}
  B: Integer = 23; {Typed constant B gets an initial value}

var
  X, Y: Integer;   {Variables X and Y are type Integer}
  J: Real;         {Variable J is type Real}

begin
  X := 7;          {Variable X is assigned a value}
  Y := 8;          {Variable Y is assigned a value}
  X := Y + Y;      {The value of variable X changes}
  B := 57;         {Typed constant B gets a new value}
  J := 0.075;     {Variable J gets a floating-point value}
end.

```

In this simple and not very useful program, *X* is assigned the value 7 originally; two statements later it is assigned a new value, $Y + Y$. As you can see, the value of a variable can vary.

A is a *constant*. The program gives it a value of 12 and this value can't change—its value remains constant throughout the program.

B is a *typed constant*. It's given a value when it's declared, but it's also given a type of *Integer*. You can think of a typed constant as a variable with an initial value. The program can later change the initial value of *B* to some other value.

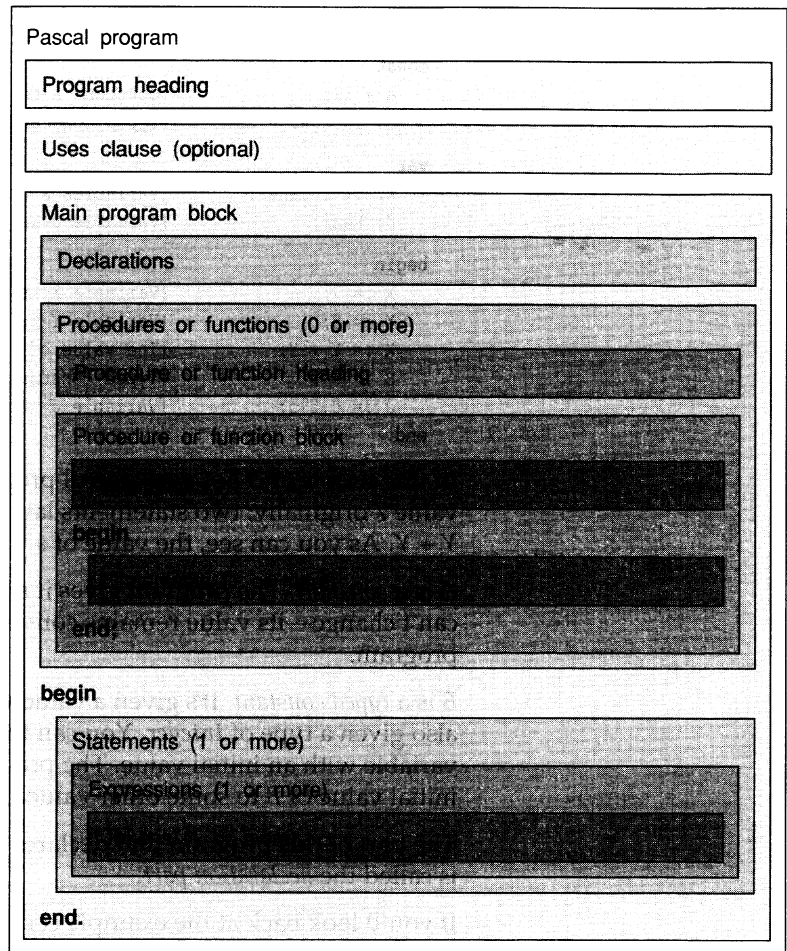
The part of this program that declares the constants and variables is called the *declaration part*.

If you'll look back at the example code on page 7, you'll see that the function *GetNumber* has a declaration part that declares a variable. Procedures and functions can contain a declaration part just as a program or unit can.

Putting it all together

Now that you've been introduced to the primary components of a Borland Pascal program, you need to see how they all fit together. Here's a diagram of a Borland Pascal program:

Figure 1.4
An expanded Pascal
program diagram



The program heading, the optional **uses** clause (we'll talk about this in the next section), and the main program block make up a Pascal program. Within the main program block can exist the smaller blocks of procedures and functions. Although the diagram doesn't show this, procedures and functions can be nested within other procedures and functions. In other words, blocks can contain other blocks.

Combined with other tokens and blank spaces, tokens make up expressions which make up statements.

In turn, statements combined with declaration parts make up blocks, either the main program block or a block in a procedure or function.

Units

A Borland Pascal program can use blocks of code in separate modules called *units*. You can think of a unit as a mini-program your application can use. Like a program, it has a heading, called a unit heading, and a main block that contains a code section delineated by **begin** and **end**.

Any Borland Pascal main program block can include a line that enables the program to use one or more units. For example, if you are writing a DOS program called *Colors* and you want to change the color of the text as it appears on your screen, you can specify that your program use the standard *Crt* unit that is part of the Borland Pascal run-time library:

```
program Colors;  
uses Crt;  
begin  
  :  
end.
```


The **uses** *Crt* line tells Borland Pascal to include the *Crt* unit in the executable program. The *Crt* unit contains all the necessary code to change the color of the text in your program, among other things. Simply by including **uses** *Crt*, your program can use the code that is in the *Crt* unit. If you put all the code required to create the functionality of the *Crt* unit within your program, it would be a lot more work, and it would sidetrack you from the main purpose of your program.

Borland Pascal's run-time libraries include several units you'll find useful. For example, use the *Dos* or *WinDos* units and your program has access to several operating system and file-handling routines.

You can also write your own units. Use them to divide large programs into logically-related modules. Code you place in a unit can be used by any program. You only have to write the code once, then you can use it many times.

Syntax diagrams

As you read Chapters 2 through 11, which define the Borland Pascal language, you'll encounter *syntax diagrams*. For example,



To read a syntax diagram, follow the arrows. Frequently, more than one path is possible. The above diagram indicates that a formal parameter list is optional in a procedure heading. You can follow the path from the identifier to the end of the procedure heading, or you can follow it to the formal parameter list before reaching the end.

The names in boxes stand for constructions. Those in circles—reserved words, operators, and punctuation—are the actual terms used in the program; they are boldfaced in the diagrams.

Tokens

Tokens are the smallest meaningful units of text in a Pascal program. They are categorized as special symbols, identifiers, labels, numbers, and string constants.

Separators can't be part of tokens except in string constants.


A Pascal program is made up of tokens and *separators*. A separator is either a blank or a comment. Two adjacent tokens must be separated by one or more separators if each token is a reserved word, an identifier, a label, or a number.


Special symbols

Borland Pascal uses the following subsets of the ASCII character set:

- **Letters**—the English alphabet, *A* through *Z* and *a* through *z*
- **Digits**—the Arabic numerals 0 through 9
- **Hex digits**—the Arabic numerals 0 through 9, the letters *A* through *F*, and the letters *a* through *f*
- **Blanks**—the space character (ASCII 32) and all ASCII control characters (ASCII 0 through 31), including the end-of-line or return character (ASCII 13)

These are the syntax diagrams for letter, digit, and hex digit:





Special symbols and reserved words are characters that have one or more fixed meanings. The following single characters are special symbols:

+ - * / = < > [] . , () : ; ^ @ { } \$ #

These character pairs are also special symbols:

<= >= := .. (* *) (. .)

A left bracket ([) is equivalent to the character pair of left parenthesis and a period—(., and a right bracket (]) is equivalent to the character pair of a period and a right parenthesis—.). Likewise, a left brace { is equivalent to the character pair of left parenthesis and an asterisk—(*, and a right brace } is equivalent to the character pair of an asterisk and a right parenthesis—*).

Reserved words and standard directives

Reserved words can't be redefined.

Reserved words appear in lowercase **boldface** throughout this manual. Borland Pascal is *not* case sensitive, however, so you can use either uppercase or lowercase letters in your programs.

Following are Borland Pascal's reserved words:

Table 2.1
Borland Pascal reserved words

and	exports	mod	shr
array	file	nil	string
asm	for	not	then
begin	function	object	to
case	goto	of	type
const	if	or	unit
constructor	implementation	packed	until
destructor	in	procedure	uses
div	inherited	program	var
do	inline	record	while
downto	interface	repeat	with
else	label	set	xor
end	library	shl	

Table 2.2
Borland Pascal directives

absolute	far	name	resident
assembler	forward	near	virtual
export	index	private	
external	interrupt	public	

private and **public** act as reserved words within object type declarations, but are otherwise treated as directives.

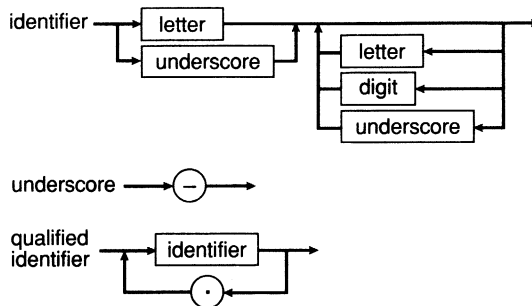
Identifiers

Identifiers denote constants, types, variables, procedures, functions, units, programs, and fields in records.

An identifier can be of any length, but only the first 63 characters are significant. An identifier must begin with a letter or an underscore character (`_`) and can't contain spaces. Letters, digits, and underscore characters (ASCII \$5F) are allowed after the first character. Like reserved words, identifiers are *not* case sensitive.

Units are described in Chapter 7 of the User's Guide and Chapter 10 of this manual.

When several instances of the same identifier exist, you may need to qualify the identifier by another identifier to select a specific instance. For example, to qualify the identifier *Ident* by the unit identifier *UnitName*, write *UnitName.Ident*. The combined identifier is called a *qualified identifier*.



Here are some examples of identifiers and qualified identifiers:

```

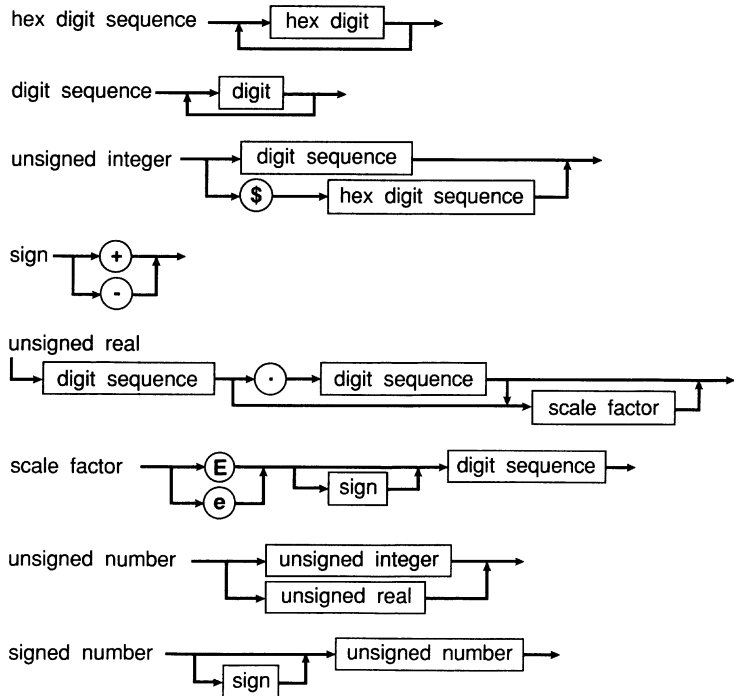
Writeln
Exit
Real2String
System.MemAvail
Strings.StrLen
WinCrt.ReadText

```

In this manual, standard and user-defined identifiers are *italicized* when they are referred to in text.

Numbers

Ordinary decimal notation is used for numbers that are constants of integer and real types. A hexadecimal integer constant uses a dollar sign (\$) as a prefix. Engineering notation (E or e, followed by an exponent) is read as “times ten to the power of” in real types. For example, 7E-2 means 7×10^{-2} ; 12.25e+6 or 12.25e6 both mean $12.25 \times 10^{+6}$. Syntax diagrams for writing numbers follow:




Numbers with decimals or exponents denote real-type constants. Other decimal numbers denote integer-type constants; they must be within the range $-2,147,483,648$ to $2,147,483,647$.

Hexadecimal numbers denote integer-type constants; they must be within the range $\$00000000$ to $\$FFFFFFF$. The resulting value's sign is implied by the hexadecimal notation.

Labels

A label is a digit sequence in the range 0 to 9999. Leading zeros are not significant. Labels are used with **goto** statements.



As an extension to Standard Pascal, Borland Pascal also allows identifiers to function as labels.

Character strings

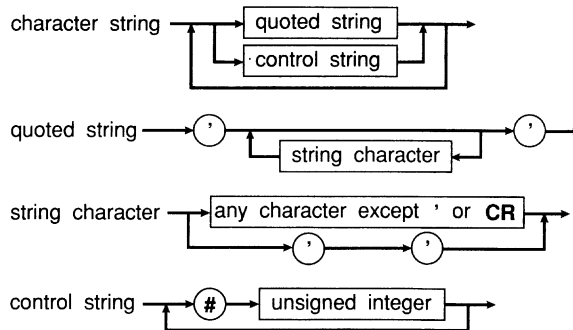
A character string is a sequence of zero or more characters from the extended ASCII character set, written on one line in the program and enclosed by apostrophes.

A character string with nothing between the apostrophes is a *null string*. Two sequential apostrophes in a character string denote a single character, an apostrophe. For example,

```
'BORLAND'      { BORLAND }
'You'll see'   { You'll see }
''             { ' }
''            { null string }
' '           { a space }
```

As an extension to Standard Pascal, Borland Pascal lets you embed control characters in character strings. The # character followed by an unsigned integer constant in the range 0 to 255 denotes a character of the corresponding ASCII value. There must be no separators between the # character and the integer constant. Likewise, if several are part of a character string, there must be no separators between them. For example,

```
#13#10
'Line 1'#13'Line2'
#7#7'Wake up!'#7#7
```



A character string's *length* is the actual number of characters in the string. A character string of any length is compatible with any string type, and with the *PChar* type when the extended syntax is enabled **{SX+}**. Also, a character string of length one is compatible with any *Char* type, and a character string of length *N*, where *N* is greater than or equal to one, is compatible with packed arrays of *N* characters.

Comments

The following constructs are comments and are ignored by the compiler:

```
{ Any text not containing right brace }
(* Any text not containing star/right parenthesis *)
```

The compiler directives are explained in Chapter 2 of the Programmer's Reference


A comment that contains a dollar sign (\$) immediately after the opening { or (* is a *compiler directive*. A mnemonic of the compiler command follows the \$ character.

Program lines

Borland Pascal program lines have a maximum length of 126 characters.

Constants

A constant is an identifier that marks a value that can't change. A *constant declaration* declares a constant within the block containing the declaration. A constant identifier can't be included in its own declaration.



Wherever Standard Pascal allows only a simple constant, Borland Pascal allows a constant expression.


As an extension to Standard Pascal, Borland Pascal allows the use of *constant expressions*. A constant expression is an expression that can be evaluated by the compiler without actually executing the program. Examples of constant expressions follow:

```

100
'A'
256 - 1
(2.5 + 1) / (2.5 - 1)
'Borland' + ' ' + 'Pascal'
Chr(32)
Ord('Z') - Ord('A') + 1

```

The simplest case of a constant expression is a simple constant, such as 100 or 'A'.



Because the compiler has to be able to completely evaluate a constant expression at compile time, the following constructs are *not* allowed in constant expressions:

- References to variables and typed constants (except in constant address expressions as described on page 59)
- Function calls (except those noted in the following text)
- The address operator (@) (except in constant address expressions as described on page 59)

For expression syntax, see Chapter 6, "Expressions."

Except for these restrictions, constant expressions follow the syntactical rules as ordinary expressions.

The following standard functions are allowed in constant expressions:

<i>Abs</i>	<i>Length</i>	<i>Ord</i>	<i>SizeOf</i>
<i>Chr</i>	<i>Lo</i>	<i>Pred</i>	<i>Succ</i>
<i>Hi</i>	<i>Low</i>	<i>Ptr</i>	<i>Swap</i>
<i>High</i>	<i>Odd</i>	<i>Round</i>	<i>Trunc</i>

Here are some examples of the use of constant expressions in constant declarations:


```

const
  Min = 0;
  Max = 100;
  Center = (Max - Min) div 2;
  Beta = Chr(225);
  NumChars = Ord('Z') - Ord('A') + 1;
  Message = 'Out of memory';
  ErrStr = ' Error: ' + Message + '. ';
  ErrPos = 80 - Length(ErrStr) div 2;
  Ln10 = 2.302585092994045684;
  Ln10R = 1 / Ln10;
  Numeric = ['0'..'9'];
  Alpha = ['A'..'Z', 'a'..'z'];
  AlphaNum = Alpha + Numeric;


```

Types

When you declare a variable, you must state its *type*. A variable's type circumscribes the set of values it can have and the operations that can be performed on it. A *type declaration* specifies the identifier that denotes a type.




When an identifier occurs on the left side of a type declaration, it's declared as a *type identifier* for the block in which the type declaration occurs. A type identifier's scope doesn't include itself except for pointer types.



There are five major type classes. They are described in the following sections.

Simple types

Simple types define ordered sets of values.



Chapter 2 explains how to denote constant integer type and real type values.


A real type identifier is one of the standard identifiers: *Real*, *Single*, *Double*, *Extended*, or *Comp*.

Ordinal types

Ordinal types are a subset of simple types. All simple types other than real types are ordinal types, which are set off by six characteristics:

- All possible values of a given ordinal type are an ordered set, and each possible value is associated with an *ordinality*, which is an integral value. Except for integer type values, the first value of every ordinal type has ordinality 0, the next has ordinality 1, and so on for each value in that ordinal type. The ordinality of an integer type value is the value itself. In any ordinal type, each value other than the first has a predecessor, and each value other than the last has a successor based on the ordering of the type.
- The standard function *Ord* can be applied to any ordinal type value to return the ordinality of the value.
- The standard function *Pred* can be applied to any ordinal type value to return the predecessor of the value. If applied to the first value in the ordinal type and if range checking is enabled **{\$R+}**, *Pred* produces a run-time error.
- The standard function *Succ* can be applied to any ordinal type value to return the successor of the value. If applied to the last value in the ordinal type and if range checking is enabled **{\$R+}**, *Succ* produces a run-time error.
- The standard function *Low* can be applied to an ordinal type identifier and to a variable reference of an ordinal type. The result is the lowest value in the range of the given ordinal type.
- The standard function *High* can be applied to an ordinal type identifier and to a variable reference of an ordinal type. The result is the highest value in the range of the given ordinal type.

The syntax of an ordinal type follows:



Borland Pascal has ten predefined ordinal types: *Integer*, *Shortint*, *Longint*, *Byte*, *Word*, *Boolean*, *ByteBool*, *WordBool*, *LongBool*, and *Char*. In addition, there are two other classes of user-defined ordinal types: enumerated types and subrange types.

Integer types

There are five predefined integer types: *Shortint*, *Integer*, *Longint*, *Byte*, and *Word*. Each type denotes a specific subset of the whole numbers, according to the following table:

Table 4.1
Predefined integer types

Type	Range	Format
<i>Shortint</i>	-128 .. 127	Signed 8-bit
<i>Integer</i>	-32768 .. 32767	Signed 16-bit
<i>Longint</i>	-2147483648 .. 2147483647	Signed 32-bit
<i>Byte</i>	0 .. 255	Unsigned 8-bit
<i>Word</i>	0 .. 65535	Unsigned 16-bit

Arithmetic operations with integer-type operands use 8-bit, 16-bit, or 32-bit precision, according to the following rules:

- The type of an integer constant is the predefined integer type with the smallest range that includes the value of the integer constant.
- For a binary operator (an operator that takes two operands), both operands are converted to their common type before the operation. The common type is the predefined integer type with the smallest range that includes all possible values of both types. For example, the common type of *Integer* and *Byte* is *Integer*, and the common type of *Integer* and *Word* is *Longint*. The operation is performed using the precision of the common type, and the result type is the common type.
- The expression on the right of an assignment statement is evaluated independently from the size or type of the variable on the left.
- Any byte-sized operand is converted to an intermediate word-sized operand that is compatible with both *Integer* and *Word* before any arithmetic operation is performed.

Typecasting is described in Chapters 5 and 6.

An integer-type value can be explicitly converted to another integer type through typecasting.

Boolean types There are four predefined Boolean types: *Boolean*, *ByteBool*, *WordBool*, and *LongBool*. Boolean values are denoted by the predefined constant identifiers *False* and *True*. Because Booleans are enumerated types, these relationships hold:

- $False < True$
- $Ord(False) = 0$
- $Ord(True) = 1$
- $Succ(False) = True$
- $Pred(True) = False$


Boolean and *ByteBool* variables occupy one byte, a *WordBool* variable occupies two bytes (one word), and a *LongBool* variable occupies four bytes (two words). *Boolean* is the preferred type and uses the least memory; *ByteBool*, *WordBool*, and *LongBool* primarily exist to provide compatibility with other languages and the Windows environment.

A *Boolean* variable can assume the ordinal values 0 and 1 only, but variables of type *ByteBool*, *WordBool*, and *LongBool* can assume other ordinal values. An expression of type *ByteBool*, *WordBool*, or *LongBool* is considered *False* when its ordinal value is zero, and *True* when its ordinal value is nonzero. Whenever a *ByteBool*, *WordBool*, or *LongBool* value is used in a context where a *Boolean* value is expected, the compiler will automatically generate code that converts any nonzero value to the value *True*.

Char type *Char*'s set of values are characters, ordered according to the extended ASCII character set. The function call $Ord(Ch)$, where *Ch* is a *Char* value, returns *Ch*'s ordinality.

A string constant of length 1 can denote a constant character value. Any character value can be generated with the standard function *Chr*.

Enumerated types Enumerated types define ordered sets of values by enumerating the identifiers that denote these values. Their ordering follows the sequence the identifiers are enumerated in.



When an identifier occurs within the identifier list of an enumerated type, it's declared as a constant for the block the enumerated type is declared in. This constant's type is the enumerated type being declared.

An enumerated constant's ordinality is determined by its position in the identifier list it's declared in. The enumerated type it's declared in becomes the constant's type. The first enumerated constant in a list has an ordinality of zero.

Here's an example of an enumerated type:


```
type  
Suit = (Club, Diamond, Heart, Spade);
```

Given these declarations, *Diamond* is a constant of type *Suit*.

When the *Ord* function is applied to an enumerated type's value, *Ord* returns an integer that shows where the value falls with respect to the other values of the enumerated type. Given the preceding declarations, *Ord(Club)* returns zero, *Ord(Diamond)* returns 1, and so on.

Subrange types

A subrange type is a range of values from an ordinal type called the *host type*. The definition of a subrange type specifies the smallest and the largest value in the subrange; its syntax follows:



Both constants must be of the same ordinal type. Subrange types of the form *A..B* require that *A* is less than or equal to *B*.

These are examples of subrange types:

```
0..99  
-128..127  
Club..Heart
```

A variable of a subrange type has all the properties of variables of the host type, but its run-time value must be in the specified interval.

One syntactic ambiguity arises from allowing constant expressions where Standard Pascal only allows simple constants. Consider the following declarations:

```

const
  X = 50;
  Y = 10;
type
  Color = (Red, Green, Blue);
  Scale = (X - Y) * 2..(X + Y) * 2;

```

Standard Pascal syntax dictates that, if a type definition starts with a parenthesis, it's an enumerated type, such as the *Color* type in the previous example. The intent of the declaration of *scale* is to define a subrange type, however. The solution is to reorganize the first subrange expression so that it doesn't start with a parenthesis, or to set another constant equal to the value of the expression and use that constant in the type definition:

```

type
  Scale = 2 * (X - Y)..(X + Y) * 2;

```

Real types

A real type has a set of values that is a subset of real numbers, which can be represented in floating-point notation with a fixed number of digits. A value's floating-point notation normally comprises three values—*M*, *B*, and *E*—such that $M \times B^E = N$, where *B* is always 2, and both *M* and *E* are integral values within the real type's range. These *M* and *E* values further prescribe the real type's range and precision.

There are five kinds of real types: *Real*, *Single*, *Double*, *Extended*, and *Comp*. The real types differ in the range and precision of values they hold as shown in the following table:

Table 4.2
Real data types

The Comp type holds only integral values within $-2^{63}+1$ to $2^{63}-1$, which is approximately -9.2×10^{18} to 9.2×10^{18} .

Type	Range	Significant digits	Size in bytes
<i>Real</i>	$2.9 \times 10^{-39} .. 1.7 \times 10^{38}$	11-12	6
<i>Single</i>	$1.5 \times 10^{-45} .. 3.4 \times 10^{38}$	7-8	4
<i>Double</i>	$5.0 \times 10^{-324} .. 1.7 \times 10^{308}$	15-16	8
<i>Extended</i>	$3.4 \times 10^{-4932} .. 1.1 \times 10^{4932}$	19-20	10
<i>Comp</i>	$-2^{63}+1 .. 2^{63}-1$	19-20	8

Borland Pascal supports two models of code generation for performing real-type operations: *software* floating point and *80x87* floating point. Use the **\$N** compiler directive to select the appropriate model. If no *80x87* is present and your target platform is DOS real or protected mode, enable the **\$E** compiler directive to provide full *80x87* emulation in software.

Software floating point In the **{ $\$N-$ }** state, which is selected by default, the generated code performs all real-type calculations in software by calling run-time library routines. For reasons of speed and code size, only operations on variables of type *Real* are allowed in this state. Any attempt to compile statements that operate on the *Single*, *Double*, *Extended*, and *Comp* types generates an error.

80x87 floating point In the **{ $\$N+$ }** state, the generated code performs all real-type calculations using 80x87 instructions and can use all five real types.

For more details on 80x87 floating-point code generation and software emulation, refer to Chapter 15, "Using the 80x87."


Borland Pascal includes run-time libraries that will automatically emulate an 80x87 in software if one isn't present for a DOS real-mode or protected-mode application. The **$\$E$** compiler directive is used to determine whether or not the 80x87 emulator should be included in a DOS program. If you're creating a DOS real-mode or protected-mode application and no 80x87 is present, enable the **$\$E$** compiler directive to provide full 80x87 emulation in the software. The **$\$E$** directive will have no effect in a Windows program, as Windows provides its own emulation routines.

String types

Operators for the string types are described in the sections "String operator" and "Relational operators" in Chapter 6.

String-type standard procedures and functions are described in "String procedures and functions" on page 149.

A string-type value is a sequence of characters with a dynamic length attribute (depending on the actual character count during program execution), and a constant size attribute from 1 to 255. A string type declared without a size attribute is given the default size attribute 255. The length attribute's current value is returned by the standard function *Length*.



The ordering between any two string values is set by the ordering relationship of the character values in corresponding positions. In two strings of unequal length, each character in the longer string without a corresponding character in the shorter string takes on a higher or greater-than value; for example, 'xs' is greater than 'x'. Null strings can be equal only to other null strings, and they hold the least string values.

Characters in a string can be accessed as components of an array. See the section “Arrays, strings, and indexes” on page 55.

The *Low* and *High* standard functions can be applied to a string-type identifier and to a variable reference of a string type. In this case, *Low* returns zero, and *High* returns the size attribute (maximum length) of the given string.


Read about open string parameters on page 113.

A variable parameter declared using the *OpenString* identifier, or using the **string** keyword in the **{SP+}** state, is an *open string parameter*. Open string parameters allow string variables of varying sizes to be passed to the same procedure or function.

Structured types

The maximum permitted size of any structured type in Borland Pascal is 65,520 bytes.


A structured type, characterized by its structuring method and by its component type(s), holds more than one value. If a component type is structured, the resulting structured type has more than one level of structuring. A structured type can have unlimited levels of structuring.




In Standard Pascal, the word **packed** in a structured type’s declaration tells the compiler to compress data storage, even at the cost of diminished access to a component of a variable of this type. In Borland Pascal, however, **packed** has no effect; instead packing occurs automatically whenever possible.

Array types

Arrays have a fixed number of components of one type—the *component type*. In the following syntax diagram, the component type follows the word **of**.





The index types, one for each dimension of the array, specify the number of elements. Valid index types are all ordinal types except *Longint* and subranges of *Longint*. The array can be indexed in each dimension by all values of the corresponding index type; therefore, the number of elements is the product of the number of values in each index type.

The following is an example of an array type:

```
array[1..100] of Real
```

If an array type's component type is also an array, you can treat the result as an array of arrays or as a single multidimensional array. For example,

```
array[Boolean] of array[1..10] of array[Size] of Real
```

is interpreted the same way by the compiler as

```
array[Boolean,1..10,Size] of Real
```

You can also express

```
packed array[1..10] of packed array[1..8] of Boolean
```

as

```
packed array[1..10,1..8] of Boolean
```

See "Arrays, strings, and indexes" on page 55.

You access an array's components by supplying the array's identifier with one or more indexes in brackets.

When applied to an array-type identifier or a variable reference of an array type, the *Low* and *High* standard functions return the low and high bounds of the index type of the array.

An array type of the form

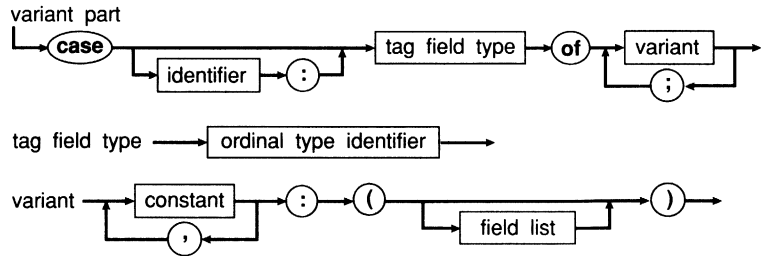
```
packed array[M..N] of Char
```

See "Identical and compatible types" on page 47.

where *M* is less than *N* is called a *packed string type* (the word **packed** can be omitted because it has no effect in Borland Pascal). A packed string type has certain properties not shared by other array types, as explained below.

An array type of the form

```
array[0..X] of Char
```

You can see from the diagram that each variant is identified by at least one constant. All constants must be distinct and of an ordinal type compatible with the tag field type. Variant and fixed fields are accessed the same way.

An optional identifier, the *tag field identifier*, can be placed in the variant part. If a tag field identifier is present, it becomes the identifier of an additional fixed field—the tag field—of the record. The program can use the tag field’s value to show which variant is active at a given time. Without a tag field, the program selects a variant by another criterion.

Some record types with variants follow:

```

type
  TPerson = record
    FirstName, LastName: string[40];
    BirthDate: TDate;
    case Citizen: Boolean of
      True: (BirthPlace: string[40]);
      False: (Country: string[20];
        EntryPort: string[20];
        EntryDate: TDate;
        ExitDate: TDate);
    end;

  TPolygon = record
    X, Y: Real;
    case Kind: Figure of
      TRectangle: (Height, Width: Real);
      TTriangle: (Side1, Side2, Angle: Real);
      TCircle: (Radius: Real);
    end;

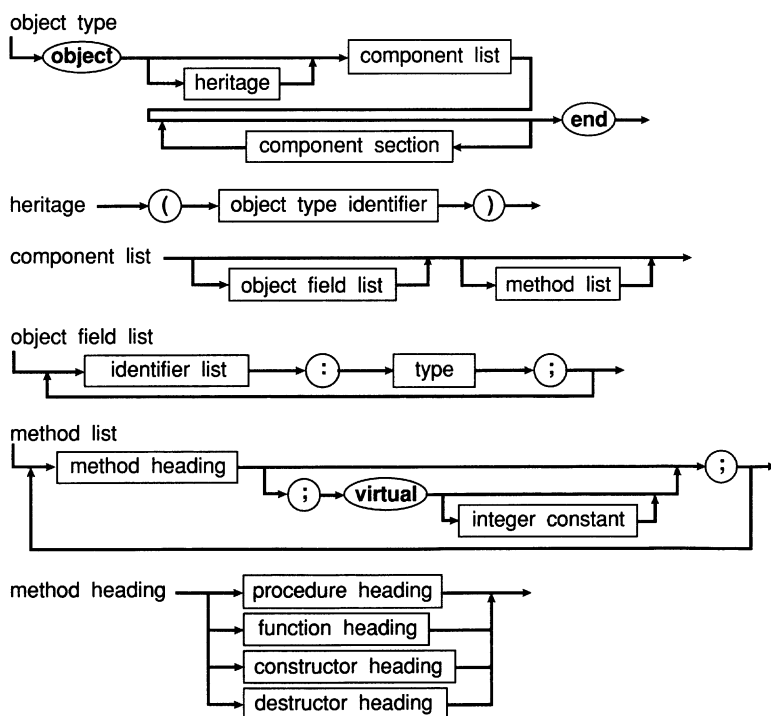
```

Object types

An object type is a structure consisting of a fixed number of components. Each component is either a *field*, which contains data of a particular type, or a *method*, which performs an operation on the object. Similar to a variable declaration, the declaration of a field specifies the field's data type and an identifier that names the field. Similar to a procedure or function declaration, the declaration of a method specifies a procedure, function, constructor, or destructor heading.

An object type can *inherit* components from another object type. If $T2$ inherits from $T1$, then $T2$ is a *descendant* of $T1$, and $T1$ is an *ancestor* of $T2$.

Inheritance is transitive; that is, if $T3$ inherits from $T2$, and $T2$ inherits from $T1$, then $T3$ also inherits from $T1$. The *domain* of an object type consists of itself and all its descendants.





The following code shows examples of object-type declarations:

These declarations are referred to by other examples throughout this chapter.

```

type
  TPoint = object
    X, Y: Integer;
  end;

  TRectangle = object
    A, B: TPoint;
    procedure Init(XA, YA, XB, YB: Integer);
    procedure Copy(var R: TRectangle);
    procedure Move(DX, DY: Integer);
    procedure Grow(DX, DY: Integer);
    procedure Intersect(var R: TRectangle);
    procedure Union(var R: TRectangle);
    function Contains(D: TPoint): Boolean;
  end;

  PString = ^String;
  PField = ^TField;

  TField = object
  private
    X, Y, Len: Integer;
    Name: String;
  public
    constructor Copy(var F: TField);
    constructor Init(FX, FY, FLen: Integer; FName: String);
    destructor Done; virtual;
    procedure Display; virtual;
    procedure Edit; virtual;
    function GetStr: String; virtual;
    function PutStr(S: String): Boolean; virtual;
  private
    procedure DisplayStr(X, Y: Integer; S: String);
  end;

  PStrField = ^TStrField;
  TStrField = object(TField)
  private
    Value: PString;
  public
    constructor Init(FX, FY, FLen: Integer; FName: String);
    destructor Done; virtual;
    function GetStr: String; virtual;

```

```

    function PutStr(S: String): Boolean; virtual;
    function Get: String;
    procedure Put(S: String);
end;

PNumField = ^TNumField;

TNumField = object(TField)
private
    Value, Min, Max: Longint;
public
    constructor Init(FX, FY, FLen: Integer; FName: String;
        FMin, FMax: Longint);
    function GetStr: String; virtual;
    function PutStr(S: String): Boolean; virtual;
    function Get: Longint;
    procedure Put(N: Longint);
end;

PZipField = ^TZipField;

TZipField = object(TNumField)
public
    function GetStr: String; virtual;
    function PutStr(S: String): Boolean; virtual;
end;

```

Contrary to other types, an object type can be declared only in a type declaration part in the outermost scope of a program or unit. Therefore, an object type can't be declared in a variable declaration part or within a procedure, function, or method block.

The component type of a file type can't be an object type, or any structured type with an object-type component.

Components and scope

The scope of a component identifier extends over the domain of its object type. Also, the scope of a component identifier extends over procedure, function, constructor, and destructor blocks that implement methods of the object type and its descendants. For this reason, the spelling of a component identifier must be unique within an object type and all its descendants and all its methods.

Component identifiers declared in the component list that immediately follows the object-type heading and component identifiers declared in **public** component sections have no special restrictions on their scope. In contrast, the scope of component identifiers declared in **private** component sections is restricted to the module (program or unit) that contains the object-type declaration. In other words, **private** component identifiers act like

normal public component identifiers within the module that contains the object-type declaration, but outside the module, any **private** component identifiers are unknown and inaccessible. By placing related object types in the same module, these object types can gain access to each other's **private** components without making the **private** components known to other modules.

Within an object-type declaration, a method heading can specify parameters of the object type being declared, even though the declaration isn't yet complete. In the previous example starting on page 35, the *Copy*, *Intersect*, and *Union* methods of the *TRectangle* type illustrate this.


Methods

Methods can be called only through an object instance variable.

Read more about methods on page 104.

The declaration of a method within an object type corresponds to a **forward** declaration of that method. This means that somewhere after the object-type declaration, and within the same scope as the object-type declaration, the method must be *implemented* by a defining declaration.

When unique identification of a method is required, a *qualified method identifier* is used. It consists of an object-type identifier, followed by a period (.), followed by a method identifier. Like any other identifier, a qualified method identifier can be prefixed with a unit identifier and a period, if required.



Virtual methods

By default, methods are *static*. With the exception of constructor methods, they can be made *virtual* by including a **virtual** directive in the method declaration. The compiler resolves calls to *static* methods at compile time. Calls to virtual methods are resolved at run time; this is known as *late binding*.

If an object type declares or inherits any virtual methods, then variables of that type must be *initialized* through a constructor call before any call to a virtual method. Therefore, any object type that declares or inherits any virtual methods must also declare or inherit at least one constructor method.

An object type can *override* (redefine) any of the methods it inherits from its ancestors. If a method declaration in a descendant specifies the same method identifier as a method declaration in an ancestor, then the declaration in the descendant overrides the declaration in the ancestor. The scope of an override method

extends over the domain of the descendant in which it's introduced, or until the method identifier is again overridden.

An override of a static method is free to change the method heading any way it pleases. In contrast, an override of a virtual method must match exactly the order, types, and names of the parameters, and the type of the function result, if any. The override must again include a **virtual** directive.

Dynamic methods Borland Pascal supports an additional class of late-bound methods called *dynamic methods*. Dynamic methods differ from virtual methods only in the way dynamic method calls are dispatched at run time. For all other purposes, a dynamic method can be considered equivalent to a virtual method.

The declaration of a dynamic method is like that of a virtual method except that a dynamic method declaration must include a *dynamic method index* right after the **virtual** keyword. The dynamic method index must be an integer constant in the range 1..65535 and it must be unique among the dynamic method indexes of any other dynamic methods contained in the object type or its ancestors. For example,

```
procedure FileOpen(var Msg: TMessage); virtual 100;
```

To learn more about dynamic methods and the differences in dispatching virtual and dynamic method calls, see page 297.

An override of a dynamic method must match the order, types, and names of the parameters and the type of the function result of the ancestral method exactly. The override must also include a **virtual** directive followed by the same dynamic method index as was specified in the ancestor object type.

Instantiating objects An object is *instantiated*, or created, through the declaration of a variable or typed constant of an object type, or by applying the *New* procedure to a pointer variable of an object type. The resulting object is called an *instance* of the object type. For example, given these variable declarations,

```
var
  F: TField;
  Z: TZipField;
  FP: PField;
  ZP: PZipField;
```

F is an instance of *TField* and *Z* is an instance of *TZipField*. Likewise, after applying *New* to *FP* and *ZP*, *FP* points to an instance of *TField* and *ZP* points to an instance of *TZipField*.

If an object type contains virtual methods, then instances of that object type must be initialized through a constructor call before any call to a virtual method. Here's an example:

```
var
  S: TStrField;
begin
  S.Init(1, 1, 25, 'Firstname');
  S.Put('Frank');
  S.Display;
  :
  S.Done;
end;
```

If *S.Init* had not been called, then the call to *S.Display* causes this example to fail.



Assignment to an instance of an object type doesn't initialize the instance.

An object is initialized by compiler-generated code that executes between the time that the constructor call takes place and when execution actually reaches the first statement of the constructor's code block.

If an object instance isn't initialized and range checking is on **{*\$R+*}**, the first call to a virtual method of the object instance results in a run-time error. If range checking is off **{*\$R-*}**, calling a virtual method of an uninitialized object instance results in undefined behavior.

The rule of required initialization also applies to instances that are components of structured types. For example,

```
var
  Comment: array[1..5] of TStrField;
  I: Integer;
begin
  for I := 1 to 5 do Comment[I].Init(1, I + 10, 40, 'Comment');
  :
  for I := 1 to 5 do Comment[I].Done;
end;
```

For dynamic instances, initialization is typically coupled with allocation, and cleanup is typically coupled with deallocation, using the extended syntax of the *New* and *Dispose* procedures. Here's an example:

```

var
  SP: PStrField;
begin
  New(SP, Init(1, 1, 25, 'Firstname'));
  SP^.Put('Frank');
  SP^.Display;
  :
  Dispose(SP, Done);
end;

```

A pointer to an object type is assignment-compatible with a pointer to any ancestor object type; therefore, during execution of a program, a pointer to an object type might point to an instance of that type or to an instance of any descendant type.


For example, a pointer of type *PZipField* can be assigned to pointers of type *PZipField*, *PNumField*, and *PField*, and during execution of a program, a pointer of type *PField* might be either **nil** or point to an instance of *TField*, *TStrField*, *TNumField*, or *TZipField*, or any other instance of a descendant of *TField*.

Pointer assignment-compatibility rules also apply to object-type variable parameters. For example, the *TField.Copy* method might be passed an instance of *TField*, *TStrField*, *TNumField*, *TZipField*, or any other instance of a descendant of *TField*.

Method activations

See "Function calls" on page 76 and "Procedure statements" on page 82.

A method is activated through a function call or procedure statement consisting of a *method designator* followed by an actual parameter list. This type of call is known as a *method activation*.



The variable reference specified in a method designator must denote an instance of an object type, and the method identifier must denote a method of that object type.

The instance denoted by a method designator becomes an implicit actual parameter of the method; it corresponds to a formal variable parameter named *Self* that possesses the object type corresponding to the activated method.

For static methods, the *declared* (compile-time) type of the instance determines which method to activate. For example, the designators *F.Init* and *FP^.Init* always activates *TField.Init* because the declared type of *F* and *FP^* is *TField*.

For virtual methods, the *actual* (run-time) type of the instance governs the selection. For example, the designator `FP^.Display` might activate `TField.Display`, `TStrField.Display`, `TNumField.Display`, or `TZipField.Display`, depending on the actual type of the instance pointed to by `FP`.


See "With statements" on page 90 and "Method declarations" on page 104.

Within a **with** statement that references an instance of an object type, the variable-reference part of a method designator can be omitted. In that case, the implicit *Self* parameter of the method activation becomes the instance referenced by the **with** statement. Likewise, within a method, the variable-reference part of a method designator can be omitted. In that case, the implicit *Self* parameter of the method activation becomes the *Self* of the method containing the call.

Qualified-method activations

See "Function calls" on page 76 and "Procedure statements" on page 82.

Within a method, a function call or procedure statement allows a *qualified-method designator* to denote activation of a specific method. This type of call is known as a *qualified-method activation*.



The object type specified in a qualified-method designator must be the same as the enclosing method's object type or an ancestor of it.



The reserved word **inherited** can be used to denote the ancestor of the enclosing method's object type; **inherited** can't be used within methods of an object type that has no ancestor.

The implicit *Self* parameter of a qualified-method activation becomes the *Self* of the method containing the call. A qualified-method activation never employs the virtual method dispatch mechanism—the call is always static and always invokes the specified method.

A qualified-method activation is generally used within an override method to activate the overridden method. Referring to the types declared earlier starting on page 35, here are some examples of qualified-method activations:

```

constructor TNumField.Init(FX, FY, FLen: Integer;
    FName: String; FMin, FMax: Longint);
begin
    inherited Init(FX, FY, FLen, FName);
    Value := 0;

```

```

    Min := FMin;
    Max := FMax;
end;

function TZipField.PutStr(S: String): Boolean;
begin
    PutStr := (Length(S) = 5) and TNumField.PutStr(S);
end;


```

As these examples demonstrate, a qualified-method activation allows an override method to “reuse” the code of the method it overrides.

Set types

A set type’s range of values is the power set of a particular ordinal type (the base type). The power set is the set of all possible subsets of values of the base type including the empty set. Therefore, each possible value of a set type is a subset of the possible values of the base type.

A variable of a set type can hold from none to all the values of the set.




Set-type operators are described in the section “Set operators” in Chapter 6. “Set constructors” in the same chapter shows how to construct set values.

The base type must not have more than 256 possible values, and the ordinal values of the upper and lower bounds of the base type must be within the range 0 to 255.

Every set type can hold the value [], which is called the *empty set*.

File types

A file type consists of a linear sequence of components of the component type, which can be of any type except a file type, any structured type with a file-type component, or an object type. The number of components isn’t set by the file-type declaration.




If the word **of** and the component type are omitted, the type denotes an untyped file. Untyped files are low-level I/O (input/output) channels primarily used for direct access to any disk file regardless of its internal format.

The standard file type *Text* signifies a file containing characters organized into lines. Text files use special I/O procedures, which are discussed in Chapter 14, “Input and output.”

Pointer types

A pointer type defines a set of values that point to dynamic variables of a specified type called the *base type*. A pointer-type variable contains the memory address of a dynamic variable.



If the base type is an undeclared identifier, it must be declared in the same type declaration part as the pointer type.

You can assign a value to a pointer variable with the *New* procedure, the @ operator, the *Ptr* function, or the *GetMem* procedure. *New* allocates a new memory area in the application heap for a dynamic variable and stores the address of that area in the pointer variable. The @ operator directs the pointer variable to the memory area containing any existing variable or procedure or function entry point, including variables that already have identifiers. *Ptr* points the pointer variable to a specific memory address. *GetMem* creates a new dynamic variable of a specified size, and puts the address of the block in the pointer variable.

The reserved word denotes a pointer-valued constant that doesn't point to anything.

Type Pointer

See Chapter 5's section entitled "Pointers and dynamic variables" on page 56 for the syntax of referencing the dynamic variable pointed to by a pointer variable.

The predefined type *Pointer* denotes an untyped pointer; that is, a pointer that doesn't point to any specific type. Variables of type *Pointer* can't be dereferenced; writing the pointer symbol ^ after such a variable is an error. Generic pointers, however, can be typecast to allow dereferencing. Like the value denoted by the word **nil**, values of type *Pointer* are compatible with all other pointer types.

Type PChar

Borland Pascal has a predefined type, *PChar*, to represent a pointer to a null-terminated string. The *System* unit declares *PChar* as


```
type PChar = ^Char;
```

Borland Pascal supports a set of *extended syntax* rules to facilitate handling of null-terminated strings using the *PChar* type. For a complete discussion of this topic, see Chapter 18, "Using null-terminated strings."

Procedural types

Standard Pascal regards procedures and functions as program parts that can be executed through procedure or function calls. Borland Pascal has a much broader view of procedures and functions: It allows procedures and functions to be treated as entities that can be assigned to variables and passed as parameters. Such actions are made possible through *procedural types*.

A procedural-type declaration specifies the parameters and, for a function, the result type.



In essence, the syntax for writing a procedural-type declaration is exactly the same as for writing a procedure or function header, except that the identifier after the **procedure** or **function** keyword is omitted. Some examples of procedural-type declarations follow:

```
type
  Proc = procedure;
  SwapProc = procedure(var X, Y: Integer);
  StrProc = procedure(S: string);
  MathFunc = function(X: Real): Real;
  DeviceFunc = function(var F: Text): Integer;
  MaxFunc = function(A, B: Real; F: MathFunc): Real;
```

The parameter names in a procedural-type declaration are purely decorative—they have no effect on the declaration's meaning.



Borland Pascal doesn't let you declare functions that return procedural-type values; a function result must be a string, real, integer, char, boolean, pointer, or user-defined enumeration-type value. But you can return the address of a procedure or function using a function result of type *Pointer* and then typecast it to the procedural type you desire.

Procedural values

A variable of a procedural type can be assigned a *procedural value*. Procedural values can be one of these:

- The value **nil**
- A variable reference of a procedural type
- A procedure or function identifier

In the context of procedural values, a procedure or function declaration can be viewed as a special kind of constant declaration, the value of the constant being the procedure or function. For example, given the following declarations:

```
var
  P: SwapProc;
  F: MathFunc;

procedure Swap(var A, B: Integer); far;
var
  Temp: Integer;
begin
  Temp := A;
  A := B;
  B := Temp;
end;

function Tan(Angle: Real); far;
begin
  Tan := Sin(Angle) / Cos(Angle);
end;
```

the variables *P* and *F* can be assigned values as follows:

```
P := Swap;
F := Tan;
```

and calls can be made using *P* and *F* as follows:

```
P(I, J);      { Equivalent to Swap(I, J) }  
X := F(X);   { Equivalent to X := Tan(X) }
```

Using a procedural variable that has been assigned the value **nil** in a procedure statement or a function call results in an error. **nil** is intended to indicate that a procedural variable is unassigned, and whenever there is a possibility that a procedural variable is **nil**, procedure statements or function calls involving that procedural variable should be guarded by a test:

```
if @P <> nil then P(I, J);
```

See "Procedural types in expressions" on page 78.

Notice the use of the **@** operator to indicate that *P* is being examined rather than being called.

Type compatibility

To be considered compatible, procedural types must have the same number of parameters, and parameters in corresponding positions must be of identical types. Finally, the result types of functions must be identical. Parameter names have no significance when determining procedural-type compatibility.

The value **nil** is compatible with any procedural type.

To be used as procedural values, procedures and functions must be declared with a **far** directive or compiled in the **(\$F+)** state. Also, standard procedures and functions, nested procedures and functions, methods, **inline** procedures and functions, and **interrupt** procedures can't be used as procedural values.

Standard procedures and functions are the ones declared by the *System* unit, such as *WriteLn*, *ReadLn*, *Chr*, and *Ord*. To use a standard procedure or function as a procedural value, write a "shell" around it. For example, the following function *FSin* is assignment-compatible with the *MathFunc* type declared above.

```
function FSin(X: Real): Real; far;  
begin  
  FSin := Sin(X);  
end;
```

A procedure or function is *nested* when it's declared within another procedure or function. Such nested procedures and functions can't be used as procedural values.

Identical and compatible types

Two types can be the same, and this sameness (identity) is mandatory in some contexts. At other times, the two types need only be compatible or merely assignment-compatible. They are identical when they are declared with, or their definitions stem from, the same type identifier.

Type identity

Type identity is required only between actual and formal variable parameters in procedure and function calls.

Two types—say, *T1* and *T2*—are identical if one of the following is true: *T1* and *T2* are the same type identifier; *T1* is declared to be equivalent to a type identical to *T2*.

The second condition connotes that *T1* doesn't have to be declared directly to be equivalent to *T2*. The type declarations

```
T1 = Integer;  
T2 = T1;  
T3 = Integer;  
T4 = T2;
```

result in *T1*, *T2*, *T3*, *T4*, and *Integer* as identical types. The type declarations

```
T5 = set of Integer;  
T6 = set of Integer;
```

don't make *T5* and *T6* identical because **set of** *Integer* isn't a type identifier. Two variables declared in the same declaration, for example,

```
V1, V2: set of Integer;
```

are of identical types—unless the declarations are separate. The declarations

```
V1: set of Integer;  
V2: set of Integer;  
V3: Integer;  
V4: Integer;
```

mean *V3* and *V4* are of identical type, but not *V1* and *V2*.

Type compatibility

Compatibility between two types is sometimes required, such as in expressions or in relational operations. Type compatibility is important, however, as a precondition of assignment compatibility.

Type compatibility exists when at least one of the following conditions is true:

- Both types are the same.
- Both types are real types.
- Both types are integer types.
- One type is a subrange of the other.
- Both types are subranges of the same host type.
- Both types are set types with compatible base types.
- Both types are packed string types with an identical number of components.
- One type is a string type and the other is either a string type, packed string type, or *Char* type.
- One type is *Pointer* and the other is any pointer type.
- One type is *PChar* and the other is a zero-based character array of the form **array**[0..X] **of** *Char*. (This applies only when extended syntax is enabled with the **{SX+}** directive.)
- Both types are pointers to identical types. (This applies only when type-checked pointers are enabled with the **{ST+}** directive.)
- Both types are procedural types with identical result types, an identical number of parameters, and a one-to-one identity between parameter types.

Assignment compatibility

Assignment compatibility is necessary when a value is assigned to something, such as in an assignment statement or in passing value parameters.

A value of type T_2 is assignment-compatible with a type T_1 (that is, $T_1 := T_2$ is allowed) if any of the following are *True*:

- T_1 and T_2 are identical types and neither is a file type or a structured type that contains a file-type component at any level of structuring.
- T_1 and T_2 are compatible ordinal types, and the values of type T_2 falls within the range of possible values of T_1 .
- T_1 and T_2 are real types, and the value of type T_2 falls within the range of possible values of T_1 .
- T_1 is a real type, and T_2 is an integer type.
- T_1 and T_2 are string types.
- T_1 is a string type, and T_2 is a *Char* type.
- T_1 is a string type, and T_2 is a packed string type.
- T_1 and T_2 are compatible, packed string types.
- T_1 and T_2 are compatible set types, and all the members of the value of type T_2 fall within the range of possible values of T_1 .
- T_1 and T_2 are compatible pointer types.
- T_1 is a *PChar* and T_2 is a string constant. (This applies only when extended syntax is enabled **{ $\$X+$ }**.)
- T_1 is a *PChar* and T_2 is a zero-based character array of the form **array[0..X] of Char**. (This applies only when extended syntax is enabled **{ $\$X+$ }**.)
- T_1 and T_2 are compatible procedural types.
- T_1 is a procedural type, and T_2 is a procedure or function with an identical result type, an identical number of parameters, and a one-to-one identity between parameter types.
- T_2 is assignment-compatible with an object type T_1 if T_2 is an object type in the domain of T_1 .
- A pointer type P_2 , pointing to an object type T_2 , is assignment-compatible with a pointer type P_1 , pointing to an object type T_1 , if T_2 is in the domain of T_1 .

A compile-time error occurs when assignment compatibility is necessary and none of the items in the preceding list are true.

The type declaration part

Programs, procedures, functions, and methods that declare types have a *type declaration part*. This is an example of a type declaration part:

```

type
  TRange = Integer;
  TNumber = Integer;
  TColor = (Red, Green, Blue);
  TCharVal = Ord('A')..Ord('Z');
  TTestIndex = 1..100;
  TTestValue = -99..99;
  TTestList = array[TTestIndex] of TTestValue;
  PTestList = ^TTestList;

  TDate = object
    Year: Integer;
    Month: 1..12;
    Day: 1..31;
    procedure SetDate(D, M, Y: Integer);
    function ShowDate: String;
  end;
  TMeasureData = record
    When: TDate;
    Count: TTestIndex;
    Data: PTestList;
  end;
  TMeasureList = array[1..50] of TMeasureData;
  TName = string[80];
  TSex = (Male, Female);
  PPersonData = ^TPersonData;
  TPersonData = record
    Name, FirstName: TName;
    Age: Integer;
    Married: Boolean;
    TFather, TChild, TSibling: PPersonData;
    case S: TSex of
      Male: (Bearded: Boolean);
      Female: (Pregnant: Boolean);
    end;
  TPersonBuf = array[0..SizeOf(TPersonData)-1] of Byte;
  TPeople = file of TPersonData;


```

In the example, *Range*, *Number*, and *Integer* are identical types. *TTestIndex* is compatible and assignment-compatible with, but not identical to, the types *Number*, *Range*, and *Integer*. Notice the use of constant expressions in the declarations of *TCharVal* and *TPersonBuf*.

Variables and typed constants

Variable declarations

A variable is an identifier that marks a value that can change. A *variable declaration* embodies a list of identifiers that designate new variables and their type.



The type given for the variable(s) can be a type identifier previously declared in a **type** declaration part in the same block, in an enclosing block, or in a unit; it can also be a new type definition.

When an identifier is specified within the identifier list of a variable declaration, that identifier is a variable identifier for the block in which the declaration occurs. The variable can then be referred to throughout the block, unless the identifier is redeclared in an enclosed block. Redeclaration creates a new variable using the same identifier, without affecting the value of the original variable.

An example of a variable declaration part follows:

```

var
  X, Y, Z: Real;
  I, J, K: Integer;
  Digit: 0..9;
  
```

```
C: Color;
Done, Error: Boolean;
Operator: (Plus, Minus, Times);
Hue1, Hue2: set of Color;
Today: Date;
Results: MeasureList;
P1, P2: Person;
Matrix: array[1..10, 1..10] of Real;
```

Variables declared outside procedures and functions are called *global variables* and they reside in the *data segment*. Variables declared within procedures and functions are called *local variables* and they reside in the *stack segment*.

The data segment

The maximum size of the data segment is 65,520 bytes. When a program is linked (this happens automatically at the end of the compilation of a program), the global variables of all units used by the program, as well as the program's own global variables, are placed in the data segment.

If you need more than 65,520 bytes of global data, you should allocate the larger structures as dynamic variables. For more details on this subject, see "Pointers and dynamic variables" on page 56.

The stack segment

The size of the stack segment is set through a **\$M** compiler directive—it can be anywhere from 1,024 to 65,520 bytes. The default stack-segment size is 8,192 bytes for a Windows application and it's 16,384 bytes for a DOS real-mode or protected-mode application.

Each time a procedure or function is activated (called), it allocates a set of local variables on the stack. On exit, the local variables are disposed of. At any time during the execution of a program, the total size of the local variables allocated by the active procedures and functions can't exceed the size of the stack segment.


The **\$S** compiler directive is used to include stack-overflow checks in the code. In the default **{\$\$+}** state, code is generated to check for stack overflow at the beginning of each procedure and function. In the **{\$\$-}** state, no such checks are performed. A stack

If you're writing a Windows application, Windows places special demands on the data and stack segments of your program, so the working maximum stack and the data-segment space can be less than the maximum data and stack segment space mentioned here.

overflow can cause a system crash, so don't turn off stack checks unless you're absolutely sure that an overflow will never occur.

Absolute variables

Variables can be declared to reside at specific memory addresses, and are then called *absolute variables*. The declaration of such variables must include an **absolute** clause following the type:



The variable declaration's identifier list can specify only one identifier when an **absolute** clause is present.

The first form of the **absolute** clause specifies the segment and offset at which the variable is to reside:

```
CrtMode : Byte absolute $0040:$0049;
```

The first constant specifies the segment base, and the second specifies the offset within that segment. Both constants must be within the range \$0000 to \$FFFF (0 to 65,535).



Use the first form of the **absolute** clause very carefully, if at all, in a Windows or DOS protected-mode program. While a Windows or DOS program is running in protected mode, it might not have access rights to memory areas outside your program. Attempting to access these memory areas will likely crash your program.

The second form of the **absolute** clause is used to declare a variable "on top" of another variable, meaning it declares a variable that resides at the same memory address as another variable:

```
var  
Str: string[32];  
StrLen: Byte absolute Str;
```

This declaration specifies that the variable *StrLen* should start at the same address as the variable *Str*, and because the first byte of a string variable contains the dynamic length of the string, *StrLen* contains the length of *Str*.


This second form of the **absolute** clause is safe to use in Windows and DOS protected-mode programming. Memory you are accessing is within your program's domain.

Variable references

A variable reference signifies one of the following:

- A variable
- A component of a structured- or string-type variable
- A dynamic variable pointed to by a pointer-type variable


This is the syntax of a variable reference:



The syntax for a variable reference allows an expression that computes a pointer type value. The expression must be followed by a qualifier that dereferences the pointer value (or indexes the pointer value if the extended syntax is enabled with the `{ $\$X+$ }` directive) to produce an actual variable reference.

Qualifiers

A variable reference can contain zero or more qualifiers that modify the meaning of the variable reference.



An array identifier with no qualifier, for example, references the entire array:

```
Results
```

An array identifier followed by an index denotes a specific component of the array—in this case, a structured variable:

```
Results[Current + 1]
```

With a component that is a record or object, the index can be followed by a field designator. Here the variable access signifies a specific field within a specific array component:

```
Results[Current + 1].Data
```

The field designator in a pointer field can be followed by the pointer symbol (^) to differentiate between the pointer field and the dynamic variable it points to:

```
Results[Current + 1].Data^
```


If the variable being pointed to is an array, indexes can be added to denote components of this array:

```
Results[Current + 1].Data^[J]
```

Arrays, strings, and indexes

A specific component of an array variable is denoted by a variable reference that refers to the array variable, followed by an index that specifies the component.

A specific character within a string variable is denoted by a variable reference that refers to the string variable, followed by an index that specifies the character position.



The index expressions select components in each corresponding dimension of the array. The number of expressions can't exceed the number of index types in the array declaration. Also, each expression's type must be assignment-compatible with the corresponding index type.

When indexing a multidimensional array, multiple indexes or multiple expressions within an index can be used interchangeably. For example,

```
Matrix[I][J]
```

is the same as

```
Matrix[I, J]
```

You can index a string variable with a single index expression, whose value must be in the range 0..N, where N is the declared size of the string. This accesses one character of the string value, with the type *Char* given to that character value.


The first character of a string variable (at index 0) contains the dynamic length of the string; that is, *Length(S)* is the same as *Ord(S[0])*. If a value is assigned to the length attribute, the compiler doesn't check whether this value is less than the declared size of the string. It's possible to index a string beyond its current dynamic length. The characters read are random and

assignments beyond the current length don't affect the actual value of the string variable.

When the extended syntax is enabled (using the `{SX+}` compiler directive), a value of type *PChar* can be indexed with a single index expression of type *Word*. The index expression specifies an *offset* to add to the character pointer before it's dereferenced to produce a *Char* type variable reference.

Records and field designators

A specific field of a record variable is denoted by a variable reference that refers to the record variable, followed by a field designator specifying the field.



These are examples of a field designator:

```
Today.Year  
Results[1].Count  
Results[1].When.Month
```

In a statement within a **with** statement, a field designator doesn't have to be preceded by a variable reference to its containing record.

Object component designators

The format of an object component designator is the same as that of a record field designator; that is, it consists of an instance (a variable reference), followed by a period and a component identifier. A component designator that designates a method is called a *method designator*. A **with** statement can be applied to an instance of an object type. In that case, the instance and the period can be omitted in referencing components of the object type.

The instance and the period can also be omitted within any method block, and when they are, the effect is the same as if *Self* and a period were written before the component reference.

Pointers and dynamic variables

The value of a pointer variable is either **nil** or the address of a dynamic variable.

The dynamic variable pointed to by a pointer variable is referenced by writing the pointer symbol (^) after the pointer variable.

You create dynamic variables and their pointer values with the procedures *New* and *GetMem*. You can use the @ (address-of)

operator and the function *Ptr* to create pointer values that are treated as pointers to dynamic variables.


nil doesn't point to any variable. The results are undefined if you access a dynamic variable when the pointer's value is **nil** or undefined.

These are examples of references to dynamic variables:

```
P1^  
P1^.Sibling^  
Results[1].Data^
```

Variable typecasts

A variable reference of one type can be changed into a variable reference of another type through a *variable typecast*.



The programmer is responsible for determining the validity of a typecast.

When a variable typecast is applied to a variable reference, the variable reference is treated as an instance of the type specified by the type identifier. The size of the variable (the number of bytes occupied by the variable) must be the same as the size of the type denoted by the type identifier. A variable typecast can be followed by one or more qualifiers, as allowed by the specified type.

Some examples of variable typecasts follow:

```
type  
TByteRec = record  
  Lo, Hi: Byte;  
end;  
TWordRec = record  
  Low, High: Word;  
end;  
TPtrRec = record  
  Ofs, Seg: Word;  
end;  
PByte = ^Byte;  
var  
B: Byte;  
W: Word;  
L: Longint;  
P: Pointer;
```

```

begin
  W := $1234;
  B := TByteRec(W).Lo;
  TByteRec(W).Hi := 0;
  L := $01234567;
  W := TWordRec(L).Low;
  B := TByteRec(TWordRec(L).Low).Hi;
  B := PByte(L)^;
  P := Ptr($40,$49);
  W := TPtrRec(P).Seg;
  Inc(TPtrRec(P).Ofs, 4);
end.

```

Notice the use of the *TByteRec* type to access the low- and high-order bytes of a word. This corresponds to the built-in functions *Lo* and *Hi*, except that a variable typecast can also be used on the left side of an assignment. Also, observe the use of the *TWordRec* and *TPtrRec* types to access the low- and high-order words of a long integer and the offset and segment parts of a pointer.

Borland Pascal fully supports variable typecasts involving procedural types. For example, given the declarations

```

type
  Func = function(X: Integer): Integer;
var
  F: Func;
  P: Pointer;
  N: Integer;

```

you can construct the following assignments:

```

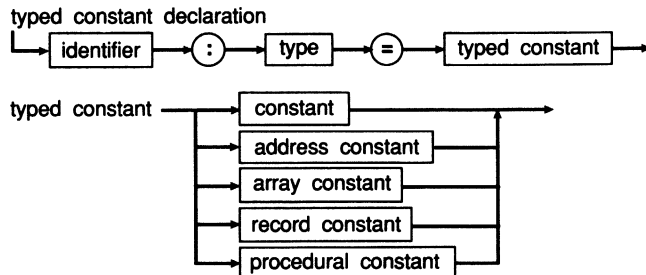
F := Func(P);           { Assign procedural value in P to F }
Func(P) := F;          { Assign procedural value in F to P }
@F := P;                { Assign pointer value in P to F }
P := @F;                { Assign pointer value in F to P }
N := F(N);              { Call function via F }
N := Func(P)(N);       { Call function via P }

```

In particular, notice that the address operator (**@**), when applied to a procedural variable, can be used on the left side of an assignment. Also, notice the typecast on the last line to call a function via a pointer variable.

Typed constants

Typed constants can be compared to initialized variables—variables whose values are defined on entry to their block. Unlike an untyped constant, the declaration of a typed constant specifies both the type and the value of the constant.



Typed constants can be used exactly like variables of the same type and they can appear on the left-hand side in an assignment statement. Note that typed constants are initialized *only once*—at the beginning of a program. Therefore, for each entry to a procedure or function, the locally declared typed constants aren't reinitialized.

In addition to a normal constant expression, the value of a typed constant can be specified using a *constant-address expression*. A constant-address expression is an expression that involves taking the address, offset, or segment of a global variable, a typed constant, a procedure, or a function. Constant-address expressions can't reference local variables (stack-based) or dynamic (heap-based) variables, because their addresses can't be computed at compile time.

Simple-type constants

Declaring a typed constant as a simple type specifies the value of the constant:

```
const  
Maximum: Integer = 9999;  
Factor: Real = -0.1;  
Breakchar: Char = #3;
```

As mentioned earlier, the value of a typed constant can be specified using a constant-address expression, that is, an expression that takes the address, offset, or segment of a global

variable, a typed constant, a procedure, or a function. For example,

```
var
  Buffer: array[0..1023] of Byte;
const
  BufferOfs: Word = OfS(Buffer);
  BufferSeg: Word = Seg(Buffer);
```

Because a typed constant is actually a variable with a constant value, it can't be interchanged with ordinary constants. For example, it can't be used in the declaration of other constants or types:

```
const
  Min: Integer = 0;
  Max: Integer = 99;
type
  Vector = array[Min..Max] of Integer;
```

The *Vector* declaration is invalid because *Min* and *Max* are typed constants.

String-type constants

The declaration of a typed constant of a string type specifies the maximum length of the string and its initial value:


```
const
  Heading: string[7] = 'Section';
  NewLine: string[2] = #13#10;
  TrueStr: string[5] = 'Yes';
  FalseStr: string[5] = 'No';
```

Structured-type constants

The declaration of a structured-type constant specifies the value of each of the structure's components. Borland Pascal supports the declaration of array, record, object, and set type constants. File type constants and constants of array, record, and object types that contain file type components aren't allowed.

Array-type constants

The declaration of an array-type constant, enclosed in parentheses and separated by commas, specifies the values of the components.



This is an example of an array-type constant:

```
type
    TStatus = (Active, Passive, Waiting);
    TStatusMap = array[TStatus] of string[7];
const
    StatStr: TStatusMap = ('Active', 'Passive', 'Waiting');
```

This example defines the array constant *StatStr*, which can be used to convert values of type *TStatus* into their corresponding string representations. These are the components of *StatStr*:

```
StatStr[Active] = 'Active'
StatStr[Passive] = 'Passive'
StatStr[Waiting] = 'Waiting'
```

The component type of an array constant can be any type except a file type. Packed string-type constants (character arrays) can be specified both as single characters and as strings. The definition

```
const
    Digits: array[0..9] of Char = ('0', '1', '2', '3', '4', '5',
    '6', '7', '8', '9');
```

can be expressed more conveniently as

```
const
    Digits: array[0..9] of Char = '0123456789';
```

When the extended syntax is enabled (using a **{\$X+}** compiler directive), a zero-based character array can be initialized with a string that is shorter than the declared length of the array. For example,

```
const
    FileName = array[0..79] of Char = 'TEST.PAS';
```

For more about null-terminated strings, see Chapter 18.

In such cases, the remaining characters are set to NULL (#0) and the array effectively contains a null-terminated string.

Multidimensional-array constants are defined by enclosing the constants of each dimension in separate sets of parentheses, separated by commas. The innermost constants correspond to the rightmost dimensions. The declaration


```
type
    TCube = array[0..1, 0..1, 0..1] of Integer;
const
    Maze: TCube = (((0, 1), (2, 3)), ((4, 5), (6, 7)));
```

provides an initialized array *Maze* with the following values:

```
Maze[0, 0, 0] = 0
Maze[0, 0, 1] = 1
Maze[0, 1, 0] = 2
Maze[0, 1, 1] = 3
Maze[1, 0, 0] = 4
Maze[1, 0, 1] = 5
Maze[1, 1, 0] = 6
Maze[1, 1, 1] = 7
```

Record-type constants

The declaration of a record-type constant specifies the identifier and value of each field, enclosed in parentheses and separated by semicolons.



Some examples of record constants follow:

```
type
TPoint = record
    X, Y: Real;
end;
TVector = array[0..1] of Point;
TMonth = (Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep, Oct,
    Nov, Dec);
TDate = record
    D: 1..31;
    M: Month;
    Y: 1900..1999;
end;
const
Origin: TPoint = (X: 0.0; Y: 0.0);
Line: TVector = ((X: -3.1; Y: 1.5), (X: 5.8; Y: 3.0));
SomeDay: TDate = (D: 2; M: Dec; Y: 1960);
```

The fields must be specified in the same order as they appear in the definition of the record type. If a record contains fields of file types, the constants of that record type can't be declared. If a record contains a variant, only fields of the selected variant can be specified. If the variant contains a tag field, then its value must be specified.

Object-type constants The declaration of an object-type constant uses the same syntax as the declaration of a record-type constant. No value is, or can be, specified for method components. Referring to the earlier object-type declarations starting on page 35, these are examples of object-type constants:

```
const
ZeroPoint: TPoint = (X: 0; Y: 0);
ScreenRect: TRect = (A: (X: 0; Y: 0); B: (X: 80; Y: 25));
CountField: TNumField = (X: 5; Y: 20; Len: 4; Name: nil;
Value: 0; Min: -999; Max: 999);
```



Constants of an object type that contains virtual methods need *not* be initialized through a constructor call—this initialization is handled automatically by the compiler.

Set-type constants Just like a simple-type constant, the declaration of a set-type constant specifies the value of the set using a constant expression. Here are some examples:

```
type
TDigits = set of 0..9;
TLetters = set of 'A'..'Z';
const
EvenDigits: TDigits = [0, 2, 4, 6, 8];
Vowels: TLetters = ['A', 'E', 'I', 'O', 'U', 'Y'];
HexDigits: set of '0'..'z' = ['0'..'9', 'A'..'F', 'a'..'f'];
```

Pointer-type constants

The declaration of a pointer-type constant uses a constant-address expression to specify the pointer value. Some examples follow:

```
type
TDirection = (Left, Right, Up, Down);
TStringPtr = ^String;
TNodePtr = ^Node;
TNode = record
    Next: TNodePtr;
    Symbol: TStringPtr;
    Value: TDirection;
end;
const
S1: string[4] = 'DOWN';
S2: string[2] = 'UP';
```

```

S3: string[5] = 'RIGHT';
S4: string[4] = 'LEFT';
N1: TNode = (Next: nil; Symbol: @S1; Value: Down);
N2: TNode = (Next: @N1; Symbol: @S2; Value: Up);
N3: TNode = (Next: @N2; Symbol: @S3; Value: Right);
N4: TNode = (Next: @N3; Symbol: @S4; Value: Left);
DirectionTable: TNodePtr = @N4;

```

When the extended syntax is enabled (using a **{SX+}** compiler directive), a typed constant of type *PChar* can be initialized with a string constant. For example,

```


const
  Message: PChar = 'Program terminated';
  Prompt: PChar = 'Enter values: ';
  Digits: array[0..9] of PChar = (
    'Zero', 'One', 'Two', 'Three', 'Four',
    'Five', 'Six', 'Seven', 'Eight', 'Nine');

```

The result is that the pointer now points to an area of memory that contains a zero-terminated copy of the string literal. See Chapter 18, "Using null-terminated strings," for more information.

Procedural-type constants

A procedural-type constant must specify the identifier of a procedure or function that is assignment-compatible with the type of the constant, or it must specify the value **nil**.



Here's an example:

```

type
  TErrorProc = procedure(ErrorCode: Integer);

procedure DefaultError(ErrorCode: Integer); far;
begin
  Writeln('Error ', ErrorCode, '.');
end;

const
  ErrorHandler: TErrorProc = DefaultError;

```


Expressions

Expressions are made up of *operators* and *operands*. Most Pascal operators are *binary*; they take two operands. The rest are *unary* and take only one operand. Binary operators use the usual algebraic form (for example, $A + B$). A unary operator always precedes its operand (for example, $-B$).

In more complex expressions, rules of precedence clarify the order in which operations are performed.

Table 6.1
Precedence of operators

Operators	Precedence	Categories
@, not	first (high)	unary operators
*, /, div, mod, and, shl, shr	second	multiplying operators
+, -, or, xor	third	adding operators
=, <>, <, >, <=, >=, in	fourth (low)	relational operators

There are three basic rules of precedence:

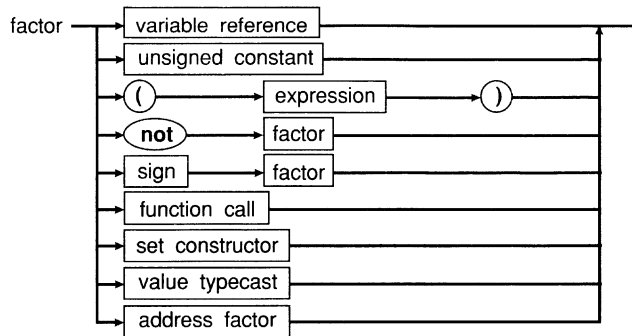
- An operand between two operators of different precedence is bound to the operator with higher precedence.
- An operand between two equal operators is bound to the one on its left.
- Expressions within parentheses are evaluated prior to being treated as a single operand.

Operations with equal precedence are normally performed from left to right, although the compiler may rearrange the operands to generate optimum code.

Expression syntax

The precedence rules follow from the syntax of expressions, which are built from factors, terms, and simple expressions.

A factor's syntax follows:




A function call activates a function and denotes the value returned by the function. See "Function calls" on page 76.

A set constructor denotes a value of a set type. See "Set constructors" on page 76.

A value typecast changes the type of a value. See "Value typecasts" on page 77.

An address factor computes the address of a variable, procedure, function, or method. See "The @ operator" on page 75.


An unsigned constant has the following syntax:



These are some examples of factors:

X	{ Variable reference }
@X	{ Pointer to a variable }
15	{ Unsigned constant }
(X + Y + Z)	{ Subexpression }
Sin(X / 2)	{ Function call }
exit['0'..'9', 'A'..'Z']	{ Set constructor }
not Done	{ Negation of a Boolean }
Char(Digit + 48)	{ Value typecast }


Terms apply the multiplying operators to factors:



Here are some examples of terms:

X * Y
 Z / (1 - Z)
 Y **shl** 2
 (X <= Y) **and** (Y < Z)

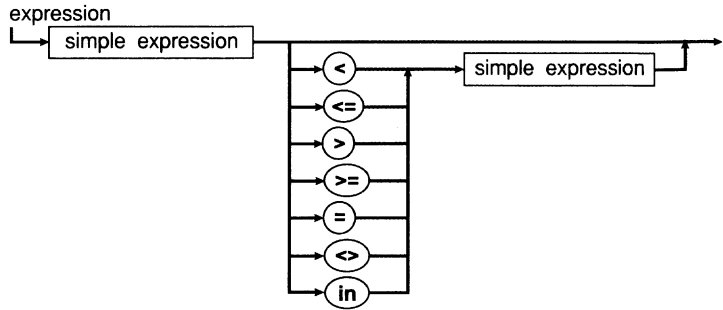
Simple expressions apply adding operators and signs to terms:



Here are some examples of simple expressions:

X + Y
 -X
 Hue1 + Hue2
 I * J + 1

An expression applies the relational operators to simple expressions:



Here are some examples of expressions:

```
X = 1.5
Done <> Error
(I < J) = (J < K)
C in Hue1
```

Operators

Operators are classified as arithmetic operators, logical operators, string operators, character-pointer operators, set operators, relational operators, and the @ operator.

Arithmetic operators

The following tables show the types of operands and results for binary and unary arithmetic operations.

Table 6.2
Binary arithmetic operations
*The + operator is also used as a string or set operator, and the +, -, and * operators are also used as set operators.*

Operator	Operation	Operand types	Result type
+	addition	integer type real type	integer type real type
-	subtraction	integer type real type	integer type real type
*	multiplication	integer type real type	integer type real type
/	division	integer type real type	real type real type
div	integer division	integer type	integer type
mod	remainder	integer type	integer type

Table 6.3
Unary arithmetic operations

Operator	Operation	Operand types	Result type
+	sign identity	integer type real type	integer type real type
-	sign negation	integer type real type	integer type real type

Any operand whose type is a subrange of an ordinal type is treated as if it were of the ordinal type.

For a definition of common types, see page 25.

If both operands of a **+**, **-**, *****, **div**, or **mod** operator are of an integer type, the result type is of the common type of the two operands.

If one or both operands of a **+**, **-**, or ***** operator are of a real type, the type of the result is *Real* in the **{\$N-}** state or *Extended* in the **{\$N+}** state.

If the operand of the sign identity or sign negation operator is of an integer type, the result is of the same integer type. If the operand is of a real type, the type of the result is *Real* or *Extended*.

The value of X / Y is always of type *Real* or *Extended* regardless of the operand types. A run-time error occurs if Y is zero.

The value of $I \text{ div } J$ is the mathematical quotient of I / J , rounded in the direction of zero to an integer-type value. A run-time error occurs if J is zero.

The **mod** operator returns the remainder obtained by dividing its two operands; that is,

$$I \text{ mod } J = I - (I \text{ div } J) * J$$

The sign of the result of **mod** is the same as the sign of I . A run-time error occurs if J is zero.

Logical operators

The types of operands and results for logical operations are shown in Table 6.4.

Table 6.4
Logical operations

The **not** operator is a unary operator.

Operator	Operation	Operand types	Result type
not	bitwise negation	integer type	<i>Boolean</i>
and	bitwise and	integer type	<i>Boolean</i>
or	bitwise or	integer type	<i>Boolean</i>
xor	bitwise xor	integer type	<i>Boolean</i>
shl	shift left	integer type	<i>Boolean</i>
shr	shift right	integer type	<i>Boolean</i>

If the operand of the **not** operator is of an integer type, the result is of the same integer type.

If both operands of an **and**, **or**, or **xor** operator are of an integer type, the result type is the common type of the two operands.

The operations *I shl J* and *I shr J* shift the value of *I* to the left right by *J* bits. The result type is the same as the type of *I*.

Boolean operators

Table 6.5
Boolean operations

The **not** operator is a unary operator.

The types of operands and results for Boolean operations are shown in Table 6.5.

Operator	Operation	Operand types	Result type
not	negation	Boolean type	<i>Boolean</i>
and	logical and	Boolean type	<i>Boolean</i>
or	logical or	Boolean type	<i>Boolean</i>
xor	logical xor	Boolean type	<i>Boolean</i>

Normal Boolean logic governs the results of these operations. For instance, *A and B* is *True* only if both *A* and *B* are *True*.

Borland Pascal supports two different models of code generation for the **and** and **or** operators: complete evaluation and short-circuit (partial) evaluation.

Complete evaluation means that every operand of a Boolean expression built from the **and** and **or** operators is guaranteed to be evaluated, even when the result of the entire expression is already known. This model is convenient when one or more operands of an expression are functions with side effects that alter the meaning of the program.

Short-circuit evaluation guarantees strict left-to-right evaluation and that evaluation stops as soon as the result of the entire expression becomes evident. This model is convenient in most cases because it guarantees minimum execution time, and usually

minimum code size. Short-circuit evaluation also makes possible the evaluation of constructs that would not otherwise be legal. For example,

```

while (I <= Length(S)) and (S[I] <> ' ') do
  Inc(I);
while (P <> nil) and (P^.Value <> 5) do
  P := P^.Next;

```

In both cases, the second test isn't evaluated if the first test is *False*.

The evaluation model is controlled through the **\$B** compiler directive. The default state is **{\$B-}**, and in this state, the compiler generates short-circuit evaluation code. In the **{\$B+}** state, the compiler generates complete evaluation.

Because Standard Pascal doesn't specify which model should be used for Boolean expression evaluation, programs dependent on either model aren't truly portable. You may decide, however, that sacrificing portability is worth the gain in execution speed and simplicity provided by the short-circuit model.

String operator

The types of operands and results for string operation are shown in Table 6.6.

Table 6.6
String operation

Operator	Operation	Operand types	Result type
+	concatenation	string type, <i>Char</i> type, or packed string type	string type

Borland Pascal allows the **+** operator to be used to concatenate two string operands. The result of the operation $S + T$, where S and T are of a string type, a *Char* type, or a packed string type, is the concatenation of S and T . The result is compatible with any string type (but not with *Char* types and packed string types). If the resulting string is longer than 255 characters, it's truncated after character 255.

Character-pointer operators

The extended syntax (enabled using a **{\$X+}** compiler directive) supports a number of character-pointer operations. The plus (**+**) and minus (**-**) operators can be used to increment and decrement

the offset part of a pointer value, and the minus operator can be used to calculate the distance (difference) between the offset parts of two character pointers. Assuming that P and Q are values of type $PChar$ and I is a value of type $Word$, these constructs are allowed:

Table 6.7
Permitted PChar constructs

Operation	Result
$P + I$	Add I to the offset part of P
$I + P$	Add I to the offset part of P
$P - I$	Subtract I from the offset part of P
$P - Q$	Subtract offset part of Q from offset part of P

The operations $P + I$ and $I + P$ adds I to the address given by P , producing a pointer that points I characters after P . The operation $P - I$ subtracts I from the address given by P , producing a pointer that points I characters before P .

The operation $P - Q$ computes the distance between Q (the lower address) and P (the higher address), resulting in a value of type $Word$ that gives the number of characters between Q and P . This operation assumes that P and Q point within the same character array. If the two character pointers point into different character arrays, the result is undefined.

Set operators

Table 6.8
Set operations

The types of operands for set operations are shown in Table 6.8.

Operator	Operation	Operand types
+	union	compatible set types
-	difference	compatible set types
*	intersection	compatible set types

The results of set operations conform to the rules of set logic:

- An ordinal value C is in $A + B$ only if C is in A or B .
- An ordinal value C is in $A - B$ only if C is in A and not in B .
- An ordinal value C is in $A * B$ only if C is in both A and B .

If the smallest ordinal value that is a member of the result of a set operation is A and the largest is B , then the type of the result is **set of $A..B$** .

Relational operators

The types of operands and results for relational operations are shown in Table 6.9.

Table 6.9
Relational operations

Operator type	Operation	Operand types	Result type
=	equal	compatible simple, pointer, set, string, or packed string types	<i>Boolean</i>
<>	not equal	compatible simple, pointer, set, string, or packed string types	<i>Boolean</i>
<	less than	compatible simple, string, packed string types, or <i>PChar</i>	<i>Boolean</i>
>	greater than	compatible simple, string, packed string types, or <i>PChar</i>	<i>Boolean</i>
<=	less than or equal to	compatible simple, string, packed string types, or <i>PChar</i>	<i>Boolean</i>
>=	greater than or equal to	compatible simple, string, or packed string types, or <i>PChar</i>	<i>Boolean</i>
<=	subset of	compatible set types	<i>Boolean</i>
>=	superset of	compatible set types	<i>Boolean</i>
in	member of	left operand, any ordinal type <i>T</i> ; right operand, set whose base is compatible with <i>T</i>	<i>Boolean</i>

Comparing simple types

When the operands =, <>, <, >, >=, or <= are of simple types, they must be compatible types; however, if one operand is of a real type, the other can be of an integer type.


Comparing strings	<p>The relational operators =, <>, <, >, >=, and <= compare strings according to the ordering of the extended ASCII character set. Any two string values can be compared because all string values are compatible.</p> <p>A character-type value is compatible with a string-type value. When the two are compared, the character-type value is treated as a string-type value with length 1. When a packed string-type value with <i>N</i> components is compared with a string-type value, it's treated as a string-type value with length <i>N</i>.</p>
Comparing packed strings	<p>The relational operators =, <>, <, >, >=, and <= can also be used to compare two packed string-type values if both have the same number of components. If the number of components is <i>N</i>, then the operation corresponds to comparing two strings, each of length <i>N</i>.</p>
Comparing pointers	<p>The operators = and <> can be used on compatible pointer-type operands. Two pointers are equal only if they point to the same object.</p>
Comparing character pointers	<p>The extended syntax (enabled using a {\$X+} compiler directive) allows the >, <, >=, and <= operators to be applied to <i>PChar</i> values. Note, however, that these relational tests assume that the two pointers being compared <i>point within the same character array</i>, and for that reason, the operators only compare the offset parts of the two pointer values. If the two character pointers point into different character arrays, the result is undefined.</p>
Comparing sets	<p>If <i>A</i> and <i>B</i> are set operands, their comparisons produce these results:</p> <ul style="list-style-type: none"> ■ <i>A</i> = <i>B</i> is <i>True</i> only if <i>A</i> and <i>B</i> contain exactly the same members; otherwise, <i>A</i> <> <i>B</i>. ■ <i>A</i> <= <i>B</i> is <i>True</i> only if every member of <i>A</i> is also a member of <i>B</i>. ■ <i>A</i> >= <i>B</i> is <i>True</i> only if every member of <i>B</i> is also a member of <i>A</i>.

Testing set membership

The **in** operator returns *True* when the value of the ordinal-type operand is a member of the set-type operand; otherwise, it returns *False*.

The @ operator

The @ operator is used in an address factor to compute the address of a variable, procedure, function, or method.



The @ operator returns the address of its operand, that is, it constructs a pointer value that points to the operand.

@ with a variable

Special rules apply to use of the @ operator with a procedural variable. For more details, see "Procedural types in expressions" on page 78.

When applied to a variable reference, @ returns a pointer to the variable. The type of the resulting pointer value is controlled through the **\$T** compiler directive: In the **{*\$T*-}** state (the default), the result type is *Pointer*. In other words, the result is an untyped pointer, which is compatible with all other pointer types. In the **{*\$T*+}** state, the type of the result is $\wedge T$, where *T* is the type of the variable reference. In other words, the result is of a type that is compatible only with other pointers to the type of the variable.

@ with a procedure, function, or method

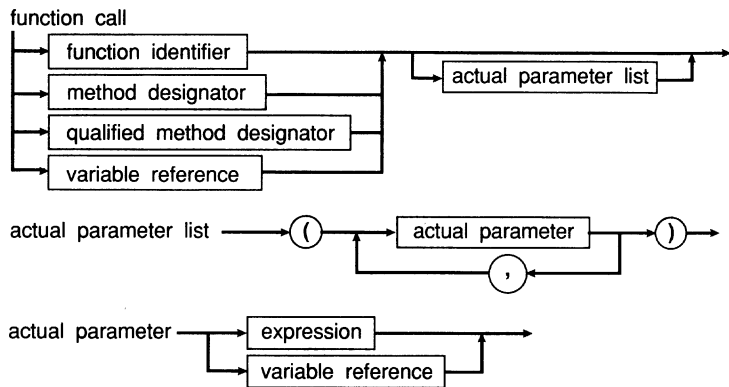
You can apply @ to a procedure, function, or method to produce a pointer to the routine's entry point. The type of the resulting pointer is always *Pointer*, regardless of the state of the **\$T** compiler directive. In other words, the result is always an untyped pointer, which is compatible with all other pointer types.

When @ is applied to a method, the method must be specified through a *qualified-method identifier* (an object-type identifier, followed by a period, followed by a method identifier).

Function calls

See "Method activations" on page 40, "Qualified-method activations" on page 41, and "Procedural types" on page 44.

A function call activates a function specified by a function identifier, a method designator, a qualified-method designator, or a procedural-type variable reference. The function call must have a list of actual parameters if the corresponding function declaration contains a list of formal parameters. Each parameter takes the place of the corresponding formal parameter according to parameter rules explained in Chapter 9, "Procedures and functions," on page 109.




These are some examples of function calls:

```
Sum(A, 63)
Maximum(147, J)
Sin(X + Y)
Eof(F)
Volume(Radius, Height)
```

In the extended syntax **{SX+}** mode, function calls can be used as statements; that is, the result of a function call can be discarded.

Set constructors

A set constructor denotes a set-type value, and is formed by writing expressions within brackets ([]). Each expression denotes a value of the set.



The notation `[]` denotes the empty set, which is assignment-compatible with every set type. Any member group `X..Y` denotes as set members all values in the range `X..Y`. If `X` is greater than `Y`, then `X..Y` doesn't denote any members and `[X..Y]` denotes the empty set.


All expression values in member groups in a particular set constructor must be of the same ordinal type.

These are some examples of set constructors:

```
[red, C, green]
[1, 5, 10..K mod 12, 23]
['A'..'Z', 'a'..'z', Chr(Digit + 48)]
```

Value typecasts

The type of an expression can be changed to another type through a value typecast.



The expression type and the specified type must both be either ordinal types or pointer types. For ordinal types, the resulting value is obtained by converting the expression. The conversion may involve truncation or extension of the original value if the size of the specified type is different from that of the expression. In cases where the value is extended, the sign of the value is always preserved; that is, the value is sign-extended.

See "Variable typecasts" on page 57.

The syntax of a value typecast is almost identical to that of a variable typecast. Value typecasts operate on values, however, not on variables, and therefore they can't participate in variable references; that is, a value typecast can't be followed by qualifiers. In particular, value typecasts can't appear on the left side of an assignment statement.

These are some examples of value typecasts:

```
Integer('A')
Char(48)
Boolean(0)
Color(2)
Longint(@Buffer)
BytePtr(Ptr($40, $49))
```

Procedural types in expressions

Usually, using a procedural variable in a statement or an expression calls the procedure or function stored in the variable. There is one exception: When the compiler sees a procedural variable on the left side of an assignment statement, it knows that the right side has to represent a procedural value. For example, consider the following program:

```
type
  IntFunc = function: Integer;

var
  F: IntFunc;
  N: Integer;

function ReadInt: Integer; far;
var
  I: Integer;
begin
  Read(I);
  ReadInt := I;
end;

begin
  F := ReadInt; { Assign procedural value }
  N := ReadInt; { Assign function result }
end.
```

The first statement in the main program assigns the procedural value (address of) *ReadInt* to the procedural variable *F*, where the second statement calls *ReadInt* and assigns the returned value to *N*. The distinction between getting the procedural value or calling the function is made by the type of the variable being assigned (*F* or *N*).

Unfortunately, there are situations where the compiler can't determine the desired action from the context. For example, in the following statement there is no obvious way the compiler can know if it should compare the procedural value in *F* to the

procedural value of *ReadInt* to determine if *F* currently points to *ReadInt*, or if it should call *F* and *ReadInt* and then compare the returned values.

```
if F = ReadInt then
  WriteLn('Equal');
```

Standard Pascal syntax, however, specifies that the occurrence of a function identifier in an expression denotes a call to that function, so the effect of the preceding statement is to call *F* and *ReadInt*, and then compare the returned values. To compare the procedural value in *F* to the procedural value of *ReadInt*, the following construct must be used:

```
if @F = @ReadInt then
  WriteLn('Equal');
```

When applied to a procedural variable or a procedure or function identifier, the address (**@**) operator prevents the compiler from calling the procedure, and at the same time converts the argument into a pointer. **@F** converts *F* into an untyped pointer variable that contains an address, and **@ReadInt** returns the address of *ReadInt*; the two pointer values can then be compared to determine if *F* currently refers to *ReadInt*.

The **@** operator is often used when assigning an untyped pointer value to a procedural variable. For example, the *GetProcAddress* function defined by Windows (in the *WinProcs* unit) returns the address of an exported function in a DLL as an untyped pointer value. Using the **@** operator, the result of a call to *GetProcAddress* can be assigned to a procedural variable:


```
type
  TStrComp = function(Str1, Str2: PChar): Integer;
var
  StrComp: TStrComp;
  :
begin
  :
  @StrComp := GetProcAddress(KernelHandle, 'lstrcmpi');
  :
end.
```



To get the memory address of a procedural variable rather than the address stored in it, use a double address (**@@**) operator. For example, where **@P** means convert *P* into an untyped pointer variable, **@@P** means return the physical address of the variable *P*.

Statements

Statements describe algorithmic actions that can be executed. Labels can prefix statements, and these labels can be referenced by **goto** statements.




A label is either a digit sequence in the range 0 to 9999 or an identifier.

There are two main types of statements: simple statements and structured statements.


Simple statements

A simple statement is a statement that doesn't contain any other statements.



Assignment statements

Assignment statements replace the current value of a variable with a new value specified by an expression. They can be used to set the return value of the function also.



See the section "Type compatibility" on page 48.

The expression must be assignment-compatible with the type of the variable or the type of the function result.

Some examples of assignment statements follow:

```
X := Y + Z;
Done := (I >= 1) and (I < 100);
Hue1 := [Blue, Succ(C)];
I := Sqr(J) - I * K;
```

Object-type assignments

The rules of object-type assignment-compatibility allow an instance of an object type to be assigned an instance of any of its descendant types. Such an assignment constitutes a *projection* of the descendant onto the space spanned by its ancestor. In the example code starting on page 35, given an instance *F* of type *TField* and an instance *Z* of type *TZipField*, the assignment *F* := *Z* copies only the fields *X*, *Y*, *Len*, and *Name*.




Assigning an instance of an object doesn't initialize the instance. Referring to the preceding example, the assignment *F* := *Z* doesn't mean that a constructor call for *F* can be omitted.

Procedure statements

See Chapter 9, "Procedures and functions."

A procedure statement activates a procedure specified by a procedure identifier, a method designator, a qualified-method designator, or a procedural-type variable reference. If the corresponding procedure declaration contains a list of formal parameters, then the procedure statement must have a matching list of actual parameters (parameters listed in definitions are *formal parameters*; in the calling statement, they are *actual*

parameters). The actual parameters are passed to the formal parameters as part of the call.




Some examples of procedure statements follow:

```
PrintHeading;  
Transpose(A, N, M);  
Find(Name, Address);
```

Goto statements

A **goto** statement transfers program execution to the statement marked by the specified label. The syntax diagram of a **goto** statement follows:




When using **goto** statements, observe the following rules:

Good programming practices recommend that you use goto statements as little as possible.

- The label referenced by a **goto** statement must be in the same block as the **goto** statement. In other words, it's not possible to jump into or out of a procedure or function.
- Jumping into a structured statement from outside that structured statement (that is, jumping to a deeper level of nesting) can have undefined effects, although the compiler doesn't indicate an error. For example, you shouldn't jump into the middle of a **for** loop.


Structured statements

Structured statements are constructs composed of other statements that are to be executed in sequentially (compound and **with** statements), conditionally (conditional statements), or repeatedly (repetitive statements).



Compound statements

The compound statement specifies that its component statements are to be executed in the same sequence as they are written. The component statements are treated as one statement, crucial in contexts where the Pascal syntax only allows one statement. **begin** and **end** bracket the statements, which are separated by semicolons.




Here's an example of a compound statement:

```
begin
  Z := X;
  X := Y;
  Y := Z;
end;
```


Conditional statements

A conditional statement selects for execution a single one (or none) of its component statements.



If statements

The syntax for an **if** statement reads like this:



The expression must yield a result of the standard type *Boolean*. If the expression produces the value *True*, then the statement following **then** is executed.

If the expression produces *False* and the **else** part is present, the statement following **else** is executed; if the **else** part isn't present, execution continues at the next statement following the **if** statement.

The syntactic ambiguity arising from the construct

```
if e1 then if e2 then s1 else s2;
```

is resolved by interpreting the construct as follows:

```
if e1 then
begin
  if e2 then
    s1
  else
    s2
end;
```

Note: No semicolon is allowed preceding an else clause.

Usually, an **else** is associated with the closest **if** not already associated with an **else**.

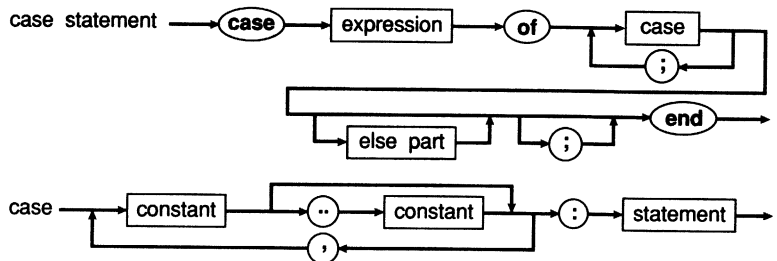
Two examples of **if** statements follow:


```
if X < 1.5 then
  Z := X + Y
else
  Z := 1.5;

if P1 <> nil then
  P1 := P1^.Father;
```

Case statements

The **case** statement consists of an expression (the selector) and a list of statements, each prefixed with one or more constants (called *case constants*) or with the word **else**. The selector must be of a byte-sized or word-sized ordinal type, so string types and the integer type *Longint* are invalid selector types. All **case** constants must be unique and of an ordinal type compatible with the selector type.





The **case** statement executes the statement prefixed by a **case** constant equal to the value of the selector or a **case** range containing the value of the selector. If no such **case** constant of the **case** range exists and an **else** part is present, the statement following **else** is executed. If there is no **else** part, execution continues with the next statement following the **if** statement.

These are examples of **case** statements:

```


case Operator of
  Plus: X := X + Y;
  Minus: X := X - Y;
  Times: X := X * Y;
end;

case I of
  0, 2, 4, 6, 8: Writeln('Even digit');
  1, 3, 5, 7, 9: Writeln('Odd digit');
  10..100: Writeln('Between 10 and 100');
else
  Writeln('Negative or greater than 100');
end;

```

Repetitive statements

Repetitive statements specify certain statements to be executed repeatedly.




If the number of repetitions is known beforehand, the **for** statement is the appropriate construct. Otherwise, the **while** or **repeat** statement should be used.

The *Break* and *Continue* standard procedures can be used to control the flow of repetitive statements: *Break* terminates a repetitive statement, and *Continue* continues with the next iteration of a repetitive statement. For more details on these standard procedures, see Chapter 1, "Library reference," in the *Programmer's Reference*.

Repeat statements

A **repeat** statement contains an expression that controls the repeated execution of a statement sequence within that **repeat** statement.



The expression must produce a result of type *Boolean*. The statements between the symbols **repeat** and **until** are executed in sequence until, at the end of a sequence, the expression yields *True*. The sequence is executed at least once because the expression is evaluated *after* the execution of each sequence.


These are examples of **repeat** statements:

```
repeat  
  K := I mod J;  
  I := J;  
  J := K;  
until J = 0;
```

```
repeat  
  Write('Enter value (0..9): ');  
  Readln(I);  
until (I >= 0) and (I <= 9);
```

While statements

A **while** statement contains an expression that controls the repeated execution of a statement (which can be a compound statement).



The expression controlling the repetition must be of type *Boolean*. It is evaluated *before* the contained statement is executed. The contained statement is executed repeatedly as long as the expression is *True*. If the expression is *False* at the beginning, the statement isn't executed at all.

These are examples of **while** statements:

```
while Data[I] <> X do I := I + 1;
```

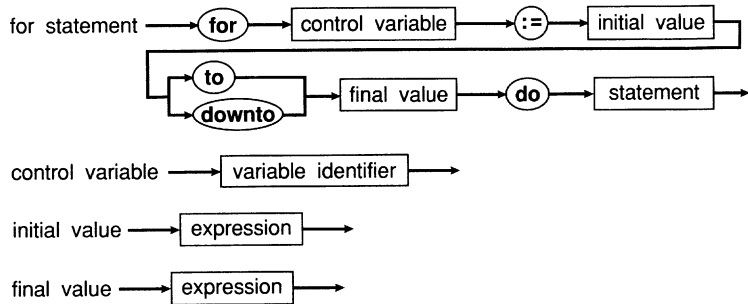
```

while I > 0 do
begin
  if Odd(I) then Z := Z * X;
  I := I div 2;
  X := Sqr(X);
end;

while not Eof(InFile) do
begin
  Readln(InFile, Line);
  Process(Line);
end;

```

For statements The **for** statement causes a statement to be repeatedly executed while a progression of values is assigned to a control variable. Such a statement can be a compound statement.



See Chapter 8 for a discussion of locality and scope.

The control variable must be a variable identifier (without any qualifier) that is local in scope to the block containing the **for** statement. The control variable must be of an ordinal type. The initial and final values must be of a type assignment-compatible with the ordinal type.

When a **for** statement is entered, the initial and final values are determined once for the remainder of the execution of the **for** statement.

The statement contained by the **for** statement is executed once for every value in the range *initial value* to *final value*. The control variable always starts off at *initial value*. When a **for** statement uses **to**, the value of the control variable is incremented by one for each repetition. If *initial value* is greater than *final value*, the contained statement isn't executed. When a **for** statement uses **downto**, the value of the control variable is decremented by one for each repetition. If *initial value* value is less than *final value*, the contained statement isn't executed.

If the contained statement alters the value of the control variable, your results will probably not be what you expect. After a **for** statement is executed, the value of the control variable value is undefined, unless execution of the **for** statement was interrupted by a **goto** from the **for** statement.

With these restrictions in mind, the **for** statement

```
for V := Expr1 to Expr2 do Body;
```

is equivalent to

```
begin
  Temp1 := Expr1;
  Temp2 := Expr2;
  if Temp1 <= Temp2 then
    begin
      V := Temp1;
      Body;
      while V <> Temp2 do
        begin
          V := Succ(V);
          Body;
        end;
      end;
    end;
end;
```

and the **for** statement

```
for V := Expr1 downto Expr2 do Body;
```

is equivalent to

```
begin
  Temp1 := Expr1;
  Temp2 := Expr2;
  if Temp1 >= Temp2 then
    begin
      V := Temp1;
      Body;
      while V <> Temp2 do
        begin
          V := Pred(V);
          Body;
        end;
      end;
    end;
end;
```

where *Temp1* and *Temp2* are auxiliary variables of the host type of the variable *V* and don't occur elsewhere in the program.

These are examples of **for** statements:

```

for I := 2 to 63 do
  if Data[I] > Max then
    Max := Data[I]


for I := 1 to 10 do
  for J := 1 to 10 do
  begin
    X := 0;
    for K := 1 to 10 do
      X := X + Mat1[I, K] * Mat2[K, J];
    Mat[I, J] := X;
  end;

for C := Red to Blue do Check(C);

```

With statements

The **with** statement is shorthand for referencing the fields of a record, and the fields and methods of an object. Within a **with** statement, the fields of one or more specific record variables can be referenced using their field identifiers only. The syntax of a **with** statement follows:



Given this type declaration,

```

type
  TDate = record
    Day : Integer;
    Month: Integer;
    Year : Integer;
  end;

var OrderDate: TDate;

```

here is an example of a **with** statement:

```
with OrderDate do
  if Month = 12 then
    begin
      Month := 1;
      Year := Year + 1
    end
  else
    Month := Month + 1;
```

This is equivalent to

```
if OrderDate.Month = 12 then
  begin
    OrderDate.Month := 1;
    OrderDate.Year := TDate.Year + 1
  end
else
  OrderDate.Month := TDate.Month + 1;
```

Within a **with** statement, each variable reference is first checked to see if it can be interpreted as a field of the record. If so, it's always interpreted as such, even if a variable with the same name is also accessible. Suppose the following declarations have been made:

```
type
  TPoint = record
    X, Y: Integer;
  end;
var
  X: TPoint;
  Y: Integer;
```

In this case, both *X* and *Y* can refer to a variable or to a field of the record. In the statement

```
with X do
  begin
    X := 10;
    Y := 25;
  end;
```

the *X* between **with** and **do** refers to the variable of type *TPoint*, but in the compound statement, *X* and *Y* refer to *X.X* and *X.Y*.

The statement

```
with V1, V2, ... Vn do S;
```

is equivalent to

```
with V1 do
  with V2 do
    :
    with Vn do
      S;
```

In both cases, if V_n is a field of both V_1 and V_2 , it's interpreted as $V_2.V_n$, not $V_1.V_n$.

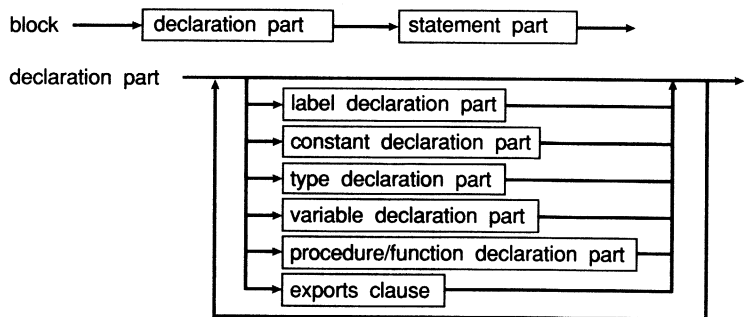
If the selection of a record variable involves indexing an array or dereferencing a pointer, these actions are executed once before the component statement is executed.

Blocks, locality, and scope

A *block* is made up of declarations, which are written and combined in any order, and statements. Each block is part of a procedure declaration, a function declaration, a method declaration, or a program or unit. All identifiers and labels declared in the declaration part are local to the block.

Blocks

The overall syntax of any block follows this format:




Labels that mark statements in the corresponding statement part are declared in the *label declaration part*. Each label must mark only one statement.



A label must be an identifier or a digit sequence in the range 0 to 9999.


The *constant declaration part* consists of constant declarations local to the block.




The *type declaration part* includes all type declarations local to the block.



The *variable declaration part* is composed of variable declarations local to the block.




The procedure and function declarations local to the block make up the *procedure and function declaration part*.



See "The exports clause" on page 133.

The **exports** clause lists all procedures and functions that are exported by the current program or dynamic-link library. An **exports** clause is allowed only in the outermost declaration part of a program or dynamic-link library—it isn't allowed in the declaration part of a procedure, function, or unit.

The *statement part* defines the statements or algorithmic actions to be executed by the block.



Rules of scope

The presence of an identifier or label in a declaration defines the identifier or label. Each time the identifier or label occurs again, it must be within the *scope* of this declaration.

Block scope

The scope of an identifier or label declared in a label, constant, type, variable, procedure, or function declaration stretches from the point of declaration to the end of the current block, and includes all blocks enclosed by the current block.

An identifier or label declared in an outer block can be *redeclared* in an inner block enclosed by the outer block. Before the point of declaration in the inner block, and after the end of the inner block, the identifier or label represents the entity declared in the outer block.

```
program Outer;           { Start of outer scope }
type
  I = Integer;          { define I as type Integer }
var
  T: I;                 { define T as an Integer variable }
procedure Inner;        { Start of inner scope }
type
  T = I;               { redefine T as type Integer }
var
  I: T;                { redefine I as an Integer variable }
begin
  I := 1;
end;                   { End of inner scope }

begin
  T := 1;
end.                   { End of outer scope }
```

Record scope

The scope of a field identifier declared in a record-type definition extends from the point of declaration to the end of the record-type definition. Also, the scope of field identifiers includes field designators and **with** statements that operate on variable references of the given record type.

See "Record types" on page 32.

Object scope

See "Object types" on page 34.

The scope of a component identifier declared in an object-type definition extends from the point of declaration to the end of the object-type definition, and extends over all descendants of the object type and the blocks of all method declarations of the object type. The scope of component identifiers includes field designators and **with** statements that operate on variable references of the given object type.

Unit scope

The scope of identifiers declared in the interface section of a unit follows the rules of block scope and extends over all *clients* of the unit. In other words, programs or units containing **uses** clauses have access to the identifiers belonging to the interface parts of the units in those **uses** clauses.

Each unit in a **uses** clause imposes a new scope that encloses the remaining units used and the program or unit containing the **uses** clause. The first unit in a **uses** clause represents the outermost scope, and the last unit represents the innermost scope. This implies that if two or more units declare the same identifier, an unqualified reference to the identifier selects the instance declared by the last unit in the **uses** clause. If you use a qualified identifier (a unit identifier, followed by a period, followed by the identifier), every instance of the identifier can be selected.

The identifiers of Borland Pascal's predefined constants, types, variables, procedures, and functions act as if they were declared in a block enclosing all used units and the entire program. In fact, these standard objects are defined in a unit called *System*, which is used by any program or unit before the units named in the **uses** clause. This means that any unit or program can redeclare the standard identifiers, but a specific reference can still be made through a qualified identifier, for example, *System.Integer* or *System.Writeln*.

Procedures and functions

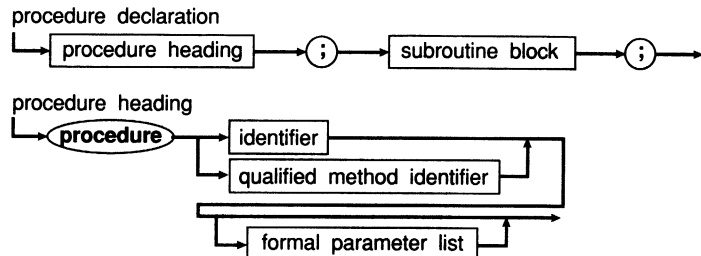
See Chapter 8, "Blocks, locality, and scope," for a definition of a block.


Procedures and functions let you nest additional blocks in the main program block. Each procedure or function declaration has a heading followed by a block. A procedure is activated by a procedure statement; a function is activated by the evaluation of an expression that contains its call and returns a value to that expression.

This chapter discusses the different types of procedure and function declarations and their parameters.

Procedure declarations

A *procedure declaration* associates an identifier with a block as a procedure; that procedure can then be activated by a procedure statement.





The syntax for a formal parameter list is shown in the section "Parameters" on page 109.

The procedure heading names the procedure's identifier and specifies the formal parameters (if any).

A procedure is activated by a procedure statement, which states the procedure's identifier and actual parameters, if any. The statements to be executed on activation are noted in the statement part of the procedure's block. If the procedure's identifier is used in a procedure statement within the procedure's block, the procedure is executed recursively (it calls itself while executing).

Here's an example of a procedure declaration:

```

procedure NumString(N: Integer; var S: string);
var
  V: Integer;
begin
  V := Abs(N);
  S := '';
  repeat
    S := Chr(N mod 10 + Ord('0')) + S;
    N := N div 10;
  until N = 0;
  if N < 0 then
    S := '-' + S;
end;

```

Near and far declarations

Near and far calls are described in Chapter 22, "Control issues."

Borland Pascal supports two procedure and function call models: near and far. In terms of code size and execution speed, the near call model is the more efficient, but **near** procedures and functions can only be called from within the module they are declared in. On the other hand, **far** procedures and functions can be called from any module, but the code for a far call is slightly less efficient.

The compiler automatically selects the correct call model based on a procedure's or function's declaration: Procedures and functions

declared in the **interface** part of a unit use the far call model—they can be called from other modules. Procedures and functions declared in a program or in the **implementation** part of a unit use the near call model—they can only be called from within that program or unit.

For some purposes, a procedure or function may be required to use the far call model. For example, if a procedure or function is to be assigned to a procedural variable, it has to use the far call model. The **\$F** compiler directive can be used to override the compiler's automatic call model selection. Procedures and functions compiled in the **{\$F+}** state always use the far call model; in the **{\$F-}** state, the compiler automatically selects the correct model. The default state is **{\$F-}**.

To force a specific call model, a procedure or function declaration can optionally specify a **near** or **far** directive before the block—if such a directive is present, it overrides the setting of the **\$F** compiler directive as well as the compiler's automatic call model selection.

Export declarations

The **export** directive makes a procedure or function exportable by forcing the routine to use the far call model and generating special procedure entry and exit code.

Procedures and functions must be exportable in these cases:

- Procedures and functions that are exported by a DLL (dynamic-link library)
- Callback procedures and functions in a Windows program

Chapter 11, “Dynamic-link libraries,” discusses how to export procedures and functions in a DLL. Even though a procedure or function is compiled with an **export** directive, the actual exporting of the procedure or function doesn't occur until the routine is listed in a library's **exports** clause.

Callback procedures and functions are routines in your application that are called by Windows and not by your application itself. Callback routines must be compiled with the **export** directive, but they don't have to be listed in an **exports** clause. Here are some examples of common callback procedures and functions:

- Window procedures
- Dialog procedures
- Enumeration callback procedures
- Memory-notification procedures
- Window-hook procedures (filters)

See "Entry and exit code" on page 298.

Borland Pascal automatically generates *smart callbacks* for procedures and functions that are exported by a Windows program. Smart callbacks alleviate the need to use the *MakeProcInstance* and *FreeProcInstance* Windows API routines when creating callback routines.

Interrupt declarations

Optionally, a procedure declaration can specify an **interrupt** directive before the block; the procedure is then considered an *interrupt procedure*. For now, note that **interrupt** procedures can't be called from procedure statements, and that every **interrupt** procedure must specify a parameter list like the following:

Don't use the *interrupt* directive when writing programs for Windows—your programs are likely to crash.

```
procedure MyInt (Flags, CS, IP, AX, BX, CX, DX, SI, DI, DS, ES,
                BP: Word);
interrupt;
```

The parameter list doesn't have to match this syntax perfectly; it can be shorter and use different names, but the register contents are passed in the order listed above.

Forward declarations

A procedure or function declaration that specifies the directive **forward** instead of a block is a **forward** declaration. Somewhere after this declaration, the procedure must be defined by a *defining declaration*. The defining declaration can omit the formal parameter list and the function result, or it can optionally repeat it. In the latter case, the defining declaration's heading must match exactly the order, types, and names of parameters, and the type of the function result, if any.

No **forward** declarations are allowed in the **interface** part of a unit.

The **forward** declaration and the defining declaration must appear in the same procedure and function declaration part. Other procedures and functions can be declared between them, and they can call the forward-declared procedure. Therefore, mutual recursion is possible.

The **forward** declaration and the defining declaration constitute a complete procedure or function declaration. The procedure or function is considered declared at the **forward** declaration.

This is an example of a **forward** declaration:

```

procedure Walter(M, N: Integer); forward;

procedure Clara(X, Y: Real);
begin
  :
  Walter(4, 5);
  :
end;

procedure Walter;
begin
  :
  Clara(8.3, 2.4);
  :
end;


```

A procedure's or function's defining declaration can be an **external** or **assembler** declaration; however, it can't be a **near**, **far**, **export**, **interrupt**, or **inline** declaration or another **forward** declaration.

External declarations

For more details about linking with assembly language, see Chapter 25.

With **external** declarations, you can interface with separately compiled procedures and functions written in assembly language. They also allow you to import procedures and functions from DLLs.



External directives consisting only of the reserved word **external** are used in conjunction with **{\$L filename}** directives to link with **external** procedures and functions implemented in **.OBJ** files.

These are examples of **external** procedure declarations:

```
procedure MoveWord(var Source, Dest; Count: Word); external;
procedure MoveLong(var Source, Dest; Count: Word); external;

procedure FillWord(var Dest; Data: Integer; Count: Word); external;
procedure FillLong(var Dest; Data: Longint; Count: Word); external;

{$L BLOCK.OBJ}
```

External directives that specify a dynamic-link library name (and optionally an import name or an import ordinal number) are used to import procedures and functions from dynamic-link libraries. For example, this **external** declaration imports a function called *GlobalAlloc* from the DLL called *KERNEL* (the Windows kernel):

```
function GlobalAlloc(Flags: Word; Bytes: Longint): THandle; far;
external 'KERNEL' index 15;
```


To read more about importing procedures and functions from a DLL, see Chapter 11.

The **external** directive takes the place of the declaration and statement parts in an imported procedure or function. Imported procedures and functions must use the far call model selected by using a **far** procedure directive or a **(\$F+)** compiler directive. Aside from this requirement, imported procedures and functions are just like regular procedures and functions.

Assembler declarations


For more details on assembler procedures and functions, see Chapter 24.

With **assembler** declarations, you can write entire procedures and functions in inline assembly language.



Inline declarations

The **inline** directive enables you to write machine code instructions in place of a block of Pascal code.



See the inline statement syntax diagram on page 344.

When a normal procedure or function is called, the compiler generates code that pushes the procedure's or function's parameters onto the stack and then generates a CALL instruction to call the procedure or function. When you call an **inline** procedure or function, the compiler generates code from the **inline** directive instead of the CALL. Therefore, an **inline** procedure or

function is expanded every time you refer to it, just like a macro in assembly language.

Here's a short example of two **inline** procedures:

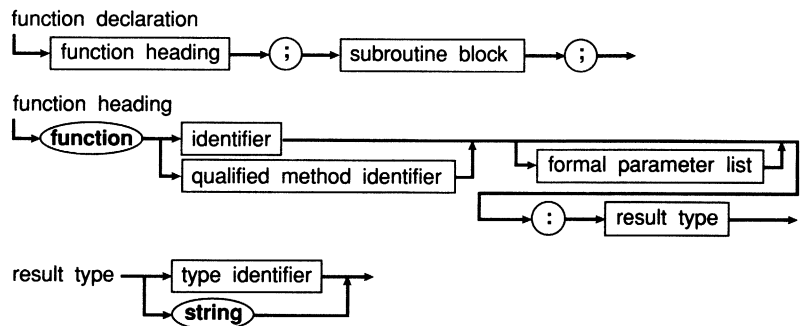
```

procedure DisableInterrupts; inline($FA); { CLI }
procedure EnableInterrupts; inline($FB); { STI }

```

Function declarations

A function declaration defines a block that computes and returns a value.



A function can't return a structured type or a procedural type.

The function heading specifies the identifier for the function, the formal parameters (if any), and the function result type.

A function is activated by the evaluation of a function call. The function call gives the function's identifier and actual parameters, if any, required by the function. A function call appears as an operand in an expression. When the expression is evaluated, the function is executed, and the value of the operand becomes the value returned by the function.

The statement part of the function's block specifies the statements to be executed upon activation of the function. The block should contain at least one assignment statement that assigns a value to the function identifier. The result of the function is the last value assigned. If no such assignment statement exists or if it isn't executed, the value returned by the function is undefined.

If the function's identifier is used in a function call within the function's block, the function is executed recursively.

Following are examples of function declarations:

```
function Max(A: Vector; N: Integer): Extended;
var
  X: Extended;
  I: Integer;
begin
  X := A[1];
  for I := 2 to N do
    if X < A[I] then X := A[I];
  Max := X;
end;

function Power(X: Extended; Y: Integer): Extended;
var
  Z: Extended;
  I: Integer;
begin
  Z := 1.0; I := Y;
  while I > 0 do
    begin
      if Odd(I) then Z := Z * X;
      I := I div 2;
      X := Sqr(X);
    end;
  Power := Z;
end;
```

Like procedures, functions can be declared as **near**, **far**, **export**, **forward**, **external**, **assembler**, or **inline**; however, **interrupt** functions are *not* allowed.

Method declarations

See page 37 for more about declaring methods in objects.

The declaration of a method within an object type corresponds to a **forward** declaration of that method. Somewhere after the object-type declaration and within the same scope as the object-type declaration, the method must be implemented by a defining declaration.

For procedure and function methods, the defining declaration takes the form of a normal procedure or function, but the procedure or function identifier is a *qualified-method identifier*. This is an object-type identifier followed by a period (.) and then by a method identifier.

For constructor methods and destructor methods, the defining declaration takes the form of a procedure method declaration, except that the **procedure** reserved word is replaced by a **constructor** or **destructor** reserved word.

Optionally, a method's defining declaration can repeat the formal parameter list of the method heading in the object type. The defining declaration's method heading must match exactly the order, types, and names of the parameters, and the type of the function result, if any.

In the defining declaration of a method, there is always an implicit parameter with the identifier *Self*, corresponding to a formal variable parameter that possesses the object type. Within the method block, *Self* represents the instance whose method component was designated to activate the method. Therefore, any changes made to the values of the fields of *Self* are reflected in the instance.

The scope of a component identifier in an object type extends over any procedure, function, constructor, or destructor block that implements a method of the object type. The effect is the same as if the entire method block was embedded in a **with** statement of the form

```
with Self do begin ... end
```

For this reason, the spellings of component identifiers, formal method parameters, *Self*, and any identifiers introduced in a method implementation must be unique.

Here are some examples of method implementations:

See the object-type declarations of these examples on page 35.

```
procedure TRectangle.Intersect(var R: TRectangle);
begin
  if A.X < R.A.X then A.X := R.A.X;
  if A.Y < R.A.Y then A.Y := R.A.Y;
  if B.X > R.B.X then B.X := R.B.X;
  if B.Y > R.B.Y then B.Y := R.B.Y;
  if (A.X >= B.X) or (A.Y >= B.Y) then Init(0, 0, 0, 0);
end;

procedure TField.Display;
begin
  GotoXY(X, Y);
  Write(Name^, ' ', GetStr);
end;
```

```

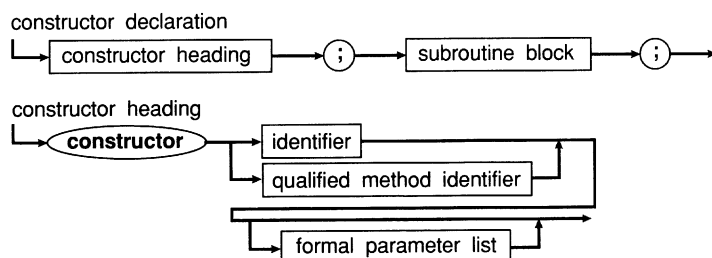
function TNumField.PutStr(S: String): Boolean;
var
  E: Integer;
begin
  Val(S, Value, E);
  PutStr := (E = 0) and (Value >= Min) and (Value <= Max);
end;

```

Constructors and destructors

Constructors and *destructors* are specialized forms of methods. Used in connection with the extended syntax of the *New* and *Dispose* standard procedures, constructors and destructors have the ability to allocate and deallocate dynamic objects. In addition, constructors have the ability to perform the required initialization of objects that contain virtual methods. Like other methods, constructors and destructors can be inherited, and an object can have any number of constructors and destructors.

Constructors are used to initialize newly created objects. Typically, the initialization is based on values passed as parameters to the constructor. Constructors can't be virtual because the virtual-method dispatch-mechanism depends on a constructor first having initialized the object.



Here are some examples of constructors:

```

constructor TField.Copy(var F: TField);
begin
  Self := F;
end;

constructor TField.Init(FX, FY, FLen: Integer; FName: String);
begin
  X := FX;
  Y := FY;
  Len := FLen;
  GetMem(Name, Length(FName) + 1);
  Name^ := FName;
end;

```

```

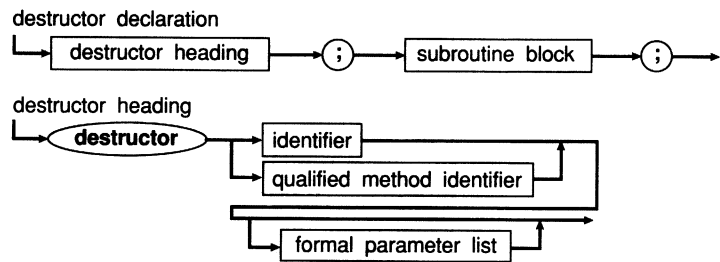
constructor TStrField.Init(FX, FY, FLen: Integer; FName: String);
begin
    inherited Init(FX, FY, FLen, FName);
    GetMem(Value, Len);
    Value^ := '';
end;

```

The first action of a constructor of a descendant type, such as the preceding *TStrField.Init*, is almost always to call its immediate ancestor's corresponding constructor to initialize the inherited fields of the object. Having done that, the constructor then initializes the fields of the object that were introduced in the descendant.

Destructors can be virtual, and often are. Destructors seldom take any parameters.

Destructors are the counterparts of constructors, and are used to clean up objects after their use. Usually, the cleanup consists of disposing of any pointer fields that were allocated by the object.



Here are some examples of destructors:

```

destructor TField.Done;
begin
    FreeMem(Name, Length(Name^) + 1);
end;

destructor TStrField.Done;
begin
    FreeMem(Value, Len);
    inherited Done;
end;

```

A destructor of a descendant type, such as the preceding *TStrField.Done*, usually disposes of the pointer fields introduced in the descendant, and then, as its last action, calls the corresponding destructor of its immediate ancestor to dispose of any inherited pointer fields of the object.

Constructor error recovery

Borland Pascal allows you to install a heap-error function through the *HeapError* variable in the *System* unit (see Chapter 21). This functionality affects the way object-type constructors work.

By default, when there isn't enough memory to allocate a dynamic instance of an object type, a constructor call using the extended syntax of the *New* standard procedure generates runtime error 203. If you install a heap-error function that returns 1 rather than the standard function result of 0, a constructor call through *New* returns **nil** when it can't complete the request (instead of aborting the program).

The code that performs allocation and virtual method table (VMT) field initialization of a dynamic instance is part of a constructor's entry sequence: When control arrives at the **begin** of the constructor's statement part, the instance has been allocated and initialized already. If allocation fails, and if the heap-error function returns 1, the constructor skips execution of the statement part and returns a **nil** pointer. Therefore, the pointer specified in the *New* construct that called the constructor is set to **nil**.

Once control arrives at the **begin** of a constructor's statement part, the object-type instance is guaranteed to have been allocated and initialized successfully. The constructor itself might attempt to allocate dynamic variables to initialize pointer fields in the instance, however, and these allocations might in turn fail. If that happens, a well-behaved constructor should reverse any successful allocations, and finally deallocate the object-type instance so that the net result becomes a **nil** pointer. To make such "backing out" possible, Borland Pascal implements a standard procedure called *Fail*, which takes no parameters and can be called only from within a constructor. A call to *Fail* causes a constructor to deallocate the dynamic instance that was allocated upon entry to the constructor and causes the return of a **nil** pointer to indicate its failure.

When dynamic instances are allocated through the extended syntax of *New*, a resulting value of **nil** in the specified pointer variable indicates that the operation failed. Unfortunately, there is no such pointer variable to inspect after the construction of a static instance or when an inherited constructor is called. Instead, Borland Pascal allows a constructor to be used as a Boolean function in an expression: A return value of *True* indicates success,


and a return value of *False* indicates failure due to a call to *Fail* within the constructor.

On disk you'll find two programs, *NORECVER.PAS* and *RECOVER.PAS*. Both implement two simple object types that contain pointers. The *NORECVER* version of the program doesn't implement constructor-error recovery.

RECOVER.PAS demonstrates how the program can be rewritten to implement error recovery. Notice how the corresponding destructors in *Base.Init* and *Derived.Init* are used to reverse any successful allocations before *Fail* is called to finally fail the operation. Also notice that in *Derived.Init*, the call to *Base.Init* is coded within an expression so that the success of the inherited constructor can be tested.

Parameters

The declaration of a procedure or function specifies a *formal parameter list*. Each parameter declared in a formal parameter list is local to the procedure or function being declared and can be referred to by its identifier in the block associated with the procedure or function.



There are four kinds of parameters: *value*, *constant*, *variable*, and *untyped*. These are characterized as follows:

- A parameter group without a preceding **var** and followed by a type is a list of value parameters.
- A parameter group preceded by **const** and followed by a type is a list of constant parameters.
- A parameter group preceded by **var** and followed by a type is a list of variable parameters.
- A parameter group preceded by **var** or **const** and *not* followed by a type is a list of untyped parameters.

Open parameters are described on page 113.

String and array-type parameters can be *open parameters*. A variable parameter declared using the *OpenString* identifier, or using the **string** keyword in the **{\$P+}** state, is an *open-string parameter*. A value, constant, or variable parameter declared using the syntax **array of T** is an *open-array parameter*.

Value parameters

A formal value parameter acts like a variable local to the procedure or function, except it gets its initial value from the corresponding actual parameter upon activation of the procedure or function. Changes made to a formal value parameter don't affect the value of the actual parameter.

A value parameter's corresponding actual parameter in a procedure statement or function call must be an expression, and its value must not be of file type or of any structured type that contains a file type.

The actual parameter must be assignment-compatible with the type of the formal value parameter. If the parameter type is **string**, then the formal parameter is given a size attribute of 255.

Constant parameters

A formal constant parameter acts like a local read-only variable that gets its value from the corresponding actual parameter upon activation of the procedure or function. Assignments to a formal constant parameter are *not* allowed, and likewise a formal constant parameter *can't* be passed as an actual variable parameter to another procedure or function.

A constant parameter's corresponding actual parameter in a procedure statement or function must follow the same rules as an actual value parameter.

In cases where a formal parameter never changes its value during the execution of a procedure or function, a constant parameter should be used instead of a value parameter. Constant parameters allow the implementor of a procedure or function to protect against accidental assignments to a formal parameter. Also, for structured- and string-type parameters, the compiler can generate more efficient code when constant parameters are used instead of value parameters.

Variable parameters

A variable parameter is used when a value must be passed from a procedure or function to the caller. The corresponding actual parameter in a procedure statement or function call must be a variable reference. The formal variable parameter represents the actual variable during the activation of the procedure or function, so any changes to the value of the formal variable parameter are reflected in the actual parameter.

File types can be passed only as variable parameters.

Within the procedure or function, any reference to the formal variable parameter accesses the actual parameter itself. The type of the actual parameter must be identical to the type of the formal variable parameter (you can bypass this restriction through untyped parameters).

The **\$P** compiler directive controls the meaning of a variable parameter declared using the **string** keyword. In the default (**\$P-**) state, **string** corresponds to a string type with a size attribute of 255. In the (**\$P+**) state, **string** indicates that the parameter is an open-string parameter. See page 113 for information on open-string parameters.

If referencing an actual variable parameter involves indexing an array or finding the object of a pointer, these actions are executed before the activation of the procedure or function.

The rules of object-type assignment compatibility also apply to object-type variable parameters: For a formal parameter of type *T1*, the actual parameter might be of type *T2* if *T2* is in the domain of *T1*. For example, given the object-type declarations found on page 35, the *TField.Copy* method might be passed an instance of *TField*, *TStrField*, *TNumField*, *TZipField*, or any other instance of a descendant of *TField*.

Untyped parameters

When a formal parameter is an untyped parameter, the corresponding actual parameter can be any variable or constant reference, regardless of its type. An untyped parameter declared using the **var** keyword can be modified, whereas an untyped parameter declared using the **const** keyword is read-only.

Within the procedure or function, the untyped parameter is typeless; that is, it is incompatible with variables of all other types, unless it is given a specific type through a variable typecast.

This is an example of untyped parameters:

```
function Equal(var Source, Dest; Size: Word): Boolean;
type
  TBytes = array[0..65534] of Byte;
var
  N: Word;
begin
  N := 0;
  while (N < Size) and (TBytes(Dest)[N] = TBytes(Source)[N]) do
    Inc(N);
  Equal := N = Size;
end;
```

This function can be used to compare any two variables of any size. For example, given these declarations,

```
type
  TVector = array[1..10] of Integer;
  TPoint = record
    X, Y: Integer;
  end;
var
  Vec1, Vec2: TVector;
  N: Integer;
  P: TPoint;
```

the function then calls

```
Equal(Vec1, Vec2, SizeOf(TVector))
Equal(Vec1, Vec2, SizeOf(Integer) * N)
Equal(Vec1[1], Vec1[6], SizeOf(Integer) * 5)
Equal(Vec1[1], P, 4)
```

which compares *Vec1* to *Vec2*, the first *N* components of *Vec1* to the first *N* components of *Vec2*, the first five components of *Vec1* to the last five components of *Vec1*, and *Vec1[1]* to *P.X* and *Vec1[2]* to *P.Y*.



While untyped parameters give you greater flexibility, they can be riskier to use. The compiler can't verify that operations on untyped variables are valid.

Open parameters

Open parameters allow strings and arrays of varying sizes to be passed to the same procedure or function.

Open-string parameters

Open-string parameters can be declared in two ways:

- Using the *OpenString* identifier
- Using the **string** keyword in the **{ $\$P+$ }** state

The *OpenString* identifier is declared in the *System* unit. It denotes a special string type that can only be used in the declaration of string parameters. For reasons of backward compatibility, *OpenString* isn't a reserved word; therefore, *OpenString* can be redeclared as a user-defined identifier.

When backward compatibility isn't an issue, a **{ $\$P+$ }** compiler directive can be used to change the meaning of the **string** keyword. In the **{ $\$P+$ }** state, a variable declared using the **string** keyword is an open-string parameter.

For an open-string parameter, the actual parameter can be a variable of any string type. Within the procedure or function, the size attribute (maximum length) of the formal parameter will be the same as that of the actual parameter.

Open-string parameters behave exactly as variable parameters of a string type, except that they can't be passed as regular variable parameters to other procedures and functions. They can, however, be passed as open-string parameters again.

In this example, the *S* parameter of the *AssignStr* procedure is an open-string parameter:

```
procedure AssignStr(var S: OpenString);
begin
  S := '0123456789ABCDEF';
end;
```

Because *S* is an open-string parameter, variables of any string type can be passed to *AssignStr*:

```
var
  S1: string[10];
  S2: string[20];
```

```

begin
  AssignStr(S1);      { S1 = '0123456789' }
  AssignStr(S2);      { S2 = '0123456789ABCDEF' }
end;

```

Within *AssignStr*, the maximum length of the *S* parameter is the same as that of the actual parameter. Therefore, in the first call to *AssignStr*, the assignment to the *S* parameter truncates the string because the declared maximum length of *S1* is 10.

When applied to an open-string parameter, the *Low* standard function returns zero, the *High* standard function returns the declared maximum length of the actual parameter, and the *SizeOf* function returns the size of the actual parameter.

In the next example, the *FillString* procedure fills a string to its maximum length with a given character. Notice the use of the *High* standard function to obtain the maximum length of an open-string parameter.

```

procedure FillString(var S: OpenString; Ch: Char);
begin
  S[0] := Chr(High(S));      { Set string length }
  FillChar(S[1], High(S), Ch); { Set string characters }
end;

```



Value and constant parameters declared using the *OpenString* identifier or the **string** keyword in the **(\$P+)** state are *not* open-string parameters. Instead, such parameters behave as if they were declared using a string type with a maximum length of 255 and the *High* standard function always returns 255 for such parameters.

Open-array
parameters

A formal parameter declared using the syntax

```
array of T
```

is an *open-array parameter*. *T* must be a type identifier, and the actual parameter must be a variable of type *T*, or an array variable whose element type is *T*. Within the procedure or function, the formal parameter behaves as if it was declared as

```
array[0..N - 1] of T
```

where *N* is the number of elements in the actual parameter. In effect, the index range of the actual parameter is mapped onto the integers 0 to *N* - 1. If the actual parameter is a simple variable of type *T*, it is treated as an array with one element of type *T*.

A formal open-array parameter can be accessed by element only. Assignments to an entire open array aren't allowed, and an open array can be passed to other procedures and functions only as an open-array parameter or as an untyped variable parameter.

Open-array parameters can be value, constant, and variable parameters and have the same semantics as regular value, constant, and variable parameters. In particular, assignments to elements of a formal open array constant parameter are not allowed, and assignments to elements of a formal open array value parameter don't affect the actual parameter.



For an open array value parameter, the compiler creates a local copy of the actual parameter within the procedure or function's stack frame. Therefore, be careful not to overflow the stack when passing large arrays as open array value parameters.

When applied to an open-array parameter, the *Low* standard function returns zero, the *High* standard function returns the index of the last element in the actual array parameter, and the *SizeOf* function returns the size of the actual array parameter.

The *Clear* procedure in the next example assigns zero to each element of an array of *Real*, and the *Sum* function computes the sum of all elements in an array of *Real*. Because the *A* parameter in both cases is an open-array parameter, the subroutines can operate on any array with an element type of *Real*.

```
procedure Clear(var A: array of Real);
var
  I: Word;
begin
  for I := 0 to High(A) do A[I] := 0;
end;

function Sum(const A: array of Real): Real;
var
  I: Word;
  S: Real;
begin
  S := 0;
  for I := 0 to High(A) do S := S + A[I];
  Sum := S;
end;
```

When the element type of an open-array parameter is *Char*, the actual parameter may be a string constant. For example, given the procedure declaration,

```

procedure PrintStr(const S: array of Char);
var
  I: Integer;
begin
  for I := 0 to High(S) do
    if S[I] <> #0 then Write(S[I]) else Break;
end;

```

the following are valid procedure statements:

```

PrintStr('Hello world');
PrintStr('A');

```

When passed as an open-character array, an empty string is converted to a string with one element containing a NULL character, so the statement *PrintStr("")* is identical to the statement *PrintStr(#0)*.

Dynamic object-type variables

The *New* and *Dispose* standard procedures allow a constructor call or destructor call as a second parameter for allocating or disposing of a dynamic object-type variable. This is the syntax:

```
New(P, Construct)
```

and

```
Dispose(P, Destruct)
```

where *P* is a pointer variable, pointing to an object type, and *Construct* and *Destruct* are calls to constructors and destructors of that object type. For *New*, the effect of the extended syntax is the same as executing

```

New(P);
P^.Construct;

```

And for *Dispose*, the effect of the extended syntax is the same as executing

```

P^.Destruct;
Dispose(P);

```

Without the extended syntax, you would frequently have to call *New* followed by a constructor call or call a destructor followed by a call to *Dispose*. The extended syntax improves readability and generates shorter and more efficient code.

The following illustrates the use of the extended *New* and *Dispose* syntax:

```
var
    SP: PStrField
    ZP: PZipField
begin
    New(SP, Init(1, 1, 25, 'Firstname'));
    New(ZP, Init(1, 2, 5, 'Zip code', 0, 99999));
    SP^.Edit;
    ZP^.Edit;
    :
    Dispose(ZP, Done);
    Dispose(SP, Done);
end;
```

You can also use *New* as a *function* that allocates and returns a dynamic variable of a specified type:

```
New(T)
```

OR

```
New(T, Construct)
```

In the first form, *T* can be any pointer type. In the second form, *T* must point to an object type and *Construct* must be a call to a constructor of that object type. In both cases, the type of the function result is *T*.

Here's an example:

```
var
    F1, F2: PField
begin
    F1 := New(PStrField, Init(1, 1, 25, 'Firstname'));
    F2 := New(PZipField, Init(1, 2, 5, 'Zip code', 0, 99999));
    :
    WriteLn(F1^.GetStr);           { Calls TStrField.GetStr }
    WriteLn(F2^.GetStr);           { Calls TZipField.GetStr }
    :
    Dispose(F2, Done);             { Calls TField.Done }
    Dispose(F1, Done);            { Calls TStrField.Done }
end;
```


Notice that even though *F1* and *F2* are of type *PField*, the extended-pointer assignment-compatibility rules allow *F1* and *F2* to be assigned a pointer to any descendant of *TField*. Because *GetStr* and *Done* are virtual methods, the virtual-method

dispatch-mechanism correctly calls *TStrField.GetStr*, *TZipField.GetStr*, *TField.Done*, and *TStrField.Done*, respectively.

Programs and units


Program syntax

A Borland Pascal program consists of a program heading, an optional **uses** clause, and a block.



The program heading


The program heading specifies the program's name and its parameters.



The program heading, if present, is ignored by the compiler.

The uses clause

The **uses** clause identifies all units used by the program.



The *System* unit is always used automatically. *System* implements all low-level, run-time routines to support such features as file input and output (I/O), string handling, floating point, dynamic memory allocation, and others.

Apart from *System*, Borland Pascal implements many standard units, such as *Dos* and *Crt*. These aren't used automatically; you must include them in your **uses** clause. For example,

```
uses Dos, Crt;                               { Can now use Dos and Crt }
```


The order of the units listed in the **uses** clause determines the order of their initialization (see "The initialization part" on page 122).



To find the unit file containing a compiled unit, the compiler truncates the unit name listed in the **uses** clause to the first eight characters and adds the file extension. If your target is DOS, the extension will be .TPU. If your target is Windows, the file extension will be .TPW. If your target is DOS protected mode, the file extension is .TPP. For example, a unit named *MathFunctions* in your Windows program is stored in a file called MATHFUNC.TPW. Even though the file name is truncated, a **uses** clause must still specify the full unit identifier.

Unit syntax

Units are the basis of modular programming in Borland Pascal. They're used to create libraries you can include in various programs without making the source code available, and to divide large programs into logically related modules.




The unit heading

The unit heading specifies the unit's name.

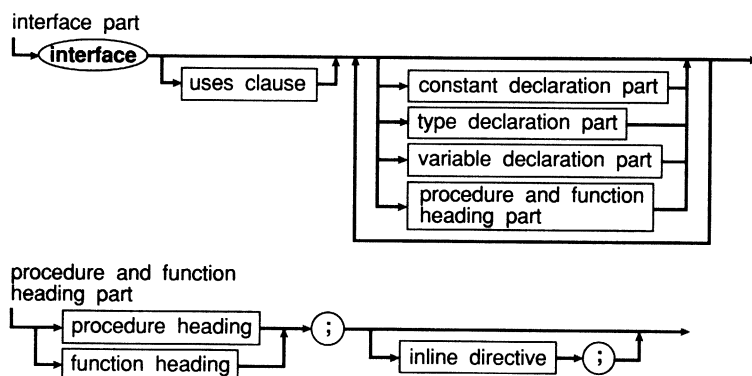


The unit name is used when referring to the unit in a **uses** clause. The name must be unique: Two units with the same name can't be used at the same time.

 The name of a unit's source file and binary file must be the same as the unit identifier, truncated to the first eight characters. If this isn't the case, the compiler can't find the source and/or binary file when compiling a program or unit that uses the unit.

The interface part


The interface part declares constants, types, variables, procedures, and functions that are *public*; that is, available to the host (the program or unit using the unit). The host can access these entities as if they were declared in a block that encloses the host.



Unless a procedure or function is **inline**, the interface part only lists the procedure or function heading. The block of the procedure or function follows in the implementation part.

The implementation part

The implementation part defines the block of all public procedures and functions. In addition, it declares constants, types, variables, procedures, and functions that are *private*; that is, not available to the host.



In effect, the procedure and function declarations in the interface part are like forward declarations, although the **forward** directive


isn't specified. Therefore, these procedures and functions can be defined and referenced in any sequence in the implementation part.



Procedure and function headings can be duplicated from the interface part. You don't have to specify the formal parameter list, but if you do, the compiler will issue a compile-time error if the interface and implementation declarations don't match.

The initialization part

The initialization part is the last part of a unit. It consists either of the reserved word **end** (in which case, the unit has no initialization code) or of a statement part to be executed to initialize the unit.



The initialization parts of units used by a program are executed in the same order that the units appear in the **uses** clause.

Indirect unit references

The **uses** clause in a module (program or unit) need only name the units used directly by that module. Consider the following:

```
program Prog;
uses Unit2;
const a = b;
begin
end.

unit Unit2;
interface
uses Unit1;
const b = c;
implementation
end.

unit Unit1;
interface
const c = 1;
implementation
const d = 2;
end.
```

In the previous example, *Unit2* is directly dependent on *Unit1* and *Prog* is directly dependent on *Unit2*. Also, *Prog* is indirectly dependent on *Unit1* (through *Unit2*), even though none of the identifiers declared in *Unit1* are available to *Prog*.

To compile a module, the compiler must be able to locate all units the module depends upon, directly or indirectly. So, to compile *Prog*, the compiler must be able to locate both *Unit1* and *Unit2*, or else an error occurs.

Note for C and other language users: The uses clauses of a Borland Pascal program provide the "make" logic information traditionally found in make or project files of other languages. With the uses clause, Borland Pascal can build all the dependency information into the module itself and reduce the chance for error.

When changes are made in the interface part of a unit, other units using the unit must be recompiled. If you use *Make* or *Build*, the compiler does this for you automatically. If changes are made only to the implementation or the initialization part, however, other units that use the unit *don't* have to be recompiled. In the previous example, if the interface part of *Unit1* is changed (for example, *c = 2*) *Unit2* must be recompiled; changing the implementation part (for example, *d = 1*) doesn't require recompilation of *Unit2*.

Borland Pascal can tell when the interface part of a unit has changed by computing a *unit version number* when the unit is compiled. In the preceding example, when *Unit2* is compiled, the current version number of *Unit1* is saved in the compiled version of *Unit2*. When *Prog* is compiled, the version number of *Unit1* is checked against the version number stored in *Unit2*. If the version numbers don't match (indicating that a change was made in the interface part of *Unit1* because *Unit2* was compiled), the compiler reports an error or recompiles *Unit2*, depending on the mode of compilation.

Circular unit references

If you place a **uses** clause in the implementation section of a unit, you hide the inner details of the unit referenced in the **uses** clause; the referenced unit is private and not available to the program or unit using the unit it's referenced in. You can use this technique to construct mutually-dependent units.

The following program demonstrates how two units can "use" each other:

```
program Circular;
{ Display text using WriteXY }

uses
  Crt, Display;
```

```

begin
  ClrScr;
  WriteXY(1, 1, 'Upper left corner of screen');
  WriteXY(1000, 1000, 'Way off the screen');
  WriteXY(81 - Length('Back to reality'), 15, 'Back to reality');
end.

```

The main program, *Circular*, uses a unit named *Display*:

```

unit Display;
{ Contains a simple video display routine }

interface

procedure WriteXY(X, Y: Integer; Message: String);

implementation
uses
  Crt, Error;

procedure WriteXY(X, Y: Integer; Message: String);
begin
  if (X in [1..80]) and (Y in [1..25]) then
    begin
      GoToXY(X, Y);
      Write(Message);
    end
  else
    ShowError('Invalid WriteXY coordinates');
  end;
end.

```

The *Display* unit declares the *WriteXY* procedure in its interface section. *WriteXY* writes a message on the screen. The program *Circular* specifies the content and screen position of the message in the parameters passed to *WriteXY*. If the screen coordinates aren't onscreen, *WriteXY* calls the *ShowError* procedure.

ShowError isn't in the *Display* unit, but in another unit, *Error*, referenced in the **uses** section of the *Display* unit's implementation section. This is the *Error* unit:

```

unit Error;
{ Contains a simple error-reporting routine }

interface

procedure ShowError(ErrMsg: String)

implementation
uses Display;

```

```

procedure ShowError(ErrMsg: String);
begin
    WriteXY(1, 25, 'Error: ' + ErrMsg);
end;

end.

```

The *Error* unit is somewhat unusual: its one declared procedure, *ShowError*, uses the *WriteXY* procedure declared in the *Display* unit, the unit that calls the *ShowError* procedure. The **uses** clause in the implementation sections of both the *Display* and *Error* units refer to each other. This is possible because Borland Pascal can compile complete interface sections for both. The compiler accepts a reference to a partially-compiled unit in the **implementation** section of another unit, as long as neither unit's interface section depends upon the other. Therefore, the units follow Pascal's strict rules for declaration order.



If the interface sections are interdependent, you get a circular unit-reference error.

Sharing other declarations

If you want to modify the *WriteXY* and *ShowError* procedures to take an additional parameter that specifies a rectangular window onscreen, you might write this:

```

procedure WriteXY(SomeWindow: WindRec; X, Y: Integer;
    Message: String);

procedure ShowError(SomeWindow: WindRec; ErrMsg: String);

```

These procedures are declared in the interface sections of different units. Because both need to use the *WindRec* type, *WindRec* can't be declared in either of the interface sections—that would make them depend on each other. The solution is to create a third unit that contains only the definition of the window record:

```

unit WindData;
interface
type
    WindRec = record
        X1, Y1, X2, Y2: Integer;
        ForeColor, BackColor: Byte;
        Active: Boolean;
    end;
implementation
end.

```

You can now add *WindData* to the **uses** clause in interface sections of both the *Display* and *Error* units. Both of these units

can use the new record type, but *Display* and *Error* still refer to each other only in their respective implementation sections.



Mutually-dependent units can be useful in special situations, but use them judiciously. If you use them when they aren't needed, they can make your program harder to maintain and more susceptible to errors.

Dynamic-link libraries

Dynamic-link libraries (DLLs) permit several Windows and DOS protected-mode applications to share code and resources. With Borland Pascal, you can use DLLs as well as write your own DLLs to be used by other applications. DOS real-mode applications can't use DLLs.

What is a DLL?

A DLL is an executable module containing code or resources for use by other applications or DLLs. Conceptually, a DLL is similar to a unit—both have the ability to provide services in the form of procedures and functions to a program. There are, however, many differences between DLLs and units. In particular, units are *statically linked*, whereas DLLs are *dynamically linked*.

When a program uses a procedure or function from a unit, a copy of that procedure or function's code is statically linked into the program's executable file. If two programs are running simultaneously and they use the same procedure or function from a unit, there will be two copies of that routine present in the system. It would be more efficient if the two programs could share a single copy of the routine. DLLs provide that ability.

In contrast to a unit, the code in a DLL isn't linked into a program that uses the DLL. Instead, a DLL's code and resources are in a separate executable file with a .DLL extension. This file must be

present when the client program runs. The procedure and function calls in the program are dynamically linked to their entry points in the DLLs used by the application.

Another difference between units and DLLs is that units can export types, constants, data, and objects whereas DLLs can export procedures and functions only.



A DLL doesn't have to be written in Borland Pascal for a Borland Pascal application to be able to use it. Also, programs written in other languages can use DLLs written in Borland Pascal. DLLs are therefore ideal for multi-language programming projects.

Using DLLs

For a module to use a procedure or function in a DLL, the module must import the procedure or function using an external declaration. For example, the following external declaration imports a function called *GlobalAlloc* from the DLL called KERNEL (the Windows kernel):

```
function GlobalAlloc(Flags: Word; Bytes: Longint): THandle; far;
external 'KERNEL' index 15;
```

See "External declarations" on page 101.

In imported procedures and functions, the **external** directive takes the place of the declaration and statement parts that would otherwise be present. Imported procedures and functions must use the far call model selected by using a **far** procedure directive or a **(\$F+)** compiler directive, but otherwise they behave no differently than normal procedures and functions.

Borland Pascal imports procedures and functions in three ways:

- By name
- By new name
- By ordinal

The format of **external** directives for each of the three methods is demonstrated in the following example.

When no **index** or **name** clause is specified, the procedure or function is imported by name. The name used is the same as the procedure or function's identifier. In this example, the *ImportByName* procedure is imported from 'TESTLIB' using the name 'IMPORTBYNAME'.

```
procedure ImportByName; external 'TESTLIB';
```


When a **name** clause is specified, the procedure or function is imported by a different name than its identifier. Here the *ImportByNewName* procedure is imported from 'TESTLIB' using the name 'REALNAME':

```
procedure ImportByNewName; external 'TESTLIB' name 'REALNAME';
```

Finally, when an **index** clause is present, the procedure or function is imported by ordinal. Importing by ordinal reduces the module's load time because the name doesn't have to be looked up in the DLL's name table. In the example, the *ImportByOrdinal* procedure is imported as the fifth entry in the 'TESTLIB' DLL:

```
procedure ImportByOrdinal; external 'TESTLIB' index 5;
```

The DLL name specified after the **external** keyword and the new name specified in a **name** clause don't have to be string literals. Any constant-string expression is allowed. Likewise, the ordinal number specified in an **index** clause can be any constant-integer expression.

```
const  
  TestLib = 'TESTLIB';  
  Ordinal = 5;  
  
procedure ImportByName; external TestLib;  
procedure ImportByNewName; external TestLib name 'REALNAME';  
procedure ImportByOrdinal; external TestLib index Ordinal;
```

Although a DLL can have variables, it's not possible to import them in other modules. Any access to a DLL's variables must take place through a procedural interface.

Import units

Declarations of imported procedures and functions can be placed directly in the program that imports them. Usually, though, they are grouped together in an *import unit* that contains declarations for all procedures and functions in a DLL, along with any constants and types required to interface with the DLL. The *WinTypes*, *WinProcs*, and *WinAPI* units supplied with Borland Pascal are examples of such import units. Import units aren't a requirement of the DLL interface, but they do simplify maintenance of projects that use multiple DLLs.

As an example, consider a DLL called DATETIME.DLL that has four routines to get and set the date and time, using a record type that contains the day, month, and year, and another record type

that contains the second, minute, and hour. Instead of specifying the associated procedure, function, and type declarations in every program that uses the DLL, you can construct an import unit to go along with the DLL. This code creates a .TPW file (assuming Windows is the compiler's target), but it doesn't contribute code or data to the programs that use it:

```
unit DateTime;

interface

type
  TTimeRec = record
    Second: Integer;
    Minute: Integer;
    Hour: Integer;
  end;

type
  TDateRec = record
    Day: Integer;
    Month: Integer;
    Year: Integer;
  end;

procedure SetTime(var Time: TTimeRec);
procedure GetTime(var Time: TTimeRec);
procedure SetDate(var Date: TDateRec);
procedure GetDate(var Date: TDateRec);

implementation

procedure SetTime; external 'DATETIME' index 1;
procedure GetTime; external 'DATETIME' index 2;
procedure SetDate; external 'DATETIME' index 3;
procedure GetDate; external 'DATETIME' index 4;

end.
```

Any program that uses DATETIME.DLL can now simply specify *DateTime* in its **uses** clause. Here is a Windows program example:

```
program ShowTime;

uses WinCrt, DateTime;

var
  Time: TTimeRec;

begin
  GetTime(Time);
  with Time do
    WriteLn('The time is ', Hour, ':', Minute, ':', Second);
end.
```

Another advantage of an import unit such as *DateTime* is that when the associated DATETIME.DLL is modified, only one unit, the *DateTime* import unit, needs updating to reflect the changes.

When you compile a program that uses a DLL, the compiler doesn't look for the DLL so it need not be present. The DLL must be present when you run the program, however.



If you write your own DLLs, they aren't automatically compiled when you choose Compile | Make on a program that uses the DLL. Instead, DLLs must be compiled separately.

Static and dynamic imports

The **external** directive provides the ability to statically import procedures and functions from a DLL. A statically-imported procedure or function always refers to the same entry point in the same DLL. Windows and Borland's protected-mode DOS extensions also support dynamic imports, whereby the DLL name and the name or ordinal number of the imported procedure or function is specified at run time. The *ShowTime* program shown here uses dynamic importing to call the *GetTime* procedure in DATETIME.DLL. Note the use of a procedural-type variable to represent the address of the *GetTime* procedure.

```
program ShowTime;
uses WinProcs, WinTypes, WinCrt;

type
  TTimeRec = record
    Second: Integer;
    Minute: Integer;
    Hour: Integer;
  end;
  TGetTime = procedure(var Time: TTimeRec);

var
  Time: TTimeRec;
  Handle: THandle;
  GetTime: TGetTime;

begin
  Handle := LoadLibrary('DATETIME.DLL');
  if Handle >= 32 then
    begin
      @GetTime := GetProcAddress(Handle, 'GETTIME');
      if @GetTime <> nil then
```


```

begin
  GetTime(Time);
  with Time do
    WriteLn('The time is ', Hour, ':', Minute, ':', Second);
  end;
  FreeLibrary(Handle);
end;
end;

```

Writing DLLs

The structure of a Borland Pascal DLL is identical to that of a program, except a DLL starts with a **library** header instead of a **program** header. The **library** header tells Borland Pascal to produce an executable file with the extension .DLL instead of .EXE, and also ensures that the executable file is marked as being a DLL.



The example here implements a very simple DLL with two exported functions, *Min* and *Max*, that calculate the smaller and larger of two integer values.

```

library MinMax;

function Min(X, Y: Integer): Integer; export;
begin
  if X < Y then Min := X else Min := Y;
end;

function Max(X, Y: Integer): Integer; export;
begin
  if X > Y then Max := X else Max := Y;
end;

exports
  Min index 1,
  Max index 2;

begin
end.

```

Note the use of the **export** procedure directive to prepare *Min* and *Max* for exporting, and the **exports** clause to actually export the

two routines, supplying an optional ordinal number for each of them.

Although the preceding example doesn't demonstrate it, libraries can and often do consist of several units. In such cases, the library source file itself is frequently reduced to a **uses** clause, an **exports** clause, and the library's initialization code. For example,

```
library Editors;
uses EdInit, EdInOut, EdFormat, EdPrint;
exports
  InitEditors index 1,
  DoneEditors index 2,
  InsertText index 3,
  DeleteSelection index 4,
  FormatSelection index 5,
  PrintSelection index 6,
  :
  SetErrorHandler index 53;
begin
  InitLibrary;
end.
```

The export procedure directive


See "Export declarations" on page 99.

If procedures and functions are to be exported by a DLL, they must be compiled with the **export** procedure directive. The **export** directive belongs to the same family of procedure directives as the **near**, **far**, **inline**, and **interrupt** directives. This means that an **export** directive, if present, must be specified upon the first introduction of procedure or function—it can't be supplied in the defining declaration of a forward declaration.

The **export** directive makes a procedure or function exportable. It forces the routine to use the far call model and prepares the routine for export by generating special procedure entry and exit code. Note, however, that the actual exporting of the procedure or function doesn't occur until the routine is listed in a library's **exports** clause.

The exports clause

A procedure or function is exported by a DLL when it's listed in the library's **exports** clause.



An **exports** clause can appear anywhere and any number of times in a program or library's declaration part. Each entry in an **exports** clause specifies the identifier of a procedure or function to be exported. That procedure or function must be declared before the **exports** clause appears, however, and its declaration must contain the **export** directive. You can precede the identifier in the **exports** clause with a unit identifier and a period; this is known as a fully qualified identifier.

The quickest way to look up a DLL entry is by index.

An exports entry can also include an **index** clause, which consists of the word **index** followed by an integer constant between 1 and 32,767. When an **index** clause is specified, the procedure or function to be exported uses the specified ordinal number. If no **index** clause is present in an exports entry, an ordinal number is automatically assigned.

An entry can also have a **name** clause, which consists of the word **name** followed by a string constant. When there is a **name** clause, the procedure or function to be exported uses the name specified by the string constant. If no **name** clause is present in an exports entry, the procedure or function is exported by its identifier and converted to all uppercase.

Finally, an exports entry can include the **resident** keyword. When **resident** is specified, the export information stays in memory while the DLL is loaded. The **resident** option significantly reduces the time it takes to look up a DLL entry by name, so if client programs that use the DLL are likely to import certain entries by name, those entries should be exported using the **resident** keyword.

A program can contain an **exports** clause, but it seldom does because Windows doesn't allow application modules to export functions for use by other applications.

Library initialization code

The statement part of a library constitutes the library's *initialization code*. The initialization code is executed once, when the library is initially loaded. When subsequent applications that use the library are loaded, the initialization code isn't executed again, but the DLL's use count is incremented.

A DLL is kept in memory as long as its use count is greater than zero. When the use count becomes zero, indicating that all applications that used the DLL have terminated, the DLL is removed from memory. At that point, the library's exit procedures are executed. Exit procedures are registered using the *ExitProc* variable, as described in Chapter 22, "Control issues."

A DLL's initialization code typically performs tasks like registering window classes for window procedures contained in the DLL and setting initial values for the DLL's global variables. The initialization code of a library can signal an error condition by setting the *ExitCode* variable to zero. (*ExitCode* is declared by the *System* unit.) *ExitCode* defaults to 1, indicating initialization was successful. If the initialization code sets *ExitCode* to zero, the DLL is unloaded from system memory and the calling application is notified of the failure to load the DLL.

When a library's exit procedures are executed, the *ExitCode* variable doesn't contain a process-termination code, as is the case with a program. Instead, *ExitCode* contains one of the values *wep_System_Exit* or *wep_Free_DLL*, which are defined in the *WinTypes* unit. *wep_System_Exit* indicates that Windows is shutting down, whereas *wep_Free_DLL* indicates that just this single DLL is being unloaded.

Here is an example of a library with initialization code and an exit procedure:

```
library Test;
{$S-}
uses WinTypes, WinProcs;
var
  SaveExit: Pointer;
procedure LibExit; far;
begin
  if ExitCode = wep_System_Exit then
```

```

begin
  :
  { System shutdown in progress }
  :
end else
begin
  :
  { DLL is being unloaded }
  :
end;
ExitProc := SaveExit;
end;

begin
  :
  { Perform DLL initialization }
  :
  SaveExit := ExitProc;      { Save old exit procedure pointer }
  ExitProc := @LibExit;     { Install LibExit exit procedure }
end.

```



In DOS protected mode, the *ExitCode* passed to a DLL exit procedure is always zero, corresponding to *wep_FREE_DLL*.

When a DLL is unloaded, an exported function called WEP in the DLL is called, if it's present. A Borland Pascal library automatically exports a WEP function, which continues to call the address stored in the *ExitProc* variable until *ExitProc* becomes *nil*. Because this works the same way exit procedures are handled in Borland Pascal programs, you can use the same exit procedure logic in both programs and libraries.



Exit procedures in a DLL *must* be compiled with stack-checking disabled (the **{S-}** state) because the operating system switches to an internal stack when terminating a DLL. Also, the operating system crashes if a run-time error occurs in a DLL exit procedure, so you must include sufficient checks in your code to prevent run-time errors.

Library programming notes

The following sections note important points you should keep in mind while working with DLLs.



Global variables in a DLL

A DLL has its own data segment and any variables declared in a DLL are private to that DLL. A DLL can't access variables declared by modules that call the DLL, and it's not possible for a DLL to export its variables for use by other modules. Such access must take place through a procedural interface.

Global memory and files in a DLL

As a rule, a DLL doesn't "own" any files that it opens or any global memory blocks that it allocates from the system. Such objects are owned by the application that (directly or indirectly) called the DLL.

When an application terminates, any open files owned by it are automatically closed, and any global memory blocks owned by it are automatically deallocated. This means that file and global memory-block handles stored in global variables in a DLL can become invalid at any time without the DLL being notified. For that reason, DLLs should refrain from making assumptions about the validity of file and global memory-block handles stored in global variables across calls to the DLL. Instead, such handles should be made parameters of the procedures and functions of the DLL, and the calling application should be responsible for maintaining them.

-  Under Windows, global memory blocks allocated with the `GMEM_DDESHARE` attribute (defined in the `WinTypes` unit) are owned by the DLL, not by the calling applications. Such memory blocks remain allocated until they are explicitly deallocated by the DLL, or until the DLL is unloaded.
-  The DOS protected-mode memory manager doesn't support shared memory blocks, and it ignores the `GMEM_DDESHARE` flag. In DOS protected mode, memory blocks allocated by a DLL are *always* owned by the application that called the DLL.

DLLs and the System unit

During a DLL's lifetime, the `HINSTANCE` variable contains the instance handle of the DLL. The `HPrevInst` and `CmdShow` variables are always zero in a DLL, as is the `PrefixSeg` variable, because a DLL doesn't have a Program Segment Prefix (PSP). `PrefixSeg` is

never zero in an application, so the test *PrefixSeg <> 0* returns *True* if the current module is an application, and *False* if the current module is a DLL.

For details about the heap manager, see Chapter 21.

To ensure proper operation of the heap manager contained in the *System* unit, the start-up code of a library sets the *HeapAllocFlags* variable to *gmem_Moveable + gmem_DDEShare*. Under Windows, this causes all memory blocks allocated via *New* and *GetMem* to be owned by the DLL instead of the applications that call the DLL.

Run-time errors in

DLLs

If a run-time error occurs in a DLL, the application that called the DLL terminates. The DLL itself isn't necessarily removed from memory at that time because other applications might still be using it.

Because a DLL has no way of knowing whether it was called from a Borland Pascal application or an application written in another programming language, it's not possible for the DLL to invoke the application's exit procedures before the application is terminated. The application is simply aborted and removed from memory. For this reason, make sure there are sufficient checks in any DLL code so such errors don't occur.

If a run-time error does occur in a DLL under Windows, the safest thing to do is to exit Windows entirely. If you simply try to modify and rebuild the faulty DLL code, when you run your program again, Windows won't load the new version if the buggy one is still in memory. Exiting Windows and then restarting Windows and Borland Pascal ensures that your corrected version of the DLL is loaded.

DLLs and stack segments

Unlike an application, a DLL doesn't have its own stack segment. Instead, it uses the stack segment of the application that called the DLL. This can create problems in DLL routines that assume that the DS and SS registers refer to the same segment, which is the case in a Windows application module.

The Borland Pascal compiler never generates code that assumes DS = SS, and none of the Borland Pascal run-time library routines make this assumption. If you write assembly language code, don't assume that SS and DS registers contain the same value.

Writing shared DLLs

Borland Pascal supports DLLs that can be shared between DOS protected mode and Windows. Shared DLLs are *binary compatible*, which means that the same .DLL file can be used by a DOS protected-mode application or a Windows application.

When compiling a shared DLL, the selected target platform must be Windows:

- In the IDE, choose Compile | Target and select Windows application in the Target dialog box.
- When using the command-line compiler, use the `/CW` switch to select Windows as the target platform.

A DLL compiled for DOS protected-mode can't be used under Windows because the DOS protected-mode run-time library uses certain DOS and DPAPI functions calls that must be avoided under Windows.

To read about the WinAPI unit, see Chapter 17, "Programming in DOS protected mode."

A shared DLL can interface with the operating system (DOS protected-mode or Windows) through the *WinAPI* unit only. The *WinAPI* unit represents the least-common denominator of DOS protected mode and Windows; other Windows interface units, such as *WinTypes* and *WinProcs*, declare a large number of API routines that aren't supported under DOS protected-mode.

It's important to note that while a shared DLL might be executing simultaneously under Windows and under DOS protected-mode in a Windows DOS box, it isn't possible to communicate between the two environments through the DLL. In reality, two copies of the DLL will be present in the system, each protected from the other in their own entirely separate memory space.

P

A

R

T

2

The run-time libraries

Overview of the run-time libraries

Borland Pascal includes run-time libraries for DOS real mode, DOS protected mode, and Windows. The most commonly used units of each run-time library are located in the TURBO.TPL (DOS real mode), TPP.TPL (DOS protected mode), and TPW.TPL (Windows) files. Additional units are distributed in separate .TPU, .TPP, and .TPW files.

- For DOS real mode, the TURBO.TPL library contains the *System*, *Overlay*, *Crt*, *Dos*, and *Printer* units. In addition, the *Graph*, *Strings*, *WinDos*, *Turbo3*, and *Graph3* units are supplied in separate .TPU files.
- For DOS protected mode, the TPP.TPL library contains the *System*, *Crt*, *Dos*, *Printer*, *Strings*, *WinDos*, and *WinAPI* units. In addition, the *Graph* unit is supplied in a separate .TPP file.
- For Windows, the TPW.TPL library contains the *System*, *Strings*, *WinTypes*, *WinProcs*, *Win31*, *WinAPI*, *WinDos*, *WinCrt*, and *WinPrn* units. A number of additional Windows support units are supplied in source code form.

In addition to the run-time libraries, Borland Pascal includes the Turbo Vision application framework for DOS real mode and DOS protected mode, and the ObjectWindows application framework for Windows. These libraries are described in the *Turbo Vision Programming Guide* and the *ObjectWindows Programming Guide*.

This chapter briefly describes each of the run-time library units.

System unit

The *System* unit implements low-level, run-time support routines for all built-in features, such as file I/O, string handling, floating point, and dynamic memory allocation. The *System* unit is used automatically by any unit or program and doesn't need to be referred to in a **uses** clause.

Dos and WinDos units

To read about the Dos and WinDos units, see Chapter 16, "Interfacing with DOS."

The *Dos* and *WinDos* units implement a number of very useful operating system and file-handling routines. None of the routines in these units are defined by Standard Pascal, so they have been placed in their own modules. For a complete description of DOS operations, refer to a DOS programmer's reference.

Crt unit

For information on the Crt unit, see page 162 in Chapter 14, "Input and output."

The *Crt* unit permits you to write programs that send their screen output directly to the BIOS or to video memory. The result is increased speed and flexibility.

WinCrt unit

For more about the WinCrt unit, see page 166 in Chapter 14, "Input and output."

The *WinCrt* unit is a text file device driver that redirects your program's output to a scrollable window. Although most of your Windows programs will create their own windows, you can use the *WinCrt* unit for quick and simple text-based programs where you are concerned with quick results and not the look of the program.

Printer unit

For more about the Printer unit, see page 161 in Chapter 14, "Input and output."

The *Printer* unit lets you send standard Pascal output to your printer using *Write* and *Writeln*.

WinPrn unit

To read about the WinPrn unit, see page 170 in Chapter 14, "Input and output."

The *WinPrn* unit permits you to send the output of your Windows program to the printer of your choice.

Overlay unit

To read about the Overlay unit, see Chapter 20, "Using overlays."

The *Overlay* unit enables you to reduce your DOS real-mode program's total run-time memory requirements. In fact, you can

write programs that are larger than the total available memory, because only parts of your program will reside in memory at any given time.

Strings unit

See page 217 in Chapter 18, "Using null-terminated strings," for information about using the Strings unit.

With Borland Pascal's extended syntax and the *Strings* unit, your programs can use null-terminated strings, so that they are more compatible with any Windows programs you write.

Graph unit

Read about the Graph unit in Chapter 19, "Using the Borland Graphics Interface."

The *Graph* unit supplies a set of fast, powerful graphics routines. It implements the device-independent Borland graphics handler that supports CGA, EGA, VGA, Hercules, AT&T 400, MCGA, 3270 PC, and 8514 graphics. The *Graph* unit isn't built into TURBO.TPL, but is on the same disk with the .BGI (Borland Graphic Interface) and .CHR files.

Turbo3 and Graph3 units

You'll find information on the Turbo3 and Graph3 units in the online file TURBO3.INT.

These units are provided for backward compatibility only. *Turbo3* contains two variables and several procedures no longer supported by Borland Pascal. *Graph3* supports the full set of graphics routines—basic, advanced, and turtlegraphics—from version 3.0.

WinTypes and WinProcs units

To learn about the details of the Windows API found in WinTypes and WinProcs, use the Borland Pascal Help system.

Together, the *WinTypes* and *WinProcs* units make up the Windows application programming interface (API). *WinTypes* contains all the constants, data structures, and styles; *WinProcs* contains all the functions and procedures.

Win31 unit

The *Win31* unit provides an interface to the additional API routines found in Windows 3.1. Applications that use the *Win31* unit don't run under Windows 3.0.

WinAPI unit

For more about the WinAPI unit, see Chapter 17, "Programming in DOS protected mode."

The *WinAPI* unit defines the subset of Windows API routines supported by both Windows and Borland's DOS protected-mode extender technology.

Windows 3.1 support units

You can find information about these units in the Borland Pascal Help system.

Borland Pascal supports the Windows 3.1 API in several units:

<i>ColorDlg</i>	<i>LZExpand</i>	<i>ShellAPI</i>
<i>CommDlg</i>	<i>MMSystem</i>	<i>Stress</i>
<i>Cpl</i>	<i>OLE</i>	<i>ToolHelp</i>
<i>DDEML</i>	<i>PenWin</i>	<i>Ver</i>
<i>Dlgs</i>	<i>Print</i>	<i>WinMem32</i>

Standard procedures and functions

This chapter briefly describes standard (built-in) procedures and functions in Borland Pascal and the predeclared variables defined in the *System* unit. For in-depth information about a particular procedure, function, or predeclared variable, look it up in the alphabetical listing in Chapter 1, “Library reference,” in the *Programmer’s Reference*.

You’ll find these topics in this chapter:

There are other standard procedures and functions also. You can read about them in Chapter 14, “Input and output.”

- Flow-control procedures
- Transfer functions
- Arithmetic functions
- Ordinal procedures and functions
- String procedures and functions
- Dynamic-allocation procedures and functions
- Pointer and address functions
- Miscellaneous procedures and functions
- Predeclared variables in the *System* unit

Standard procedures and functions are predeclared. Because all predeclared entities act as if they were declared in a block surrounding the program, you can redefine the same identifier within the program.

Flow-control procedures

These are the procedures that change the flow of logic in your program:

Table 13.1
Flow-control procedures

Procedure	Description
<i>Break</i>	Terminates a for , while , or repeat statement.
<i>Continue</i>	Continues with another iteration of a for , while , or repeat statement.
<i>Exit</i>	Exits immediately from the current block.
<i>Halt</i>	Stops program execution and returns to the operating system.
<i>RunError</i>	Stops program execution and generates a run-time error.

Transfer functions

The Transfer functions are listed here:

Table 13.2
Transfer functions

Function	Description
<i>Chr</i>	Returns a character of a specified ordinal number.
<i>High</i>	Returns the highest value in the range of the argument.
<i>Low</i>	Returns the lowest value in the range of the argument.
<i>Ord</i>	Returns the ordinal number of an ordinal-type value.
<i>Round</i>	Rounds a real-type value to a type <i>Longint</i> value.
<i>Trunc</i>	Truncates a real-type value to a type <i>Longint</i> value.

The transfer procedures *Pack* and *Unpack*, as defined in *Standard Pascal*, aren't implemented by *Borland Pascal*.

Arithmetic functions

These functions are useful in performing arithmetic operations. When you're compiling in numeric processing mode, **{SN+}**, the return values of the floating-point routines in the *System* unit (*Sqrt*, *Pi*, *Sin*, and so on) are of type *Extended* instead of *Real*.

Table 13.3
Arithmetic functions

Function	Description
<i>Abs</i>	Returns the absolute value of the argument.
<i>ArcTan</i>	Returns the arctangent of the argument.
<i>Cos</i>	Returns the cosine of the argument.
<i>Exp</i>	Returns the exponential part of the argument.
<i>Frac</i>	Returns the fractional part of the argument.
<i>Int</i>	Returns the integer part of the argument.
<i>Ln</i>	Returns the natural logarithm of the argument.

Table 13.3: Arithmetic functions (continued)

<i>Pi</i>	Returns the value of <i>Pi</i> (3.1415926535897932385).
<i>Sin</i>	Returns the sine of the argument.
<i>Sqr</i>	Returns the square of the argument.
<i>Sqrt</i>	Returns the square root of the argument.

Ordinal procedures and functions

The ordinal routines operate on the ordinality of a variable:

Table 13.4
Ordinal procedures and functions

Procedure or function	Description
<i>Dec</i>	Decrements a variable.
<i>Inc</i>	Increments a variable.
<i>Odd</i>	Tests if the argument is an odd number.
<i>Pred</i>	Returns the predecessor of the argument.
<i>Succ</i>	Returns the successor of the argument.

String procedures and functions

These procedures and functions are used on the traditional Pascal-style strings.

Table 13.5
String procedures and functions

Procedure or function	Description
<i>Concat</i>	Concatenates a sequence of strings.
<i>Copy</i>	Returns a substring of a string.
<i>Delete</i>	Deletes a substring from a string.
<i>Insert</i>	Inserts a substring into a string.
<i>Length</i>	Returns the dynamic length of a string.
<i>Pos</i>	Searches for a substring in a string.
<i>Str</i>	Converts a numeric value to its string representation.
<i>Val</i>	Converts the string value to its numeric representation.

Dynamic-allocation procedures and functions

The dynamic-allocation procedures and functions are used to manage the *heap*—a memory area that occupies all or some of the free memory left when a program is executed. Heap-management techniques are discussed on page 259 for DOS programs and on page 273 for Windows programs.

Table 13.6
Dynamic-allocation
procedures and functions

Procedure or function	Description
<i>Dispose</i>	Disposes of a dynamic variable.
<i>FreeMem</i>	Disposes of a dynamic variable of a given size.
<i>GetMem</i>	Creates a new dynamic variable of a given size and sets a pointer variable to point to it.
<i>MaxAvail</i>	Returns the size of the largest contiguous free block in the heap, indicating the size of the largest dynamic variable that can be allocated at the time of the call to <i>MaxAvail</i> .
<i>MemAvail</i>	Returns the number of free bytes of heap storage available.
<i>New</i>	Creates a new dynamic variable and sets a pointer variable to point to it.

Pointer and address
functions

Table 13.7
Pointer and address
functions

The pointer and address functions are listed in this table:

Function	Description
<i>Addr</i>	Returns the address of a specified object.
<i>Assigned</i>	Tests to determine if a pointer or procedural variable is nil .
<i>CSeg</i>	Returns the current value of the CS register.
<i>DSeg</i>	Returns the current value of the DS register.
<i>Ofs</i>	Returns the offset of a specified object.
<i>Ptr</i>	Converts a segment base and an offset address to a pointer-type value.
<i>Seg</i>	Returns the segment of a specified object.
<i>SPtr</i>	Returns the current value of the SP register.
<i>SSeg</i>	Returns the current value of the SS register.

Miscellaneous
standard procedures
and functions

These standard procedures and functions don't fit neatly in any other category.

Table 13.8
Miscellaneous standard
procedures and functions

Procedure or function	Description
<i>Exclude</i>	Excludes an element from a set.
<i>FillChar</i>	Fills a specified number of contiguous bytes with a specified value.
<i>Hi</i>	Returns the high-order byte of the argument.
<i>Include</i>	Includes an element in a set.
<i>Lo</i>	Returns the low-order byte of the argument.
<i>Move</i>	Copies a specified number of contiguous bytes from a source range to a destination range.
<i>ParamCount</i>	Returns the number of parameters passed to the program on the command line.
<i>ParamStr</i>	Returns a specified command-line parameter.
<i>Random</i>	Returns a random number.
<i>Randomize</i>	Initializes built-in random generator with a random value.
<i>SizeOf</i>	Returns number of bytes occupied by the argument.
<i>Swap</i>	Swaps the high- and low-order bytes of the argument.
<i>TypeOf</i>	Points to an object type's virtual method table.
<i>UpCase</i>	Converts a character to uppercase.

Predeclared variables

The *System* unit also supplies several predeclared variables. The list varies depending on which run-time library the *System* unit belongs to.

These predeclared variables are in the *System* unit in TURBO.TPL, the run-time library for DOS real-mode applications:

Table 13.9
Variables in the DOS
real-mode System unit

Variable	Type	Description
<i>ErrorAddr</i>	<i>Pointer</i>	Run-time error address
<i>ExitCode</i>	<i>Integer</i>	Exit code
<i>ExitProc</i>	<i>Pointer</i>	Exit procedure
<i>FileMode</i>	<i>Byte</i>	File open mode
<i>FreeList</i>	<i>Pointer</i>	Free heap block list
<i>FreeZero</i>	<i>Pointer</i>	Must be zero
<i>HeapEnd</i>	<i>Pointer</i>	Heap end
<i>HeapError</i>	<i>Pointer</i>	Heap-error function
<i>HeapOrg</i>	<i>Pointer</i>	Heap origin
<i>HeapPtr</i>	<i>Pointer</i>	Heap pointer
<i>Input</i>	<i>Text</i>	Input standard file
<i>InOutRes</i>	<i>Integer</i>	I/O result buffer

Table 13.9: Variables in the DOS real-mode System unit (continued)

<i>Output</i>	<i>Text</i>	Output standard file
<i>OvrCodeList</i>	<i>Word</i>	Overlay code segment list
<i>OvrDebugPtr</i>	<i>Pointer</i>	Overlay debugger hook
<i>OvrDosHandle</i>	<i>Word</i>	Overlay DOS handle
<i>OvrEmsHandle</i>	<i>Word</i>	Overlay EMS handle
<i>OvrHeapEnd</i>	<i>Word</i>	Overlay buffer end
<i>OvrHeapOrg</i>	<i>Word</i>	Overlay buffer origin
<i>OvrHeapPtr</i>	<i>Word</i>	Overlay buffer pointer
<i>OvrHeapSize</i>	<i>Word</i>	Initial overlay buffer size
<i>OvrLoadList</i>	<i>Word</i>	Loaded overlays list
<i>PrefixSeg</i>	<i>Word</i>	Program Segment Prefix
<i>RandSeed</i>	<i>Longint</i>	Random seed
<i>SaveInt00</i>	<i>Pointer</i>	Saved exception \$00
<i>SaveInt02</i>	<i>Pointer</i>	Saved interrupt \$02
<i>SaveInt1B</i>	<i>Pointer</i>	Saved interrupt \$1B
<i>SaveInt21</i>	<i>Pointer</i>	Saved interrupt \$21
<i>SaveInt23</i>	<i>Pointer</i>	Saved interrupt \$23
<i>SaveInt24</i>	<i>Pointer</i>	Saved interrupt \$24
<i>SaveInt34</i>	<i>Pointer</i>	Saved interrupt \$34
<i>SaveInt35</i>	<i>Pointer</i>	Saved interrupt \$35
<i>SaveInt36</i>	<i>Pointer</i>	Saved interrupt \$36
<i>SaveInt37</i>	<i>Pointer</i>	Saved interrupt \$37
<i>SaveInt38</i>	<i>Pointer</i>	Saved interrupt \$38
<i>SaveInt39</i>	<i>Pointer</i>	Saved interrupt \$39
<i>SaveInt3A</i>	<i>Pointer</i>	Saved interrupt \$3A
<i>SaveInt3B</i>	<i>Pointer</i>	Saved interrupt \$3B
<i>SaveInt3C</i>	<i>Pointer</i>	Saved interrupt \$3C
<i>SaveInt3D</i>	<i>Pointer</i>	Saved interrupt \$3D
<i>SaveInt3E</i>	<i>Pointer</i>	Saved interrupt \$3E
<i>SaveInt3F</i>	<i>Pointer</i>	Saved interrupt \$3F
<i>SaveInt75</i>	<i>Pointer</i>	Saved interrupt \$75
<i>Seg0040</i>	<i>Word</i>	Selector for segment \$0040
<i>SegA000</i>	<i>Word</i>	Selector for segment \$A000
<i>SegB000</i>	<i>Word</i>	Selector for segment \$B000
<i>SegB800</i>	<i>Word</i>	Selector for segment \$B800
<i>SelectorInc</i>	<i>Word</i>	Selector increment
<i>StackLimit</i>	<i>Word</i>	Minimum stack pointer
<i>Test8086</i>	<i>Byte</i>	80x86 test result
<i>Test8087</i>	<i>Byte</i>	80x87 test result

These variables are in the *System* unit in TPW.TPL, the run-time library for Windows applications:

Table 13.10
Variables in the Windows
System unit

Variable	Type	Description
<i>CmdLine</i>	<i>PChar</i>	Command-line pointer
<i>CmdShow</i>	<i>Integer</i>	<i>CmdShow</i> parameter for <i>CreateWindow</i>
<i>ErrorAddr</i>	<i>Pointer</i>	Run-time error address
<i>ExitCode</i>	<i>Integer</i>	Exit code
<i>ExitProc</i>	<i>Pointer</i>	Exit procedure

Table 13.10: Variables in the Windows System unit (continued)

<i>FileMode</i>	<i>Byte</i>	File open mode
<i>Input</i>	<i>Text</i>	Input standard file
<i>HeapAllocFlags</i>	<i>Word</i>	Heap-block allocation flags
<i>HeapBlock</i>	<i>Word</i>	Heap-block size
<i>HeapError</i>	<i>Pointer</i>	Heap-error function
<i>HeapLimit</i>	<i>Word</i>	Heap small-block limit
<i>HeapList</i>	<i>Word</i>	Heap-segment list
<i>HInstance</i>	<i>Word</i>	Handle of this instance
<i>HPrevInst</i>	<i>Word</i>	Handle of previous instance
<i>InOutRes</i>	<i>Integer</i>	I/O result buffer
<i>Output</i>	<i>Text</i>	Output standard file
<i>PrefixSeg</i>	<i>Word</i>	Program Segment Prefix
<i>RandSeed</i>	<i>Longint</i>	Random seed
<i>SelectorInc</i>	<i>Word</i>	Selector increment
<i>Test8086</i>	<i>Byte</i>	80x86 test result

These variables are in the *System* unit in TPP.TPL, the run-time library for DOS protected-mode applications:

Table 13.11
Variables in the DOS
protected-mode System unit

Variable	Type	Description
<i>ErrorAddr</i>	<i>Pointer</i>	Run-time error address
<i>ExitCode</i>	<i>Integer</i>	Exit code
<i>ExitProc</i>	<i>Pointer</i>	Exit procedure
<i>FileMode</i>	<i>Byte</i>	File open mode
<i>HeapAllocFlags</i>	<i>Word</i>	Heap-allocation flags, <i>gmem_Moveable</i>
<i>HeapBlock</i>	<i>Word</i>	Heap-block size
<i>HeapError</i>	<i>Pointer</i>	Heap-error function
<i>HeapLimit</i>	<i>Word</i>	Heap small-block limit
<i>HeapList</i>	<i>Word</i>	Heap-segment list
<i>HInstance</i>	<i>Word</i>	Module-instance handler
<i>Input</i>	<i>Text</i>	Input standard file
<i>InOutRes</i>	<i>Integer</i>	I/O result buffer
<i>Output</i>	<i>Text</i>	Output standard file
<i>PrefixSeg</i>	<i>Word</i>	Program segment prefix
<i>RandSeed</i>	<i>Longint</i>	Random seed
<i>RealModeRegs</i>	array[0..49] of byte	Real mode registers
<i>SaveInt00</i>	<i>Pointer</i>	Saved exception \$00
<i>SaveInt02</i>	<i>Pointer</i>	Saved interrupt \$02
<i>SaveInt0C</i>	<i>Pointer</i>	Saved exception \$0C
<i>SaveInt0D</i>	<i>Pointer</i>	Saved exception \$0D
<i>SaveInt1B</i>	<i>Pointer</i>	Saved interrupt \$1B
<i>SaveInt21</i>	<i>Pointer</i>	Saved interrupt \$21
<i>SaveInt23</i>	<i>Pointer</i>	Saved real-mode interrupt \$23
<i>SaveInt24</i>	<i>Pointer</i>	Saved real-mode interrupt \$24
<i>SaveInt34</i>	<i>Pointer</i>	Saved interrupt \$34
<i>SaveInt35</i>	<i>Pointer</i>	Saved interrupt \$35
<i>SaveInt36</i>	<i>Pointer</i>	Saved interrupt \$36
<i>SaveInt37</i>	<i>Pointer</i>	Saved interrupt \$37

Table 13.11: Variables in the DOS protected-mode System unit (continued)

<i>SaveInt38</i>	<i>Pointer</i>	Saved interrupt \$38
<i>SaveInt39</i>	<i>Pointer</i>	Saved interrupt \$39
<i>SaveInt3A</i>	<i>Pointer</i>	Saved interrupt \$3A
<i>SaveInt3B</i>	<i>Pointer</i>	Saved interrupt \$3B
<i>SaveInt3C</i>	<i>Pointer</i>	Saved interrupt \$3C
<i>SaveInt3D</i>	<i>Pointer</i>	Saved interrupt \$3D
<i>SaveInt3E</i>	<i>Pointer</i>	Saved interrupt \$3E
<i>SaveInt3F</i>	<i>Pointer</i>	Saved interrupt \$3F
<i>SaveInt75</i>	<i>Pointer</i>	Saved interrupt \$75
<i>Seg0040</i>	<i>Word</i>	Selector for segment \$0040
<i>SegA000</i>	<i>Word</i>	Selector for segment \$A000
<i>SegB000</i>	<i>Word</i>	Selector for segment \$B000
<i>SegB800</i>	<i>Word</i>	Selector for segment \$B800
<i>Test8086</i>	<i>Byte</i>	80x86 test result
<i>Test8087</i>	<i>Byte</i>	80x87 test result

For more information about these variables, look them up in the alphabetical listing in Chapter 1, “Library reference,” in the *Programmer’s Reference*.

Input and output

This chapter describes the standard (or built-in) input and output (I/O) procedures and functions of Borland Pascal. You'll find them in the *System* unit.

Table 14.1
Input and output procedures
and functions

Procedure or function	Description
<i>Append</i>	Opens an existing text file for appending.
<i>Assign</i>	Assigns the name of an external file to a file variable.
<i>BlockRead</i>	Reads one or more records from an untyped file.
<i>BlockWrite</i>	Writes one or more records into an untyped file.
<i>ChDir</i>	Changes the current directory.
<i>Close</i>	Closes an open file.
<i>Eof</i>	Returns the end-of-file status of a file.
<i>Eoln</i>	Returns the end-of-line status of a text file.
<i>Erase</i>	Erases an external file.
<i>FilePos</i>	Returns the current file position of a typed or untyped file.
<i>FileSize</i>	Returns the current size of a file; not used for text files.
<i>Flush</i>	Flushes the buffer of an output text file.
<i>GetDir</i>	Returns the current directory of a specified drive.
<i>IOResult</i>	Returns an integer value that is the status of the last I/O function performed.
<i>MkDir</i>	Creates a subdirectory.

Table 14.1: Input and output procedures and functions (continued)

<i>Read</i>	Reads one or more values from a file into one or more variables.
<i>ReadLn</i>	Does what a <i>Read</i> does and then skips to the beginning of the next line in the text file.
<i>Rename</i>	Renames an external file.
<i>Reset</i>	Opens an existing file.
<i>Rewrite</i>	Creates and opens a new file.
<i>RmDir</i>	Removes an empty subdirectory.
<i>Seek</i>	Moves the current position of a typed or untyped file to a specified component. Not used with text files.
<i>SeekEof</i>	Returns the end-of-file status of a text file.
<i>SeekEoln</i>	Returns the end-of-line status of a text file.
<i>SetTextBuf</i>	Assigns an I/O buffer to a text file.
<i>Truncate</i>	Truncates a typed or untyped file at the current file position.
<i>Write</i>	Writes one or more values to a file.
<i>Writeln</i>	Does the same as a <i>Write</i> , and then writes an end-of-line marker to the text file.

File input and output

A Pascal file variable is any variable whose type is a file type. There are three classes of Pascal files: *typed*, *text*, and *untyped*.

The syntax for writing file types is given on page 42.

Before a file variable can be used, it must be associated with an external file through a call to the *Assign* procedure. An external file is typically a named disk file, but it can also be a device, such as the keyboard or the display. The external file stores the information written to the file or supplies the information read from the file.

Once the association with an external file is established, the file variable must be “opened” to prepare it for input or output. An existing file can be opened via the *Reset* procedure, and a new file can be created and opened via the *Rewrite* procedure. Text files opened with *Reset* are read-only and text files opened with *Rewrite* and *Append* are write-only. Typed files and untyped files always allow both reading and writing regardless of whether they were opened with *Reset* or *Rewrite*.

Every file is a linear sequence of components, each of which has the component type (or record type) of the file. Each component has a component number. The first component of a file is considered to be component zero.

Files are normally accessed *sequentially*; that is, when a component is read using the standard procedure *Read* or written using the standard procedure *Write*, the current file position moves to the next numerically ordered file component. Typed files and untyped files can also be accessed randomly, however, via the standard procedure *Seek*, which moves the current file position to a specified component. The standard functions *FilePos* and *FileSize* can be used to determine the current file position and the current file size.

When a program completes processing a file, the file must be closed using the standard procedure *Close*. After a file is closed, its associated external file is updated. The file variable can then be associated with another external file.

If you're writing a Windows program, and you don't want Windows to handle disk or other I/O errors for you, call `SetErrorMode(1)`.

By default, all calls to standard I/O procedures and functions are automatically checked for errors: If an error occurs, the program terminates, displaying a run-time error message. This automatic checking can be turned on and off using the **{SI+}** and **{SI-}** compiler directives. When I/O checking is off—that is, when a procedure or function call is compiled in the **{SI-}** state—an I/O error doesn't cause the program to halt. To check the result of an I/O operation, you must call the standard function *IOResult* instead.



You must call the *IOResult* function to clear whatever error may have occurred, even if you aren't interested in the error. If you don't and **{SI+}** is the current state, the next I/O function call fails with the lingering *IOResult* error.

Text files

In Borland Pascal, the type `Text` is distinct from the type `file of Char`.

This section summarizes I/O using file variables of the standard type *Text*.

When a text file is opened, the external file is interpreted in a special way: It is considered to represent a sequence of characters formatted into lines, where each line is terminated by an end-of-line marker (a carriage-return character, possibly followed by a linefeed character).

For text files, there are special forms of *Read* and *Write* that let you read and write values that are not of type *Char*. Such values are

automatically translated to and from their character representation. For example, *Read(F, I)*, where *I* is a type *Integer* variable, reads a sequence of digits, interprets that sequence as a decimal integer, and stores it in *I*.

Borland Pascal defines two standard text-file variables, *Input* and *Output*. The standard file variable *Input* is a read-only file associated with the operating system's standard input file (typically the keyboard). The standard file variable *Output* is a write-only file associated with the operating system's standard output file (typically the display). Before a DOS program begins running, *Input* and *Output* are automatically opened, as if the following statements were executed:

```
Assign(Input, '');  
Reset(Input);  
Assign(Output, '');  
Rewrite(Output);
```

See page 166 for more about the *WinCrt* unit.

Because Windows doesn't directly support text-oriented I/O, the I/O files are unassigned in a Windows application, and any attempt to read or write to them will produce an I/O error. If a Windows application uses the *WinCrt* unit, however, *Input* and *Output* refer to a scrollable text window. *WinCrt* contains the complete control logic required to emulate a text screen in the Windows environment, and no Windows-specific programming is required in an application that uses *WinCrt*.

Some of the standard I/O routines that work on text files don't need to have a file variable explicitly given as a parameter. If the file parameter is omitted, *Input* or *Output* is assumed by default, depending on whether the procedure or function is input- or output-oriented. For example, *Read(X)* corresponds to *Read(Input, X)* and *Write(X)* corresponds to *Write(Output, X)*.

If you do specify a file when calling one of the input or output routines that work on text files, the file must be associated with an external file using *Assign*, and opened using *Reset*, *Rewrite*, or *Append*. A run-time error occurs if you pass a file that was opened with *Reset* to an output-oriented procedure or function. It's also an error to pass a file that was opened with *Rewrite* or *Append* to an input-oriented procedure or function.

Untyped files

Untyped files are low-level I/O channels primarily used for direct access to any disk file regardless of type and structuring. An untyped file is declared with the word **file** and nothing more. For example,

```
var
  DataFile: file;
```

For untyped files, the *Reset* and *Rewrite* procedures allow an extra parameter to specify the record size used in data transfers. For historical reasons, the default record size is 128 bytes. A record size of 1 is the only value that correctly reflects the exact size of any file (no partial records are possible when the record size is 1).

Except for *Read* and *Write*, all typed-file standard procedures and functions are also allowed on untyped files. Instead of *Read* and *Write*, two procedures called *BlockRead* and *BlockWrite* are used for high-speed data transfers.

The FileMode variable

*New files created using
Rewrite are always opened
in read/write mode,
corresponding to
FileMode = 2.*

The *FileMode* variable defined by the *System* unit determines the access code to pass to DOS when typed and untyped files (not text files) are opened using the *Reset* procedure.

The default *FileMode* is 2, which allows both reading and writing. Assigning another value to *FileMode* causes all subsequent *Resets* to use that mode.

The range of valid *FileMode* values depends on the version of DOS in use. For all versions, the following modes are defined:

- 0 Read only
- 1 Write only
- 2 Read/Write

DOS version 3.x and higher defines additional modes, which are primarily concerned with file-sharing on networks. (For more details, see your DOS programmer's reference manual.)

Devices in Borland Pascal

Borland Pascal and the DOS operating system regard external hardware, such as the keyboard, the display, and the printer, as *devices*. From the programmer's point of view, a device is treated as a file, and is operated on through the same standard procedures and functions as files.

Borland Pascal supports two kinds of devices: DOS devices and text file devices.

DOS devices

Use DOS devices in DOS programs only. Use the device I/O functions of the Windows API for Windows programs.

DOS devices are implemented through reserved file names that have a special meaning attached to them. DOS devices are completely transparent—in fact, Borland Pascal isn't even aware when a file variable refers to a device instead of a disk file. For example, the DOS program

```
var
  Lst: Text;
begin
  Assign(Lst, 'LPT1');
  Rewrite(Lst);
  Writeln(Lst, 'Hello World...');
  Close(Lst);
end.
```

writes the string "Hello World..." on the printer, even though the syntax for doing so is exactly the same as for a disk file.

The devices implemented by DOS are used for obtaining input or for presenting legible output for a DOS program. Therefore, DOS devices are normally used only in connection with text files. On rare occasions, untyped files can also be useful for interfacing with DOS devices.

The CON device

CON refers to the CONsole device, in which output is sent to the display, and input is obtained from the keyboard. The *Input* and *Output* standard files and all files assigned an empty name refer to the CON device when input or output isn't redirected.

Input from the CON device is line-oriented and uses the line-editing facilities described in your DOS manual. Characters are read from a line buffer, and when the buffer becomes empty, a new line is input.

An end-of-file character is generated by pressing *Ctrl+Z*, after which the *Eof* function will return *True*.

The LPT1, LPT2, and LPT3 devices

The line printer devices are the three possible printers you can use. If only one printer is connected, it is usually referred to as LPT1, for which the synonym PRN can also be used.

The line printer devices are output-only devices—an attempt to *Reset* a file assigned to one of these generates an immediate end-of-file.

To read about printing from a Windows program, see page 170.

The standard unit *Printer* declares a text-file variable called *Lst*, and makes it refer to the LPT1 device. To easily write something on the printer from one of your DOS programs, include *Printer* in the program's **uses** clause, and use *Write(Lst,...)* and *Writeln(Lst,...)* to produce your output.

The COM1 and COM2 devices

The communication port devices are the two serial communication ports. The synonym AUX can be used instead of COM1.

The NUL device

The NUL device ignores anything written to it, and generates an immediate end-of-file when read from. You should use this when you don't want to create a particular file, but the program requires an input or output file name.



In general, you should avoid using DOS devices under Windows and you should use the device I/O functions provided by the Windows API instead. Some devices, such as CON, won't function properly. Other devices might work, but the results might not be as you expect. For example, if you use LPT1, your printout might appear in the middle of another print job. It's always safest to use the Windows API.

Text-file devices

Text-file devices are used to implement devices unsupported by DOS or to provide another set of features similar to those supplied by another DOS device. A good example of a text-file device is the CRT device implemented by the *Crt* standard unit. It provides an interface to the display and the keyboard, like the CON device in DOS, but the CRT device is much faster and supports such invaluable features as color and windows.

Unlike DOS devices, text-file devices have no reserved file names; in fact, they have no file names at all. Instead, a file is associated

with a text-file device through a customized *Assign* procedure. For example, the *Crt* standard unit implements an *AssignCrt* procedure that associates text files with the CRT device.

Input and output with the Crt unit

DOS real- and protected-mode programs only

The *Crt* unit implements a range of powerful routines that give your DOS real-mode and DOS protected-mode programs full control of your PC's features, such as screen mode control, extended keyboard codes, colors, windows, and sound. *Crt* can only be used in programs that run on IBM PCs, ATs, PS/2s, and true compatibles.

One of the major advantages to using *Crt* is the added speed and flexibility of screen output operations. Programs that don't use the *Crt* unit send their screen output through DOS, which adds a lot of overhead. With the *Crt* unit, output is sent directly to the BIOS or, for even faster operation, directly to video memory.

Using the Crt unit

To use the *Crt* unit, include it in your program's **uses** clause as you would any other unit:

```
uses Crt;
```

The initialization code of the *Crt* unit assigns the *Input* and *Output* standard text files to refer to the CRT instead of to DOS's standard input and output files. These statements execute at the beginning of a program:

```
AssignCrt(Input); Reset(Input);  
AssignCrt(Output); Rewrite(Output);
```

This means that I/O redirection of the *Input* and *Output* files is no longer possible unless these files are explicitly assigned back to standard input and output by executing this:

```
Assign(Input, ''); Reset(Input);  
Assign(Output, ''); Rewrite(Output);
```

CRT windows

Crt supports a simple yet powerful form of windows. The *Window* procedure lets you define a window anywhere on the screen. When you write in such a window, the window behaves exactly as if you were using the entire screen, leaving the rest of the screen untouched. In other words, the screen outside the window isn't accessible. Inside the window, lines can be inserted and deleted, the cursor wraps around at the right edge, and the text scrolls when the cursor reaches the bottom line.

All screen coordinates, except the ones used to define a window, are relative to the current window, and screen coordinates (1,1) correspond to the upper left corner of the window.

The default window is the entire screen.

Special characters

When writing to *Output* or a file that has been assigned to the CRT, the following control characters have special meanings:

Table 14.2
Crt unit special characters

Char	Name	Description
#7	BELL	Emits a beep from the internal speaker.
#8	BS	Moves the cursor left one column. If the cursor is already at the left edge of the current window, nothing happens.
#10	LF	Moves the cursor down one line. If the cursor is already at the bottom of the current window, the window is scrolled up one line.
#13	CR	Returns the cursor to the left edge of the current window.

Line input

When reading from *Input* or from a text file that has been assigned to *Crt*, text is input one line at a time. The line is stored in the text file's internal buffer, and when variables are read, this buffer is used as the input source. Whenever the buffer has been emptied, a new line is input.

When entering lines, the following editing keys are available:

Table 14.3
Crt unit editing keys

Editing key	Description
<i>Backspace</i>	Deletes the last character entered.
<i>Esc</i>	Deletes the entire input line.
<i>Enter</i>	Terminates the input line and stores the end-of-line marker (carriage return/linefeed) in the buffer.
<i>Ctrl+S</i>	Same as <i>Backspace</i> .
<i>Ctrl+D</i>	Recalls one character from the last input line.
<i>Ctrl+A</i>	Same as <i>Esc</i> .
<i>Ctrl+F</i>	Recalls the last input line.
<i>Ctrl+Z</i>	Terminates the input line and generates an end-of-file marker.

Ctrl+Z will only generate an end-of-file marker if the *CheckEOF* variable has been set to *True*; it is *False* by default.

To test keyboard status and input single characters under program control, use the *KeyPressed* and *ReadKey* functions.

Crt procedures and functions

Table 14.4
Crt unit procedures and functions

The following table lists the procedures and functions defined in the *Crt* unit.

Procedure or function	Description
<i>AssignCrt</i>	Associates a text file with the CRT window.
<i>ClrEol</i>	Clears all the characters from the cursor position to the end of the line.
<i>ClrScr</i>	Clears the screen and returns cursor to the upper left-hand corner.
<i>Delay</i>	Delays a specified number of milliseconds.
<i>DelLine</i>	Deletes the line containing the cursor and moves all lines below that line one line up. The bottom line is cleared.
<i>GotoXY</i>	Positions the cursor. X is the horizontal position. Y is the vertical position.
<i>HighVideo</i>	Selects high-intensity characters.
<i>InsLine</i>	Inserts an empty line at the cursor position.
<i>KeyPressed</i>	Returns <i>True</i> if a key has been pressed on the keyboard.

See the *Programmer's Reference* for more details about using the *Crt* procedures and functions.

Table 14.4: Crt unit procedures and functions (continued)

<i>LowVideo</i>	Selects low-intensity characters.
<i>NormVideo</i>	Selects normal characters.
<i>NoSound</i>	Turns off the internal speaker.
<i>Sound</i>	Starts the internal speaker.
<i>TextBackground</i>	Selects the background color.
<i>TextColor</i>	Selects the foreground character color.
<i>TextMode</i>	Selects a specific text mode.
<i>Window</i>	Defines a text window onscreen.
<i>ReadKey</i>	Reads a character from the keyboard.
<i>WhereX</i>	Returns the x-coordinate of the current cursor location, relative to the current window.
<i>WhereY</i>	Returns the y-coordinate of the current cursor location, relative to the current window.

Crt unit constants and variables

The *Crt* unit has several constants that your programs can use. To learn more about using them, look them up in Chapter 1 of the *Programmer's Reference*. You'll find them grouped like this:

Table 14.5
Crt unit constants

Constant group	Description
Crt mode constants	Graphics-mode constants used as parameters for the <i>TextMode</i> procedure.
Text-color constants	Constants used to set the colors of the CRT window using the <i>TextColor</i> and <i>TextBackground</i> procedures.

For example, to find the value of a constant that will color the text in your program red, look up Text Color Constants, and you'll discover that the constant *Red* has a value of 4.

These are the variables in the *Crt* unit and their functions:

Table 14.6
Crt unit variables

Variable	Description
<i>CheckBreak</i>	Enables and disables checks for <i>Ctrl+Break</i> .
<i>CheckEOF</i>	Enables and disables the end-of-file character.
<i>CheckSnow</i>	Enables and disables "snow checking."
<i>DirectVideo</i>	Enables and disables direct memory access for <i>Write</i> and <i>Writeln</i> statements that output to the screen.

Table 14.6: Crt unit variables (continued)

<i>LastMode</i>	Stores the current video mode when each time <i>TextMode</i> is called.
<i>TextAttr</i>	Stores the currently-selected text attributes.
<i>WindMin</i>	Stores the screen coordinates of the upper-left corner of the current window.
<i>WindMax</i>	Stores the screen coordinates of the lower-right corner of the current window.

Input and output with the WinCrt unit

Windows programs only

The *WinCrt* unit implements a terminal-like text screen in a window. With *WinCrt*, you can easily create a Windows program that uses the *Read*, *Readln*, *Write*, and *Writeln* standard procedures to perform input and output operations just as you would in a traditional text-mode application. *WinCrt* contains the complete control logic required to emulate a text screen in the Windows environment. You don't need to write "Windows-specific" code if your program uses *WinCrt*.

Using the WinCrt unit

To use the *WinCrt* unit, simply include it in your program's **uses** clause, just as you would any other unit:

```
uses WinCrt;
```

By default, the *Input* and *Output* standard text files defined in the *System* unit are unassigned, and any *Read*, *Readln*, *Write*, or *Writeln* procedure call without a file variable causes a run-time error. But when a program uses the *WinCrt* unit, the initialization code of the unit assigns the *Input* and *Output* standard text files to refer to a window that emulates a text screen. It's as if the following statements are executed at the beginning of your program:

```
AssignCrt (Input); Reset (Input);
AssignCrt (Output); Rewrite (Output);
```

When the first *Read*, *Readln*, *Write*, or *Writeln* call executes in the program, a CRT window opens on the Windows desktop. The default title of a CRT window is the full path of the program's .EXE file. When the program finishes (when control reaches the final **end** reserved word), the title of the CRT window is changed

to “(Inactive *nnnnn*)”, where *nnnnn* is the title of the window in its active state.



Even though the program has finished, the window stays up so that the user can examine the program’s output. Just like any other Windows application, the program doesn’t completely terminate until the user closes the window.

The *InitWinCrt* and *DoneWinCrt* routines give you greater control over the CRT window’s life cycle. A call to *InitWinCrt* immediately creates the CRT window rather than waiting for the first call to *Read*, *Readln*, *Write*, or *Writeln*. Likewise, calling *DoneWinCrt* immediately destroys the CRT window instead of when the user closes it.

The CRT window is a scrollable panning window on a virtual text screen. The default dimensions of the virtual text screen are 80 columns by 25 lines, but the actual size of the CRT window may be less. If the size is less, the user can use the window’s scroll bars or the cursor keys to move this panning window over the larger text screen. This is particularly useful for scrolling back to examine previously written text. By default, the panning window tracks the text screen cursor. In other words, the panning window automatically scrolls to ensure that the cursor is always visible. You can disable the autotracking feature by setting the *AutoTracking* variable to *False*.

The dimensions of the virtual text screen are determined by the *ScreenSize* variable. You can change the virtual screen dimensions by assigning new dimensions to *ScreenSize* before your program creates the CRT window. When the window is created, a screen buffer is allocated in dynamic memory. The size of this buffer is *ScreenSize.X* multiplied by *ScreenSize.Y*, and it can’t be larger than 65,520 bytes. It’s up to you to ensure that the values you assign to *ScreenSize.X* and *ScreenSize.Y* don’t overflow this limit. If, for example, you assign 64 to *ScreenSize.X*, the largest allowable value for *ScreenSize.Y* is 1,023.

At any time while running a program that uses the *WinCrt* unit, the user can terminate the application by choosing the Close command on the CRT window’s Control menu, double-clicking the Control-menu box, or pressing *Alt+F4*. The user can also press *Ctrl+C* or *Ctrl+Break* at any time to halt the application and force the window into its inactive state; you can disable these features by setting the *CheckBreak* variable to *False* at the beginning of the program.

Special characters When writing to *Output* or a file that has been assigned to the CRT window, the following control characters have special meanings:

Table 14.7
Special characters in the
WinCrt window

Char	Name	Description
#7	BELL	Emits a beep from the internal speaker.
#8	BS	Moves the cursor left one column and erases the character at that position. If the cursor is already at the left edge of the screen, nothing happens.
#10	LF	Moves the cursor down one line. If the cursor is already at the bottom of the virtual screen, the screen is scrolled up one line.
#13	CR	Returns the cursor to the left edge of the screen.

Line input When your program reads from *Input* or a file that has been assigned to the CRT window, text is input one line at a time. The line is stored in the text file's internal buffer, and when variables are read, this buffer is used as the input source. When the buffer empties, a new line is read and stored in the buffer.

When entering lines in the CRT window, the user can use the *Backspace* key to delete the last character entered. Pressing *Enter* terminates the input line and stores an end-of-line marker (CR/LF) in the buffer. In addition, if the *CheckEOF* variable is set to *True*, a *Ctrl+Z* also terminates the input line and generates an end-of-file marker. *CheckEOF* is *False* by default.

To test keyboard status and input single characters under program control, use the *KeyPressed* and *ReadKey* functions.

WinCrt procedures and functions

Table 14.8
WinCrt procedures and
functions

The following tables list the procedures and functions defined by the *WinCrt* unit.

Procedure or function	Description
<i>AssignCrt</i>	Associates a text file with the CRT window.
<i>ClrEol</i>	Clears all the characters from the cursor position to the end of the line.

Table 14.8: WinCrt procedures and functions (continued)

See the Programmer's Reference for more details about using the WinCrt procedures and functions.

<i>ClrScr</i>	Clears the screen and returns cursor to the upper left-hand corner.
<i>CursorTo</i>	Moves the cursor to the given coordinates within the virtual screen. The origin coordinates are 0,0.
<i>DoneWinCrt</i>	Destroys the CRT window.
<i>GotoXY</i>	Moves the cursor to the given coordinates within the virtual screen. The origin coordinates are 1,1.
<i>InitWinCrt</i>	Creates the CRT window.
<i>KeyPressed</i>	Returns <i>True</i> if a key has been pressed on the keyboard.
<i>ReadBuf</i>	Inputs a line from the CRT window.
<i>ReadKey</i>	Reads a character from the keyboard.
<i>ScrollTo</i>	Scrolls the CRT window to show a screen location.
<i>TrackCursor</i>	Scrolls the CRT window to keep the cursor visible.
<i>WhereX</i>	Returns the x-coordinate of the current cursor location. The origin coordinates are 1,1.
<i>WhereY</i>	Returns the y-coordinate of the current cursor location. The origin coordinates are 1,1.
<i>WriteBuf</i>	Writes a block of characters to the CRT window.
<i>WriteChar</i>	Writes a single character to the CRT window.

WinCrt unit variables

Table 14.9
WinCrt variables

See the Programmer's Reference for more details about using the WinCrt variables.

The *WinCrt* unit declares several variables:

Variable	Description
<i>WindowOrg</i>	Determines the initial location of the CRT window.
<i>WindowSize</i>	Determines the initial size of the CRT window.
<i>ScreenSize</i>	Determines the width and height in characters of the virtual screen within the CRT window.
<i>Cursor</i>	Contains the current position of the cursor within the virtual screen. <i>Cursor</i> is 0,0 based.
<i>Origin</i>	Contains the virtual screen coordinates of the character cell displayed in the upper left corner of the CRT window. <i>Origin</i> is 0,0 based.
<i>InactiveTitle</i>	Points to a null-terminated string to use when constructing the title of an inactive CRT window.
<i>AutoTracking</i>	Enables and disables the automatic scrolling of the window to keep the cursor visible.

Table 14.9: WinCrt variables (continued)

<i>CheckEOF</i>	Enables and disables the end-of-file character.
<i>CheckBreak</i>	Enables and disables user termination of an application.
<i>WindowTitle</i>	Determines the title of the CRT window.

Printing from a Windows program

The *WinPrn* unit allows you to print text from your Windows programs. To use *WinPrn*, place *WinPrn* in your program's **uses** clause:

```
uses WinPrn;
```

Before you begin printing, you need to assign a text-file variable to the printer. Do this in one of two ways—assign the default printer, or select a specific printer, driver, and port. To write to the default printer, call *AssignDefPrn* followed by a call to *Rewrite*. To select a specific printer, call *AssignPrn*. Any writes to the assigned text-file variable are then directed to the printer.

Changing titles

By default, the Windows Print Manager will display all printing jobs started with *WinPrn* as “Untitled.” You can provide a title by calling the *TitlePrn* procedure followed by a call to *Rewrite*. For example,

```
AssignDefPrn(Prn);
TitlePrn(Prn, 'End of the year report');
Rewrite(Prn);
```

sets up the default printer as the one to write to, changes the title to “End of the year report,” and writes to the default printer. If *TitlePrn* is called after *Rewrite*, it has no effect.

Changing fonts

WinPrn uses the default font that is returned by the device driver. To change the font, call the *SetPrnFont* function, passing it the handle to the font you want to use. *SetPrnFont* returns the current font in use. Use the returned font in a future call to *SetPrnFont* or

pass it to *DeleteObject*. Here is an example program that shows how to change fonts:

```
program Test;
uses WinTypes, WinProcs, WinCrt, WinPrn;

var
  Prn: Text;
  OldFont: HFont;

begin
  Writeln('Printing...');
  AssignDefPrn(Prn);
  Rewrite(Prn);

  Writeln(Prn, 'Some text');
  OldFont := SetPrnFont(Prn, CreateFont(100, 0, 0, 0, 0, 0, 0, 0,
    1, Out_Default_Precis, Clip_Default_Precis, Default_Quality,
    ff_Roman, nil));
  Writeln(Prn, 'Some text in a new font');
  DeleteObject(SetPrnFont(Prn, OldFont));
  Writeln(Prn, 'Back to the old font');

  Close(Prn);
  Writeln('Done');
end.
```

Stopping a print job

To stop a print job started with *WinPrn*, call the *AbortPrn* procedure. This terminates printing and resets the device to prepare for the next print job.

Special characters

These characters have a special meaning while your program uses the *WinPrn* unit.

Table 14.10
Special characters in the
WinPrn unit

Char	Name	Description
#9	TAB	Begins printing characters at the next tab stop, which is spaced 8 times the average font width apart from the next tab stop.
#10	LF	Begins printing on a new line.
#12	FF	Forces a page break.
#13	CR	Begins printing at the beginning of the line.

WinPrn procedures and functions

Table 14.11
WinPrn procedures and
functions

The following tables list the procedures and functions available in the *WinPrn* unit.

Procedure or function	Description
<i>AbortPrn</i>	Terminates a print job dropping all unprinted data.
<i>AssignPrn</i>	Assigns a file to a printer.
<i>AssignDefPrn</i>	Assigns a file to the default printer.
<i>SetPrnFont</i>	Begins printing a file with a chosen font.
<i>TitlePrn</i>	Titles the file being printed.

Text-file device drivers

Borland Pascal lets you define your own text-file device drivers for both your DOS and Windows programs. A text-file device driver is a set of four functions that completely implement an interface between Borland Pascal's file system and some device.

The four functions that define each device driver are *Open*, *InOut*, *Flush*, and *Close*. The function header of each function is

```
function DeviceFunc(var F: TextRec): Integer;
```

where *TextRec* (or *TTextRec* for a Windows program) is the text-file record-type defined on page 276. Each function must be compiled in the **{SF+}** state to force it to use the far call model. The return value of a device-interface function becomes the value returned by *IOResult*. If the return value is zero, the operation was successful.

To associate the device-interface functions with a specific file, you must write a customized *Assign* procedure (like the *AssignCrt* procedure in the *Crt* unit). The *Assign* procedure must assign the addresses of the four device-interface functions to the four function pointers in the text-file variable. In addition, it should store the *fnClosed* "magic" constant in the *Mode* field, store the size of the text-file buffer in *BufSize*, store a pointer to the text-file buffer in *BufPtr*, and clear the *Name* string.

Assuming, for example, that the four device-interface functions are called *DevOpen*, *DevInOut*, *DevFlush*, and *DevClose*, the *Assign* procedure might look like this:

```
procedure AssignDev(var F: Text);
begin
  with TextRec(F) do
  begin
    Mode := fmClosed;
    BufSize := SizeOf(Buffer);
    BufPtr := @Buffer;
    OpenFunc := @DevOpen;
    InOutFunc := @DevInOut;
    FlushFunc := @DevFlush;
    CloseFunc := @DevClose;
    Name[0] := #0;
  end;
end;
```

The device-interface functions can use the *UserData* field in the file record to store private information. This field isn't modified by the Borland Pascal file system at any time.

The Open function

The *Open* function is called by the *Reset*, *Rewrite*, and *Append* standard procedures to open a text file associated with a device. On entry, the *Mode* field contains *fmInput*, *fmOutput*, or *fmInOut* to indicate whether the *Open* function was called from *Reset*, *Rewrite*, or *Append*.

The *Open* function prepares the file for input or output, according to the *Mode* value. If *Mode* specified *fmInOut* (indicating that *Open* was called from *Append*), it must be changed to *fmOutput* before *Open* returns.

Open is always called before any of the other device-interface functions. For that reason, *AssignDev* only initializes the *OpenFunc* field, leaving initialization of the remaining vectors up to *Open*. Based on *Mode*, *Open* can then install pointers to either input- or output-oriented functions. This saves the *InOut*, *Flush* functions and the *Close* from determining the current mode.

The InOut function

The *InOut* function is called by the *Read*, *Readln*, *Write*, *Writeln*, *Eof*, *Eoln*, *SeekEof*, *SeekEoln*, and *Close* standard procedures and functions whenever input or output from the device is required.

When *Mode* is *fmInput*, the *InOut* function reads up to *BufSize* characters into *BufPtr*[^], and returns the number of characters read in *BufEnd*. In addition, it stores zero in *BufPos*. If the *InOut* function returns zero in *BufEnd* as a result of an input request, *Eof* becomes *True* for the file.

When *Mode* is *fmOutput*, the *InOut* function writes *BufPos* characters from *BufPtr*[^], and returns zero in *BufPos*.

The Flush function

The *Flush* function is called at the end of each *Read*, *Readln*, *Write*, and *Writeln*. It can optionally flush the text-file buffer.

If *Mode* is *fmInput*, the *Flush* function can store zero in *BufPos* and *BufEnd* to flush the remaining (unread) characters in the buffer. This feature is seldom used.

If *Mode* is *fmOutput*, the *Flush* function can write the contents of the buffer exactly like the *InOut* function, which ensures that text written to the device appears on the device immediately. If *Flush* does nothing, the text doesn't appear on the device until the buffer becomes full or the file is closed.

The Close function

The *Close* function is called by the *Close* standard procedure to close a text file associated with a device. (The *Reset*, *Rewrite*, and *Append* procedures also call *Close* if the file they are opening is already open.) If *Mode* is *fmOutput*, then before calling *Close*, Borland Pascal's file system calls the *InOut* function to ensure that all characters have been written to the device.

Using the 80x87

There are two kinds of numbers you can work with in Borland Pascal: integers (*Shortint, Integer, Longint, Byte, Word*) and reals (*Real, Single, Double, Extended, Comp*). Reals are also known as floating-point numbers. The 80x86 family of processors is designed to handle integer values easily, but handling reals is considerably more difficult. To improve floating-point performance, the 80x86 family of processors has a corresponding family of math coprocessors, the 80x87s.

The 80x87 is a special hardware numeric processor that can be installed in your PC. It executes floating-point instructions very quickly, so if you use floating point often, you'll probably want a numeric coprocessor or a 80486 processor, which has a numeric coprocessor built in.

Borland Pascal provides optimal floating-point performance whether or not you have an 80x87:

- For programs running on any PC, with or without an 80x87, Borland Pascal provides the *Real* type and an associated library of software routines that handle floating-point operations. The *Real* type occupies 6 bytes of memory, providing a range of 2.9×10^{-39} to 1.7×10^{38} with 11 to 12 significant digits. The software floating-point library is optimized for speed and size, trading in some of the fancier features provided by the 80x87 processor.
- If you need the added precision and flexibility of the 80x87, you can instruct Borland Pascal to produce code that uses the 80x87 chip. This gives you access to four additional real types (*Single,*

Double, Extended, and Comp), and an *Extended* floating-point range of 3.4×10^{-4951} to 1.1×10^{4932} with 19 to 20 significant digits.

You switch between the two different models of floating-point code generation using the **\$N** compiler directive or the 80x87 Code check box in the Options | Compiler dialog box. The default state is **{\$N-}**, and in this state, the compiler uses the 6-byte floating-point library, allowing you to operate only on variables of type *Real*. In the **{\$N+}** state, the compiler generates code for the 80x87, giving you increased precision and access to the four additional real types.

In a DOS real-mode or DOS protected-mode application, even if you don't have an 80x87 in your machine, you can instruct Borland Pascal to include a run-time library that *emulates* the numeric coprocessor. Then, if an 80x87 is present, it's used. If an 80x87 isn't present, the run-time library emulates an 80x87, although your program runs a bit slower.

The **\$E** compiler directive and the Emulation check box in the Options | Compiler dialog box are used to enable and disable 80x87 emulation. The default state is **{\$E+}**, and in this state, the full 80x87 emulator is automatically included in programs that use the 80x87. In the **{\$E-}** state, a substantially smaller floating-point library is used, and the resulting application can run only on machines with an 80x87.



The **\$E** compiler directive has no effect in a Windows application. Likewise, **\$E** has no effect if used in a unit; it only applies to the compilation of a DOS real-mode or DOS protected-mode program. Also, if the program is compiled in the **{\$N-}** state, and if all the units used by the program were compiled with **{\$N-}**, then an 80x87 run-time library isn't required, and the **\$E** compiler directive is ignored.

WIN87EM.DLL is part of Windows, not Borland Pascal.

Windows applications don't need an 80x87 run-time library—instead, they use the WIN87EM.DLL support library that comes with Windows. It provides the necessary interface between the 80x87 processor, Windows, and your application. If an 80x87 processor isn't present in your system, WIN87EM.DLL will *emulate* it in software. Emulation is substantially slower than the real 80x87 processor, but it does guarantee that an application using the 80x87 can be run on any machine.



When you run a Windows application compiled in the **{\$N+}** state, be sure that Windows can find the WIN87EM.DLL file on your

system. Even if you do have an 80x87 processor in your system, the WIN87EM.DLL emulator library must still be present when running Windows programs compiled in the **{\$N+}** state.

Important! When you're compiling in 80x87 Code mode, **{\$N+}**, the return values of the floating-point routines in the *System* unit (*Sqrt*, *Pi*, *Sin*, and so on) are of type *Extended* instead of *Real*:

```
{$N+}
begin
  Writeln(Pi);                               { 3.14159265358979E+0000 }
end.

{$N-}
begin
  Writeln(Pi)                                { 3.1415926536E+00 }
end.
```

The remainder of this chapter discusses special issues concerning Borland Pascal programs that use the 80x87 coprocessor.

The 80x87 data types

For programs that use the 80x87, Borland Pascal provides four floating-point types in addition to the type *Real*.

- The *Single* type is the smallest format you can use with floating-point numbers. It occupies 4 bytes of memory, providing a range of 1.5×10^{-45} to 3.4×10^{38} with 7 to 8 significant digits.
- The *Double* type occupies 8 bytes of memory, providing a range of 5.0×10^{-324} to 1.7×10^{308} with 15 to 16 significant digits.
- The *Extended* type is the largest floating-point type supported by the 80x87. It occupies 10 bytes of memory, providing a range of 3.4×10^{-4932} to 1.1×10^{4932} with 19 to 20 significant digits. Any arithmetic involving real-type values is performed with the range and precision of the *Extended* type.
- The *Comp* type stores integral values in 8 bytes, providing a range of $-2^{63}+1$ to $2^{63}-1$, which is approximately -9.2×10^{18} to 9.2×10^{18} . *Comp* may be compared to a double-precision *Longint*, but it's considered a real type because all arithmetic done with *Comp* uses the 80x87 coprocessor. *Comp* is well suited for representing monetary values as integral values of cents or mills (thousandths) in business applications.

Whether or not you have an 80x87 processor, the 6-byte *Real* type is always available, so don't have to modify your source code when switching to the 80x87, and you can still read data files generated by programs that use software floating point.

Note, however, that 80x87 floating-point calculations on variables of type *Real* are slightly slower than on other types. Because the 80x87 can't directly process the *Real* format, calls must be made to library routines to convert *Real* values to *Extended* before operating on them. If you're concerned with optimum speed and always run on a system with an 80x87, you might want to use the *Single*, *Double*, *Extended*, and *Comp* types exclusively.

Extended range arithmetic

The *Extended* type is the basis of all floating-point computations with the 80x87. Borland Pascal uses the *Extended* format to store all non-integer numeric constants and evaluates all non-integer numeric expressions using extended precision. The entire right side of the following assignment, for example, is computed in *Extended* before being converted to the type on the left side:

```
{ $N+ }  
var  
  X , A , B , C: Real;  
begin  
  X := ( B + Sqrt ( B * B - A * C ) ) / A;  
end;
```

Borland Pascal automatically performs computations using the precision and range of the *Extended* type. The added precision means smaller round-off errors, and the additional range means overflow and underflow are less common.

You can go beyond Borland Pascal's automatic *Extended* capabilities. For example, you can declare variables used for intermediate results to be of type *Extended*. The following example computes a sum of products:

```
var  
  Sum: Single;  
  X, Y: array[1..100] of Single;  
  I: Integer;  
  T: Extended; { For intermediate results }  
begin  
  T := 0.0;
```

```

for I := 1 to 100 do
  begin
    X[I] := I;
    Y[I] := I;
    T := T + X[I] * Y[I];
  end;
  Sum := T;
end;

```

Had *T* been declared *Single*, the assignment to *T* would have caused a round-off error at the limit of single precision at each loop entry. But because *T* is *Extended*, all round-off errors are at the limit of extended precision, except for the one resulting from the assignment of *T* to *Sum*. Fewer round-off errors mean more accurate results.

You can also declare formal value parameters and function results to be of type *Extended*. This avoids unnecessary conversions between numeric types, which can result in loss of accuracy. For example,

```

function Area(Radius: Extended): Extended;
begin
  Area := Pi * Radius * Radius;
end;

```

Comparing reals

Because real-type values are approximations, the results of comparing values of different real types aren't always as expected. For example, if *X* is a variable of type *Single* and *Y* is a variable of type *Double*, then these statements are *False*:

```

X := 1 / 3;
Y := 1 / 3;
Writeln(X = Y);

```

This is because *X* is accurate only to 7 to 8 digits, where *Y* is accurate to 15 to 16 digits, and when both are converted to *Extended*, they will differ after 7 to 8 digits. Similarly, these statements,

```

X := 1 / 3;
Writeln(X = 1 / 3);

```

are *False*, because the result of $1/3$ in the *Writeln* statement is calculated with 20 significant digits.

The 80x87 evaluation stack

The 80x87 coprocessor has an internal evaluation stack that can be as deep as eight levels. Accessing a value on the 80x87 stack is much faster than accessing a variable in memory. To achieve the best possible performance, Borland Pascal uses the 80x87's stack for storing temporary results.

In theory, very complicated real-type expressions can overflow the 80x87 stack, but this isn't likely to occur because the expression would need to generate more than eight temporary results.

A more tangible danger lies in recursive function calls. If such constructs aren't coded correctly, they can overflow the 80x87 stack.

Consider the following procedure that calculates Fibonacci numbers using recursion:

```
function Fib(N: Integer): Extended;
begin
  if N = 0 then
    Fib := 0.0
  else
    if N = 1 then
      Fib := 1.0
    else
      Fib := Fib(N - 1) + Fib(N - 2);
end;
```

A call to this version of *Fib* will cause an 80x87 stack overflow for values of *N* larger than 8. This is because the calculation of the last assignment requires a temporary on the 80x87 stack to store the result of *Fib(N-1)*. Each recursive invocation allocates one such temporary, causing an overflow the ninth time. This is the correct construct:

```
function Fib(N: Integer): Extended;
var
  F1, F2: Extended;
```

```

begin
  if N = 0 then
    Fib := 0.0
  else
    if N = 1 then
      Fib := 1.0
    else
      begin
        F1 := Fib(N - 1);
        F2 := Fib(N - 2);
        Fib := F1 + F2;
      end;
    end;
  end;
end;

```

The temporary results are now stored in variables allocated on the 8086 stack. (The 8086 stack can also overflow, but this would usually require many more recursive calls.)

Writing reals with the 80x87

In the **{\$N+}** state, the *Write* and *Writeln* standard procedures output four digits, not two, for the exponent in a floating-point decimal string to provide for the extended numeric range. The *Str* standard procedure also returns a four-digit exponent when floating-point format is selected.

Units using the 80x87

Units that use the 80x87 can be used only by other units or programs that are compiled in the **{\$N+}** state.

The fact that a unit uses the 80x87 is determined by whether it contains 80x87 instructions—not by the state of the **\$N** compiler directive at the time of its compilation. This makes the compiler more forgiving in cases where you accidentally compile a unit that doesn't use the 80x87 in the **{\$N+}** state.



When you compile in numeric processing mode (**{\$N+}**), the return values of the floating-point routines in the *System* unit—*Sqrt*, *Pi*, *Sin*, and so on—are of type *Extended* instead of *Real*.

Detecting the 80x87 in a DOS program

The Borland Pascal 80x87 run-time library built into your DOS program compiled with **{\$N+}** includes startup code that automatically detects the presence of an 80x87 chip. If an 80x87 is available, the program will use it. If not, the program will use the emulation run-time library. If the program was compiled in the **{\$E-}** state, and an 80x87 can't be detected at startup, the program displays "Numeric coprocessor required," and ends.

You might want to override this default autodetection behavior occasionally. For example, your own system might have an 80x87, but you want to verify that your program will work as intended on systems without a coprocessor. Or your program might need to run on a PC-compatible system, but that particular system returns incorrect information to the autodetection logic (saying that an 80x87 is present when it's not, or vice versa).

Borland Pascal provides an option for overriding the startup code's default autodetection logic: the *87* environment variable.

You set the *87* environment variable at the DOS prompt with the SET command, like this:

```
SET 87 = Y
```

or

```
SET 87 = N
```

Setting the *87* environment variable to *N* (for no) tells the startup code that you don't want to use the 80x87, even though it might be present in the system. Conversely, setting the *87* environment variable to *Y* (for yes) means that the coprocessor is there, and you want the program to use it.



If you set *87* = *Y* when there is no 80x87 available, your program will either crash or hang!

If the *87* environment variable has been defined (to any value) but you want to undefine it, enter this at the DOS prompt:

```
SET 87 =
```

If an *87* = *Y* entry is present in the DOS environment, or if the autodetection logic succeeds in detecting a coprocessor, the startup code executes additional checks to determine what kind of coprocessor it's (8087, 80287, or 80387). This is required so that

Borland Pascal can correctly handle certain incompatibilities that exist between the different coprocessors.

The result of the autodetection and the coprocessor classification is stored in the *Test8087* variable (which is declared by the *System* unit). The following values are defined:

Table 15.1
Test8087 variable values

Value	Definition
0	No coprocessor detected
1	8087 detected
2	80287 detected
3	80387 or 80486 detected

Your program can examine the *Test8087* variable to determine the characteristics of the system it's running on. In particular, *Test8087* can be examined to determine if floating-point instructions are being emulated or truly executed.

Detecting the 80x87 in a Windows program

The Windows environment and the WIN87EM.DLL emulator library automatically detect the presence of an 80x87 chip. If an 80x87 is available in your system, it's used. If not, WIN87EM.DLL emulates it in software. You can use the *GetWinFlags* function (defined in the *WinProcs* unit) and the *wf_80x87* bit mask (defined in the *WinTypes* unit) to determine whether an 80x87 processor is present in your system. For example,

```
if GetWinFlags and wf_80x87 <> 0 then
  WriteLn('80x87 is present') else
  WriteLn('80x87 is not present');
```

Emulation in assembly language

When linking in object files using **{*\$L filename*}** directives, make sure that these object files were compiled with the 80x87 emulation enabled. For example, if you're using 80x87 instructions in assembly language **external** procedures, enable emulation when you assemble the .ASM files into .OBJ files. Otherwise, the 80x87 instructions can't be emulated on machines without an 80x87. Use Turbo Assembler's **/E** command-line switch to enable emulation.

Interfacing with DOS

The *Dos* and *WinDos* units implement a number of operating system and file-handling routines. None of the routines in the *Dos* and *WinDos* units are defined by Standard Pascal, so they have been placed in their own modules.

For a complete description of DOS operations, refer to a DOS programmer's reference manual.

Read more about the differences between standard Pascal-style and null-terminated strings on page 217.

The primary difference between the *Dos* and *WinDos* units is that the procedures and functions of the *Dos* unit use standard Pascal-style strings and the *WinDos* procedures and functions use null-terminated strings. A standard Pascal-style string is a length byte followed by a sequence of characters. A null-terminated string is a sequence of non-null characters followed by a NULL (#0) character.

If you write Windows applications, use the *WinDos* unit.

To read about the Strings unit, see Chapter 18, "Using null-terminated strings."

If you develop only DOS applications, you'll probably want to use the *Dos* unit for the programs you write, as most Pascal programs traditionally use the Pascal-style strings. If you also develop applications for the Windows environment, however, you'll be able to write code you can more easily share between the DOS and Windows platforms if you use the *WinDos* unit along with the *Strings* unit; Windows requires the use of null-terminated strings. You also might want to use the *WinDos* and *Strings* units if you have a C data file you want to use or convert. The C language uses null-terminated strings.

This chapter discusses the *Dos* unit first. To read about the *WinDos* unit, turn to page 189.

Dos unit procedures and functions

These are the procedures and functions in the *Dos* unit. To use them, you must refer to the *Dos* unit with the **uses** statement in your program. See Chapter 1, "Library reference" in the *Programmer's Reference*.

Table 16.1
Dos date and time
procedures

Procedure	Description
<i>GetDate</i>	Returns the current date set in the operating system.
<i>GetFTime</i>	Returns the date and time a file was last modified.
<i>GetTime</i>	Returns the current time set in the operating system.
<i>PackTime</i>	Converts a <i>DateTime</i> record into a 4-byte, packed date-and-time <i>Longint</i> used by <i>SetFTime</i> .
<i>SetDate</i>	Sets the current date in the operating system.
<i>SetFTime</i>	Sets the date and time a file was last modified.
<i>SetTime</i>	Sets the current time in the operating system.
<i>UnpackTime</i>	Converts a 4-byte, packed date-and-time <i>Longint</i> returned by <i>GetFTime</i> , <i>FindFirst</i> , or <i>FindNext</i> into an unpacked <i>DateTime</i> record.

Table 16.2
Dos interrupt support
procedures

Procedure	Description
<i>GetIntVec</i>	Returns the address stored in a specified interrupt vector.
<i>Intr</i>	Executes a specified software interrupt with a specified <i>Registers</i> package.
<i>MsDos</i>	Executes a DOS function call with a specified <i>Registers</i> package.
<i>SetIntVec</i>	Sets a specified interrupt vector to a specified address.

Table 16.3
Dos disk status functions

Functions	Description
<i>DiskFree</i>	Returns the number of free bytes of a specified disk drive.
<i>DiskSize</i>	Returns the total size in bytes of a specified disk drive.

Table 16.4
Dos file-handling procedures and functions

Procedure or function	Description
<i>FExpand</i>	Takes a file name and returns a fully qualified file name (drive, directory, name, and extension).
<i>FSearch</i>	Searches for a file in a list of directories.
<i>FSplit</i>	Splits a file name into its three component parts (drive and directory, file name, and extension).
<i>FindFirst</i>	Searches the specified directory for the first entry matching the specified file name and set of attributes.
<i>FindNext</i>	Returns the next entry that matches the name and attributes specified in a previous call to <i>FindFirst</i> .
<i>GetFAttr</i>	Returns the attributes of a file.
<i>SetFAttr</i>	Sets the attributes of a file.

Table 16.5
Dos environment-handling functions

Functions	Description
<i>EnvCount</i>	Returns the number of strings contained in the DOS environment.
<i>EnvStr</i>	Returns a specified environment string.
<i>GetEnv</i>	Returns the value of a specified environment variable.

Table 16.6
Dos process-handling procedures

Procedure	Description
<i>Exec</i>	Executes a specified program with a specified command line.
<i>Keep</i>	<i>Keep</i> (or Terminate Stay Resident) terminates the program and makes it stay in memory.
<i>SwapVectors</i>	Swaps all saved interrupt vectors with the current vectors.

Table 16.7
Dos miscellaneous
procedures and functions

Procedure or function	Description
<i>DosVersion</i>	Returns the DOS version number.
<i>GetCBreak</i>	Returns the state of <i>Ctrl+Break</i> checking in DOS.
<i>GetVerify</i>	Returns the state of the verify flag in DOS.
<i>SetCBreak</i>	Sets the state of <i>Ctrl-Break</i> checking in DOS.
<i>SetVerify</i>	Sets the state of the verify flag in DOS.

Dos unit constants, types, and variables

Each of the constants, types, and variables defined by the *Dos* unit are briefly discussed in this section. For more information, look them up in Chapter 1, "Library reference," in the *Programmer's Reference*.

Constants

The *Dos* unit defines uses several constants. These constants can be grouped by their function. To learn more about these constants, look them up as part of the group they belong to. For example, to find the value of *FParity*, look up "Flag constants" in the *Programmer's Reference*.

Table 16.8
Dos constants

Constant group	Description
Flag	Used to test individual flag bits in the Flags register after a call to <i>Intr</i> or <i>MsDos</i> : <i>FCarry</i> , <i>FParity</i> , <i>FAuxiliary</i> , <i>FZero</i> , <i>FSign</i> , <i>FOverflow</i>
fmXXXX	Defines the allowable values for <i>Mode</i> field of a <i>TextRec</i> text file record: <i>fmClosed</i> , <i>fmInput</i> , <i>fmOutput</i> , <i>fmInOut</i>
File attribute	Used to construct file attributes for use with <i>FindFirst</i> , <i>GetFAttr</i> , and <i>SetFAttr</i> : <i>ReadOnly</i> , <i>Hidden</i> , <i>SysFile</i> , <i>VolumeID</i> , <i>Directory</i> , <i>Archive</i> , <i>AnyFile</i>

Types

The *Dos* unit defines these types:

Table 16.9
Dos types

Types	Description
File record types	<i>FileRec</i> defines the internal data format for both typed and untyped files; <i>TextRec</i> is the internal format of a variable of type <i>Text</i> .
<i>Registers</i>	Variables of this type are used by <i>Intr</i> and <i>MsDos</i> to specify the input register contents and examine the output register contents of a software interrupt.
<i>DateTime</i>	Variables of this type are used to examine and construct 4-byte, packed date-and-time values for <i>GetFTime</i> , <i>SetFTime</i> , <i>FindFirst</i> , and <i>FindNext</i> .
<i>SearchRec</i>	Variables of this type are used by <i>FindFirst</i> and <i>FindNext</i> to scan directories.
File-handling string types	String types used by various procedures and functions in the <i>Dos</i> unit: <i>ComStr</i> , <i>PathStr</i> , <i>DirStr</i> , <i>NameStr</i> , <i>ExtStr</i> .

Variables *DosError* is used by many of the routines in the *Dos* unit to report errors.

WinDos unit procedures and functions

These are the procedures and functions in the *WinDos* unit. To use them, you must refer to the *WinDos* unit with the **uses** statement in your program:

Table 16.10
WinDos date and time
procedures

Procedure	Description
<i>GetDate</i>	Returns the current date set in the operating system.
<i>GetFTime</i>	Returns the date and time a file was last modified.
<i>GetTime</i>	Returns the current time set in the operating system.
<i>PackTime</i>	Converts a <i>TDateTime</i> record into a 4-byte, packed date-and-time <i>Longint</i> used by <i>SetFTime</i> .
<i>SetDate</i>	Sets the current date in the operating system.
<i>SetFTime</i>	Sets the date and time a file was last modified.

Table 16.10: WinDos date and time procedures (continued)

<i>SetTime</i>	Sets the current time in the operating system.
<i>UnpackTime</i>	Converts a 4-byte, packed date-and-time <i>Longint</i> returned by <i>GetFTime</i> , <i>FindFirst</i> , or <i>FindNext</i> into an unpacked <i>TDateTime</i> record.

Table 16.11
WinDos interrupt support
procedures

*Don't use these functions
when running Windows in
protected mode.*

Procedure	Description
<i>GetIntVec</i>	Returns the address stored in a specified interrupt vector.
<i>Intr</i>	Executes a specified software interrupt with a specified <i>TRegisters</i> package.
<i>MsDos</i>	Executes a DOS function call with a specified <i>TRegisters</i> package.
<i>SetIntVec</i>	Sets a specified interrupt vector to a specified address.

Table 16.12
WinDos disk status functions

Functions	Description
<i>DiskFree</i>	Returns the number of free bytes of a specified disk drive.
<i>DiskSize</i>	Returns the total size in bytes of a specified disk drive.

Table 16.13
File-handling procedures and
functions

Procedure or function	Description
<i>FileExpand</i>	Takes a file name and returns a fully qualified file name (drive, directory, name, and extension).
<i>FileSearch</i>	Searches for a file in a list of directories.
<i>FileSplit</i>	Splits a file name into its three component parts (directory, file name, and extension).
<i>FindFirst</i>	Searches the specified directory for the first entry matching the specified file name and set of attributes.
<i>FindNext</i>	Returns the next entry that matches the name and attributes specified in a previous call to <i>FindFirst</i> .
<i>GetFAttr</i>	Returns the attributes of a file.
<i>SetFAttr</i>	Sets the attributes of a file.

Table 16.14
WinDos directory-handling
procedures and functions

Procedure or function	Description
<i>CreateDir</i>	Creates a new subdirectory.
<i>GetCurDir</i>	Returns the current directory of a specified drive.
<i>RemoveDir</i>	Removes a subdirectory.
<i>SetCurDir</i>	Changes the current directory.

Table 16.15
WinDos environment-
handling functions

Functions	Description
<i>GetArgCount</i>	Returns the number of parameters passed to the program on the command line.
<i>GetArgStr</i>	Returns a specified command-line argument.
<i>GetEnvVar</i>	Returns a pointer to the value of a specified environment variable.

Table 16.16
WinDos miscellaneous
procedures and functions

Procedure or function	Description
<i>DosVersion</i>	Returns the DOS version number.
<i>GetCBreak</i>	Returns the state of <i>Ctrl-Break</i> checking in DOS.
<i>GetVerify</i>	Returns the state of the verify flag in DOS.
<i>SetCBreak</i>	Sets the state of <i>Ctrl-Break</i> checking in DOS.
<i>SetVerify</i>	Sets the state of the verify flag in DOS.

WinDos unit constants, types, and variables

Each of the constants, types, and variables defined by the *WinDos* unit are briefly discussed in this section.

Constants

The *WinDos* unit uses several constants. These constants can be grouped by their function. To learn more about these constants, look them up as part of the constant group they belong to. For example, to find the value of *fParity*, look up “Flag constants” in the *Programmer’s Reference*.

Table 16.17
WinDos constants

Constant group	Description
Flag	Test individual flag bits in the Flags register after a call to <i>Intr</i> or <i>MsDos</i> : <i>fCarry</i> , <i>fParity</i> , <i>fAuxiliary</i> , <i>fZero</i> , <i>fSign</i> , <i>fOverflow</i>
fmXXXX	Used by file-handling procedures when opening and closing disk files: <i>fmClosed</i> , <i>fmInput</i> , <i>fmOutput</i> , <i>fmInOut</i>
faXXXX	Test, set, and clear file attribute bits in connection with the file-handling procedures: <i>faReadOnly</i> , <i>faHidden</i> , <i>faSysFile</i> , <i>faVolumeID</i> , <i>faDirectory</i> , <i>faArchive</i> , <i>faAnyFile</i>
fsXXXX	Maximum file name component string lengths used by <i>FileSearch</i> and <i>FileExpand</i> : <i>fsPathName</i> , <i>fsDirectory</i> , <i>fsFileName</i> , <i>fsExtension</i>
fcXXXX	Return flags used by the <i>FileSplit</i> function: <i>fcExtension</i> , <i>fcFileName</i> , <i>fcDirectory</i> , <i>fcWildcards</i>

Types

The *WinDos* unit defines these types:

Table 16.18
WinDos types

Types	Description
File record types	<i>TFileRec</i> is used for both typed and untyped files; <i>TTextRec</i> is the internal format of a variable of type text.
<i>TRegisters</i>	Variables of this type are used by <i>Intr</i> and <i>MsDos</i> to specify the input register contents and examine the output register contents of a software interrupt.
<i>TDateTime</i>	Variables of this type are used to examine and construct 4-byte, packed date-and-time values for <i>GetFTime</i> , <i>SetFTime</i> , <i>FindFirst</i> , and <i>FindNext</i> .
<i>TSearchRec</i>	Variables of this type are used by <i>FindFirst</i> and <i>FindNext</i> to scan directories.

Variables

DosError is used by many of the routines in the *WinDos* unit to report errors.

Programming in DOS protected mode

The 80286 microprocessor introduces a new way of addressing memory: protected virtual address mode, or simply protected mode. This new addressing mode has three primary benefits:

- Access of up to 16 megabytes of memory
- A logical address space larger than the physical address space
- A way to isolate programs from each other so that one program would not corrupt another running at the same time

With Borland Pascal, you can easily write DOS applications that will run in protected mode without the need for a third-party DOS extender. You may find that many of your real-mode programs run just fine in protected mode. This chapter will help you modify those programs that don't, and it explains some of the basic protected-mode versus real-mode issues.

What is protected mode?

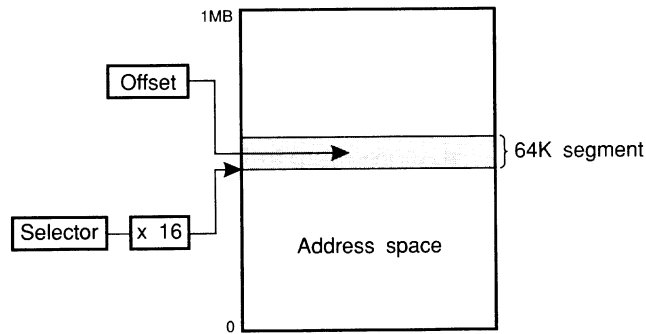
The 80286 and later processors support two modes of operation: *real mode* and *protected mode*. Real mode is compatible with the original 8086 processor and allows an application to address up to one megabyte of memory. Protected mode extends the addressing range to 16 megabytes.

The main difference between real and protected mode lies in the way the processor converts *logical addresses* into *physical addresses*.

Logical addresses are the addresses used in an application. In both real and protected mode, a *logical address* is a 32-bit value consisting of a 16-bit selector (segment address) and a 16-bit offset. Physical addresses are the addresses that the processor uses to exchange data with memory components in the system. In real mode, a physical address is a 20-bit value, whereas in protected mode, a physical address is a 24-bit value.

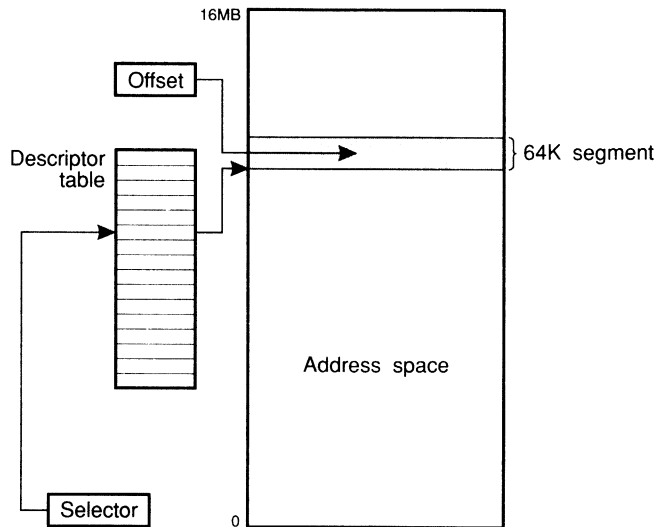
Whenever the processor accesses memory—to fetch an instruction or to load or store a variable—it generates a physical address from a logical address. In real mode, the physical-address generation consists of shifting the selector (segment address) left by four bits (multiplying by 16) and adding the offset. The resulting 20-bit address is then used to access memory.

Figure 17.1
Real-mode physical address generation



To generate a physical address in protected mode, the selector part of the logical address is used as an index into a *descriptor table*. The entry in the descriptor table contains a 24-bit base address to which the offset part of the logical address is then added to form the physical address.

Figure 17.2
Protected-mode physical
address generation



Each entry in the descriptor table is called a *descriptor* and defines a segment in memory. A descriptor-table entry occupies eight bytes, and the information stored in a descriptor includes the *base address*, the *limit*, and the *access rights* flags for the segment.

The segment limit and access rights entries in a descriptor define the size and type of the segment. Segments can be anywhere from 1 to 65536 bytes in size and are either *code segments* or *data segments*. Code segments can contain executable machine instructions and read-only data, and data segments can contain data that can be both read and written. It isn't possible to write data to code segments or to execute instructions in data segments. Any attempt to do so, or any attempt to access data beyond a segment's limit, causes a *general protection fault* (or GP fault for short)—hence the name protected mode.

Given a logical address in real mode, an application can determine the corresponding physical address. Usually this isn't true in protected mode, because the selector part of a logical address is an index into a descriptor table, and the selector itself has no direct bearing on the physical address calculation. This is an advantage in that true virtual-memory management can be implemented without affecting applications. For example, by simply updating the base address field of a segment's descriptor, the operating system can move a segment around in physical

memory without affecting the application that uses the segment. The application refers to the segment's selector only, and the selector isn't affected by changes to the descriptor.

Applications seldom, if ever, need to deal with descriptors. The operating system and the memory manager create and destroy descriptors as required, and an application knows only about the corresponding selectors. Selectors are similar to file handles—from the application's point of view they are "magic cookies" that are managed by the operating system, but within the operating system, they function as indexes into tables that contain additional information.

Borland protected-mode DOS extensions

Borland Pascal's protected-mode extensions are implemented through two components: the DPMServer (the DPMServer16BI.OVL file), and the run-time manager (the RTM.EXE file).

The DPMServer

The DOS Protected Mode Interface (DPMServer) is an industry standard that allows DOS programs to access the advanced features of the 80286-, 80386-, and 80486-based PCs in a well-behaved, hardware-independent fashion. DPMServer functions are defined to manage descriptor tables, perform mode switching, allocate extended memory, allocate DOS memory, control the interrupt subsystem, and communicate with real-mode programs.

Borland Pascal's protected-mode extensions are based on the DPMServer 0.9 specification. While the DPMServer specification doesn't support DOS calls from protected-mode applications, Borland's DPMServer server and servers from many other companies, including enhanced-mode Windows 3.x, do support INT 21H and other standard DOS and BIOS interrupts commonly used in DOS protected-mode applications.

The run-time manager

The run-time manager (RTM.EXE) builds on top of the DPMServer server to provide a number of services to protected-mode applications. The run-time manager contains a protected-mode

program loader and a protected-mode memory manager, and it allows multiple protected-mode clients to exist under DPMI.

Borland protected-mode applications use the same executable file format as Windows 3.x and OS/2 1.x. The run-time manager's program loader can load both executables (.EXE files) and dynamic-link libraries (.DLL files).

The protected-mode memory manager allows protected-mode applications to allocate blocks of dynamic memory. The memory manager supports fixed, moveable, and discardable blocks, and also manages an application's code and data segments. Using features unique to protected mode, the memory manager also functions as an overlay manager by automatically loading and discarding code segments (for this reason, a protected-mode application doesn't need to use the *Overlay* unit).

Applications can access the run-time manager through the *WinAPI* unit. *WinAPI*, described in a following section, implements a subset of the Windows API (application programming interface) to support memory management, module management, resource management, loading of dynamic-link libraries, and low-level access to selectors. Because the run-time manager API is a subset of the Windows API, you can write binary-compatible dynamic-link libraries that can be used in both DOS protected mode and Windows.

Developing a DOS protected-mode application

Writing a protected-mode application is quite simple. You don't have to worry about selectors and memory addresses. The operating system with Borland's extensions does it all for you. In fact, most of your real-mode applications will run just fine in protected mode. The following sections outline some of the differences between real and protected mode you should be aware of when writing a protected-mode application.

Safe programming in protected mode

There are a few practices used occasionally in real-mode programs that result in a general protection fault, or GP fault, if they are used in a protected-mode program. Borland Pascal

GP faults help protect your system from poor programming practices.

reports a GP fault as run-time error 216. GP faults occur when you try to access memory that isn't legal for your application to access. The operating system halts your application and tells you that a GP fault occurred, but your system won't crash. Although GP faults terminate your program, your system is "protected" from crashing. These practices result in GP faults:

- Loading invalid values into segment registers
- Accessing memory beyond a segment's limit
- Writing to code segments
- Dereferencing **nil** pointers

Loading invalid values into segment registers

When the processor is running in protected mode, segment registers (CS, DS, ES, and SS) can hold only selectors. Because selectors are indexes into descriptor tables, they have no physical relationship with the memory they reference. If you attempt to load an arbitrary value into a segment register, you'll probably get a GP fault, because the value might not denote a valid descriptor.

The *Ptr* function and the *Mem* arrays

The compiler generates code to load a segment register whenever you dereference a pointer. If you construct pointers using the *Ptr* standard function, you must ensure that the segment part of the pointer is a valid selector. Similarly, when using the *Mem*, *MemW*, and *MemL* arrays, you must use selectors instead of physical addresses. For example, when accessing the ROM BIOS workspace area (segment \$0040) or the video RAM areas (segments \$A000, \$B000, and \$B800), you should use the *SegXXXX* variables instead of absolute values. The *SegXXXX* variables are described on page 200.

Absolute variables

In protected mode, you can't specify an absolute address for a variable. You need to rewrite any code that specifies a segment and an offset in an **absolute** clause. For example, you might need to construct a pointer using the *SegXXXX* variables.

Using segment arithmetic

Adding or subtracting values from the selector part of a pointer usually isn't valid in protected mode. For example, adding \$1000 to the selector part of a pointer in real mode increments the pointer by 64K, but in protected mode, the resulting pointer is invalid. Instead, you must use the *GlobalXXXX* functions in the *WinAPI* unit to allocate and manage memory blocks.

In Borland Pascal, there is a way to perform segment arithmetic on selectors with the *SelectorInc* variable; see page 201 for more information.

Using segment registers as temporary variables

In real mode, some older assembly language code may use segment registers as “scratchpads” for temporary variables. This doesn’t work in protected mode, because usually the temporary values stored in the segment registers aren’t valid selectors.

Accessing memory beyond a segment’s limit

In real mode, each segment is a full 64K bytes in length. In protected mode, a segment’s descriptor contains a field that specifies the limit of the segment, and if you attempt to access data beyond a segment’s limit, you get a GP fault. When loading your application, the run-time manager sets the appropriate limits for the application’s code, data, and stack segments. Also, memory blocks allocated using the *GlobalAlloc* function in the *WinAPI* unit have a segment limit corresponding to the size of the memory block.

Writing to a code segment

In real mode, it’s possible to store variables in a code segment because real mode doesn’t specify what can and can’t exist in a segment. This doesn’t work in protected mode. Protected-mode selectors have a flag for read/write or read-only access, and code selectors are always flagged as read-only. A GP fault occurs if you attempt to write to a code segment selector. You can use an *alias* to write self-modifying code, however; see page 203.

Dereferencing nil pointers

When you convert a real-mode application to protected mode, it’s common for certain bugs to suddenly appear in code that has run without errors for years. For example, you might be accidentally dereferencing a *nil* pointer, or you might find that your application contains “dangling” pointers that are dereferenced after they have been disposed of. In real mode, such bugs don’t necessarily show up right away, but in protected mode, they usually cause a GP fault. Protected mode, as the name suggests, does a much better job than real mode of protecting you from pointer-related bugs.

Code and data segments

Like a real-mode Borland Pascal application, a protected-mode application consists of a number of code segments, a data

For more information on code segments and the **\$C** compiler directive, see page 266 in this book and also Chapter 2, "Compiler directives," in the *Programmer's Reference*.

segment, and a stack segment. When a protected-mode program is loaded, the run-time manager automatically allocates selectors for the code, data, and stack segments. For code segments, certain attributes of the segments can be controlled through a **\$C** compiler directive. In particular, code segments can be either moveable or fixed in physical memory, they can be preloaded or demand loaded, and they can be discardable or permanent.

The code segment attributes effectively allow you to designate segments as static (moveable, preload, permanent) or dynamic (moveable, demandload, discardable). Therefore, in protected mode, you don't need to use the *Overlay* unit with the **\$O** compiler directive, and the *OvrCodeList*, *OvrHeapSize*, *OvrDebugPtr*, *OvrHeapOrg*, *OvrHeapPtr*, *OvrHeapEnd*, *OvrLoadList*, *OvrDosHandle*, and *OvrEmsHandle* variables don't exist in the protected-mode version of the *System* unit.

Heap management

The Borland Pascal run-time library protected-mode heap manager is quite different from the Borland Pascal run-time library real-mode heap manager. In particular, the *HeapOrg*, *HeapPtr*, *HeapEnd*, and *FreeList* variables aren't defined in the protected-mode version of the *System* unit. The Borland Pascal run-time library protected-mode heap manager (which is identical to the Borland Pascal for Windows run-time library heap manager) relies on the run-time manager to provide basic segment allocation and deallocation, and includes a segment sub-allocator to optimize performance for small blocks. For more information on the protected-mode heap manager, see on page 268.

Predefined selectors

The *System* unit provides a few predefined selectors for commonly used real-mode addresses. They are named for the physical segment they are assigned to and are provided for compatibility between DOS real and protected mode.

Table 17.1
Predefined selectors

Selector	Description
<i>Seg0040</i>	Used to access the \$40 low-memory BIOS data region
<i>SegA000</i>	Used to access the EGA and VGA graphics memory pages at segment address \$A000.

Table 17.1: Predefined selectors (continued)

<i>SegB000</i>	Used to access the Monochrome Adapter video memory at segment address \$B000.
<i>SegB800</i>	Used to access the Color Graphics Adapter video memory at segment address \$B800.

In real mode, the *SegXXXX* variables always contain \$0040, \$A000, \$B000, and \$B800 respectively. In protected mode, the run-time library's start-up code creates four selectors that reference the particular real-mode memory regions. You should use the *SegXXXX* variables whenever you reference these memory regions. For example, in the past you might have written code that looked like this:

```
CrtMode := Mem[$40: $49];
```

Now you should write this instead:

```
CrtMode := Mem[Seg0040: $49];
```

By using the *SegXXXX* variables you can ensure that your code runs unchanged in both real and protected modes.

The SelectorInc variable

The *SelectorInc* variable in the *System* unit contains the value that must be added to or subtracted from a selector to produce the next or previous selector in the descriptor table. *SelectorInc* is useful when dealing with huge memory blocks (blocks larger than 64K bytes) and when accessing aliased segments.

The *GlobalAlloc* and *GlobalAllocPtr* functions in the *WinAPI* unit can be used to allocate blocks larger than 64K bytes—such blocks are also known as huge memory blocks. Huge memory blocks are contiguous in physical memory, but due to the processor's 16-bit architecture, they can't be accessed as a whole in an application. For a huge memory block, the memory manager allocates multiple contiguous selectors that each (except for the last selector) refer to a 64K byte portion of the huge memory block. For example, to allocate a 220K huge block, the memory manager creates four selectors—the first three selectors each refer to a 64K block, and the last selector refers to a 28K memory block. By adding *SelectorInc* to a selector belonging to a huge block, you can produce the selector for the next segment, and by subtracting *SelectorInc*, you can produce the selector for the previous segment.

When allocating a huge block, the *GlobalAlloc* function always returns the handle of the first segment, and the *GlobalAllocPtr* function always returns a pointer to the first segment.

The *GetPtr* function below takes a huge block pointer (as returned by *GlobalAllocPtr*) and a 32-bit offset and returns a pointer to the given offset within the block.

```
function GetPtr(P: Pointer; Offset: Longint): Pointer;
type
  Long = record
    Lo, Hi: Word;
  end;
begin
  GetPtr := Ptr(
    Long(P).Hi + Long(Offset).Hi * SelectorInc,
    Long(P).Lo + Long(Offset).Lo);
end;
```

Notice how the high word of the *Offset* parameter is used to determine how many times to increment the selector part of *P* in order to find the correct segment. For example, if *Offset* is \$24000, the selector part of *P* will be incremented by $2 * SelectorInc$, and the offset part of *P* will be incremented by \$4000.

The following function *LoadFile* loads an entire file into a memory block and returns a pointer to the block. If the file is larger than 64K bytes, a huge block is allocated.

```
function LoadFile(const FileName: string): Pointer;
var
  Buffer: Pointer;
  Size, Offset, Count: Longint;
  F: file;
begin
  Buffer := nil;
  Assign(F, FileName);
  Reset(F, 1);
  Size := FileSize(F);
  Buffer := GlobalAllocPtr(gmem_Moveable, Size);
  if Buffer <> nil then
  begin
    Offset := 0;
    while Offset < Size do
    begin
      Count := Size - Offset;
      if Count > $8000 then Count := $8000;
      BlockRead(F, GetPtr(Buffer, Offset)^, Count);
      Inc(Offset, Count);
    end;
  end;
```

```

        end;
    end;
    LoadFile := Buffer;
end;

```

SelectorInc is also defined in the real-mode version of the *System* unit. In real mode, it always contains the value \$1000, which, when added to the segment part of a real-mode pointer, increments the pointer by 64K bytes.

You can use *SelectorInc* in another way in DOS protected-mode programs only. Use *SelectorInc* to access the *aliased selectors* allocated by the run-time manager when an application is loaded. For each code segment in an application, the run-time manager creates an aliased selector that refers to the same segment, but has data selector access rights. Aliased selectors are not created for data and stack segments.

To access the aliased selector for a particular segment, add *Selector* to the segment's selector. For example, assuming that *P* is a variable of type *Pointer* and *Foo* is a procedure or function, the assignment

```
P := Addr(Foo)
```

causes *P* to point to the executable, read-only entry point of *Foo*, whereas

```
P := Ptr(Seg(Foo) + SelectorInc, Of(Foo));
```

causes *P* to point to the same location, but with read/write access rights.

The WinAPI unit

The *WinAPI* unit gives you direct access to Borland's protected-mode DOS extensions. To make it easier to write portable applications and binary-compatible DLLs, the *WinAPI* interface is designed as a subset of the Windows API.

The *WinAPI* unit gives you access to memory management, resource management, module management, selector management, and various other API functions. A brief description of each of these functional areas follows. For complete details on the

constants, types, procedures, and functions in the *WinAPI* unit, see the *Programmer's Reference*.

Under Windows, the API routines supported through *WinAPI* are located in the *KERNEL.DLL* and *USER.DLL* dynamic-link libraries. In DOS protected mode, these DLLs aren't required. Instead, the protected-mode run-time manager contains the implementation of the *KERNEL* and *USER* routines, and modules that import from *KERNEL* and *USER* are automatically redirected to use the routines built into the run-time manager.

Memory management

When writing applications that work with dynamic memory, you usually use the *New*, *Dispose*, *GetMem*, and *FreeMem* standard procedures. You can access Borland's protected-mode memory manager directly through the *GlobalXXXX* functions in the *WinAPI* unit, however.

Notice that the *GlobalXXXXPtr* functions are "wrappers" that combine common function call sequences, such as *GlobalAlloc* followed by *GlobalLock*, or *GlobalUnlock* followed by *GlobalFree*, into one API routine.

Table 17.2
Memory management API
routines

Function	Description
<i>GetFreeSpace</i>	Retrieves the amount of free memory in the heap.
<i>GlobalAlloc</i>	Allocates a block of memory from the heap.
<i>GlobalAllocPtr</i>	Allocates and locks a memory block (using <i>GlobalAlloc</i> and <i>GlobalLock</i>).
<i>GlobalCompact</i>	Rearranges memory currently allocated to the heap so the specified amount of memory is free.
<i>GlobalDiscard</i>	Discards the given memory object.
<i>GlobalDosAlloc</i>	Allocates memory that can be accessed by DOS running in real mode; the memory will exist in the first megabyte of linear address space.
<i>GlobalDosFree</i>	Frees memory previously allocated with <i>GlobalDosAlloc</i> .
<i>GlobalFix</i>	Prevents the given memory object from moving in linear memory.
<i>GlobalFlags</i>	Retrieves information about the memory block.

Table 17.2: Memory management API routines (continued)

<i>GlobalFree</i>	Frees an unlocked memory block and invalidates its handle.
<i>GlobalFreePtr</i>	Unlocks and frees a memory block (using <i>GlobalUnlock</i> and <i>GlobalFree</i>).
<i>GlobalHandle</i>	Retrieves the handle of a memory object with the specified segment address.
<i>GlobalLock</i>	Increments the reference count of a memory block and returns a pointer to it.
<i>GlobalLockPtr</i>	Same as <i>GlobalLock</i> , but takes a pointer instead of a handle.
<i>GlobalLRUNewest</i>	Moves a memory object into the newest least-recently-used memory position, thereby minimizing the likelihood that the object will be discarded.
<i>GlobalLRUOldest</i>	Moves a memory object into the oldest least-recently-used memory position, thereby maximizing the likelihood that the object will be discarded.
<i>GlobalNotify</i>	Calls the instance address of a notification callback procedure, passing the handle of the memory block to be discarded.
<i>GlobalPageLock</i>	Increments the page-lock count for the memory associated with the given selector.
<i>GlobalPageUnlock</i>	Decrements the page-lock count for the memory associated with the specified selector.
<i>GlobalPtrHandle</i>	Given a pointer to a memory block, returns the handle of that block.
<i>GlobalReAlloc</i>	Reallocates a memory block.
<i>GlobalReAllocPtr</i>	Unlocks, reallocates, and locks a memory block (using <i>GlobalUnlock</i> , <i>GlobalReAlloc</i> , and <i>GlobalLock</i>).
<i>GlobalSize</i>	Retrieves the current size of a memory block.
<i>GlobalUnfix</i>	Allows a memory object to be moved in linear memory canceling the effects of <i>GlobalFix</i> .
<i>GlobalUnlock</i>	Unlocks a memory block previously locked by <i>GlobalLock</i> .
<i>GlobalUnlockPtr</i>	Same as <i>GlobalUnlock</i> , but takes a pointer instead of a handle.
<i>LockSegment</i>	Locks the specified discardable segment.
<i>UnlockSegment</i>	Unlocks a segment.

The *GlobalAlloc* function is used to allocate memory blocks and the *GlobalFree* function is used to free them. The memory manager supports three types of memory blocks: fixed, moveable, and discardable. A fixed memory block is guaranteed to stay at the same location in physical memory. A moveable block can be moved in physical memory to make room for other allocation requests, and a discardable block can be discarded by the memory manager to make room for other blocks. You select one of the three styles through the flags passed to *GlobalAlloc*:

- *gmem_Fixed*
- *gmem_Moveable*
- *gmem_Moveable* + *gmem_Discardable*

An application usually allocates moveable blocks only. You'll seldom need to allocate fixed or discardable blocks.

The memory manager operates on "handles", which are represented by the *THandle* type in the *WinAPI* unit. A memory handle is a word-sized value that identifies a memory block, much like a file handle is a word-sized value that identifies an open file.

Before you can access a memory block, you must lock it using the *GlobalLock* function, and when you're done accessing it, you must unlock it using the *GlobalUnlock* function. *GlobalLock* returns a full 32-bit pointer to the first byte in the block. The offset part of the pointer always zero. In DOS protected-mode, the selector part of the pointer is the same as the block's handle, but this isn't always the case in Windows.

The following example shows the proper sequence of allocating, locking, unlocking, and freeing a block. In the example, *H* is a variable of type *THandle* and *P* is a pointer.

```
H := GlobalAlloc(gmem_Moveable, 8192);           { Allocate block }
if H <> 0 then                                   { If we got the memory }
begin
  P := GlobalLock(H);                            { Lock the block }
  :                                             { Access block through P }
  GlobalUnlock(H);                               { Unlock the block }
  GlobalFree(H);                                 { Free the block }
end;
```

Locking and unlocking a memory block each time it is accessed is tedious and error prone, and in reality it is only required with discardable blocks and in Windows applications running in real mode. In all other situations, a better solution is to lock the block

right after it is allocated, and not unlock it until right before it is disposed. For that purpose, the *WinAPI* unit includes the *GlobalXXXXPtr* family of “wrapper” routines. Of particular interest are *GlobalAllocPtr* which allocates and locks a memory block, and *GlobalFreePtr* which unlocks and frees a memory block. Using these routines, the above example can be simplified:

```
P := GlobalAllocPtr(gmem_Moveable, 8192);           { Allocate block }
if P <> nil then                                    { If we got the memory }
begin
  :                                                 { Access the block }
  GlobalFreePtr(P);                                { Free the block }
end;
```

You can change the size or attributes of a memory block while preserving its contents by calling the *GlobalReAlloc* function. *GlobalReAlloc* returns the new handle of the block, which may be different from the handle that was passed to the function if the old size or the new size of the block is larger than 64K. Note that in cases where the old size and the new size of the block are both less than 64K bytes, *GlobalReAlloc* can always resize the block without changing its handle.

GlobalReAlloc can also be used to change a block’s attributes. You do that by specifying the *gmem_Modify* flag along with *gmem_Moveable* or *gmem_Discardable*.

The *GlobalReAllocPtr* function performs the same actions as *GlobalReAlloc*, but uses pointers instead of handles.

There are a number of other less frequently used *GlobalXXXX* functions, all of which are described in detail in Chapter 1, “Library reference,” in the *Programmer’s Reference*.

Module management

Table 17.3
Module management API
routines

The run-time manager supports the following module management routines:

Routine	Description
<i>FreeLibrary</i>	Invalidates the loaded library module, and frees associated memory if the module is no longer referenced.
<i>GetModuleFileName</i>	Retrieves the full path and file name of the executable file which the specified module was loaded from.

Table 17.3: Module management API routines (continued)

<i>GetModuleHandle</i>	Retrieves the handle of the specified module.
<i>GetModuleUsage</i>	Retrieves a module's reference count.
<i>GetProcAddress</i>	Retrieves the address of an exported library function.
<i>LoadLibrary</i>	Loads the named library module.

Several of these routines require a module-handle parameter. The module handle of the application itself is stored in the *HInstance* variable declared in the *System* unit.

Resource management

Table 17.4
Resource management API
functions

The run-time manager supports the following resource management routines:

Function	Description
<i>AccessResource</i>	Opens the given executable file and moves the file pointer to the beginning of the given resource.
<i>FindResource</i>	Determines the location of a resource in a specified resource file.
<i>FreeResource</i>	Decrements the reference count of a loaded resource; when the count reaches zero, the memory used by a resource is freed.
<i>LoadResource</i>	Loads the specified resource in memory.
<i>LoadString</i>	Loads the specified string resource.
<i>LockResource</i>	Locks the given resource in memory and increments its reference count.
<i>SizeOfResource</i>	Returns the size, in bytes, of the given resource.
<i>UnlockResource</i>	Unlocks the specified resource and decreases the reference count of the resource by one.

Turbo Vision resources don't follow the same conventions as Windows resources and can't be accessed using these API routines.

Resources can be linked into an application using `{$R filename}` compiler directives. The named files must be Windows resource files (.RES files). Usually a DOS protected-mode application links in only string resources and user-defined resources; other Windows resource types usually aren't applicable to a DOS application.

The instance handle required by a number of the resource-management routines is normally the application's instance handle, which can be found in the *HInstance* variable in the *System* unit.

Selector management

An application usually doesn't need to directly manipulate selectors, but in certain situations the following selector management routines can be useful:

Table 17.5
Selector management API
functions

Function	Description
<i>AllocDStoCSAlias</i>	Maps a data-segment selector to a code-segment selector.
<i>AllocSelector</i>	Allocates a new selector.
<i>ChangeSelector</i>	Generates a code selector that corresponds to a given data selector, or generates a given data selector that corresponds to a code selector.
<i>FreeSelector</i>	Frees a selector originally allocated by <i>AllocDStoCSAlias</i> or <i>AllocSelector</i> .
<i>GetSelectorBase</i>	Retrieves the base address of a selector.
<i>GetSelectorLimit</i>	Returns the limit of the specified selector.
<i>PrestoChangoSelector</i>	Generates a code selector that corresponds to a given data selector, or generates a given data selector that corresponds to a code selector.
<i>SetSelectorBase</i>	Sets the base address of a selector.
<i>SetSelectorLimit</i>	Sets the limit of a selector.

Other API routines

The run-time manager supports the following additional API routines:

Table 17.6
Other API routines

Function	Description
<i>DOS3Call</i>	Calls a DOS interrupt 21h function; called only from assembly language routines.
<i>FatalExit</i>	Sends the current state of the protected-mode environment to the debugger and prompts for instructions on how to proceed.

Table 17.6: Other API routines (continued)

<i>GetDOSEnvironment</i>	Retrieves the current task's DOS environment string.
<i>GetVersion</i>	Retrieves the current version of the Windows environment and DOS operating system.
<i>GetWinFlags</i>	Retrieves the memory configuration flags used by Windows.
<i>MessageBox</i>	Creates, displays, and operates and message-box window.

A shared DLL can use the *GetWinFlags* function and the *wf_DPMI* flag to determine if it is running in DOS protected mode or under Windows. For example,

```

if GetWinFlags and wf_DPMI <> 0 then
    Message('Running in DOS protected mode.')
else
    Message('Running under Windows.');
```

Accessing the DPMI server directly

You can also directly access the DPMI server through interrupt \$31, which calls the DPMI server directly, bypassing the run-time manager. This is a very dangerous practice, however. DPMI isn't consistent about cleaning up resources such as memory interrupt vectors; the run-time manager was designed to correct these problems. You must have a deep understanding of protected-mode issues and be aware of the significant risks you are taking before you attempt this method of accessing protected mode.

Compiling a protected-mode application

Most of the time you don't have to do anything special to produce a protected-mode application. Simply compile your application specifying protected mode as your target platform in one of these two ways:

- In the IDE, choose **Compile | Target** and select **Protected-mode Application** in the **Target Platform** dialog box.
- When using the command-line compiler, use the **/CP** switch to select protected mode as the target platform.

Running a DOS protected-mode application

Your license agreement permits you to distribute DPMI16BI.OVL and RTM.EXE along with your program.

When you run a DOS protected-mode application, you must ensure that DPMI16BI.OVL (the DPMI server), RTM.EXE (the run-time manager), and any DLLs that your application uses are present in the current directory or on the DOS path.

A DOS protected-mode .EXE file uses the same executable file format as Windows 3.x and OS/2 1.x. The file format is a superset of the normal DOS .EXE format and consists of a regular .EXE file image, called the *stub*, followed by an extended header and the protected-mode code, data, and resources. The flow of events upon executing a DOS protected-mode application is outlined here:

1. DOS loads the real-mode stub and passes control to it.
2. If no DPMI services are present, the stub loads the Borland DPMI server from the DPMI16BI.OVL file. Some newer 386 memory managers support DPMI services, as does a Windows 3.x enhanced-mode DOS box. In such configurations, the stub doesn't load the DPMI server, but uses the one that is already present instead.
3. Next, if the run-time manager isn't already loaded in memory, the stub loads it from the RTM.EXE file. If a protected-mode application executes another protected-mode application, both use the same copy of the run-time manager.
4. Once DPMI services and the run-time manager are present, the stub switches from real to protected mode and passes control to the extended .EXE loader in the run-time manager.
5. The loader first loads the DLLs used by the application (if any), and then loads the application's code and data segments. Finally, the loader passes control to the application's entry point.



When you run your DOS protected-mode program, it's always possible that a non-Borland DPMI server is already present. Because there might be slight differences among servers, especially in their handling of DOS interrupts, you should test your program to be sure it runs with all the possible servers your program might encounter.



When a DOS protected-mode application is executed in a Windows enhanced-mode DOS box, you can control how much

extended memory the run-time manager allocates by specifying an XMS memory limit in the application's .PIF file.

Controlling the amount of memory RTM uses

By default, the run-time manager consumes all available memory for itself when it loads. It then allocates memory to its clients as they request it through the memory manager API routines.

In protected mode, there is no difference between conventional memory (memory below the 1MB address) and extended memory (memory above the 1MB address); both types of memory are available to protected-mode applications. The run-time manager *favors* extended memory, however. Only when all extended memory has been allocated, or when an application specifically requests conventional memory (using the *GlobalDosAlloc* function, for example) does the run-time manager allocate blocks in conventional memory.

The reason the run-time manager favors extended memory is that a protected-mode application may spawn other applications using the *Exec* routine in the *Dos* unit. Spawned applications aren't necessarily protected-mode applications; therefore, they need as much conventional memory as possible. In fact, spawned protected-mode applications start up as real-mode programs and switch to protected mode only after the stub has successfully located DPMI services and the run-time manager.

The run-time manager attempts to free as much conventional memory as possible (by moving moveable memory blocks into extended memory, for example) before spawning an application. No attempt is made to release extended memory, however. Therefore, if protected-mode applications that don't use the run-time manager are to be spawned, you need a way to control the run-time manager's allocation of memory.

To control how much memory the run-time manager can use, at the DOS command line, add the RTM environment variable to your system's DOS environment. Here is the syntax:

```
SET RTM=[option nnnn]
```

the following table lists the options you can use. *nnnn* can be a decimal number or a hex number in the form of *xAB54* or *xab54*.

Table 17.7
Environment variable options
to control RTM's memory
allocation

Option	Description
EXTLEAVE nnnn	Always leave at least nnnn kilobytes of extended memory available. The default value is 640K.
EXTMAX nnnn	Don't allocate more than nnnn kilobytes of extended memory. The default value is 4 gigabytes. In Windows, the default value is one-half the available memory.
EXTMIN nnnn	If fewer than nnnn kilobytes are available after applying EXTMAX and EXTLEAVE limits, terminate with an Out of Memory message. The default value is zero.
REALLEAVE nnnn	Always leave at least nnnn paragraphs of real memory available. The default value is 64K or 4096 paragraphs.
REALMAX nnnn	Don't allocate more than nnnn paragraphs of real memory. The default value is 1 megabyte or 65,535 paragraphs.
REALMIN nnnn	If fewer than nnnn paragraphs are available after applying REALMAX and REALLEAVE, terminate with an Out of Memory message. The default value is zero.

The following DOS command limits RTM to 2M bytes of extended memory, and ensures that 128K bytes of real memory are left unallocated.

```
SET RTM=EXTMAX 2048 REALLEAVE 8192
```


Using null-terminated strings

Borland Pascal supports a class of character strings called *null-terminated strings*. With Borland Pascal's extended syntax and the *Strings* unit, both your DOS and Windows programs can use null-terminated strings by simply referring to the *Strings* unit with the **uses** statement in your program.

What is a null-terminated string?

The compiler stores a traditional Borland Pascal **string** type as a length byte followed by a sequence of characters. The maximum length of a Pascal string is 255 characters, and a Pascal string occupies from 1 to 256 bytes of memory.

A null-terminated string has no length byte; instead, it consists of a sequence of non-null characters followed by a NULL (#0) character. There is no inherent restriction on the length of a null-terminated string, but the 16-bit architecture of DOS does impose an upper limit of 65,535 characters.

Strings unit functions

Borland Pascal has no built-in routines specifically for null-terminated string handling. Instead you'll find all such functions in the *Strings* unit. Among them are *StrPCopy*, which you can use

to copy a Pascal string to a null-terminated string, and *StrPas*, which you can use to convert a null-terminated string to a Pascal string. Here's a brief description of each function:

Table 18.1
Strings unit functions

Function	Description
<i>StrCat</i>	Appends a source string to the end of a destination string and returns a pointer to the destination string.
<i>StrComp</i>	Compares two strings, <i>S1</i> and <i>S2</i> , and returns a value less than zero if $S1 < S2$, zero if $S1 = S2$, or greater than zero if $S1 > S2$.
<i>StrCopy</i>	Copies a source string to a destination string and returns a pointer to the destination string.
<i>StrECopy</i>	Copies a source string to a destination string and returns a pointer to the end of the destination string.
<i>StrLComp</i>	Compares two strings without case sensitivity.
<i>StrLCat</i>	Appends a source string to the end of a destination string, ensuring that the length of the resulting string doesn't exceed a given maximum, and returns a pointer to the destination string.
<i>StrLComp</i>	Compares two strings for a given maximum length.
<i>StrLCopy</i>	Copies up to a given number of characters from a source string to a destination string and returns a pointer to the destination string.
<i>StrEnd</i>	Returns a pointer to the end of a string (that is, a pointer to the null character that terminates a string).
<i>StrDispose</i>	Disposes of a previously allocated string.
<i>StrLen</i>	Returns the length of a string.
<i>StrLComp</i>	Compares two strings for a given maximum length without case sensitivity.
<i>StrLower</i>	Converts a string to lowercase and returns a pointer to the string.
<i>StrMove</i>	Moves a block of characters from a source string to a destination string, and returns a pointer to the destination string. The two blocks may overlap.
<i>StrNew</i>	Allocates a string on the heap.
<i>StrPas</i>	Converts a null-terminated string to a Pascal string.
<i>StrPCopy</i>	Copies a Pascal string to a null-terminated string and returns a pointer to the null-terminated string.
<i>StrPos</i>	Returns a pointer to the first occurrence of a given substring within a string, or nil if the substring doesn't occur within the string.

Table 18.1: Strings unit functions (continued)

<i>StrRScan</i>	Returns a pointer to the last occurrence of a given character within a string, or nil if the character doesn't occur within the string.
<i>StrScan</i>	Returns a pointer to the first occurrence of a given character within a string, or nil if the character doesn't occur within the string.
<i>StrUpper</i>	Converts a string to uppercase and returns a pointer to the string.

Using null-terminated strings

Null-terminated strings are stored as arrays of characters with a zero-based integer index type; that is, an array of the form

```
array[0..X] of Char
```

where *X* is a positive nonzero integer. These arrays are called *zero-based character arrays*. Here are some examples of declarations of zero-based character arrays that can be used to store null-terminated strings:

```
type
  TIdentifier = array[0..15] of Char;
  TFileName = array[0..79] of Char;
  TMemoText = array[0..1023] of Char;
```

The biggest difference between using Pascal strings and null-terminated strings is the extensive use of pointers in the manipulation of null-terminated strings. Borland Pascal performs operations on these pointers with a set of *extended syntax* rules.

Character pointers and string literals

When extended syntax is enabled, a string literal is *assignment-compatible* with the *PChar* type. This means that a string literal can be assigned to a variable of type *PChar*. For example,

```
var
  P: PChar;
  :
begin
  P := 'Hello world...';
end;
```

The effect of such an assignment is that the pointer points to an area of memory that contains a null-terminated copy of the string literal. This example accomplishes the same thing as the previous example:

```
const
  TempString: array[0..14] of Char = 'Hello world...'#0;
var
  P: PChar;
  ;
begin
  P := @TempString;
end;
```

You can use string literals as actual parameters in procedure and function calls when the corresponding formal parameter is of type *PChar*. For example, given a procedure with the declaration

```
procedure PrintStr(Str: PChar);
```

the following procedure calls are valid:

```
PrintStr('This is a test');
PrintStr(#10#13);
```

Just as it does with an assignment, the compiler generates a null-terminated copy of the string literal. The compiler passes a pointer to that memory area in the *Str* parameter of the *PrintStr* procedure.

Finally, you can initialize a typed constant of type *PChar* with a string constant. You can do this with structured types as well, such as arrays of *PChar* and records and objects with *PChar* fields.

```
const
  Message: PChar = 'Program terminated';
  Prompt: PChar = 'Enter values: ';
  Digits: array[0..9] of PChar = (
    'Zero', 'One', 'Two', 'Three', 'Four',
    'Five', 'Six', 'Seven', 'Eight', 'Nine');
```

A string constant expression is always evaluated as a Pascal-style string even if it initializes a typed constant of type *PChar*; therefore, a string constant expression is always limited to 255 characters in length.

Character pointers and character arrays

When you enable the extended syntax with **\$X**, a zero-based character array is *compatible* with the *PChar* type. This means that whenever a *PChar* is expected, you can use a zero-based character array instead. When you use a character array in place of a *PChar* value, the compiler converts the character array to a pointer *constant* whose value corresponds to the address of the first element of the array. For example,

```
var
  A: array[0..63] of Char;
  P: PChar;
  :
begin
  P := A;
  PrintStr(A);
  PrintStr(P);
end;
```

Because of this assignment statement, *P* now points to the first element of *A*, so *PrintStr* is called twice with the same value.

You can initialize a typed constant of a zero-based character array type with a string literal that is shorter than the declared length of the array. The remaining characters are set to NULL (#0) and the array effectively contains a null-terminated string.

```
type
  TFileName = array[0..79] of Char;
const
  FileNameBuf: TFileName = 'TEST.PAS';
  FileNamePtr: PChar = FileNameBuf;
```

Character pointer indexing

Just as a zero-based character array is compatible with a character pointer, so can a character pointer be indexed as if it were a zero-based character array:

```
var
  A: array[0..63] of Char;
  P: PChar;
  Ch: Char;
  :
```

```

begin
  P := A;
  Ch := A[5];
  Ch := P[5];
end;

```

Both of the last two statements assign *Ch* the value contained in the sixth character element of *A*.

When you index a character pointer, the index specifies an unsigned *offset* to add to the pointer before it is dereferenced. Therefore, *P[0]* is equivalent to *P^* and specifies the character pointed to by *P*. *P[1]* specifies the character right after the one pointed to by *P*, *P[2]* specifies the next character, and so on. For purposes of indexing, a *PChar* behaves as if it were declared as this:

```

type
  TCharArray = array[0..65535] of Char;
  PChar = ^TCharArray;

```

The compiler performs no range checks when indexing a character pointer because it has no type information available to determine the maximum length of the null-terminated string pointed to by the character pointer. Your program must perform any such range checking.

The *StrUpper* function shown here illustrates the use of character pointer indexing to convert a null-terminated string to uppercase.

```

function StrUpper(Str: PChar): PChar;
var
  I: Word;
begin
  I := 0;
  while Str[I] <> #0 do
  begin
    Str[I] := UpCase(Str[I]);
    Inc(I);
  end;
  StrUpper := Str;
end;

```

Notice that *StrUpper* is a function, not a procedure, and that it always returns the value that it was passed as a parameter. Because the extended syntax allows the result of a function call to be ignored, *StrUpper* can be treated as if it were a procedure:

```
StrUpper(A);
PrintStr(A);
```

However, as *StrUpper* always returns the value it was passed, the preceding statements can be combined into one:

```
PrintStr(StrUpper(A));
```

Nesting calls to null-terminated string-handling functions can be very convenient when you want to indicate a certain interrelationship between a set of sequential string manipulations.



See page 71 for information about *PChar* operations.

Null-terminated strings and standard procedures

Borland Pascal's extended syntax allows the *Read*, *Readln*, *Str*, and *Val* standard procedures to be applied to zero-based character arrays, and allows the *Write*, *Writeln*, *Val*, *Assign*, and *Rename* standard procedures to be applied to both zero-based character arrays and character pointers. For more details, see the descriptions of these standard procedures Chapter 1, "Library reference," in the *Programmer's Reference*.

An example using string-handling functions

Here's a code example that shows how we used some of the string-handling functions when we wrote the *FileSplit* function in the *WinDos* unit:

```
{ Maximum file name component string lengths }

const
  fsPathName   = 79;
  fsDirectory  = 67;
  fsFileName   = 8;
  fsExtension  = 4;

{ FileSplit return flags }

const
  fcExtension  = $0001;
  fcFileName   = $0002;
  fcDirectory  = $0004;
  fcWildcards  = $0008;

{ FileSplit splits the file name specified by Path into its
{ three components. Dir is set to the drive and directory path }
{ with any leading and trailing backslashes, Name is set to the }
{ file name, and Ext is set to the extension with a preceding }
```

```

{ period. If a component string parameter is NIL, the      }
{ corresponding part of the path is not stored. If the path }
{ does not contain a given component, the returned component }
{ string is empty. The maximum lengths of the strings returned }
{ in Dir, Name, and Ext are defined by the fsDirectory,      }
{ fsFileName, and fsExtension constants. The returned value is }
{ a combination of the fcDirectory, fcFileName, and fcExtension }
{ bit masks, indicating which components were present in the   }
{ path. If the name or extension contains any wildcard       }
{ characters (* or ?), the fcWildcards flag is set in the    }
{ returned value.                                           }

```

```

function FileSplit(Path, Dir, Name, Ext: PChar): Word;
var
    DirLen, NameLen, Flags: Word;
    NamePtr, ExtPtr: PChar;
begin
    NamePtr := StrRScan(Path, '\');
    if NamePtr = nil then NamePtr := StrRScan(Path, ':');
    if NamePtr = nil then NamePtr := Path else Inc(NamePtr);
    ExtPtr := StrScan(NamePtr, '.');
    if ExtPtr = nil then ExtPtr := StrEnd(NamePtr);
    DirLen := NamePtr - Path;
    if DirLen > fsDirectory then DirLen := fsDirectory;
    NameLen := ExtPtr - NamePtr;
    if NameLen > fsFilename then NameLen := fsFilename;
    Flags := 0;
    if (StrScan(NamePtr, '?') <> nil) or
        (StrScan(NamePtr, '*') <> nil) then
        Flags := fcWildcards;
    if DirLen <> 0 then Flags := Flags or fcDirectory;
    if NameLen <> 0 then Flags := Flags or fcFilename;
    if ExtPtr[0] <> #0 then Flags := Flags or fcExtension;
    if Dir <> nil then StrLCopy(Dir, Path, DirLen);
    if Name <> nil then StrLCopy(Name, NamePtr, NameLen);
    if Ext <> nil then StrLCopy(Ext, ExtPtr, fsExtension);
    FileSplit := Flags;
end;

```

Using the Borland Graphics Interface

The *Graph* unit features a complete library of more than 50 graphics routines that range from high-level calls such as *SetViewPort*, *Circle*, *Bar3D*, and *DrawPoly*, to bit-oriented routines such as *GetImage* and *PutImage*. It supports several fill and line styles and there are several fonts that may be magnified, justified, and oriented horizontally or vertically.

Your license agreement permits you to distribute the .CHR and .BGI files along with your programs.

Table 19.1
Library and Graph unit file names

To compile a program that uses the *Graph* unit, you need your program's source code, the compiler, and access to the standard units in the run-time library and the *Graph* unit:

Type of program	Library	Graph unit name
Real mode	TURBO.TPL	GRAPH.TPU
Protected mode	TPP.TPL	GRAPH.TPP

To run a program that uses the *Graph* unit, you'll need one or more of the graphics drivers (.BGI files listed in the next section) in addition to your .EXE program. Also, if your program uses any stroked fonts, you'll need one or more font (.CHR) files as well.

Drivers

Graphics drivers are provided for the following graphics adapters (and true compatibles):

- CGA
- MCGA
- EGA
- VGA

- Hercules
- 3270 PC
- AT&T 400 line
- IBM 8514

Each driver contains code and data and is stored in a separate file on disk. At run time, the *InitGraph* procedure identifies the graphics hardware, loads and initializes the appropriate graphics driver, puts the system into graphics mode, and then returns control to the calling routine. The *CloseGraph* procedure unloads the driver from memory and restores the previous video mode. You can switch back and forth between text and graphics modes using the *RestoreCrtMode* and *SetGraphMode* routines. To load the driver files yourself or link them into your .EXE file, refer to *RegisterBGIdriver* in Chapter 1, "Library reference," in the *Programmer's Reference*.

Graph supports computers with dual monitors. When *Graph* is initialized by calling *InitGraph*, the correct monitor will be selected for the graphics driver and mode requested. When terminating a graphics program, the previous video mode will be restored. If autodetection of graphics hardware is requested on a dual monitor system, *InitGraph* will select the monitor and graphics card that will produce the highest quality graphics output.

Table 19.2
BGI drivers

Driver	Equipment
ATT.BGI	AT&T 6300 (400 line)
CGA.BGI	IBM CGA, MCGA
EGAVGA.BGI	IBM EGA, VGA
HERC.BGI	Hercules monochrome
IBM8514.BGI	IBM 8514
PC3270.BGI	IBM 3270 PC

IBM 8514 support

Borland Pascal supports the IBM 8514 graphics card, a high-resolution graphics card capable of resolutions up to 1024 × 768 pixels and a color palette of 256 colors from a list of 256K colors. The driver file name is IBM8514.BGI.

Borland Pascal can't properly autodetect the IBM 8514 graphics card (the autodetection logic recognizes it as VGA). Therefore, to use the IBM 8514 card, the *GraphDriver* variable must be assigned the value IBM8514 (which is defined in the *Graph* unit) when *InitGraph* is called. You should not use *DetectGraph* (or *Detect* with *InitGraph*) with the IBM 8514 unless you want the emulated VGA mode.

The supported modes of the IBM 8514 card are IBM8514LO (640 × 480 pixels), and IBM8514HI (1024 × 768 pixels). Both mode constants are defined in the interface for GRAPH.TPU or GRAPH.TPP.

The IBM 8514 uses three 6-bit values to define colors. There is a 6-bit Red, Green, and Blue component for each defined color. The *SetRGBPalette* procedure allows you to define colors for the IBM 8514; it's defined in the *Graph* unit as:

```
procedure SetRGBPalette(ColorNum, Red, Green, Blue: Word);
```

The argument *ColorNum* defines the palette entry to be loaded. *ColorNum* is an integer from 0 to 255 (decimal). The arguments *Red*, *Green*, and *Blue* define the component colors of the palette entry. Only the lower byte of these values is used, and out of this byte, only the 6 most-significant bits are loaded in the palette.


The other palette manipulation routines of the graphics library can't be used with the IBM 8514 driver (that is, *SetAllPalette*, *SetPalette*, and *GetPalette*).

For compatibility with the balance of the IBM graphics adapters, the BGI driver defines the first 16 palette entries of the IBM 8514 to the default colors of the EGA/VGA. These values can be used as is, or changed using the *SetRGBPalette* routine.

Coordinate system

By convention, the upper left corner of the graphics screen is (0,0). The *x* values, or columns, increment to the right. The *y* values, or rows, increment downward. In 320×200 mode on a CGA, the screen coordinates for each of the four corners with a specified point in the middle of the screen would look like this:

Figure 19.1
Screen with xy-coordinates



Current pointer

Many graphics systems support the notion of a current pointer (CP). The CP is similar in concept to a text mode cursor except that the CP isn't visible.

In text mode, this *Write* statement

```
Write('ABC');
```

leaves the cursor in the column immediately following the letter C. If the C is written in column 80, then the cursor wraps around to column 1 of the next line. If the C is written in column 80 on the 25th line, the entire screen scrolls up one line, and the cursor is in column 1 of line 25.

```
MoveTo(0,0)  
LineTo(20,20)
```

In graphics mode, the preceding *LineTo* statement leaves the CP at the last point referenced (20,20). The actual line output is clipped to the current viewport if clipping is active. Note that the CP is never clipped.

The *MoveTo* command is the equivalent of *GoToXY*. Its only purpose is to move the CP. Only the commands that *use* the CP *move* the CP: *InitGraph*, *MoveTo*, *MoveRel*, *LineTo*, *LineRel*, *OutText*, *SetGraphMode*, *GraphDefaults*, *ClearDevice*, *SetViewPort*, and *ClearViewPort*. The latter five commands move the CP to (0,0).

Text

An 8x8 bit-mapped font and several stroked fonts are included for text output while in graphics mode. A bit-mapped character is defined by an 8x8 matrix of pixels. A stroked font is defined by a series of vectors that tell the graphics system how to draw the font.

The advantage of using a stroked font is apparent when you start to draw large characters. Because a stroked font is defined by vectors, it retains good resolution and quality when the font is enlarged.

When a bit-mapped font is enlarged, the matrix is multiplied by a scaling factor and, as the scaling factor becomes larger, the

characters' resolution becomes coarser. For small characters, the bit-mapped font is usually sufficient, but for larger text you will want to select a stroked font.

The justification of graphics text is controlled by the *SetTextJustify* procedure. Scaling and font selection is done with the *SetTextStyle* procedure. Graphics text is output by calling either the *OutText* or *OutTextXY* procedures. Inquiries about the current text settings are made by calling the *GetTextSettings* procedure. The size of stroked fonts can be customized by the *SetUserCharSize* procedure.

Each stroked font is kept in its own file on disk with a .CHR file extension. Font files can be loaded from disk automatically by the *Graph* unit at run time (as described), or they can be linked in or loaded by the user program and "registered" with the *Graph* unit.

Borland Pascal provides a special utility, BINOBJ.EXE, that converts a font file (or any binary data file, for that matter) to an .OBJ file that can be linked into a unit or program using the **{\$L}** compiler directive. This makes it possible for a program to have all its font files built into the .EXE file. (Read the comments at the beginning of the BGILINK.PAS sample program.)

Figures and styles

All kinds of support routines are provided for drawing and filling figures, including points, lines, circles, arcs, ellipses, rectangles, polygons, bars, 3-D bars, and pie slices. Use *SetLineStyle* to control whether lines are thick or thin, or whether they are solid, dotted, or built using your own pattern.

Use *SetFillStyle* and *SetFillPattern*, *FillPoly* and *FloodFill* to fill a region or a polygon with cross-hatching or other intricate patterns.

Viewports and bit images

The *SetViewPort* procedure makes all output commands operate in a rectangular region onscreen. Plots, lines, figures—all graphics output—are viewport-relative until the viewport is changed. Other routines are provided to clear a viewport and read the current viewport definitions. If clipping is active, all graphics

output is clipped to the current port. Note that the CP is never clipped.

GetPixel and *PutPixel* are provided for reading and plotting pixels. *GetImage* and *PutImage* can be used to save and restore rectangular regions onscreen. They support the full complement of *BitBlt* operations (copy, **xor**, **or**, **and**, **not**).

Paging and colors

There are many other routines that support palettes, colors, multiple graphic pages (EGA, VGA, and Hercules only), and so on.

Error handling

Internal errors in the *Graph* unit are returned by the function *GraphResult*. *GraphResult* returns an error code that reports the status of the last graphics operation. Find the error return codes under *GraphResult* Errors in Chapter 1, “Library reference,” in the *Programmer’s Reference*.

The following routines set *GraphResult*:

<i>Bar</i>	<i>ImageSize</i>	<i>SetFillPattern</i>
<i>Bar3D</i>	<i>InitGraph</i>	<i>SetFillStyle</i>
<i>ClearViewPort</i>	<i>InstallUserDriver</i>	<i>SetGraphBufSize</i>
<i>CloseGraph</i>	<i>InstallUserFont</i>	<i>SetGraphMode</i>
<i>DetectGraph</i>	<i>PieSlice</i>	<i>SetLineStyle</i>
<i>DrawPoly</i>	<i>RegisterBGIDriver</i>	<i>SetPalette</i>
<i>FillPoly</i>	<i>RegisterBGIfont</i>	<i>SetTextJustify</i>
<i>FloodFill</i>	<i>SetAllPalette</i>	<i>SetTextStyle</i>
<i>GetGraphMode</i>		



GraphResult is reset to zero after it has been called. Therefore, the user should store the value of *GraphResult* into a temporary variable and then test it.

Getting started

Here's a simple graphics program:

```
program GrafTest;
uses
  Graph;
const
  S = 'Borland Graphics Interface (BGI)';
var
  GraphDriver: Integer;
  GraphMode: Integer;
  ErrorCode: Integer;
  Size: Word;
begin
  GraphDriver := Detect;           { Set flag: do detection }
  InitGraph(GraphDriver, GraphMode, 'C:\BP\BGI');
  ErrorCode := GraphResult;
  if ErrorCode <> GrOk then           { Error? }
    begin
      Writeln('Graphics error: ', GraphErrorMsg(ErrorCode));
      Writeln('Program aborted...');
      Halt(1);
    end;
  Rectangle(0, 0, GetMaxX, GetMaxY);   { Draw full screen box }
  SetTextJustify(CenterText, CenterText);   { Center text }
  Size := 3;
  repeat
    SetTextStyle(DefaultFont, HorizDir, Size);
    Dec(Size);
  until (Size = 0) or (TextWidth(S) < GetMaxX);
  if Size <> 0 then
    OutTextXY(GetMaxX div 2, GetMaxY div 2, S);   { Center of screen }
  Readln;
  CloseGraph;
end.{ GrafTest }
```

The program begins with a call to *InitGraph*, which autodetects the hardware and loads the appropriate graphics driver. For the program to run correctly, the driver and fonts must be in the same directory as the executable program, or the program must specify an explicit directory. In this example, the directory is C:\BP\BGI. If the program fails to recognize graphics hardware or an error occurs during initialization, the program displays an error message and terminates. Otherwise, the program draws a box along the edge of the screen and displays text in the center of the screen.



Neither the AT&T 400 line card nor the IBM 8514 graphics adapter is autodetected. You can still use these drivers by overriding autodetection and passing *InitGraph* the driver code and a valid graphics mode. To use the AT&T driver, for example, replace the ninth and tenth lines in the preceding example with the following three lines of code:

```
GraphDriver := ATT400;
GraphMode := ATT400Hi;
InitGraph(GraphDriver, GraphMode, 'C:\BP\BGI');
```

This instructs the graphics system to load the AT&T 400 line driver located in C:\BP\BGI and set the graphics mode to 640 by 400.

Here's another example that demonstrates how to switch back and forth between graphics and text modes:

```
program GrafTst2;
uses
  Graph;
var
  GraphDriver: Integer;
  GraphMode: Integer;
  ErrorCode: Integer;
begin
  GraphDriver := Detect;                { Set flag: do detection }
  InitGraph(GraphDriver, GraphMode, 'C:\BP\BGI');
  ErrorCode := GraphResult;
  if ErrorCode <> grOk then           { Error? }
  begin
    Writeln('Graphics error: ', GraphErrorMsg(ErrorCode));
    Writeln('Program aborted...');
    Halt(1);
  end;
  OutText('In Graphics mode. Press <RETURN>');
  Readln;
  RestoreCRTMode;
  Write('Now in text mode. Press <RETURN>');
  Readln;
  SetGraphMode(GraphMode);
  OutText('Back in Graphics mode. Press <RETURN>');
  Readln;
  CloseGraph;
end. { GrafTst2 }
```

Note that the *SetGraphMode* call near the end of the example resets all the graphics parameters (palette, current pointer, foreground, and background colors, and so on) to the default values.

The call to *CloseGraph* restores the video mode that was detected initially by *InitGraph* and frees the heap memory that was used to hold the graphics driver.

Heap management routines

Two heap management routines are used by the *Graph* unit: *GraphGetMem* and *GraphFreeMem*. *GraphGetMem* allocates memory for graphics device drivers, stroked fonts, and a scan buffer. *GraphFreeMem* deallocates the memory allocated to the drivers. The standard routines take the following form:

```
procedure GraphGetMem(var P: Pointer; Size: Word);  
{ Allocate memory for graphics }  
  
procedure GraphFreeMem(var P: Pointer; Size: Word);  
{ Deallocate memory for graphics }
```

Two pointers are defined by *Graph* that, by default, point to the two standard routines described here. The pointers are defined as follows:

```
var  
  GraphGetMemPtr: Pointer; { Pointer to memory allocation routine }  
  GraphFreeMemPtr: Pointer { Pointer to memory deallocation routine }
```

The *Graph* unit calls the heap management routines referenced by *GraphGetMemPtr* and *GraphFreeMemPtr* to allocate and deallocate memory for three different purposes:

- A multi-purpose graphics buffer whose size can be set by a call to *SetGraphBufSize* (default equals 4K)
- A device driver that is loaded by *InitGraph* (*.BGI files)
- A stroked font file that is loaded by *SetTextStyle* (*.CHR files)

The graphics buffer is always allocated on the heap. The device driver is allocated on the heap unless your program loads or links one in and calls *RegisterBGIDriver*. The font file is allocated on the heap when you select a stroked font using *SetTextStyle*—unless your program loads or links one in and calls *RegisterBGIfont*.

When the *Graph* unit is initialized, these pointers point to the standard graphics allocation and deallocation routines that are defined in the implementation section of the *Graph* unit. You can

insert your own memory-management routines by assigning these pointers the address of your routines. The user-defined routines must have the same parameter lists as the standard routines and must be *far* procedures. The following is an example of user-defined allocation and deallocation routines; notice the use of *MyExitProc* to automatically call *CloseGraph* when the program terminates:

```

program UserHeapManagement;
{ Illustrates how the user can steal the heap }
{ management routines used by the Graph unit. }

uses
    Graph;

var
    GraphDriver, GraphMode: Integer;
    ErrorCode: Integer;           { Stores GraphResult return code }
    PreGraphExitProc: Pointer;    { Saves original exit proc }

procedure MyGetMem(var P: Pointer; Size: Word); far;
{ Allocate memory for graphics device drivers, fonts, and scan buffer }
begin
    GetMem(P, Size)
end; { MyGetMem }

procedure MyFreeMem(var P: Pointer; Size: Word); far;
{ Deallocate memory for graphics device drivers, fonts, and scan
  buffer }
begin
    if P <> nil then                               { Don't free nil pointers! }
    begin
        FreeMem(P, Size);
        P := nil;
    end;
end; { MyFreeMem }

procedure MyExitProc; far;
{ Always gets called when program terminates }
begin
    ExitProc := PreGraphExitProc;                    { Restore original exit proc }
    CloseGraph;                                       { Do heap clean up }
end; { MyExitProc }

begin                                               { Install clean-up routine }
    PreGraphExitProc := ExitProc;
    ExitProc := @MyExitProc;

    GraphGetMemPtr := @MyGetMem;                     { Control memory allocation }
    GraphFreeMemPtr := @MyFreeMem;                   { Control memory deallocation }

    GraphDriver := Detect;

```



```

InitGraph(GraphDriver, GraphMode, '');
ErrorCode := GraphResult;
if ErrorCode <> grOk then
begin
  Writeln('Graphics error: ', GraphErrorMsg(ErrorCode));
  Readln;
  Halt(1);
end;
Line(0, 0, GetMaxX, GetMaxY);
OutTextXY(1, 1, 'Press <Return>:');
Readln;
end. { UserHeapManagement }

```

If your target is DOS protected mode, you'll be able to use programs like the previous example if you remember one thing: You must guarantee that any pointer returned by *GetMem* have a zero offset. You can do this with the *GlobalAllocPtr* function:

```

procedure MyGetMem(var P: Pointer; Size: Word); far;
var
  P : Pointer;
begin
  P := GlobalAllocPtr(HeapAllocFlags, Size);
  GetMem(P, 4096);
end; { MyGetMem }

```

Graph procedures and functions

The *Graph* unit provides many procedures and functions for use in your programs:

Table 19.3: Graph unit procedures and functions

<i>Arc</i>	Draws a circular arc from start angle to end angle using (x,y) as the center point.
<i>Bar</i>	Draws a bar using the current fill style and color.
<i>Bar3D</i>	Draws a 3-D bar using the current fill style and color.
<i>Circle</i>	Draws a circle using (x,y) as the center point.
<i>ClearDevice</i>	Clears the currently selected output device and homes the current pointer.
<i>ClearViewPort</i>	Clears the current viewport.
<i>CloseGraph</i>	Shuts down the graphics system.
<i>DetectGraph</i>	Checks the hardware and determines which graphics driver and mode to use.
<i>DrawPoly</i>	Draws the outline of a polygon using the current line style and color.
<i>Ellipse</i>	Draws an elliptical arc from start angle to end angle, using (x,y) as the center point.

Table 19.3: Graph unit procedures and functions (continued)

<i>FillEllipse</i>	Draws a filled ellipse using (x,y) as a center point and <i>XRadius</i> and <i>YRadius</i> as the horizontal and vertical axes.
<i>FillPoly</i>	Fills a polygon, using the scan converter.
<i>FloodFill</i>	Fills a bounded region using the current fill pattern and fill color.
<i>GetArcCoords</i>	Allows the user to inquire about the coordinates of the last <i>Arc</i> command.
<i>GetAspectRatio</i>	Returns the effective resolution of the graphics screen from which the aspect ratio (<i>Xasp:Yasp</i>) can be computed.
<i>GetBkColor</i>	Returns the current background color.
<i>GetColor</i>	Returns the current drawing color.
<i>GetDefaultPalette</i>	Returns the default hardware palette in a record of <i>PaletteType</i> .
<i>GetDriverName</i>	Returns a string containing the name of the current driver.
<i>GetFillPattern</i>	Returns the last fill pattern set by a call to <i>SetFillPattern</i> .
<i>GetFillSettings</i>	Allows the user to inquire about the current fill pattern and color as set by <i>SetFillStyle</i> or <i>SetFillPattern</i> .
<i>GetGraphMode</i>	Returns the current graphics mode.
<i>GetImage</i>	Saves a bit image of the specified region into a buffer.
<i>GetLineSettings</i>	Returns the current line style, line pattern, and line thickness as set by <i>SetLineStyle</i> .
<i>GetMaxColor</i>	Returns the highest color that can be passed to <i>SetColor</i> .
<i>GetMaxMode</i>	Returns the maximum mode number for the currently loaded driver.
<i>GetMaxX</i>	Returns the rightmost column (<i>x</i> resolution) of the current graphics driver and mode.
<i>GetMaxY</i>	Returns the bottommost row (<i>y</i> resolution) of the current graphics driver and mode.
<i>GetModeName</i>	Returns a string containing the name of the specified graphics mode.
<i>GetModeRange</i>	Returns the lowest and highest valid graphics mode for a given driver.
<i>GetPaletteSize</i>	Returns the size of the palette color lookup table.
<i>GetPixel</i>	Gets the pixel value at (x,y) .
<i>GetPalette</i>	Returns the current palette and its size.
<i>GetTextSettings</i>	Returns the current text font, direction, size, and justification as set by <i>SetTextStyle</i> and <i>SetTextJustify</i> .
<i>GetViewSettings</i>	Allows the user to inquire about the current viewport and clipping parameters.
<i>GetX</i>	Returns the x-coordinate of the current position (CP).
<i>GetY</i>	Returns the y-coordinate of the current position (CP).
<i>GraphDefaults</i>	Homes the current pointer (CP) and resets the graphics system.
<i>GraphErrorMsg</i>	Returns an error message string for the specified <i>ErrorCode</i> .

Table 19.3: Graph unit procedures and functions (continued)

<i>GraphResult</i>	Returns an error code for the last graphics operation.
<i>ImageSize</i>	Returns the number of bytes required to store a rectangular region of the screen.
<i>InstallUserDriver</i>	Installs a vendor-added device driver to the BGI device driver table.
<i>InstallUserFont</i>	Installs a new font file that isn't built into the BGI system.
<i>InitGraph</i>	Initializes the graphics system and puts the hardware into graphics mode.
<i>Line</i>	Draws a line from the (x1, y1) to (x2, y2).
<i>LineRel</i>	Draws a line to a point that is a relative distance from the current pointer (CP).
<i>LineTo</i>	Draws a line from the current pointer to (x,y).
<i>MoveRel</i>	Moves the current pointer (CP) a relative distance from its current position.
<i>MoveTo</i>	Moves the current graphics pointer (CP) to (x,y).
<i>OutText</i>	Sends a string to the output device at the current pointer.
<i>OutTextXY</i>	Sends a string to the output device.
<i>PieSlice</i>	Draws and fills a pie slice, using (x,y) as the center point and drawing from start angle to end angle.
<i>PutImage</i>	Puts a bit image onto the screen.
<i>PutPixel</i>	Plots a pixel at (x,y).
<i>Rectangle</i>	Draws a rectangle using the current line style and color.
<i>RegisterBGIdriver</i>	Registers a valid BGI driver with the graphics system.
<i>RegisterBGIfont</i>	Registers a valid BGI font with the graphics system.
<i>RestoreCrtMode</i>	Restores the original screen mode before graphics is initialized.
<i>Sector</i>	Draws and fills an elliptical sector.
<i>SetActivePage</i>	Sets the active page for graphics output.
<i>SetAllPalette</i>	Changes all palette colors as specified.
<i>SetAspectRatio</i>	Changes the default aspect ratio.
<i>SetBkColor</i>	Sets the current background color using the palette.
<i>SetColor</i>	Sets the current drawing color using the palette.
<i>SetFillPattern</i>	Selects a user-defined fill pattern.
<i>SetFillStyle</i>	Sets the fill pattern and color.
<i>SetGraphBufSize</i>	Lets you change the size of the buffer used for scan and flood fills.
<i>SetGraphMode</i>	Sets the system to graphics mode and clears the screen.
<i>SetLineStyle</i>	Sets the current line width and style.
<i>SetPalette</i>	Changes one palette color as specified by <i>ColorNum</i> and <i>Color</i> .
<i>SetRGBPalette</i>	Lets you modify palette entries for the IBM 8514 and the VGA drivers.
<i>SetTextJustify</i>	Sets text justification values used by <i>OutText</i> and <i>OutTextXY</i> .

Table 19.3: Graph unit procedures and functions (continued)

<i>SetTextStyle</i>	Sets the current text font, style, and character magnification factor.
<i>SetUserCharSize</i>	Lets you change the character width and height for stroked fonts.
<i>SetViewPort</i>	Sets the current output viewport or window for graphics output.
<i>SetVisualPage</i>	Sets the visual graphics page number.
<i>SetWriteMode</i>	Sets the writing mode (copy or xor) for lines drawn by <i>DrawPoly</i> , <i>Line</i> , <i>LineRel</i> , <i>LineTo</i> , and <i>Rectangle</i> .
<i>TextHeight</i>	Returns the height of a string in pixels.
<i>TextWidth</i>	Returns the width of a string in pixels.

For a detailed description of each procedure or function, refer to Chapter 1, "Library reference," in the *Programmer's Reference*.

Graph unit constants, types, and variables

The *Graph* unit defines several constants, types, and variables that your programs can use.

Constants

The *Graph* constants can be grouped by their function. To learn more about these constants, see Chapter 1, "Library reference," in the *Programmer's Reference*. Look up the constant under the group it belongs to. This table will help you identify the group you want:

Table 19.4
Graph unit constant groups

Constant group	Description
Driver and mode	Constants that specify video drivers and modes; used with <i>InitGraph</i> , <i>DetectGraph</i> , and <i>GetModeRange</i> .
grXXXX	Constants that identify the type of error returned from <i>GraphResult</i> .
Color	Constants that specify colors; used with <i>SetPalette</i> and <i>SetAllPalette</i> .
Color for <i>SetRGBPalette</i>	Constants used with <i>SetRGBPalette</i> to select standard EGA colors on an IBM 8514.
Line style	Constants used to determine a line style and thickness; used with <i>GetLineSettings</i> and <i>SetLineStyle</i> .

Table 19.4: Graph unit constant groups (continued)

Font control	Constants that identify fonts; used with <i>GetTextSettings</i> and <i>SetTextStyle</i> .
Justification	Constants that control horizontal and vertical justification; used with <i>SetTextJustify</i> .
Clipping	Constants that control clipping; used with <i>SetViewPort</i> .
Bar	Constants that control the drawing of a 3-D top on a bar; used with <i>Bar3D</i> .
Fill pattern	Constants that determine the pattern used to fill an area; used with <i>GetFillSettings</i> and <i>SetFillStyle</i> .
BitBlt operators	Operators (copy, xor , or , and not) used with <i>PutImage</i> and, <i>SetWriteMode</i> .
MaxColors	The constant that defines the maximum number of colors used with <i>GetPalette</i> , <i>GetDefaultPalette</i> , and <i>SetAllPalette</i> .

For example, to find the constant you need to change the background screen color to green, look under Color Constants in Chapter 1, "Library reference," in the *Programmer's Reference*.

Types

The *Graph* unit defines these types:

Table 19.5
Graph unit types

Type	Description
<i>PaletteType</i>	The record that defines the size and colors of the palette; used by <i>GetPalette</i> , <i>GetDefaultPalette</i> , and <i>SetAllPalette</i> .
<i>LineSettingsType</i>	The record that defines the style, pattern, and thickness of a line; used by <i>GetLineSettings</i> .
<i>TextSettingsType</i>	The record that defines the text; used by <i>GetTextSettings</i> .
<i>FillSettingsType</i>	The record that defines the pattern and color used to fill an area; used by <i>GetFillSettings</i> .
<i>FillPatternType</i>	The record that defines a user-defined fill pattern; used by <i>GetFillPattern</i> and <i>SetFillPattern</i> .
<i>PointType</i>	A type defined for your convenience.

Table 19.5: Graph unit types (continued)

<i>ViewPortType</i>	A record that reports the status of the current viewport; used by <i>GetViewSettings</i> .
<i>ArcCoordsType</i>	A record that retrieves information about the last call to <i>Arc</i> or <i>Ellipse</i> ; used by <i>GetArcCoords</i> .

Variables

The *Graph* unit has two variables you can use: *GraphGetMemPtr* and *GraphFreeMemPtr*. They are used by heap-management routines. Read about them in Chapter 1, "Library reference," in the *Programmer's Reference*.

Using overlays

Overlays are parts of a program that share a common memory area. Only the parts of the program that are required for a given function reside in memory at the same time; they can overwrite each other during execution.



Overlays are useful in DOS real-mode programs only. Because Windows manages memory for Windows applications and the run-time manager (RTM.EXE) manages memory for DOS protected-mode applications, and these memory managers include full sets of overlay-management services, you won't need to use overlays in your Windows or protected-mode programs.

Overlays can significantly reduce a program's total run-time memory requirements. In fact, with overlays you can execute programs that are much larger than the total available memory because only parts of the program reside in memory at any given time.

Borland Pascal manages overlays at the unit level; this is the smallest part of a program that can be made into an overlay. When an overlaid program is compiled, Borland Pascal generates an overlay file (extension .OVR) in addition to the executable file (extension .EXE). The .EXE file contains the static (nonoverlaid) parts of the program, and the .OVR file contains all the overlaid units that will be swapped in and out of memory during program execution.

Except for a few programming rules, an overlaid unit is identical to a nonoverlaid unit. In fact, as long as you observe these rules,

you don't even need to recompile a unit to make it into an overlay. The decision of whether or not to overlay a unit is made by the program that uses the unit.

When an overlay is loaded into memory, it's placed in the overlay buffer, which resides in memory between the stack segment and the heap. By default, the size of the overlay buffer is as small as possible, but it may be easily increased at run time by allocating additional space from the heap. Like the data segment and the minimum heap size, the default overlay-buffer size is allocated when the .EXE is loaded. If enough memory isn't available, an error message will be displayed by DOS ("Program too big to fit in memory") or by the IDE ("Not enough memory to run program").

One very important option of the overlay manager is the ability to load the overlay file into expanded memory when sufficient space is available. Borland Pascal supports version 3.2 or later of the Lotus/Intel/Microsoft Expanded Memory Specification (EMS) for this purpose. Once placed into EMS, the overlay file is closed, and subsequent overlay loads are reduced to fast in-memory transfers.

The overlay manager

Borland Pascal's overlay manager is implemented by the *Overlay* standard unit. The buffer-management techniques used by the *Overlay* unit are very advanced, and always guarantee optimal performance in the available memory. For example, the overlay manager always keeps as many overlays as possible in the overlay buffer to reduce the chance of having to read an overlay from disk. Once an overlay is loaded, a call to one of its routines executes just as fast as a call to a nonoverlaid routine. Also, when the overlay manager needs to dispose of an overlay to make room for another, it attempts to first dispose of overlays that are inactive (ones that have no active routines at that time).

To implement its advanced overlay-management techniques, Borland Pascal requires that you observe two important rules when writing overlaid programs:

- All overlaid units must include a **{SO+}** directive, which causes the compiler to ensure that the generated code can be overlaid.

- Whenever a call is made to an overlaid procedure or function, you must ensure that all currently active procedures and functions use the far call model.

Both rules are explained further in a section entitled “Designing overlaid programs,” beginning on page 245. For now, just note that you can easily satisfy these requirements by placing a `{O+,F+}` compiler directive at the beginning of all overlaid units, and a `{F+}` compiler directive at the beginning of all other units and the main program.



Failing to observe the far call requirement in an overlaid program causes unpredictable and possibly catastrophic results when the program is executed.

The `{O unitname}` compiler directive is used in a program to indicate which units to overlay. This directive must be placed after the program’s **uses** clause, and the **uses** clause must name the *Overlay* standard unit before any of the overlaid units. Here is an example:

```
program Editor;
{F+}           { Force FAR calls for all procedures & functions }

uses
  Overlay, Crt, Dos, EdInOut, EdFormat, EdPrint, EdFind, EdMain;

{$O EdInOut}
{$O EdFormat}
{$O EdPrint}
{$O EdFind}
{$O EdMain}
```



The compiler reports an error if you attempt to overlay a unit that wasn’t compiled in the `{O+}` state. Of the standard units, the only one that can be overlaid is *Dos*; the other standard units, can’t be overlaid. If you’re using the real-mode IDE (TURBO.EXE), you must compile programs containing overlaid units to disk; the compiler reports an error if you attempt to compile such programs to memory.

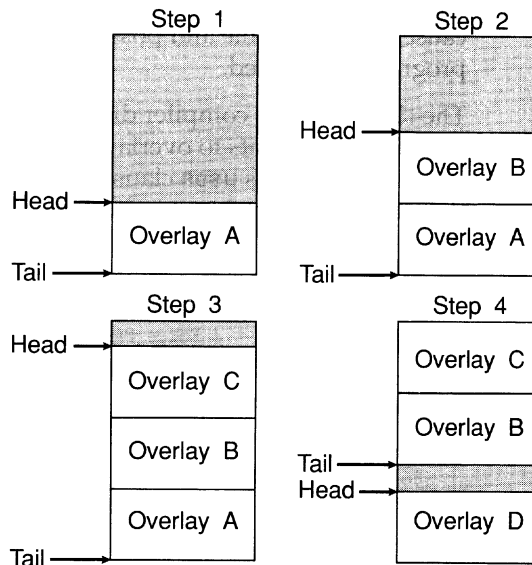
Overlay-buffer management

The Borland Pascal overlay buffer is best described as a ring buffer that has a head pointer and a tail pointer. Overlays are always loaded at the head of the buffer, pushing “older” ones toward the tail. When the buffer becomes full (that is, when there

isn't enough free space between the head and the tail), overlays are disposed of at the tail to make room for new ones.

Because ordinary memory isn't circular in nature, the actual implementation of the overlay buffer involves a few more steps to make the buffer appear to be a ring. Figure 20.1 illustrates the process. The figure shows a progression of overlays being loaded into an initially empty overlay buffer. Overlay *A* is loaded first, followed by *B*, then *C*, and finally *D*. Shaded areas indicate free buffer space.

Figure 20.1
Loading and disposing of
overlays



As you can see, a couple of interesting things happen in the transition from step 3 to step 4. First, the head pointer wraps around to the bottom of the overlay buffer, causing the overlay manager to slide all loaded overlays (and the tail pointer) upward. This sliding is required to keep the area between the head pointer and the tail pointer free. Second, to load overlay *D*, the overlay manager has to dispose of overlay *A* from the tail of the buffer. Overlay *A* in this case is the least recently loaded overlay and, therefore, the best choice for disposal when something has to go. The overlay manager continues to dispose of overlays at the tail to make room for new ones at the head, and each time the head pointer wraps around, the sliding operation repeats.

Although this is the default mode of operation for Borland Pascal's overlay manager, you can use an optional optimization of the overlay-management algorithm.

Imagine that overlay *A* contains a number of frequently used routines. Even though these routines are used all the time, *A* will still be thrown out of the overlay buffer occasionally, only to be reloaded again shortly thereafter. The overlay manager knows nothing about the *frequency* of calls to routines in *A*—only that a call is made to a routine in *A* and *A* isn't in memory, so it has to load *A*. One solution to this problem might be to trap every call to routines in *A* and then, at each call, move *A* to the head of the overlay buffer to reflect its new status as the most recently used overlay. Intercepting calls this way is very costly in terms of execution speed and, in some cases, can slow down the application even more than the additional overlay load operations will.

Borland Pascal offers a solution that incurs almost no performance overhead and yet successfully identifies frequently used overlays that shouldn't be unloaded: When an overlay gets close to the tail of the overlay buffer, it's put on "probation." If, during this probationary period, a call is made to a routine in the overlay, it's "reprieved," and isn't disposed of when it reaches the tail of the overlay buffer. Instead, it's moved to the head of the buffer and gets another free ride around the overlay-buffer ring. On the other hand, if no calls are made to an overlay during its probationary period, thereby indicating infrequent use, the overlay is disposed of when it reaches the tail of the overlay buffer.

The net effect of the probation/reprieve scheme is that frequently used overlays are kept in the overlay buffer at the cost of intercepting just *one* call every time the overlay gets close to the tail of the overlay buffer.

Two overlay-manager routines, *OvrSetRetry* and *OvrGetRetry*, control the probation/reprieve mechanism. *OvrSetRetry* sets the size of the area in the overlay buffer to keep on probation and *OvrGetRetry* returns the current setting. If an overlay falls within the last *OvrGetRetry* bytes before the overlay-buffer tail, it's automatically put on probation. Any free space in the overlay buffer is considered part of the probation area.

Overlay procedures and functions

The *Overlay* unit defines a few procedures and functions; find their definitions as well as more details in Chapter 1, “Library reference,” in the *Programmer’s Reference*.

Table 20.1
Overlay unit procedures and functions

Procedure or function	Description
<i>OvrClearBuf</i>	Clears the overlay buffer.
<i>OvrGetBuf</i>	Returns the current size of the overlay buffer.
<i>OvrGetRetry</i>	Returns the current size of the probation area, the value last set with <i>OvrSetRetry</i> .
<i>OvrInit</i>	Initializes the overlay manager and opens the overlay file.
<i>OvrInitEMS</i>	Loads the overlay file into EMS.
<i>OvrSetBuf</i>	Sets the size of the overlay buffer.
<i>OvrSetRetry</i>	Sets the size of the “probation area” in the overlay buffer.

Variables and constants

The *Overlay* unit defines five variables:

Table 20.2
Overlay unit variables

Variable	Description
<i>OvrFileMode</i>	Determines the access code to pass to DOS when the overlay file is opened.
<i>OvrLoadCount</i>	The variable incremented each time an overlay is loaded.
<i>OvrReadBuf</i>	The procedure variable that lets you intercept overlay load operations.
<i>OvrResult</i>	The variable that holds the result code when an <i>Overlay</i> procedure executes.
<i>OvrTrapCount</i>	The variable incremented each time an overlaid routine is intercepted by the overlay manager.

Find the values of these variables in Chapter 1, “Library reference,” in the *Programmer’s Reference*.

Result codes

Errors in the *Overlay* unit are reported through the *OvrResult* variable. Look up “ovrXXXX constants” in Chapter 1, “Library reference,” in the *Programmer’s Reference* to find *OvrResult* values.

Designing overlaid programs

This section provides some important information on designing programs with overlays. Look it over carefully, because a number of the issues discussed are vital to well-behaved overlaid applications.

Overlay code generation

Borland Pascal allows a unit to be overlaid only if it was compiled with **{SO+}**. In this state, the code generator takes special precautions when passing string and set constant parameters from one overlaid procedure or function to another. For example, if *UnitA* contains a procedure with the following header:

```
procedure WriteStr(S: string);
```

and if *UnitB* contains the statement

```
WriteStr('Hello world...');
```

then Borland Pascal places the string constant ‘Hello world...’ in *UnitB*’s code segment, and passes a pointer to it to the *WriteStr* procedure. If both units are overlaid, this doesn’t work because, at the call to *WriteStr*, *UnitB*’s code segment may be overwritten by *UnitA*’s and the string pointer becomes invalid. The **{SO+}** directive is used to avoid such problems; whenever Borland Pascal detects a call from one unit compiled with **{SO+}** to another unit compiled with **{SO+}**, the compiler copies all code-segment-based constants into stack temporaries before passing pointers to them.

The use of **{SO+}** in a unit doesn’t force you to overlay that unit. It just instructs Borland Pascal to ensure that the unit can be overlaid, if so desired. If you develop units that you plan to use in overlaid as well as nonoverlaid applications, compiling them with

{\$O+} ensures that you can do both with just one version of the unit.

The far call requirement

At any call to an overlaid procedure or function in another module, you *must* guarantee that all currently active procedures or functions use the far call model.

This is best illustrated by example: Assume that *OvrA* is a procedure in an overlaid unit, and that *MainB* and *MainC* are procedures in the main program. If the main program calls *MainC*, which calls *MainB*, which then calls *OvrA*, then at the call to *OvrA*, *MainB* and *MainC* are active (they have not yet returned) and they are required to use the far call model. Being declared in the main program, *MainB* and *MainC* would normally use the near call model. In this case, a **{\$F+}** compiler directive must be used to force the far call model into effect.

The easiest way to satisfy the far call requirement is to place a **{\$F+}** directive at the beginning of the main program and each unit. Alternatively, you can change the default **\$F** setting to **{\$F+}** using a **/\$F+** command-line directive or the Force Far Calls check box in the Options | Compiler dialog box. Compared to the cost of mixing near and far calls, using far calls exclusively costs little—one extra word of stack space per active procedure and one extra byte per call.

Initializing the overlay manager

Here we'll take a look at some examples of how to initialize the overlay manager. Place the initialization code before the first call to an overlaid routine. Typically you would do this at the beginning of the program's statement part.

The following code shows just how little you need to initialize the overlay manager:

```
begin
  OvrInit('EDITOR.OVR');
end;
```

No error checks are made. If there isn't enough memory for the overlay buffer or if the overlay file was not found, run-time error 208 ("Overlay manager not installed") occurs when you attempt to call an overlaid routine.

Here's another simple example that expands on the previous one:

```
begin
  OvrInit('EDITOR.OVR');
  OvrInitEMS;
end;
```

In this case, provided there is enough memory for the overlay buffer and that the overlay file can be located, the overlay manager checks to see if EMS memory is available. If it is, it loads the overlay file into EMS.

The initial overlay-buffer size is as small as possible, or in other words, just big enough to contain the largest overlay. This may be adequate for some applications, but imagine a situation where a particular function of a program is implemented through two or more units, each of which is overlaid. If the total size of those units is larger than the largest overlay, a substantial amount of swapping will occur if the units make frequent calls to each other.

The solution is to increase the size of the overlay buffer so that enough memory is available at any given time to contain all overlays that make frequent calls to each other. The following code demonstrates the use of *OvrSetBuf* to increase the overlay-buffer size:

```
const
  OvrMaxSize = 80000;
begin
  OvrInit('EDITOR.OVR');
  OvrInitEMS;
  OvrSetBuf(OvrMaxSize);
end;
```

There is no general formula for determining the ideal overlay-buffer size. Only an intimate knowledge of the application and a bit of experimenting results in a suitable value.



Using *OvrInitEMS* to place the overlay file in EMS doesn't eliminate the need for an overlay buffer. Overlays must still be copied from EMS into "normal" memory in the overlay buffer before they can be executed, but because such in-memory transfers are significantly faster than disk reads, there is less need to increase the size of the overlay buffer.

Remember, *OvrSetBuf* expands the overlay buffer by shrinking the heap. Therefore, the heap must be empty or *OvrSetBuf* has no

effect. If you are using the *Graph* unit, call *OvrSetBuf* before you call *InitGraph*, which allocates memory on the heap.

Here's a rather elaborate example of overlay-manager initialization with full error-checking:

```
const
  OvrMaxSize = 80000;
var
  OvrName: string[79];
  Size: Longint;
begin
  OvrName := 'EDITOR.OVR';
  repeat
    OvrInit(OvrName);
    if OvrResult = OvrNotFound then
      begin
        Writeln('Overlay file not found: ', OvrName, '.');
        Write('Enter correct overlay file name: ');
        Readln(OvrName);
      end;
  until OvrResult <> OvrNotFound;
  if OvrResult <> OvrOk then
    begin
      Writeln('Overlay manager error. ');
      Halt(1);
    end;
  OvrInitEMS;
  if OvrResult <> OvrOk then
    begin
      case OvrResult of
        ovrIOError: Write('Overlay file I/O error');
        ovrNoEMSDriver: Write('EMS driver not installed');
        ovrNoEMSMemory: Write('Not enough EMS memory');
      end;
      Write('. Press Enter...');
      Readln;
    end;
  OvrSetBuf(OvrMaxSize);
end;
```

First, if the default overlay file name isn't correct, the user is repeatedly prompted for a correct file name.

Next, a check is made for other errors that might have occurred during initialization. If an error is detected, the program halts, because errors in *OvrInit* are fatal. (If they are ignored, a run-time error occurs upon the first call to an overlaid routine.)

Assuming successful initialization, a call to *OvrInitEMS* is made to load the overlay file into EMS if possible. In case of error, a diagnostic message is displayed, but the program isn't halted. Instead, it continues to read overlays from disk.

Finally, *OvrSetBuf* is called to set the overlay buffer size to a suitable value, determined through analysis and experimentation with the particular application. Errors from *OvrSetBuf* are ignored, although *OvrResult* might return an error code of -3 (*OvrNoMemory*). If there isn't enough memory, the overlay manager continues to use the minimum buffer that was allocated when the program started.

Initialization sections

Like static units, overlaid units can have an initialization section. Although overlaid initialization code is no different from normal overlaid code, the overlay manager must be initialized first so it can load and execute overlaid units.

Referring to the earlier *Editor* program, assume that the *EdInOut* and *EdMain* units have initialization code. This requires that *OvrInit* is called before *EdInOut*'s initialization code. The only way to do that is to create an additional nonoverlaid unit that goes before *EdInOut* and calls *OvrInit* in its initialization section:

```
unit EdInit;
interface
implementation
uses Overlay;
const
    OvrMaxSize = 80000;
begin
    OvrInit('EDITOR.OVR');
    OvrInitEMS;
    OvrSetBuf(OvrMaxSize);
end.
```

The *EdInit* unit must be listed in the program's **uses** clause before any of the overlaid units:

```
program Editor;
{$F+}

uses Overlay, Crt, Dos, EdInit, EdInOut, EdFormat, EdPrint, EdFind,
    EdMain;
```

```
{ $O EdInOut }
{ $O EdFormat }
{ $O EdPrint }
{ $O EdFind }
{ $O EdMain }
```

In general, although initialization code in overlaid units is indeed possible, you should avoid it for a number of reasons.

First, the initialization code, even though it's executed only once, is a part of the overlay, and occupies overlay-buffer space whenever the overlay is loaded. Second, if a number of overlaid units have initialization code, each of them will have to be read into memory when the program starts.

A much better approach is to gather all the initialization code into an overlaid initialization unit, which is called once at the beginning of the program, and then never referenced again.

What not to overlay

Certain units can't be overlaid. In particular, don't try to overlay the following:

- Units compiled in the **{ \$O- }** state. The compiler reports an error if you attempt to overlay a unit that wasn't compiled with **{ \$O+ }**. Such nonoverlay units include *System*, *Overlay*, *Crt*, *Graph*, *Turbo3*, and *Graph3*.
- Units that contain interrupt handlers. Due to the non-reentrant nature of the DOS operating system, units that implement **interrupt** procedures should not be overlaid. An example of such a unit is the *Crt* standard unit, which implements a *Ctrl+Break* interrupt handler.
- BGI drivers or fonts registered with calls to *RegisterBGIdriver* or *RegisterBGIfont*.

Calling overlaid routines via procedure pointers is fully supported by Borland Pascal's overlay manager. Examples of the use of procedure pointers include exit procedures and text-file device drivers.

The overlay manager also supports passing overlaid procedures and functions as procedural parameters and assigning overlaid procedures and functions to procedural type variables.

Debugging overlays

Most debuggers have very limited overlay debugging capabilities, if any at all. This isn't so with Borland Pascal and Turbo Debugger. The integrated debugger fully supports single-stepping and breakpoints in overlays in a manner completely transparent to you. By using overlays, you can easily engineer and debug huge applications—all from inside the IDE or by using Turbo Debugger.

External routines in overlays

Like normal Pascal procedures and functions, **external** assembly language routines must observe certain programming rules to work correctly with the overlay manager.

If an assembly language routine makes calls to *any* overlaid procedures or functions, the assembly language routine must use the far model, and it must set up a stack frame using the BP register. For example, assuming that *OtherProc* is an overlaid procedure in another unit, and that the assembly language routine *ExternProc* calls it, then *ExternProc* must use the FAR model and set up a stack frame. For example,

```
ExternProc      PROC      FAR
                PUSH     BP                ;Save BP
                MOV     BP,SP            ;Set up stack frame
                SUB     SP,LocalSize    ;Allocate local variables
                :
                CALL    OtherProc       ;Call another overlaid unit
                :
                MOV     SP,BP            ;Dispose of local variables
                POP     BP                ;Restore BP
                RET     ParamSize       ;Return
ExternProc      ENDP
```

LocalSize is the size of the local variables and *ParamSize* is the size of the parameters. If *LocalSize* is zero, you can omit the two lines to allocate and dispose of local variables.

These requirements are the same if *ExternProc* makes *indirect* references to overlaid procedures or functions. For example, if *OtherProc* makes calls to overlaid procedures or functions, but

isn't itself overlaid, *ExternProc* must still use the FAR model and still has to set up a stack frame.

When an assembly language routine doesn't make any direct or indirect references to overlaid procedures or functions, there are no special requirements; the assembly language routine is free to use the near model and it doesn't have to set up a stack frame.

Overlaid assembly language routines should *not* create variables in the code segment, because any modifications made to an overlaid code segment are lost when the overlay is disposed of. Likewise, pointers to objects based in an overlaid code segment can't be expected to remain valid across calls to other overlays, because the overlay manager freely moves around and disposes of overlaid code segments.

Installing an overlay read function

The *OvrReadBuf* procedure variable lets you intercept overlay load operations. For example, you can implement error handling or check that a removable disk is present. Whenever the overlay manager needs to read an overlay, it calls the function whose address is stored in *OvrReadBuf*. If the function returns zero, the overlay manager assumes that the operation was successful; if the function result is nonzero, the compiler generates run-time error 209. The *OvrSeg* parameter indicates what overlay to load, but as you'll see later, you won't need to access this information.

Don't attempt to call any overlaid routines from within your overlay read function—such calls will crash the system.

To install your own overlay read function, you must first save the previous value of *OvrReadBuf* in a variable of type *OvrReadFunc*, and then assign your overlay read function to *OvrReadBuf*. Within your read function, you should call the saved read function to perform the actual load operation. Any validations you want to perform, such as checking that a removable disk is present, should go before the call to the saved read function, and any error checking should go after the call.

The code to install an overlay read function should go right after the call to *OvrInit*; at this point, *OvrReadBuf* will contain the address of the default disk read function.

If you also call *OvrInitEMS*, it uses your read function to read overlays from disk into EMS memory, and if no errors occur, it stores the address of the default EMS read function in

OvrReadBuf. If you also want to override the EMS read function, simply repeat the installation process after the call to *OvrInitEMS*.

The default disk read function returns zero if it succeeds, or a DOS error code if it fails. Likewise, the default EMS read function returns 0 if it succeeds, or an EMS error code (ranging from \$80 through \$FF) if it fails. For details on DOS error codes, refer to the *Programmer's Reference*. For details on EMS error codes, refer to your Expanded Memory Specification documentation.

The following code fragment demonstrates how to write and install an overlay-read function. The new overlay-read function repeatedly calls the saved overlay-read function until no errors occur.

Any errors are passed to the *DOSError* or *EMSError* procedures (not shown here) so that they can present the error to the user. Notice how the *OvrSeg* parameter is just passed on to the saved overlay-read function and never directly handled by the new overlay-read function.

```
uses Overlay;
var
  SaveOvrRead: OvrReadFunc;
  UsingEMS: Boolean;

function MyOvrRead(OvrSeg: Word): Integer; far;
var
  E: Integer;
begin
  repeat
    E := SaveOvrRead(OvrSeg);
    if E <> 0 then
      if UsingEMS then
        EMSError(E) else DOSError(E);
    until E = 0;
  MyOvrRead := 0;
end;

begin
  OvrInit('MYPROG.OVR');
  SaveOvrRead := OvrReadBuf;           { Save disk default }
  OvrReadBuf := MyOvrRead;            { Install ours }
  UsingEMS := False;
  OvrInitEMS;
  if OvrResult = OvrOK then
```

```

begin
  SaveOvrRead := OvrReadBuf;           { Save EMS default }
  OvrReadBuf := MyOvrRead;             { Install ours }
  UsingEMS := True;
end;
:
end.

```

Overlays in .EXE files

Borland Pascal allows you to store your overlays at the end of your application's .EXE file rather than in a separate .OVR file. To attach an .OVR file to the end of an .EXE file, use the DOS COPY command with a **/B** command-line switch, for example,

```
COPY/B MYPROG.EXE + MYPROG.OVR
```

You must make sure that the .EXE file was compiled *without* Turbo Debugger debug information. In the IDE, make sure the Standalone option isn't checked in Options | Debugger. With the command-line version of the compiler, make sure you don't specify a **/V** switch.

To read overlays from the end of an .EXE file instead of from a separate .OVR file, simply specify the .EXE file name in the call to *OvrInit*. If you are running under DOS 3.x or greater, you can use the *ParamStr* standard function to obtain the name of the .EXE file; for example,

```
OvrInit(ParamStr(0));
```

P

A

R

T

3

Inside Borland Pascal

Memory issues

This chapter is about Borland Pascal and memory. Borland Pascal can produce DOS real-mode, DOS protected-mode, and Windows applications; each type of application uses memory differently. This chapter explains how each of these types of programs use memory, and it also describes the internal data formats used in Borland Pascal.

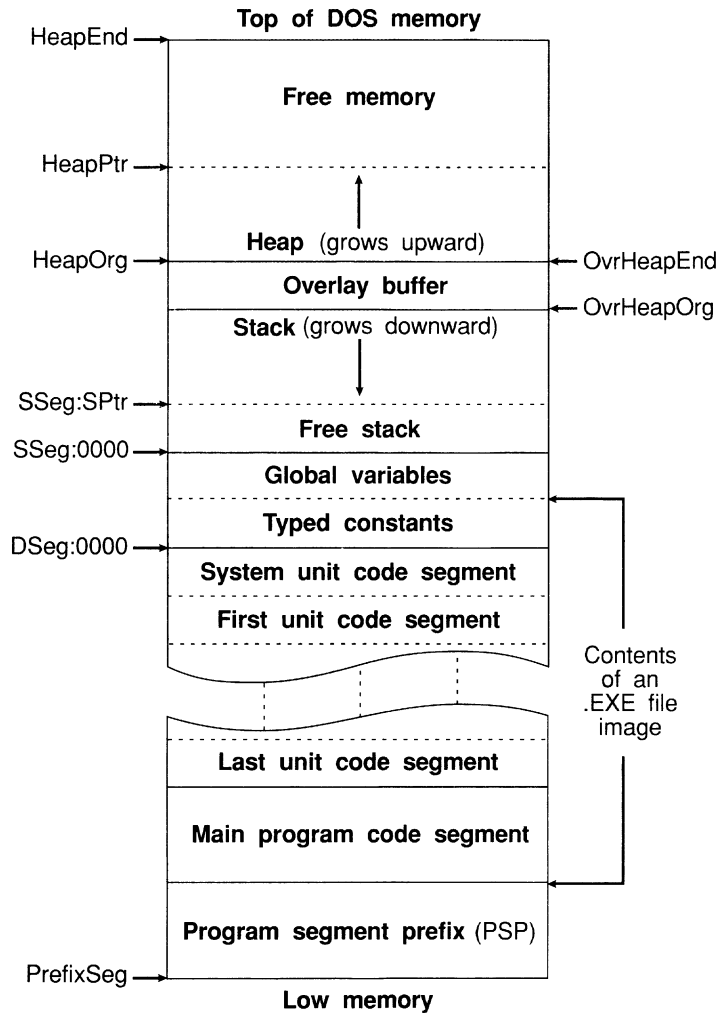
Memory issues for DOS real-mode programs

Figure 21.1 depicts the memory map of a DOS real-mode program written in Borland Pascal.

The Program Segment Prefix (PSP) is a 256-byte area built by DOS when the .EXE file is loaded. The segment address of PSP is stored in the predeclared variable *PrefixSeg*.

Each module, which includes the main program and each unit, has its own code segment. The main program occupies the first code segment. The code segments that follow it are occupied by the units in reverse order from how they are listed in the **uses** clause. The last code segment is occupied by the *System* unit. The size of a single code segment can't exceed 64K, but the total size of the code is limited only by the available memory.

Figure 21.1
Memory map for a DOS real-mode program



The data segment (addressed through DS) contains all typed constants followed by all global variables. The DS register is never changed during program execution. The size of the data segment can't exceed 64K.

On entry to the program, the stack segment register (SS) and the stack pointer (SP) are loaded so that SS:SP points to the first byte past the stack segment. The SS register is never changed during program execution, but SP can move downward until it reaches

the bottom of the segment. The size of the stack segment can't exceed 64K; the default size is 16K, but this can be changed with a **\$M** compiler directive.

The *Overlay* standard unit uses the overlay buffer to store overlaid code. The default size of the overlay buffer corresponds to the size of the largest overlay in the program; if the program has no overlays, the size of the overlay buffer is zero. The size of the overlay buffer can be increased through a call to the *OvrSetBuf* routine in the *Overlay* unit; in that case, the size of the heap is decreased accordingly, by moving *HeapOrg* upwards.

The heap stores *dynamic variables*, that is, variables allocated through calls to the *New* and *GetMem* standard procedures. It occupies all or some of the free memory left when a program is executed. The actual size of the heap depends on the minimum and maximum heap values, which can be set with the **\$M** compiler directive. Its size is guaranteed to be at least the minimum heap size and never more than the maximum heap size. If the minimum amount of memory isn't available, the program doesn't execute. The default heap minimum is 0 bytes, and the default heap maximum is 640K; this means that the heap occupies all remaining memory by default.

As you might expect, the heap manager (which is part of Borland Pascal's run-time library) manages the heap. It's described in detail in the following section.

The DOS heap manager

The heap is a stack-like structure that grows from low memory in the heap segment. The bottom of the heap is stored in the variable *HeapOrg*, and the top of the heap, corresponding to the bottom of free memory, is stored in the variable *HeapPtr*. Each time a dynamic variable is allocated on the heap (via *New* or *GetMem*), the heap manager moves *HeapPtr* upward by the size of the variable, in effect stacking the dynamic variables on top of each other.

HeapPtr is always normalized after each operation, forcing the offset part into the range \$0000 to \$000F. The maximum size of a single variable that can be allocated on the heap is 65,519 bytes (corresponding to \$10000 minus \$000F), because every variable must be completely contained in a single segment.


Disposal methods

The dynamic variables stored on the heap are disposed of in one of two ways: (1) through *Dispose* or *FreeMem* or (2) through *Mark* and *Release*. The simplest scheme is that of *Mark* and *Release*; for example, if the following statements are executed:

```
New(Ptr1);  
New(Ptr2);  
Mark(P);  
New(Ptr3);  
New(Ptr4);  
New(Ptr5);
```

the layout of the heap will then look like the this figure:


Figure 21.2
Disposal method using Mark
and Release



Executing Release(HeapOrg) completely disposes of the entire heap because HeapOrg points to the bottom of the heap.

The *Mark(P)* statement marks the state of the heap just before *Ptr3* is allocated (by storing the current *HeapPtr* in *P*). If the statement *Release(P)* is executed, the heap layout becomes like that of Figure 21.3, effectively disposing of all pointers allocated since the call to *Mark*.


Figure 21.3
Heap layout with `Release(P)`
executed



For applications that dispose of pointers in exactly the reverse order of allocation, the *Mark* and *Release* procedures are very efficient. Yet most programs tend to allocate and dispose of pointers in a more random manner, requiring the more sophisticated management technique implemented by *Dispose* and *FreeMem*. These procedures allow an application to dispose of any pointer at any time.


When a dynamic variable that's not the topmost variable on the heap is disposed of through *Dispose* or *FreeMem*, the heap becomes fragmented. Assuming that the same statement sequence has been executed, then after executing *Dispose(Ptr3)*, a "hole" is created in the middle of the heap (see Figure 21.4).

Figure 21.4
Creating a "hole" in the heap




If *New(Ptr3)* had been executed now, it would again occupy the same memory area. On the other hand, executing *Dispose(Ptr4)* enlarges the free block, because *Ptr3* and *Ptr4* were neighboring blocks (see Figure 21.5).

Figure 21.5
Enlarging the free block



Finally, executing *Dispose(Ptr5)* first creates an even bigger free block, and then lowers *HeapPtr*. This, in effect, releases the free block, because the last valid pointer is now *Ptr2* (see Figure 21.6).

Figure 21.6
Releasing the free block



The heap is now in the same state as it would be after executing *Release(P)*, as shown in Figure 21.3. The free blocks created and destroyed in the process were tracked for possible reuse, however.

The free list The addresses and sizes of the free blocks generated by *Dispose* and *FreeMem* operations are kept on a *free list*. Whenever a dynamic variable is allocated, the free list is checked before the heap is expanded. If a free block of adequate size exists (it's greater than or equal to the size of the requested block size), it's used.



The *Release* procedure always clears the free list, thereby causing the heap manager to “forget” about any free blocks that might exist below the heap pointer. If you mix calls to *Mark* and *Release* with calls to *Dispose* and *FreeMem*, you must ensure that no such free blocks exist.

The *FreeList* variable in the *System* unit points to the first free block in the heap. This block contains a pointer to the next free block, which contains a pointer to the following free block, and so on. The last free block contains a pointer to the top of the heap (that is, to the location given by *HeapPtr*). If there are no free blocks on the free list, *FreeList* will be equal to *HeapPtr*.

The format of the first 8 bytes of a free block are given by the *TFreeRec* type as follows:

```

type
  PFreeRec = ^TFreeRec;
  TFreeRec = record
    Next: PFreeRec;
    Size: Pointer;
  end;

```

The *Next* field points to the next free block, or to the same location as *HeapPtr* if the block is the last free block. The *Size* field encodes the size of the free block. The value in *Size* isn't a normal 32-bit value; rather, it's a "normalized" pointer value with a count of free paragraphs (16-byte blocks) in the high word, and a count of free bytes (between 0 and 15) in the low word. The following *BlockSize* function converts a *Size* field value to a normal *Longint* value:

```

function BlockSize(Size: Pointer): Longint;
type
  PtrRec = record Lo, Hi: Word end;
begin
  BlockSize := Longint(PtrRec(Size).Hi) * 16 + PtrRec(Size).Lo;
end;

```

To guarantee that there will always be room for a *TFreeRec* at the beginning of a free block, the heap manager rounds the size of every block allocated by *New* or *GetMem* upwards to an 8-byte boundary. Eight bytes are allocated for blocks of size 1..8, 16 bytes are allocated for blocks of size 9..16, and so on. This might seem an excessive waste of memory at first, and indeed it would be if every block was just 1 byte in size. Blocks are typically larger, however, and so the relative size of the unused space is less.

The 8-byte granularity factor ensures that a number of random allocations and deallocations of blocks of varying small sizes, such as would be typical for variable-length line records in a text-processing program, do not heavily fragment the heap. For example, say a 50-byte block is allocated and disposed of, thereby becoming an entry on the free list. The block would have been rounded to 56 bytes (7*8), and a later request to allocate anywhere from 49 to 56 bytes would completely reuse the block, instead of leaving 1 to 7 bytes of free (but most likely unusable) space, which would fragment the heap.

The `HeapError` variable The `HeapError` variable allows you to install a heap-error function, which gets called whenever the heap manager can't complete an allocation request. `HeapError` is a pointer that points to a function with the following header:

```
function HeapFunc(Size: Word): Integer; far;
```

Note that the `far` directive forces the heap-error function to use the FAR call model.

The heap-error function is installed by assigning its address to the `HeapError` variable:

```
HeapError := @HeapFunc;
```

The heap-error function gets called whenever a call to `New` or `GetMem` can't complete the request. The `Size` parameter contains the size of the block that couldn't be allocated, and the heap-error function should attempt to free a block of at least that size.

Depending on its success, the heap-error function should return 0, 1, or 2. A return of 0 indicates failure, causing a run-time error to occur immediately. A return of 1 also indicates failure, but instead of a run-time error, it causes `New` or `GetMem` to return a `nil` pointer. Finally, a return of 2 indicates success and causes a retry (which could also cause another call to the heap-error function).

The standard heap-error function always returns 0 and causes a run-time error whenever a call to `New` or `GetMem` can't be completed. For many applications, however, the simple heap-error function that follows is more appropriate:

```
function HeapFunc(Size: Word): Integer; far;
begin
  HeapFunc := 1;
end;
```

When installed, this function causes `New` or `GetMem` to return `nil` when they can't complete the request, instead of aborting the program.



A call to the heap-error function with a `Size` parameter of 0 means that the heap manager has just expanded the heap by moving `HeapPtr` upwards. This occurs whenever there are no free blocks on the free list, or when all free blocks are too small for the allocation request. Such a call doesn't indicate an error condition, because there was still adequate room for expansion between

HeapPtr and *HeapEnd*. Instead, the call indicates that the unused area above *HeapPtr* has shrunk, and the heap manager ignores the return value.

Memory issues for DOS protected-mode programs

This section explains how protected-mode programs written in Borland Pascal use memory.

Code segments

*For more information about the `$$` and `$G` *unitname* directives, see Chapter 2 of the Programmer's Reference.*

The main program and each library in an application or DLL has its own code segment. By default, units with like segment attributes are grouped in code segments; you can control this grouping with the `$$` and `$G` *unitname* directives. The size of a single code segment can't exceed 64K, but the total size of the code is limited only by the available memory.

Segment attributes

Each code segment has a set of attributes that determine the behavior of the code segment when it's loaded into memory.

MOVEABLE or FIXED

When a code segment is **MOVEABLE**, the protected-mode memory manager can move the segment around in physical memory to satisfy other memory allocation requests. When a code segment is **FIXED**, it *never* moves in physical memory. The preferred attribute is **MOVEABLE**, and unless it's absolutely necessary to keep a code segment at the same address in physical memory (such as if it contains an interrupt handler), you should use the **MOVEABLE** attribute. When you do need a fixed code segment, keep that code segment as small as possible.

PRELOAD or DEMANDLOAD

A code segment that has the **PRELOAD** attribute is automatically loaded when the application or library is activated. The **DEMANDLOAD** attribute delays the loading of the segment until a routine in the segment is actually called. Although this takes longer, it allows an application to execute in less space.

DISCARDABLE or PERMANENT

When a segment is DISCARDABLE, the protected-mode memory manager can free the memory occupied by the segment when it needs to allocate additional memory. When a segment is PERMANENT, it's kept in memory at all times. When an application makes a call to a DISCARDABLE segment that's not in memory, the protected-mode memory manager first loads it from the .EXE file. This takes longer than if the segment were PERMANENT, but it allows an application to execute in less space.

A DISCARDABLE segment in a DOS protected-mode application is much like an overlaid segment in a DOS program, while a PERMANENT segment in a DOS protected-mode application is like that of a segment that's not overlaid in a DOS program.

The data and stack segments

Each DOS protected-mode application or library has a data segment that can be up to 64K in size. The data segment is always pointed to by the data segment register (DS), and it contains all typed constants and global variables.

In addition to the data segment, a DOS protected-mode application has a stack segment that is used to store local variables allocated by procedures and functions. On entry to an application, the stack segment register (SS) and the stack pointer (SP) are loaded so that SS:SP points to the first byte past the stack segment. When procedures and functions are called, SP is moved down to allocate space for parameters, the return address, and local variables. When a routine returns, the process is reversed by incrementing SP to the value it had before the call. The default size of the stack segment is 16K, but this can be changed with a **\$M** compiler directive.

Unlike an application, a DOS protected-mode DLL has no stack segment. When a call is made to a procedure or function in a DLL, the DS register is changed to point to the DLL's data segment, but the SS:SP register pair isn't modified. Therefore, a DLL always uses the stack of the calling application.

Changing attributes

The default attributes of a code segment are `MOVEABLE`, `DEMANDLOAD`, and `DISCARDABLE`, but you can change this with a **\$C** compiler directive. For example,

```
{ $C MOVEABLE PRELOAD PERMANENT }
```

For details about the \$C compiler directive, see Chapter 2, "Compiler directives," in the Programmer's Reference.

There is no need for a separate overlay manager in a DOS protected-mode application. The DOS protected-mode memory manager includes a full set of overlay management services, controlled through code segment attributes. These services are available to any DOS protected-mode application.

The DOS protected-mode heap manager

For more details on Borland's DOS protected-mode extensions, see Chapter 17, "Programming in DOS protected mode."

Borland Pascal for DOS protected mode doesn't support the Mark and Release allocation scheme provided in DOS real mode.

To read more about using the heap manager in a DLL, see page 136 in Chapter 11, "Dynamic-link libraries."

Borland's DOS protected-mode extensions include a complete protected-mode memory manager. When a DOS protected-mode application is executed, all available memory is changed into a *global heap* which is managed by the protected-mode memory manager. An application can access the global heap through the `GlobalXXXX` routines in the `WinAPI` unit. Although global memory blocks of any size can be allocated, the global heap is intended only for large memory blocks (1024 bytes or more). Each global memory block carries an overhead of at least 32 bytes, and there is a system-wide limit of 8192 global memory blocks.

Borland Pascal includes a *heap manager* which implements the `New`, `Dispose`, `GetMem`, and `FreeMem` standard procedures. The heap manager uses the global heap for all allocations. Because the global heap has a system-wide limit of 8192 memory blocks (which certainly is less than what some applications might require), Borland Pascal's heap manager includes a *segment sub-allocator* algorithm to enhance performance and allow a substantially larger number of blocks to be allocated.

This is how the segment sub-allocator works: When allocating a large block, the heap manager simply allocates a global memory block using the `GlobalAlloc` routine. When allocating a small block, the heap manager allocates a larger global memory block and then divides (sub-allocates) that block into smaller blocks as required. Allocations of small blocks reuse all available sub-allocation space before the heap manager allocates a new global memory block, which, in turn, is further sub-allocated.

The `HeapLimit` variable defines the threshold between small and large heap blocks. The default value is 1024 bytes. The `HeapBlock`

variable defines the size the heap manager uses when allocating blocks to be assigned to the sub-allocator. The default value of *HeapBlock* is 8192 bytes. You should have no reason to change the values of *HeapLimit* and *HeapBlock*, but if you do, make sure that *HeapBlock* is at least four times the size of *HeapLimit*.

The *HeapAllocFlags* variable defines the attribute flags value passed to *GlobalAlloc* when the heap manager allocates global blocks. The default value is *GMEM_MOVEABLE*.

The HeapError variable

The *HeapError* variable allows you to install a heap-error function that's called whenever the heap manager can't complete an allocation request. *HeapError* is a pointer that points to a function with this header:

```
function HeapFunc (Size: Word): Integer; far;
```

Note that the **far** directive forces the heap-error function to use the FAR call model.

The heap-error function is installed by assigning its address to the *HeapError* variable:

```
HeapError := @HeapFunc;
```

The heap-error function is called whenever a call to *New* or *GetMem* can't complete the request. The *Size* parameter contains the size of the block that couldn't be allocated, and the heap-error function should attempt to free a block of at least that size.

Before calling the heap-error function, the heap manager attempts to allocate the block within its sub-allocation free space as well as through a direct call to the *GlobalAlloc* function.

Depending on its success, the heap-error function should return 0, 1, or 2. A return of 0 indicates failure, causing a run-time error to occur immediately. A return of 1 also indicates failure, but instead of a run-time error, it causes *New* or *GetMem* to return a **nil** pointer. Finally, a return of 2 indicates success and causes a retry (which could also cause another call to the heap-error function).

The standard heap-error function always returns 0, thereby causing a run-time error whenever a call to *New* or *GetMem* can't be completed. For many applications, however, the simple heap-error function that follows is more appropriate:

```

function HeapFunc(Size: Word): Integer; far;
begin
    HeapFunc := 1;
end;

```

When installed, this function causes *New* or *GetMem* to return **nil** when they can't complete the request, instead of aborting the program.

Memory issues for Windows programs

This section explains how Windows programs written in Borland Pascal use memory.

Code segments

Each module (the main program or library and each unit) in a Borland Pascal application or DLL has its own code segment. The size of a single code segment can't exceed 64K, but the total size of the code is limited only by the available memory.

Segment attributes Each code segment has a set of attributes that determine the behavior of the code segment when it's loaded into memory.

MOVEABLE or FIXED

When a code segment is **MOVEABLE**, Windows can move the segment around in physical memory to satisfy other memory allocation requests. When a code segment is **FIXED**, it *never* moves in physical memory. The preferred attribute is **MOVEABLE**, and unless it's absolutely necessary to keep a code segment at the same address in physical memory (such as if it contains an interrupt handler), you should use the **MOVEABLE** attribute. When you do need a fixed code segment, keep that code segment as small as possible.

PRELOAD or DEMANDLOAD

A code segment that has the **PRELOAD** attribute is automatically loaded when the application or library is activated. The **DEMANDLOAD** attribute delays the loading of the segment until

a routine in the segment is actually called. Although this takes longer, it allows an application to execute in less space.

DISCARDABLE or PERMANENT

When a segment is DISCARDABLE, Windows can free the memory occupied by the segment when it needs to allocate additional memory. When a segment is PERMANENT, it's kept in memory at all times.

When an application makes a call to a DISCARDABLE segment that's not in memory, Windows first loads it from the .EXE file. This takes longer than if the segment were PERMANENT, but it allows an application to execute in less space.

A DISCARDABLE segment in a Windows application is much like an overlaid segment in a DOS program, while a PERMANENT segment in a Windows application is like that of a segment that's not overlaid in a DOS program.

Changing attributes

The default attributes of a code segment are MOVEABLE, DEMANDLOAD, and DISCARDABLE, but you can change this with a **\$C** compiler directive. For example,

```
{ $C MOVEABLE PRELOAD PERMANENT }
```


For details about the \$C compiler directive, see Chapter 2, "Compiler directives," in the Programmer's Reference.

There is no need for a separate overlay manager in a Windows application. The Windows memory manager includes a full set of overlay management services, controlled through code segment attributes. These services are available to any Windows application.

The automatic data segment

Each application or library has one data segment called the "automatic data segment," which can be up to 64K in size. The automatic data segment is always pointed to by the data segment register (DS). It's divided into four sections:

Figure 21.7
Automatic data segment



The first 16 bytes of the automatic data segment always contain the *task header* in which Windows stores various system information.

The *static data* area contains all global variables and typed constants declared by the application or library.

The *stack* is used to store local variables allocated by procedures and functions. On entry to an application, the stack segment register (SS) and the stack pointer (SP) are loaded so that SS:SP points to the first byte past the stack area in the automatic data segment. When procedures and functions are called, SP is moved down to allocate space for parameters, the return address, and local variables. When a routine returns, the process is reversed by incrementing SP to the value it had before the call. The default size of the stack area in the automatic data segment is 8K, but this can be changed with a **\$M** compiler directive.

Unlike an application, a library has no stack area in its automatic data segment. When a call is made to a procedure or function in a DLL, the DS register points to the library's automatic data segment, but the SS:SP register pair isn't modified. Therefore, a library always uses the stack of the calling application.

The last section in the automatic data segment is the *local heap*. It contains all local dynamic data that was allocated using the *LocalAlloc* function in Windows. The default size of the local heap section is 8K, but this can be changed with a **\$M** compiler directive.

The local heap is used by Windows for the temporary storage of things such as edit control and list box buffers. Never set the local heap to zero.

Windows allows the automatic data segment to be *movable*, but programs written in Borland Pascal don't support this. The automatic data segment of a Borland Pascal application or library is always *locked*, thereby ensuring that the selector (segment address) of the automatic data segment never changes. This has no adverse effects when running in standard or extended mode because a segment retains the same selector even when it's moved in physical memory. In real mode, however, if Windows is required to expand the local heap, it probably won't be able to do so, because the automatic data segment can't be moved. If your application uses the local heap and must run in real mode, you should make sure that the initial size of the local heap (set with a **\$M** directive) is large enough to accommodate all local heap allocations.

The heap manager

See the Programmer's Reference for more details on the global and local heaps.

Windows supports dynamic memory allocations on two different heaps: The *global heap* and the *local heap*.

The global heap is a pool of memory available to all applications. Although global memory blocks of any size can be allocated, the global heap is intended only for large memory blocks (256 bytes or more). Each global memory block carries an overhead of at least 20 bytes, and under the Windows standard and 386 enhanced modes, there is a system-wide limit of 8192 global memory blocks, only some of which are available to any given application.

The local heap is a pool of memory available only to your application or library. It exists in the upper part of an application's or library's data segment. The total size of local memory blocks that can be allocated on the local heap is 64K minus the size of the application's stack and static data. For this reason, the local heap is best suited for small memory blocks (256 bytes or less). The default size of the local heap is 8K, but you can change this with the **\$M** compiler directive.

When Windows is your target, you can't use the Mark and Release allocation scheme.

Borland Pascal includes a *heap manager* which implements the *New*, *Dispose*, *GetMem*, and *FreeMem* standard procedures. The heap manager uses the global heap for all allocations. Because the global heap has a system-wide limit of 8192 memory blocks

(which certainly is less than what some applications might require), Borland Pascal's heap manager includes a *segment sub-allocator* algorithm to enhance performance and allow a substantially larger number of blocks to be allocated.

To read more about using the heap manager in a DLL, see page 136 in Chapter 11, "Dynamic-link libraries."

This is how the segment sub-allocator works: When allocating a large block, the heap manager simply allocates a global memory block using the Windows *GlobalAlloc* routine. When allocating a small block, the heap manager allocates a larger global memory block and then divides (sub-allocates) that block into smaller blocks as required. Allocations of small blocks reuse all available sub-allocation space before the heap manager allocates a new global memory block, which, in turn, is further sub-allocated.

The *HeapLimit* variable defines the threshold between small and large heap blocks. The default value is 1024 bytes. The *HeapBlock* variable defines the size the heap manager uses when allocating blocks to be assigned to the sub-allocator. The default value of *HeapBlock* is 8192 bytes. You should have no reason to change the values of *HeapLimit* and *HeapBlock*, but if you do, make sure that *HeapBlock* is at least four times the size of *HeapLimit*.

The *HeapAllocFlags* variable defines the attribute flags value passed to *GlobalAlloc* when the heap manager allocates global blocks. In a program, the default value is *GMEM_MOVEABLE*, and in a library the default value is *GMEM_MOVEABLE + GMEM_DDESHARE*.

Global blocks are always locked (using *GlobalLock*) immediately after they're allocated, and not unlocked until immediately before they're deallocated. This ensures that the selectors (segment addresses) of the blocks don't change. In Windows standard and 386 enhanced modes, fixed blocks can still be moved around in physical memory to make room for other memory allocation requests, so there is no performance penalty associated with using the Borland Pascal heap manager. In Windows real mode, however, a fixed block must remain fixed in physical memory. This precludes the Windows memory manager from moving it so it can allocate other blocks. If your application is to run in real mode, consider using the memory-management services provided by Windows when allocating dynamic memory blocks.

The HeapError variable

The *HeapError* variable allows you to install a heap-error function that's called whenever the heap manager can't complete an allocation request. *HeapError* is a pointer that points to a function with this header:

```
function HeapFunc (Size: Word): Integer; far;
```

Note that the **far** directive forces the heap-error function to use the FAR call model.

The heap-error function is installed by assigning its address to the *HeapError* variable:

```
HeapError := @HeapFunc;
```

The heap-error function is called whenever a call to *New* or *GetMem* can't complete the request. The *Size* parameter contains the size of the block that couldn't be allocated, and the heap-error function should attempt to free a block of at least that size.

Before calling the heap-error function, the heap manager attempts to allocate the block within its sub-allocation free space as well as through a direct call to the Windows *GlobalAlloc* function.

Depending on its success, the heap-error function should return 0, 1, or 2. A return of 0 indicates failure, causing a run-time error to occur immediately. A return of 1 also indicates failure, but instead of a run-time error, it causes *New* or *GetMem* to return a **nil** pointer. Finally, a return of 2 indicates success and causes a retry (which could also cause another call to the heap-error function).

The standard heap-error function always returns 0, thereby causing a run-time error whenever a call to *New* or *GetMem* can't be completed. For many applications, however, the simple heap-error function that follows is more appropriate:

```
function HeapFunc(Size: Word): Integer; far;
begin
  HeapFunc := 1;
end;
```

When installed, this function causes *New* or *GetMem* to return **nil** when they can't complete the request, instead of aborting the program.

Internal data formats

The next several pages discuss the internal data formats of Borland Pascal.

Integer types

The format selected to represent an integer-type variable depends on its minimum and maximum bounds:

- If both bounds are within the range $-128..127$ (*Shortint*), the variable is stored as a signed byte.
- If both bounds are within the range $0..255$ (*Byte*), the variable is stored as an unsigned byte.
- If both bounds are within the range $-32768..32767$ (*Integer*), the variable is stored as a signed word.
- If both bounds are within the range $0..65535$ (*Word*), the variable is stored as an unsigned word.
- Otherwise, the variable is stored as a signed double word (*Longint*).

Char types

A *Char* or a subrange of a *Char* type is stored as an unsigned byte.

Boolean types

A *Boolean* type is stored as a *Byte*, a *ByteBool* is stored as a *Byte*, a *WordBool* type is stored as a *Word*, and a *LongBool* is stored as a *Longint*.

A *Boolean* can assume the values 0 (*False*) and 1 (*True*). *ByteBool*, *WordBool*, and *LongBool* types can assume the value of 0 (*False*) or nonzero (*True*).

Enumerated types

An enumerated type is stored as an unsigned byte if the enumeration has 256 or fewer values; otherwise, it's stored as an unsigned word.


The value v of the number is determined by the following:

```

if 0 < e < 255,      then v = (-1)s * 2(e-127) * (1.f).
if e = 0 and f <> 0, then v = (-1)s * 2(-126) * (0.f).
if e = 0 and f = 0, then v = (-1)s * 0.
if e = 255 and f = 0, then v = (-1)s * Inf.
if e = 255 and f <> 0, then v is a NaN.

```

The Double type An 8-byte (64-bit) *Double* number is divided into three fields:




The value v of the number is determined by the following:

```

if 0 < e < 2047,      then v = (-1)s * 2(e-1023) * (1.f).
if e = 0 and f <> 0, then v = (-1)s * 2(-1022) * (0.f).
if e = 0 and f = 0, then v = (-1)s * 0.
if e = 2047 and f = 0, then v = (-1)s * Inf.
if e = 2047 and f <> 0, then v is a NaN.

```

The Extended type A 10-byte (80-bit) *Extended* number is divided into four fields:




The value v of the number is determined by the following:

```

if 0 <= e < 32767,    then v = (-1)s * 2(e-16383) * (i.f).
if e = 32767 and f = 0, then v = (-1)s * Inf.
if e = 32767 and f <> 0, then v is a NaN.

```

The Comp type An 8-byte (64-bit) *Comp* number is divided into two fields:



The value v of the number is determined by the following:

if $s = 1$ **and** $d = 0$, **then** v is a NaN

Otherwise, v is the two's complement 64-bit value.

Pointer types

A *Pointer* type is stored as two words (a double word), with the offset part in the low word and the segment part in the high word. The pointer value **nil** is stored as a double-word zero.

String types

A string occupies as many bytes as its maximum length plus one. The first byte contains the current dynamic length of the string, and the following bytes contain the characters of the string.

The length byte and the characters are considered unsigned values. Maximum string length is 255 characters plus a length byte (**string**[255]).

Set types

A set is a bit array where each bit indicates whether an element is in the set or not. The maximum number of elements in a set is 256, so a set never occupies more than 32 bytes. The number of bytes occupied by a particular set is calculated as

$$\text{ByteSize} = (\text{Max} \text{ div } 8) - (\text{Min} \text{ div } 8) + 1$$

where *Min* and *Max* are the lower and upper bounds of the base type of that set. The byte number of a specific element E is

$$\text{ByteNumber} = (E \text{ div } 8) - (\text{Min} \text{ div } 8)$$

and the bit number within that byte is

$$\text{BitNumber} = E \text{ mod } 8$$

where E denotes the ordinal value of the element.

Array types

An array is stored as a contiguous sequence of variables of the component type of the array. The components with the lowest indexes are stored at the lowest memory addresses. A multi-dimensional array is stored with the rightmost dimension increasing first.

Record types

The fields of a record are stored as a contiguous sequence of variables. The first field is stored at the lowest memory address. If the record contains variant parts, then each variant starts at the same memory address.

Object types

The internal data format of an object resembles that of a record. The fields of an object are stored in order of declaration, as a contiguous sequence of variables. Any fields inherited from an ancestor type are stored before the new fields defined in the descendant type.

If an object type defines virtual methods, constructors, or destructors, the compiler allocates an extra field in the object type. This 16-bit field, called the *virtual method table (VMT) field*, is used to store the offset of the object type's VMT in the data segment. The VMT field immediately follows after the ordinary fields in the object type. When an object type inherits virtual methods, constructors, or destructors, it also inherits a VMT field, so an additional one isn't allocated.

Initialization of the VMT field of an instance is handled by the object type's constructor(s). A program never explicitly initializes or accesses the VMT field.

The following examples illustrate the internal data formats of object types:


```

type
PLocation = ^TLocation;
TLocation = object
  X, Y: Integer;
  procedure Init(PX, PY: Integer);
  function GetX: Integer;
  function GetY: Integer;
end;


PPoint = ^TPoint;
TPoint = object(TLocation)
  Color: Integer;
  constructor Init(PX, PY, PColor: Integer);
  destructor Done; virtual;
  procedure Show; virtual;
  procedure Hide; virtual;
  procedure MoveTo(PX, PY: Integer); virtual;
end;

PCircle = ^TCircle;
TCircle = object(TPoint)
  Radius: Integer;
  constructor Init(PX, PY, PColor, PRadius: Integer);
  procedure Show; virtual;
  procedure Hide; virtual;
  procedure Fill; virtual;
end.

```

Figure 21.8 shows layouts of instances of *TLocation*, *TPoint*, and *TCircle*; each box corresponds to one word of storage.

Figure 21.8
Layouts of instances of
TLocation, *TPoint*, and *TCircle*



Virtual method tables

Each object type that contains or inherits virtual methods, constructors, or destructors has a VMT associated with it, which is stored in the initialized part of the program's data segment. There is only one VMT per object type (not one per instance), but two distinct object types never share a VMT, no matter how identical they appear to be. VMTs are built automatically by the compiler, and are never directly manipulated by a program. Likewise, pointers to VMTs are automatically stored in object type instances

by the object type's constructor(s) and are never directly manipulated by a program.

The first word of a VMT contains the size of instances of the associated object type; this information is used by constructors and destructors to determine how many bytes to allocate or dispose of, using the extended syntax of the *New* and *Dispose* standard procedures.

The second word of a VMT contains the negative size of instances of the associated object type; this information is used by the virtual method call validation mechanism to detect uninitialized objects (instances for which no constructor call has been made), and to check the consistency of the VMT. When virtual call validation is enabled (using the **{SR+}** compiler directive, which has been expanded to include virtual method checking), the compiler generates a call to a VMT validation routine before each virtual call. The VMT validation routine checks that the first word of the VMT isn't zero, and that the sum of the first and the second word is zero. If either check fails, the compiler generates run-time error 210.



Enabling range checking and virtual method call checking slows down your program and makes it somewhat larger, so use the **{R+}** state only when debugging, and switch to the **{SR-}** state for the final version of the program.

See page 283 for an explanation of dynamic method tables.

The third word of a VMT contains the data segment offset of the object type's dynamic method table (DMT), or zero if the object type has no dynamic methods.

The fourth word of a VMT is reserved and always contains zero.

Finally, starting at offset 8 in the VMT, is a list of 32-bit method pointers, one per virtual method in the object type, in order of declaration. Each slot contains the address of the corresponding virtual method's entry point.

Figure 21.9 shows the layouts of the VMTs of the *TPoint* and *TCircle* types; each small box corresponds to one word of storage, and each large box corresponds to two words of storage.

Figure 21.9
TPoint and TCircle's VMT
layouts

TPoint VMT	TCircle VMT
8	10
-8	-10
0	0
0	0
@TPoint.Done	@TPoint.Done
@TPoint.Show	@TCircle.Show
@TPoint.Hide	@TCircle.Hide
@TPoint.MoveTo	@TPoint.MoveTo
	@TCircle.Fill

Notice how *TCircle* inherits the *Done* and *MoveTo* methods from *TPoint*, and how it overrides the *Show* and *Hide* methods.

As mentioned already, an object type's constructors contain special code that stores the offset of the object type's VMT in the instance being initialized. For example, given an instance *P* of type *Pointer*, and an instance *C* of type *TCircle*, a call to *P.Init* automatically stores the offset of *TPoint*'s VMT in *P*'s VMT field, and a call to *C.Init* likewise stores the offset of *TCircle*'s VMT in *C*'s VMT field. This automatic initialization is part of a constructor's entry code. When control arrives at the **begin** of the constructor's statement part, the VMT field *Self* is already set up. Therefore, if the need arises, a constructor can make calls to virtual methods.

Dynamic method tables

The VMT for an object type contains a four-byte entry (a method pointer) for each virtual method declared in the object type and any of its ancestors. In cases where ancestral type(s) define a large number of virtual methods, the process of creating derived types can use up quite a lot of memory, especially if many derived types are created. Even though the derived types can override only a

few of the inherited methods, the VMT of each derived type contains method pointers for all inherited virtual methods, even if they haven't changed.

Dynamic methods provide an alternative in such situations. Instead of encoding a pointer for *all* late-bound methods in an object type, a dynamic method table (DMT) encodes only the methods that were *overridden* in the object type. When descendant types override only a few of a large number of inherited late-bound methods, the dynamic method table format uses less space than the format used by VMTs.

The following two object types illustrate DMT formats:

```
type
  TBase = object
    X: Integer;
    constructor Init;
    destructor Done; virtual;
    procedure P10; virtual 10;
    procedure P20; virtual 20;
    procedure P30; virtual 30;
    procedure P40; virtual 40;
  end;

type
  TDerived = object(TBase)
    Y: Integer;
    constructor Init;
    destructor Done; virtual;
    procedure P10; virtual 10;
    procedure P30; virtual 30;
    procedure P50; virtual 50;
  end;
```

Figures 21.10 and 21.11 shows the layouts of the VMTs and DMTs of *TBase* and *TDerived*. Each small box corresponds to one word of storage, and each large box corresponds to two words of storage.

Figure 21.10
TBase's VMT and DMT layouts

TBase VMT	TBase DMT
4	0
-4	Cached index
Offset of TBase DMT	Cached entry offset
0	4
	10
	20
	30
	40
	@TBase.P10
	@TBase.P20
	@TBase.P30
	@TBase.P40

An object type has a DMT only if it introduces or overrides dynamic methods. If an object type inherits dynamic methods, but doesn't override any of them or introduce new ones, it simply inherits the DMT of its ancestor.

As is the case for VMTs, DMTs are stored in the initialized part of the application's data segment.

Figure 21.11
TDerived's VMT and DMT
layouts

TDerived VMT	TDerived DMT
6	Offset of TBase DMT
-6	Cached index
Offset of TDerived DMT	Cached entry offset
0	3
@TDerived.Done	10
	30
	50
	@TDerived.P10
	@TDerived.P30
	@TDerived.P50

The first word of a DMT contains the data segment offset of the *parent DMT*, or zero if there is no parent DMT.

The second and third words of a DMT are used to cache dynamic method lookups, as is described on page 297.

The fourth word of a DMT contains the *DMT entry count*. It's immediately followed by a list of words, each of which contain a dynamic method index, and then followed by a list of corresponding method pointers. The length of each list is given by the DMT entry count.

File types

File types are represented as records. Typed files and untyped files occupy 128 bytes, which are laid out as follows:

```

type
  TFileRec = record
    Handle: Word;
    Mode: Word;
    RecSize: Word;
    Private: array[1..26] of Byte;
  
```

```

    UserData: array[1..16] of Byte;
    Name: array[0..79] of Char;
end;

```

Text files occupy 256 bytes, which are laid out as follows:

```

type
TTextBuf = array[0..127] of Char;
TTextRec = record
    Handle: Word;
    Mode: Word;
    BufSize: Word;
    Private: Word;
    BufPos: Word;
    BufEnd: Word;
    BufPtr: ^TTextBuf;
    OpenFunc: Pointer;
    InOutFunc: Pointer;
    FlushFunc: Pointer;
    CloseFunc: Pointer;
    UserData: array[1..16] of Byte;
    Name: array[0..79] of Char;
    Buffer: TTextBuf;
end;

```

Handle contains the file's handle (when the file is open) as returned by DOS.

The *Mode* field can assume one of the following values:

```

const
fmClosed = $D7B0;
fmInput  = $D7B1;
fmOutput = $D7B2;
fmInOut  = $D7B3;

```

fmClosed indicates that the file is closed. *fmInput* and *fmOutput* indicate that the file is a text file that has been reset (*fmInput*) or rewritten (*fmOutput*). *fmInOut* indicates that the file variable is a typed or an untyped file that has been reset or rewritten. Any other value indicates that the file variable hasn't been assigned (and thereby not initialized).

The *UserData* field is never accessed by Borland Pascal and is free for user-written routines to store data in.

Name contains the file name, which is a sequence of characters terminated by a null character (#0).

For typed files and untyped files, *RecSize* contains the record length in bytes, and the *Private* field is unused but reserved.

For text files, *BufPtr* is a pointer to a buffer of *BufSize* bytes, *BufPos* is the index of the next character in the buffer to read or write, and *BufEnd* is a count of valid characters in the buffer. *OpenFunc*, *InOutFunc*, *FlushFunc*, and *CloseFunc* are pointers to the I/O routines that control the file. The section entitled "Text file device drivers" in Chapter 14 provides information on that subject.

Procedural types

A procedural type is stored as a double word, with the offset part of the referenced procedure in the low word and the segment part in the high word.

Direct memory access

Borland Pascal implements three predefined arrays, *Mem*, *MemW*, and *MemL*, which are used to directly access memory. Each component of *Mem* is a byte, each component of *MemW* is a *Word*, and each component of *MemL* is a *Longint*.

The *Mem* arrays use a special syntax for indexes: Two expressions of the integer type *Word*, separated by a colon, are used to specify the segment base and offset of the memory location to access.

Here are two examples:

```
Mem[Seg0040:$0049] := 7;  
Data := MemW[Seg(V):Ofs(V)];
```

The first statement stores the value 7 in the byte at \$0040:\$0049. The second statement moves the *Word* value stored in the first 2 bytes of the variable *V* into the variable *Data*.

Direct port access

For access to the 80x86 CPU data ports, Borland Pascal implements two predefined arrays, *Port* and *PortW*. Both are one-dimensional arrays, and each element represents a data port, whose port address corresponds to its index. The index type is the integer type *Word*. Components of the *Port* array are of type *Byte* and components of the *PortW* array are of type *Word*.

When a value is assigned to a component of *Port* or *PortW*, the value is output to the selected port. When a component of *Port* or *PortW* is referenced in an expression, its value is input from the selected port.

Use of the *Port* and *PortW* arrays is restricted to assignment and reference in expressions only; that is, components of *Port* and *PortW* can't be used as variable parameters. Also, references to the entire *Port* or *PortW* array (reference without index) aren't allowed.

Control issues

This chapter describes in detail the various ways that Borland Pascal implements program control. Included are calling conventions and exit procedures.

Calling conventions

Parameters are transferred to procedures and functions via the stack. Before calling a procedure or function, the parameters are pushed onto the stack in their order of declaration. Before returning, the procedure or function removes all parameters from the stack.

The skeleton code for a procedure or function call looks like this:

```
PUSH  Param1
PUSH  Param2
      ⋮
PUSH  ParamX
CALL  ProcOrFunc
```

Parameters are passed either by *reference* or by *value*. When a parameter is passed by reference, a pointer that points to the actual storage location is pushed onto the stack. When a parameter is passed by value, the actual value is pushed onto the stack.

Variable parameters

Variable parameters (**var** parameters) are always passed by reference—a pointer that points to the actual storage location.

Value parameters

Value parameters are passed by value or by reference depending on the type and size of the parameter. In general, if the value parameter occupies 1, 2, or 4 bytes, the value is pushed directly onto the stack. Otherwise a pointer to the value is pushed, and the procedure or function then copies the value into a local storage location.



The 8086 does not support byte-sized PUSH and POP instructions, so byte-sized parameters are always transferred onto the stack as words. The low-order byte of the word contains the value, and the high-order byte is unused (and undefined).

An integer type or parameter is passed as a byte, a word, or a double word, using the same format as an integer-type variable. (For double words, the high-order word is pushed before the low-order word so that the low-order word ends up at the lowest address.)

A *Char* parameter is passed as an unsigned byte.

A *Boolean* parameter is passed as a byte with the value 0 or 1.

An enumerated-type parameter is passed as an unsigned byte if the enumeration has 256 or fewer values; otherwise, it is passed as an unsigned word.

A floating-point type parameter (*Real*, *Single*, *Double*, *Extended*, and *Comp*) is passed as 4, 6, 8, or 10 bytes on the stack. This is an exception to the rule that only 1-, 2-, and 4-byte values are passed directly on the stack.

A pointer-type parameter is passed as two words (a double word). The segment part is pushed before the offset part so that the offset part ends up at the lowest address.

A string-type parameter is passed as a pointer to the value.

For a set type parameter, if the bounds of the element type of the set are both within the range 0 to 7, the set is passed as a byte. If

the bounds are both within the range 0 to 15, the set is passed as a word. Otherwise, the set is passed as a pointer to an unpacked set that occupies 32 bytes.

Arrays and records with 1, 2, or 4 bytes are passed directly onto the stack. Other arrays and records are passed as pointers to the value.

Open parameters

Open string parameters are passed by first pushing a pointer to the string and then pushing a word containing the size attribute (maximum length) of the string.

Open array parameters are passed by first pushing a pointer to the array and then pushing a word containing the number of elements in the array less one.

When using the built-in assembler, the value that the *High* standard function returns for an open parameter can be accessed by loading the word just below the open parameter. In this example, the *FillString* procedure, which fills a string to its maximum length with a given character, demonstrates this.

```
procedure FillString(var Str: OpenString; Chr: Char); assembler;
asm
    LES    DI,Str           { ES:DI = @Str }
    MOV    CX,Str.Word[-2] { CX = High(Str) }
    MOV    AL,CL
    CLD
    STOSB                   { Set Str[0] }
    MOV    AL,Chr
    REP    STOSB            { Set Str[1..High] }
end;
```

Function results

Ordinal-type function results are returned in the CPU registers: Bytes are returned in *AL*, words are returned in *AX*, and double words are returned in *DX:AX* (high-order word in *DX*, low-order word in *AX*).

Real-type function results (type *Real*) are returned in the *DX:BX:AX* registers (high-order word in *DX*, middle word in *BX*, low-order word in *AX*).

80x87-type function results (type *Single*, *Double*, *Extended*, and *Comp*) are returned in the 80x87 coprocessor's top-of-stack register (ST(0)).

Pointer-type function results are returned in DX:AX (segment part in DX, offset part in AX).

For a string-type function result, the caller pushes a pointer to a temporary storage location before pushing any parameters, and the function returns a string value in that temporary location. The function must not remove the pointer.

NEAR and FAR calls

The 80x86 family of CPUs supports two kinds of call and return instructions: near and far. The near instructions transfer control to another location within the same code segment, and the far instructions allow a change of code segment.

A NEAR CALL instruction pushes a 16-bit return address (offset only) onto the stack, and a FAR CALL instruction pushes a 32-bit return address (both segment and offset). The corresponding RET instructions pop only an offset or both an offset and a segment.

Borland Pascal will automatically select the correct call model based on the procedure's declaration. Procedures declared in the interface section of a unit are far—they can be called from other units. Procedures declared in a program or in the **implementation** section of a unit are near—they can only be called from within that program or unit.

For some specific purposes, a procedure can be required to be far. For example, if a procedure or function is to be assigned to a procedural variable, it must far. The **\$F** compiler directive is used to override the compiler's automatic call model selection. Procedures and functions compiled in the **{\$F+}** state are always far; in the **{\$F-}** state, Borland Pascal automatically selects the correct model. The default state is **{\$F-}**.

Nested procedures and functions

A procedure or function is said to be nested when it is declared within another procedure or function. By default, nested procedures and functions always use the near call model, because they are visible only within a specific procedure or function in the same code segment. In an overlaid application, however, a **{\$F+}**

directive is generally used to force all procedures and functions to be far, including those that are nested.

When calling a nested procedure or function, the compiler generates a PUSH BP instruction just before the CALL, in effect passing the caller's BP as an additional parameter. Once the called procedure has set up its own BP, the caller's BP is accessible as a word stored at [BP + 4], or at [BP + 6] if the procedure is far. Using this link at [BP + 4] or [BP + 6], the called procedure can access the local variables in the caller's stack frame. If the caller itself is also a nested procedure, it also has a link at [BP + 4] or [BP + 6], and so on. The following example demonstrates how to access local variables from an **inline** statement in a nested procedure:

Nested procedures and functions cannot be declared with the external directive, and they cannot be procedural parameters.

```
procedure A; near;
var
  IntA: Integer;

procedure B; far;
var
  IntB: Integer;

procedure C; near;
var
  IntC: Integer;
begin
  asm
    MOV     AX,1
    MOV     IntC,AX           { IntC := 1 }
    MOV     BX,[BP+4]        { B's stack frame }
    MOV     SS:[BX+OFFSET IntB],AX { IntB := 1 }
    MOV     BX,[BP+4]        { B's stack frame }
    MOV     BX,SS:[BX+6]     { A's stack frame }
    MOV     SS:[BX+OFFSET IntA],AX { IntA := 1 }
  end;
end;

begin C end;

begin B end;
```

Method calling conventions

Methods use the same calling conventions as ordinary procedures and functions, except that every method has an additional implicit parameter, *Self*, that corresponds to a **var** parameter of the same type as the method's object type. The *Self* parameter is always passed as the last parameter, and always takes the form of a 32-bit pointer to the instance through which the method is

called. For example, given a variable *PP* of type *PPoint* as defined on page 281, the call *PP^.MoveTo(10, 20)* is coded as follows:

```
MOV    AX,10                ;Load 10 into AX
PUSH   AX                  ;Pass as PX parameter
MOV    AX,20               ;Load 20 into AX
PUSH   AX                  ;Pass as PY parameter
LES    DI,PP               ;Load PP into ES:DI
PUSH   ES                  ;Pass as Self parameter
PUSH   DI
MOV    DI,ES:[DI+6]        ;Pick up VMT offset from VMT field
CALL   DWORD PTR [DI+20]   ;Call VMT entry for MoveTo
```

Upon returning, a method must remove the *Self* parameter from the stack, just as it must remove any normal parameters.

Methods always use the far call model, regardless of the setting of the **\$F** compiler directive.

Virtual method calls

To call a virtual method, the compiler generates code that picks up the VMT address from the VMT field in the object, and then calls via the slot associated with the method. For example, given a variable *PP* of type *Point* (see page 281), the call *PP^.Show* generates the following code:

```
LES    DI,PP               ;Load PP into ES:DI
PUSH   ES                  ;Pass as Self parameter
PUSH   DI
MOV    DI,ES:[DI+6]        ;Pick up VMT offset from VMT field
CALL   DWORD PTR [DI+12]   ;Call VMT entry for Show
```

The type compatibility rules of object types allow *PP* to point at a *Point* or a *TCircle*, or at any other descendant of *TPoint*. And if you examine the VMTs shown on page 283, you'll see that for a *TPoint*, the entry at offset 12 in the VMT points to *TPoint.Show*; whereas for a *TCircle*, it points to *TCircle.Show*. Therefore, depending upon the *actual* run-time type of *PP*, the CALL instruction calls *TPoint.Show* or *TCircle.Show*, or the *Show* method of any other descendant of *TPoint*.

If *Show* had been a static method, this compiler would have generated this for the call to *PP^.Show*:


```

LES     DI,PP           ;Load PP into ES:DI
PUSH   ES              ;Pass as Self parameter
PUSH   DI
CALL   TPoint.Show     ;Directly call TPoint.Show

```

Here, no matter what *PP* points to, the code always calls the *TPoint.Show* method.

Dynamic method

calls

Dispatching a dynamic method call is somewhat more complicated and time consuming than dispatching a virtual method call. Instead of using a CALL instruction to call through a method pointer at a static offset in the VMT, the object type's DMT and parent DMTs must be *scanned* to find the topmost occurrence of a particular dynamic method index, and then a call must be made through the corresponding method pointer. This process involves far more instructions than can be coded in-line, so the Borland Pascal run-time library (RTL) contains a dispatch-support routine which is used when making dynamic method calls.

Had the *Show* method of the preceding type *TPoint* been declared as a dynamic method (with a dynamic method index of 200), the call *PP^.Show*, where *PP* is of type *Point*, would generate the following code:

```

LES     DI,PP           ;Load PP into ES:DI
PUSH   ES              ;Pass as Self parameter
PUSH   DI
MOV     DI,EX:[DI+6]    ;Pick up VMT offset from VMT field
MOV     AX,200         ;Load dynamic method index into AX
CALL   Dispatch        ;Call RTL routine to dispatch call

```

The RTL dispatcher first picks up the DMT offset from the VMT pointed to by the DI register. Then, using the "cached index" field of the DMT, the dispatcher checks if the dynamic method index of the method being called is the same as the last one that was called. If so, it immediately transfers control to the method, by jumping indirectly through the method pointer stored at the offset given by the "cached entry offset" field.

If the dynamic index of the method being called is not the same as the one stored in the cache, the dispatcher scans the DMT and the parent DMTs (by following the parent links in the DMTs) until it locates an entry with the given dynamic method index. The index and the offset of the corresponding method pointer is then stored in the DMT's cache fields, and control is transferred to the

method. If, for some reason, the dispatcher can't find an entry with the given dynamic method index, indicating that the DMTs have been destroyed, it terminates the application with a run-time error 210.

In spite of caching and a highly optimized RTL dispatch support routine, the dispatching of a dynamic method call takes substantially longer than a virtual method call. When the actions performed by the dynamic methods themselves take up a lot of time, however, the amount of space saved by using DMTs may outweigh this penalty.

Constructors and destructors

Constructors and destructors use the same calling conventions as other methods, except that an additional word-sized parameter, called the *VMT* parameter, is passed on the stack just before the *Self* parameter.

For constructors, the *VMT* parameter contains the *VMT* offset to store in *Self*'s *VMT* field to initialize *Self*.

When a constructor is called to allocate a dynamic object using the extended syntax of the *New* standard procedure, a **nil** pointer is passed in the *Self* parameter. The constructor allocates a new dynamic object, the address of which is passed back to the caller in *DX:AX* when the constructor returns. If the constructor can't allocate the object, a **nil** pointer is returned in *DX:AX*.

See "Constructor error recovery" on page 108.

Finally, when a constructor is called using a qualified-method identifier (that is, an object type identifier, followed by a period and a method identifier), a value of zero is passed in the *VMT* parameter. This indicates to the constructor that it should *not* initialize the *VMT* field of *Self*.

For destructors, a 0 in the *VMT* parameter indicates a normal call, and a nonzero value indicates that the destructor was called using the extended syntax of the *Dispose* standard procedure. This causes the destructor to deallocate *Self* just before returning (the size of *Self* is found by looking at the first word of *Self*'s *VMT*).

Entry and exit code

This is the standard entry and exit code for a procedure or function using the near call model:

```

PUSH    BP                ;Save BP
MOV     BP,SP             ;Set up stack frame
SUB     SP,LocalSize     ;Allocate locals (if any)
:
MOV     SP,BP             ;Deallocate locals (if any)
POP     BP                ;Restore BP
RETN   ParamSize         ;Remove parameters and return

```

For information on using exit procedures in a DLL, see Chapter 11, "Dynamic-link libraries."

For DOS programs, the entry and exit code for a routine using the far call model is the same as that of a routine using the near call model, except that a far-return instruction (RETF) is used to return from the routine. This is also true for a Windows program if the routine is compiled in the **{SW-}** state.

In the **{SW+}** state (the default), this is the entry and exit code for a Windows routine using the far call model:

```

INC     BP                ;Indicate FAR frame
PUSH   BP                ;Save odd BP
MOV    BP,SP             ;Set up stack frame
PUSH   DS                ;Save DS
SUB    SP,LocalSize     ;Allocate locals (if any)
:
MOV    SP,BP             ;Remove locals and saved DS
POP    BP                ;Restore odd BP
DEC    BP                ;Adjust BP
RETF   ParamSize         ;Remove parameters and return

```

This is the entry and exit code for an exportable routine (a procedure or function compiled with the **export** compiler directive):

```

MOV    AX,DS             ;Load DS selector into AX
NOP
INC    BP                ;Indicate FAR frame
PUSH   BP                ;Save odd BP
MOV    BP,SP             ;Set up stack frame
PUSH   DS                ;Save DS
MOV    DS,AX             ;Initialize DS
SUB    SP,LocalSize     ;Allocate locals (if any)
PUSH   SI                ;Save SI
PUSH   DI                ;Save DI
:
POP    DI                ;Restore DI
POP    SI                ;Restore SI
LEA   SP,[BP-2]         ;Deallocate locals (if any)
POP    DS                ;Restore DS
POP    BP                ;Restore odd BP

```

```

DEC     BP           ;Adjust BP
RETF   ParamSize   ;Remove parameters and return

```

For all call models, the instructions required to allocate and deallocate local variables are omitted if the routine has no local variables.

When running in real mode, Windows requires that all far stack frames (including stack frames of exported routines) have an odd saved BP value stored in the word at [BP+0] to distinguish them from near stack frames. In addition, Windows real mode requires that the word at [BP-2] in a far stack frame contains the data segment selector of the routine's caller. This explains the INC BP, PUSH DS, and DEC BP instructions generated in the {**\$W+**} state in the entry and exit code of **far** and **export** routines.

Note that only Windows real mode requires the use of {**\$W+**}. If you aren't supporting real mode, use {**\$W-**}. You'll create a smaller executable and gain a little speed in your Windows program.

Occasionally you might find {**\$W+**} useful while developing a Windows protected-mode application. Some non-Borland debugging tools require this state to work properly.

By default, Borland Pascal automatically generates *smart callbacks* for procedures and functions that are exported by an application. When linking an application in the {**\$K+**} state (the default), the linker looks for a MOV AX,DS instruction followed by a NOP instruction at every exported entry point and, for each such sequence it finds, it changes the MOV AX,DS to a MOV AX,SS. This change alleviates the need to use the Windows *MakeProcInstance* and *FreeProcInstance* API routines when creating callback routines (although it isn't harmful to do so), and also makes it possible to call exported entry points from within the application itself.

In the {**\$K-**} state or when creating a dynamic-link library, the Borland Pascal linker makes no modifications to the entry code of exported entry points. Unless a callback routine in an application is to be called from another application (which isn't recommended anyway), you shouldn't have to ever select the {**\$K-**} state.

When loading an application or dynamic-link library, Windows looks for a MOV AX,DS followed by a NOP at each exported entry point. For applications, the sequence is changed into three NOP instructions to prepare the routine for use with

MakeProcInstance. For libraries, the sequence is changed into a `MOV AX,xxxx` instruction, where `xxxx` is the selector (segment address) of the library's automatic data segment. Because smart callback entry points start with a `MOV AX,SS` instruction, they are left untouched by the Windows program loader.

Register-saving conventions

Procedures and functions should preserve the BP, SP, SS, and DS registers. All other registers can be modified. In addition, exported routines should preserve the SI and DI registers.

Exit procedures

By installing an exit procedure, you can gain control over a program's termination process. This is useful when you want to make sure specific actions are carried out before a program terminates; a typical example is updating and closing files.

The *ExitProc* pointer variable allows you to install an exit procedure. The exit procedure is always called as a part of a program's termination, whether it's a normal termination, a termination through a call to *Halt*, or a termination due to a run-time error.

An exit procedure takes no parameters and must be compiled with a **far** procedure directive to force it to use the far call model.

When implemented properly, an exit procedure actually becomes part of a chain of exit procedures. This chain makes it possible for units as well as programs to install exit procedures. Some units install an exit procedure as part of their initialization code and then rely on that specific procedure to be called to clean up after the unit. Closing files is such an example. The procedures on the exit chain are executed in reverse order of installation. This ensures that the exit code of one unit isn't executed before the exit code of any units that depend upon it.

To keep the exit chain intact, you must save the current contents of *ExitProc* before changing it to the address of your own exit procedure. Also, the first statement in your exit procedure must reinstall the saved value of *ExitProc*. The following program demonstrates a skeleton method of implementing an exit procedure:

```

program Testexit;
var
    ExitSave: Pointer;

procedure MyExit; far;
begin
    ExitProc := ExitSave;           { Always restore old vector first }
    :
end;

begin
    ExitSave := ExitProc;
    ExitProc := @MyExit;
    :
end.

```

On entry, the program saves the contents of *ExitProc* in *ExitSave*, and then installs the *MyExit* exit procedure. After having been called as part of the termination process, the first thing *MyExit* does is reinstall the previous exit procedure.

See page 135 for information on DLL exit procedures.

The termination routine in the run-time library keeps calling exit procedures until *ExitProc* becomes **nil**. To avoid infinite loops, *ExitProc* is set to **nil** before every call, so the next exit procedure is called only if the current exit procedure assigns an address to *ExitProc*. If an error occurs in an exit procedure, it won't be called again.

An exit procedure can learn the cause of termination by examining the *ExitCode* integer variable and the *ErrorAddr* pointer variable.

In case of normal termination, *ExitCode* is zero and *ErrorAddr* is **nil**. In case of termination through a call to *Halt*, *ExitCode* contains the value passed to *Halt*, and *ErrorAddr* is **nil**. Finally, in case of termination due to a run-time error, *ExitCode* contains the error code and *ErrorAddr* contains the address of the statement in error.

The last exit procedure (the one installed by the run-time library) closes the *Input* and *Output* files. If *ErrorAddr* is not **nil**, it outputs a run-time error message.

If you wish to present run-time error messages yourself, install an exit procedure that examines *ErrorAddr* and outputs a message if it's not **nil**. In addition, before returning, make sure to set *ErrorAddr* to **nil**, so that the error is not reported again by other exit procedures.

Once the run-time library has called all exit procedures, it returns to DOS or Windows, passing the value stored in *ExitCode* as a return code.

Interrupt handling

The Borland Pascal libraries and the code generated by the compiler are fully interruptible. Also, most the routines in the run-time libraries are reentrant, which allows you to write interrupt service routines in Borland Pascal.



You should not write interrupt service routines for Windows. If you do, your system is likely to crash when you run them.

Writing interrupt procedures

Interrupt procedures are declared with the **interrupt** directive. Every **interrupt** procedure must specify the following procedure header (or a subset of it, as explained later):

```
procedure IntHandler(Flags, CS, IP, AX, BX, CX, DX, SI, DI, DS, ES,  
    BP: Word);  
interrupt;  
begin  
    :  
    :  
end;
```

As you can see, all the registers are passed as pseudoparameters so you can use and modify them in your code. You can omit some or all of the parameters, starting with *Flags* and moving towards *BP*. It's an error to declare more parameters than are listed in the preceding example or to omit a specific parameter without also omitting the ones before it (although no error is reported). For example,

```
procedure IntHandler(DI, ES, BP: Word);           { Invalid header }  
procedure IntHandler(SI, DI, DS, ES, BP: Word);  { Valid header }
```

On entry, an **interrupt** procedure automatically saves all registers (regardless of the procedure header) and initializes the DS register:

```
PUSH  AX  
PUSH  BX  
PUSH  CX  
PUSH  DX
```

```

PUSH  SI
PUSH  DI
PUSH  DS
PUSH  ES
PUSH  BP
MOV   BP,SP
SUB   SP,LocalSize
MOV   AX,SEG DATA
MOV   DS,AX

```

Notice the lack of a STI instruction to enable additional interrupts. You should code this yourself (if required) using an **inline** statement. The exit code restores the registers and executes an interrupt-return instruction:

```

MOV   SP, BP
POP   BP
POP   ES
POP   DS
POP   DI
POP   SI
POP   DX
POP   CX
POP   BX
POP   AX
IRET

```

An **interrupt** procedure can modify its parameters. Changing the declared parameters modifies the corresponding register when the interrupt handler returns. This can be useful when you are using an interrupt handler as a user service, much like the DOS INT 21H services.

Interrupt procedures that handle hardware-generated interrupts shouldn't use any of Borland Pascal's input and output or dynamic memory allocation routines, because they aren't reentrant. Likewise, no DOS functions can be used because DOS is not reentrant.

Optimizing your code

Borland Pascal performs several different types of code optimizations, ranging from constant folding and short-circuit Boolean expression evaluation, all the way up to smart linking. The following sections describe some of the types of optimizations performed and how you can benefit from them in your programs.

Constant folding

If the operand(s) of an operator are constants, Borland Pascal evaluates the expression at compile time. For example,

```
X := 3 + 4 * 2
```

generates the same code as `X := 11`, and

```
S := 'In' + 'Out'
```

generates the same code as `S := 'InOut'`.

Likewise, if an operand of an *Abs*, *Chr*, *Hi*, *Length*, *Lo*, *Odd*, *Ord*, *Pred*, *Ptr*, *Round*, *Succ*, *Swap*, or *Trunc* function call is a constant, the function is evaluated at compile time.

If an array index expression is a constant, the address of the component is evaluated at compile time. For example, accessing `Data[5, 5]` is just as efficient as accessing a simple variable.

Constant merging

Using the same string constant two or more times in a statement part generates only one copy of the constant. For example, two or more `Write('Done')` statements in the same statement part references the same copy of the string constant 'Done'.

Short-circuit evaluation

Borland Pascal implements short-circuit Boolean evaluation, which means that evaluation of a Boolean expression stops as soon as the result of the entire expression becomes evident. This guarantees minimum execution time and usually minimum code size. Short-circuit evaluation also makes possible the evaluation of constructs that would not otherwise be legal. For example,

```
while (I <= Length(S)) and (S[I] <> ' ') do
  Inc(I);
while (P <> nil) and (P^.Value <> 5) do
  P := P^.Next;
```

In both cases, the second test isn't evaluated if the first test is *False*.

The opposite of short-circuit evaluation is complete evaluation, which is selected through a **{\$B+}** compiler directive. In this state, every operand of a Boolean expression is guaranteed to be evaluated.

Constant parameters

Read more about constant parameters on page 110.

Whenever possible, you should use constant parameters instead of value parameters. Constant parameters are at least as efficient as value parameters and, in many cases, more efficient. In particular, constant parameters generate less code and execute faster than value parameters for structured and string types.

Constant parameters are more efficient than value parameters because the compiler doesn't have to generate copies of the actual parameters upon entry to procedures or functions. Value parameters have to be copied into local variables so that modifications made to the formal parameters won't modify the

actual parameters. Because constant formal parameters can't be modified, the compiler has no need to generate copies of the actual parameters and code and stack space is saved.

Redundant pointer-load elimination

In certain situations, Borland Pascal's code generator can eliminate redundant pointer-load instructions, shrinking the size of the code and allowing for faster execution. When the code generator can guarantee that a particular pointer remains *constant* over a stretch of linear code (code with no jumps into it), and when that pointer is already loaded into a register pair (such as ES:DI), the code generator eliminates more redundant pointer load instructions in that block of code.

A pointer is considered constant if it's obtained from a variable parameter (variable parameters are always passed as pointers) or from the variable reference of a **with** statement. Because of this, using **with** statements is often more efficient (but never less efficient) than writing the fully qualified variable for each component reference.

Constant set inlining

When the right operand of the **in** operator is a set constant, the compiler generates the inclusion test using inline CMP instructions. Such inlined tests are more efficient than the code that would be generated by a corresponding boolean expression using relational operators. For example, the statement

```
if ((Ch >= 'A') and (Ch <= 'Z')) or
    ((Ch >= 'a') and (Ch <= 'z')) then ... ;
```

is less readable and also less efficient than

```
if Ch in ['A'..'Z', 'a'..'z'] then ... ;
```

Because constant folding applies to set constants as well as to constants of other types, it's possible to use **const** declarations without any loss of efficiency:

```
const
  Upper = ['A'..'Z'];
  Lower = ['a'..'z'];
  Alpha = Upper + Lower;
```

Given these declarations, this **if** statement generates the same code as the previous **if** statement:

```
if Ch in Alpha then ... ;
```

Small sets

The compiler generates very efficient code for operations on *small sets*. A small set is a set with a lower bound ordinal value in the range 0..7 and an upper bound ordinal value in the range 0..15. For example, the following *TByteSet* and *TWordSet* are both small sets.

```
type
  TByteSet = set of 0..7;
  TWordSet = set of 0..15;
```

Small set operations, such as union (+), difference (-), intersection (*), and inclusion tests (**in**) are generated inline using AND, OR, NOT, and TEST machine code instructions instead of calls to runtime library routines. Likewise, the *Include* and *Exclude* standard procedures generate inline code when applied to small sets.

Order of evaluation

As permitted by the Pascal standards, operands of an expression are frequently evaluated differently from the left to right order in which they are written. For example, the statement

```
I := F(J) div G(J);
```

where *F* and *G* are functions of type *Integer*, causes *G* to be evaluated before *F*, because this enables the compiler to produce better code. For this reason, it's important that an expression never depend on any specific order of evaluation of the embedded functions. Referring to the previous example, if *F* must be called before *G*, use a temporary variable:

```
T := F(J);  
I := T div G(J);
```



As an exception to this rule, when short-circuit evaluation is enabled (the **{\$B-}** state), Boolean operands grouped with **and** or **or** are *always* evaluated from left to right.

Range checking

Assignment of a constant to a variable and use of a constant as a value parameter is range-checked at compile time; no run-time range-check code is generated. For example, `X := 999`, where `X` is of type *Byte*, causes a compile-time error.

Shift instead of multiply or divide

The operation `X * C`, where `C` is a constant and a power of 2, is coded using a SHL instruction. The operation `X div C`, where `X` is an unsigned integer (*Byte* or *Word*) and `C` is a constant and a power of 2, is coded using a SHR instruction.

Likewise, when the size of an array's components is a power of 2, a SHL instruction (not a MUL instruction) is used to scale the index expression.

Automatic word alignment

By default, Borland Pascal aligns all variables and typed constants larger than 1 byte on a machine-word boundary. On all 16-bit 80x86 CPUs, word alignment means faster execution, because word-sized items on even addresses are accessed faster than words on odd addresses.

For more details about data alignment, see Chapter 2, "Compiler directives," in the Programmer's Reference.

Data alignment is controlled through the **\$A** compiler directive. In the default **{\$A+}** state, variables and typed constants are aligned as described above. In the **{\$A-}** state, no alignment measures are taken.

Eliminating dead code

Statements that never execute don't generate any code. For example, these constructs don't generate code:

```
if False then
  statement
while False do
  statement
```

Smart linking

When compiling to memory with TURBO.EXE, Borland Pascal's smart linker is disabled. This explains why some programs become smaller when compiled to disk.

Borland Pascal's built-in linker automatically removes unused code and data when building an .EXE file. Procedures, functions, variables, and typed constants that are part of the compilation, but are never referenced, are removed from the .EXE file. The removal of unused code takes place on a per procedure basis; the removal of unused data takes place on a per declaration section basis.

Consider the following program:

```
program SmartLink;

const
  H: array[0..15] of Char = '0123456789ABCDEF';

var
  I, J: Integer;
  X, Y: Real;

var
  S: string[79];

var
  A: array[1..10000] of Integer;

procedure P1;
begin
  A[1] := 1;
end;

procedure P2;
begin
  I := 1;
end;
```

```

procedure P3;
begin
  S := 'Borland Pascal';
  P2;
end;

begin
  P3;
end.

```

The main program calls *P3*, which calls *P2*, so both *P2* and *P3* are included in the .EXE file. Because *P2* references the first **var** declaration section, and *P3* references the second **var** declaration, *I*, *J*, *X*, *Y*, and *S* are also included in the .EXE file. No references are made to *P1*, however, and none of the included procedures reference *H* and *A*, so these objects are removed.

Smart linking is especially valuable in connection with units that implement procedure/function libraries. An example of such a unit is the *Dos* standard unit: It contains a number of procedures and functions, all of which are seldom used by the same program. If a program uses only one or two procedures from *Dos*, then only these procedures are included in the final .EXE file, and the remaining ones are removed, greatly reducing the size of the .EXE file.

P

A

R

T

4

Using Borland Pascal with assembly language

The built-in assembler

Borland Pascal's built-in assembler allows you to write 8086/8087 and 80286/80287 assembler code directly inside your Pascal programs. Of course, you can still convert assembler instructions to machine code manually for use in **inline** statements, or link in .OBJ files that contain **external** procedures and functions when you want to mix Pascal and assembler.

The built-in assembler implements a large subset of the syntax supported by Turbo Assembler and Microsoft's Macro Assembler. The built-in assembler supports all 8086/8087 and 80286/80287 opcodes, and all but a few of Turbo Assembler's expression operators.

Except for DB, DW, and DD (define byte, word, and double word), none of Turbo Assembler's directives, such as EQU, PROC, STRUC, SEGMENT, and MACRO, are supported by the built-in assembler. Operations implemented through Turbo Assembler directives, however, are largely matched by corresponding Borland Pascal constructs. For example, most EQU directives correspond to **const**, **var**, and **type** declarations in Borland Pascal, the PROC directive corresponds to **procedure** and **function** declarations, and the STRUC directive corresponds to Borland Pascal **record** types. In fact, Borland Pascal's built-in assembler can be thought of as an assembler language compiler that uses Pascal syntax for all declarations.

The asm statement

The built-in assembler is accessed through **asm** statements. This is the syntax of an **asm** statement:

```
asm AsmStatement [ Separator AsmStatement ] end
```

where *AsmStatement* is an assembler statement, and *Separator* is a semicolon, a new-line, or a Pascal comment.

Multiple assembler statements can be placed on one line if they are separated by semicolons. A semicolon isn't required between two assembler statements if the statements are on separate lines. A semicolon doesn't indicate that the rest of the line is a comment—comments must be written in Pascal style using { and } or (* and *).

Register use

In general, the rules of register use in an **asm** statement are the same as those of an **external** procedure or function. An **asm** statement must preserve the BP, SP, SS, and DS registers, but can freely modify the AX, BX, CX, DX, SI, DI, ES, and Flags registers. On entry to an **asm** statement, BP points to the current stack frame, SP points to the top of the stack, SS contains the segment address of the stack segment, and DS contains the segment address of the data segment. Except for BP, SP, SS, and DS, an **asm** statement can assume nothing about register contents on entry to the statement.

Assembler statement syntax

This is the syntax of an assembler statement:

```
[ Label ":" ] < Prefix > [ Opcode [ Operand < "," Operand > ] ]
```

Label is a label identifier, *Prefix* is an assembler prefix opcode (operation code), *Opcode* is an assembler instruction opcode or directive, and *Operand* is an assembler expression.

Comments are allowed between assembler statements, but not within them. For example, this is allowed:

```
asm
    MOV AX,1 {Initial value}
    MOV CX,100 {Count}
end;
```

but this is an error:

```
asm
    MOV {Initial value} AX,1;
    MOV CX, {Count} 100
end;
```

Labels

Only the first 32 characters of an identifier are significant in the built-in assembler.

Labels are defined in assembler as they are in Pascal—by writing a label identifier and a colon before a statement. And as they are in Pascal, labels defined in assembler must be declared in a **label** declaration part in the block containing the **asm** statement. There is one exception to this rule: *local labels*.

Local labels are labels that start with an at-sign (@). Because an at-sign can't be part of a Pascal identifier, such local labels are automatically restricted to use within **asm** statements. A local label is known only within the **asm** statement that defines it (that is, the scope of a local label extends from the **asm** keyword to the **end** keyword of the **asm** statement that contains it).



Unlike a normal label, a local label doesn't have to be declared in a **label** declaration part before it's used.

The exact composition of a local label identifier is an at-sign (@) followed by one or more letters (A..Z), digits (0..9), underscores (_), or at-signs. As with all labels, the identifier is followed by a colon (:).

Instruction opcodes

The built-in assembler supports all 8086/8087 and 80286/80287 instruction opcodes. 8087 opcodes are available only in the **{\$N+}** state (numeric processor enabled), 80286 opcodes are available only in the **{\$G+}** state (80286 code generation enabled), and 80287 opcodes are available only in the **{\$G+,N+}** state.

For a complete description of each instruction, refer to your 80x86 and 80x87 reference manuals.

RET instruction sizing The RET instruction opcode generates a near return or a far return machine code instruction depending on the call model of the current procedure or function.

```
procedure NearProc; near;
begin
  asm
    RET    { Generates a near return }
  end;
end;

procedure FarProc; far;
begin
  asm
    RET    { Generates a far return }
  end;
end;
```

On the other hand, the RETN and RETF instructions always generate a near return and a far return, regardless of the call model of the current procedure or function.

Automatic jump sizing Unless otherwise directed, the built-in assembler optimizes jump instructions by automatically selecting the shortest, and therefore most efficient form of a jump instruction. This automatic jump sizing applies to the unconditional jump instruction (JMP), and all conditional jump instructions, when the target is a label (not a procedure or function).

For an unconditional jump instruction (JMP), the built-in assembler generates a short jump (one byte opcode followed by a one byte displacement) if the distance to the target label is within -128 to 127 bytes; otherwise a near jump (one byte opcode followed by a two byte displacement) is generated.

For a conditional jump instruction, a short jump (1 byte opcode followed by a 1 byte displacement) is generated if the distance to the target label is within -128 to 127 bytes; otherwise, the built-in assembler generates a short jump with the inverse condition, which jumps over a near jump to the target label (5 bytes in total). For example, the assembler statement

```
JC    Stop
```

where *Stop* isn't within reach of a short jump is converted to a machine code sequence that corresponds to this:

```
JNC    Skip
JMP    Stop
Skip:
```

Jumps to the entry points of procedures and functions are always either near or far, but never short, and conditional jumps to procedures and functions are not allowed. You can force the built-in assembler to generate an unconditional near jump or far jump to a label by using a NEAR PTR or FAR PTR construct. For example, the assembler statements

```
JMP    NEAR PTR Stop
JMP    FAR PTR Stop
```

always generates a near jump and a far jump, respectively, even if *Stop* is a label within reach of a short jump.

Assembler directives

Borland Pascal's built-in assembler supports three assembler directives: DB (define byte), DW (define word), and DD (define double word). They each generate data corresponding to the comma-separated operands that follow the directive.

The DB directive generates a sequence of bytes. Each operand can be a constant expression with a value between -128 and 255 , or a character string of any length. Constant expressions generate one byte of code, and strings generate a sequence of bytes with values corresponding to the ASCII code of each character.

The DW directive generates a sequence of words. Each operand can be a constant expression with a value between $-32,768$ and $65,535$, or an address expression. For an address expression, the built-in assembler generates a near pointer, that is, a word that contains the offset part of the address.

The DD directive generates a sequence of double words. Each operand can be a constant expression with a value between $-2,147,483,648$ and $4,294,967,295$, or an address expression. For an address expression, the built-in assembler generates a far pointer, that is, a word that contains the offset part of the address, followed by a word that contains the segment part of the address.

The data generated by the DB, DW, and DD directives is always stored in the code segment, just like the code generated by other built-in assembler statements. To generate uninitialized or initialized data in the data segment, you should use Pascal **var** or **const** declarations.

Some examples of DB, DW, and DD directives follow:

```

asm
DB 0FFH { One byte }
DB 0,99 { Two bytes }
DB 'A' { Ord('A') }
DB 'Hello world...',0DH,0AH { String followed by CR/LF }
DB 12,"Borland Pascal" { Pascal style string }
DW 0FFFFH { One word }
DW 0,9999 { Two words }
DW 'A' { Same as DB 'A',0 }
DW 'BA' { Same as DB 'A','B' }
DW MyVar { Offset of MyVar }
DW MyProc { Offset of MyProc }
DD 0FFFFFFFFH { One double-word }
DD 0,999999999 { Two double-words }
DD 'A' { Same as DB 'A',0,0,0 }
DD 'DCBA' { Same as DB 'A','B','C','D' }
DD MyVar { Pointer to MyVar }
DD MyProc { Pointer to MyProc }
end;

```



In Turbo Assembler, when an identifier precedes a DB, DW, or DD directive, it causes the declaration of a byte, word, or double-word sized variable at the location of the directive. For example, Turbo Assembler allows the following:

```

ByteVar    DB    ?
WordVar    DW    ?
:
:
MOV        AL,ByteVar
MOV        BX,WordVar

```

The built-in assembler doesn't support such variable declarations. In Borland Pascal, the only kind of symbol that can be defined in an built-in assembler statement is a label. All variables must be declared using Pascal syntax, and the preceding construct corresponds to this:

```

var
  ByteVar: Byte;
  WordVar: Word;
  :
  :
asm
  MOV    AL,ByteVar
  MOV    BX,WordVar
end;

```


Operands

Built-in assembler operands are expressions that consist of a combination of constants, registers, symbols, and operators. Although built-in assembler expressions are built using the same basic principles as Pascal expressions, there are a number of important differences, as will be explained later in this chapter.

Within operands, the following reserved words have a predefined meaning to the built-in assembler:

Table 24.1
Built-in assembler reserved
words

AH	CS	LOW	SI
AL	CX	MOD	SP
AND	DH	NEAR	SS
AX	DI	NOT	ST
BH	DL	OFFSET	TBYTE
BL	DS	OR	TYPE
BP	DWORD	PTR	WORD
BX	DX	QWORD	XOR
BYTE	ES	SEG	
CH	FAR	SHL	
CL	HIGH	SHR	

The reserved words always take precedence over user-defined identifiers. For example, the code fragment,

```
var
    ch: Char;
    :
asm
    MOV     CH, 1
end;
```

loads 1 into the CH register, *not* into the CH variable. To access a user-defined symbol with the same name as a reserved word, you must use the ampersand (&) identifier override operator:

```
asm
    MOV     &ch, 1
end;
```

It's strongly suggested that you avoid user-defined identifiers with the same names as built-in assembler reserved words, because such name confusion can easily lead to obscure and hard-to-find bugs.

Expressions

The built-in assembler evaluates all expressions as 32-bit integer values. It doesn't support floating-point and string values, except string constants.

Built-in assembler expressions are built from *expression elements* and *operators*, and each expression has an associated *expression class* and *expression type*. These concepts are explained in the following sections.

Differences between Pascal and Assembler expressions

The most important difference between Pascal expressions and built-in assembler expressions is that all built-in assembler expressions must resolve to a *constant value*, a value that can be computed at compile time. For example, given these declarations:

```
const
  X = 10;
  Y = 20;
var
  Z: Integer;
```

the following is a valid built-in assembler statement:

```
asm
  MOV     Z, X+Y
end;
```

Because both *X* and *Y* are constants, the expression *X + Y* is merely a more convenient way of writing the constant 30, and the resulting instruction becomes a move immediate of the value 30 into the word-sized variable *Z*. But if you change *X* and *Y* to be variables,

```
var
  X, Y: Integer;
```

the built-in assembler can no longer compute the value of *X + Y* at compile time. The correct built-in assembler construct to move the sum of *X* and *Y* into *Z* is this:

```
asm
  MOV     AX, X
  ADD     AX, Y
  MOV     Z, AX
end;
```

Another important difference between Pascal and built-in assembler expressions is the way variables are interpreted. In a Pascal expression, a reference to a variable is interpreted as the *contents* of the variable, but in an built-in assembler expression, a variable reference denotes the *address* of the variable. For example, in Pascal, the expression $X + 4$, where X is a variable, means the contents of X plus 4, whereas in the built-in assembler, it means the contents of the word at an address four bytes higher than the address of X . So, even though you're allowed to write

```
asm
MOV    AX, X+4
end;
```

the code doesn't load the value of X plus 4 into AX , but it loads the value of a word stored four bytes beyond X instead. The correct way to add 4 to the contents of X is like this:

```
asm
MOV    AX, X
ADD    AX, 4
end;
```

Expression elements

The basic elements of an expression are *constants*, *registers*, and *symbols*.

Constants The built-in assembler supports two types of constants: *numeric constants* and *string constants*.

Numeric constants

Numeric constants must be integers, and their values must be between $-2,147,483,648$ and $4,294,967,295$.

By default, numeric constants use decimal (base 10) notation, but the built-in assembler supports binary (base 2), octal (base 8), and hexadecimal (base 16) notations as well. Binary notation is selected by writing a *B* after the number, octal notation is selected by writing a letter *O* after the number, and hexadecimal notation is selected by writing an *H* after the number or a $\$$ before the number.



The *B*, *O*, and *H* suffixes aren't supported in Pascal expressions. Pascal expressions allow only decimal notation (the default) and hexadecimal notation (using a \$ prefix).

Numeric constants must start with one of the digits 0 through 9 or a \$ character; therefore when you write a hexadecimal constant using the *H* suffix, an extra zero in front of the number is required if the first significant digit is one of the hexadecimal digits *A* through *F*. For example, 0BAD4H and \$BAD4 are hexadecimal constants, but BAD4H is an identifier because it starts with a letter and not a digit.

String constants

String constants must be enclosed in single or double quotes. Two consecutive quotes of the same type as the enclosing quotes count as only one character. Here are some examples of string constants:

```
'z'  
'Borland Pascal'  
"That's all folks"  
'"That"'s all folks," he said.'  
'100'  
'"  
"'"
```

Notice in the fourth string the use of two consecutive single quotes to denote one single quote character.

String constants of any length are allowed in DB directives, and cause allocation of a sequence of bytes containing the ASCII values of the characters in the string. In all other cases, a string constant can be no longer than four characters, and denotes a numeric value which can participate in an expression. The numeric value of a string constant is calculated as

$$\text{Ord}(\text{Ch1}) + \text{Ord}(\text{Ch2}) \text{ shl } 8 + \text{Ord}(\text{Ch3}) \text{ shl } 16 + \text{Ord}(\text{Ch4}) \text{ shl } 24$$

where *Ch1* is the rightmost (last) character and *Ch4* is the leftmost (first) character. If the string is shorter than four characters, the leftmost (first) character(s) are assumed to be zero. Here are some examples of string constants and their corresponding numeric values:

Table 24.2
String examples and their
values

String	Value
'a'	00000061H
'ba'	00006261H
'cba'	00636261H
'dcba'	64636261H
'a '	00006120H
' a'	20202061H
'a' * 2	000000E2H
'a'-'A'	00000020H
not 'a'	FFFFFF9EH

Registers

The following reserved symbols denote CPU registers:

Table 24.3
CPU registers

16-bit general purpose	AX	BX	CX	DX
8-bit low registers	AL	BL	CL	DL
8-bit high registers	AH	BH	CH	DH
16-bit pointer or index	SP	BP	SI	DI
16-bit segment registers	CS	DS	SS	ES
8087 register stack	ST			

When an operand consists solely of a register name, it's called a register operand. All registers can be used as register operands. In addition, some registers can be used in other contexts.

The base registers (BX and BP) and the index registers (SI and DI) can be written within square brackets to indicate indexing. Valid base/index register combinations are [BX], [BP], [SI], [DI], [BX+SI], [BX+DI], [BP+SI], and [BP+DI].

The segment registers (ES, CS, SS, and DS) can be used in conjunction with the colon (:): segment override operator to indicate a different segment than the one the processor selects by default.

The symbol ST denotes the topmost register on the 8087 floating-point register stack. Each of the eight floating-point registers can be referred to using ST(*x*), where *x* is a constant between 0 and 7 indicating the distance from the top of the register stack.

Symbols

The built-in assembler allows you to access almost all Pascal symbols in assembler expressions, including labels, constants, types, variables, procedures, and functions. In addition, the built-in assembler implements the following special symbols:

@Code

@Data

@Result

The `@Code` and `@Data` symbols represent the current code and data segments. They should only be used in conjunction with the `SEG` operator:

```
asm
MOV   AX, SEG @Data
MOV   DS, AX
end;
```

The `@Result` symbol represents the function result variable within the statement part of a function. For example, in this function,

```
function Sum(X, Y: Integer): Integer;
begin
  Sum := X + Y;
end;
```

the statement that assigns a function result value to `Sum` uses the `@Result` variable if it is written in built-in assembler:

```
function Sum(X, Y: Integer): Integer;
begin
  asm
    MOV   AX, X
    ADD   AX, Y
    MOV   @Result, AX
  end;
end;
```

The following symbols can't be used in built-in assembler expressions:

- Standard procedures and functions (for example, `WriteLn`, `Chr`)
- The `Mem`, `MemW`, `MemL`, `Port`, and `PortW` special arrays
- String, floating-point, and set constants
- Procedures and functions declared with the `inline` directive
- Labels that aren't declared in the current block
- The `@Result` symbol outside a function

Table 24.4 summarizes the value, class, and type of the different kinds of symbols that can be used in built-in assembler expressions. (Expression classes and types are described in a following section.)

Table 24.4
Values, classes, and types of
symbols

Symbol	Value	Class	Type
Label	Address of label	Memory	SHORT
Constant	Value of constant	Immediate	0
Type	0	Memory	Size of type
Field	Offset of field	Memory	Size of type
Variable	Address of variable	Memory	Size of type
Procedure	Address of procedure	Memory	NEAR or FAR
Function	Address of function	Memory	NEAR or FAR
Unit	0	Immediate	0
@Code	Code segment address	Memory	0FFF0H
@Data	Data segment address	Memory	0FFF0H
@Result	Result var offset	Memory	Size of type

Local variables (variables declared in procedures and functions) are always allocated on the stack and accessed relative to SS:BP, and the value of a local variable symbol is its signed offset from SS:BP. The assembler automatically adds [BP] in references to local variables. For example, given these declarations,

```

procedure Test;
var
    Count: Integer;

```

the instruction

```

asm
    MOV     AX, Count
end;

```

assembles into `MOV AX, [BP-2]`.

The built-in assembler always treats a **var** parameter as a 32-bit pointer, and the size of a **var** parameter is always 4 (the size of a 32-bit pointer). In Pascal, the syntax for accessing a **var** parameter and a value parameter is the same—this isn't the case in code you write for the built-in assembler. Because **var** parameters are really pointers, you have to treat them as such. So, to access the contents of a **var** parameter, you first have to load the 32-bit pointer and then access the location it points to. For example, if the *X* and *Y* parameters of the above function *Sum* were **var** parameters, the code would look like this:

```

function Sum(var X, Y: Integer): Integer;
begin
    asm
        LES     BX, X
        MOV     AX, ES: [BX]
        LES     BX, Y

```

```

        ADD    AX,ES:[BX]
        MOV    @Result,AX
    end;
end;

```

Some symbols, such as record types and variables, have a scope that can be accessed using the period (.) structure member selector operator. For example, given these declarations,

```

type
  TPoint = record
    X, Y: Integer;
  end;
  TRect = record
    A, B: TPoint;
  end;
var
  P: TPoint;
  R: TRect;

```

the following constructs can be used to access fields in the *P* and *R* variables:

```

asm
  MOV    AX,P.X
  MOV    DX,P.Y
  MOV    CX,R.A.X
  MOV    BX,R.B.Y
end;

```

A type identifier can be used to construct variables on the fly. Each of the following instructions generates the same machine code, which loads the contents of ES:[DI+4] into AX:

```

asm
  MOV    AX,(TRect PTR ES:[DI]).B.X
  MOV    AX,TRect(ES:[DI]).B.X
  MOV    AX,ES:TRect[DI].B.X
  MOV    AX,TRect[ES:DI].B.X
  MOV    AX,ES:[DI].TRect.B.X
end;

```

A scope is provided by type, field, and variable symbols of a record or object type. In addition, a unit identifier opens the scope of a particular unit, just like a fully qualified identifier in Pascal.

Expression classes

The built-in assembler divides expressions into three classes: *registers*, *memory references*, and *immediate values*.

An expression that consists solely of a register name is a register expression. Examples of register expressions are AX, CL, DI, and ES. Used as operands, register expressions direct the assembler to generate instructions that operate on the CPU registers.

Expressions that denote memory locations are memory references; Pascal's labels, variables, typed constants, procedures, and functions belong to this category.

Expressions that aren't registers and aren't associated with memory locations are immediate values; this group includes Pascal's untyped constants and type identifiers.

Immediate values and memory references cause different code to be generated when used as operands. For example,

```
const
  Start = 10;
var
  Count: Integer;
  :
asm
  MOV   AX, Start           { MOV AX, xxxx }
  MOV   BX, Count           { MOV BX, [xxxx] }
  MOV   CX, [Start]         { MOV CX, [xxxx] }
  MOV   DX, OFFSET Count    { MOV DX, xxxx }
end;
```

Because *Start* is an immediate value, the first MOV is assembled into a move immediate instruction. The second MOV, however, is translated into a move memory instruction, as *Count* is a memory reference. In the third MOV, the square brackets operator is used to convert *Start* into a memory reference (in this case, the word at offset 10 in the data segment), and in the fourth MOV, the OFFSET operator is used to convert *Count* into an immediate value (the offset of *Count* in the data segment).

As you can see, the square brackets and the OFFSET operators complement each other. In terms of the resulting machine code, the following **asm** statement is identical to the first two lines of the previous **asm** statement:

```

asm
  MOV   AX,OFFSET [Start]
  MOV   BX,[OFFSET Count]
end;

```

Memory references and immediate values are further classified as either *relocatable expressions* or *absolute expressions*. A relocatable expression denotes a value that requires *relocation* at link time, and an absolute expression denotes a value that requires no such relocation. Typically, an expression that refers to a label, variable, procedure, or function is relocatable, and an expression that operates solely on constants is absolute.

Relocation is the process by which the linker assigns absolute addresses to symbols. At compile time, the compiler doesn't know the final address of a label, variable, procedure, or function; it doesn't become known until link time, when the linker assigns a specific absolute address to the symbol.

The built-in assembler allows you to carry out any operation on an absolute value, but it restricts operations on relocatable values to addition and subtraction of constants.

Expression types

Every built-in assembler expression has an associated type—or more correctly, an associated size, because the built-in assembler regards the type of an expression simply as the size of its memory location. For example, the type (size) of an *Integer* variable is two, because it occupies 2 bytes.

The built-in assembler performs type checking whenever possible, so in the instructions

```

var
  QuitFlag: Boolean;
  OutBufPtr: Word;
  :
asm
  MOV   AL,QuitFlag
  MOV   BX,OutBufPtr
end;

```

the built-in assembler checks that the size of *QuitFlag* is one (a byte), and that the size of *OutBufPtr* is two (a word). An error results if the type check fails; for example, this isn't allowed:

```

asm
    MOV     DL, OutBufPtr
end;

```

The problem is DL is a byte-sized register and *OutBufPtr* is a word. The type of a memory reference can be changed through a typecast; these are correct ways of writing the previous instruction:

```

asm
    MOV     DL, BYTE PTR OutBufPtr
    MOV     DL, Byte (OutBufPtr)
    MOV     DL, OutBufPtr.Byte
end;

```

These MOV instructions all refer to the first (least significant) byte of the *OutBufPtr* variable.

In some cases, a memory reference is untyped; that is, it has no associated type. One example is an immediate value enclosed in square brackets:

```

asm
    MOV     AL, [100H]
    MOV     BX, [100H]
end;

```

The built-in assembler permits both of these instructions, because the expression [100H] has no associated type—it just means “the contents of address 100H in the data segment,” and the type can be determined from the first operand (byte for AL, word for BX). In cases where the type can’t be determined from another operand, the built-in assembler requires an explicit typecast:

```

asm
    INC     BYTE PTR [100H]
    IMUL   WORD PTR [100H]
end;

```

Table 24.5 summarizes the predefined type symbols that the built-in assembler provides in addition to any currently declared Pascal types.

Table 24.5
Predefined type symbols

Symbol	Type
BYTE	1
WORD	2
DWORD	4
QWORD	8
TBYTE	10
NEAR	0FFFEH
FAR	0FFFFH

Notice in particular the NEAR and FAR pseudotypes, which are used by procedure and function symbols to indicate their call model. You can use NEAR and FAR in typecasts just like other symbols. For example, if *FarProc* is a FAR procedure,

```
procedure FarProc; far;
```

and if you are writing built-in assembler code in the same module as *FarProc*, you can use the more efficient NEAR call instruction to call it:

```
asm
  PUSH  CS
  CALL  NEAR PTR FarProc
end;
```

Expression operators

The built-in assembler provides a variety of operators, divided into 12 classes of precedence. Table 24.6 lists the built-in assembler's expression operators in decreasing order of precedence.

Table 24.6
Summary of built-in
assembler expression
operators

Built-in assembler operator precedence is different from Pascal. For example, in a built-in assembler expression, the AND operator has lower precedence than the plus (+) and minus (-) operators, whereas in a Pascal expression, it has higher precedence.

Operator(s)	Comments
&	Identifier override operator
() , [] , *	Structure member selector
HIGH, LOW	
+, -	Unary operators
:	Segment override operator
OFFSET, SEG, TYPE, PTR, *, /, MOD, SHL, SHR	
+, -	Binary addition/ subtraction operators
NOT, AND, OR, XOR	Bitwise operators

Table 24.7: Definitions of built-in assembler expression operators

Operator	Description
&	Identifier override. The identifier immediately following the ampersand is treated as a user-defined symbol, even if the spelling is the same as a built-in assembler reserved symbol.
(...)	Subexpression. Expressions within parentheses are evaluated completely prior to being treated as a single expression element. Another expression can optionally precede the expression within the parentheses; the result in this case becomes the sum of the values of the two expressions, with the type of the first expression.
[...]	Memory reference. The expression within brackets is evaluated completely prior to being treated as a single expression element. The expression within brackets can be combined with the BX, BP, SI, or DI registers using the plus (+) operator, to indicate CPU register indexing. Another expression can optionally precede the expression within the brackets; the result in this case becomes the sum of the values of the two expressions, with the type of the first expression. The result is always a memory reference.
.	Structure member selector. The result is the sum of the expression before the period and the expression after the period, with the type of the expression after the period. Symbols belonging to the scope identified by the expression before the period can be accessed in the expression after the period.
HIGH	Returns the high-order 8 bits of the word-sized expression following the operator. The expression must be an absolute immediate value.
LOW	Returns the low-order 8 bits of the word-sized expression following the operator. The expression must be an absolute immediate value.
+	Unary plus. Returns the expression following the plus with no changes. The expression must be an absolute immediate value.
-	Unary minus. Returns the negated value of the expression following the minus. The expression must be an absolute immediate value.
:	Segment override. Instructs the assembler that the expression after the colon belongs to the segment given by the segment register name (CS, DS, SS, or ES) before the colon. The result is a memory reference with the value of the expression after the colon. When a segment override is used in an instruction operand, the instruction will be prefixed by an appropriate segment override prefix instruction to ensure that the indicated segment is selected.
OFFSET	Returns the offset part (low-order word) of the expression following the operator. The result is an immediate value.
SEG	Returns the segment part (high-order word) of the expression following the operator. The result is an immediate value.
TYPE	Returns the type (size in bytes) of the expression following the operator. The type of an immediate value is 0.
PTR	Typecast operator. The result is a memory reference with the value of the expression following the operator and the type of the expression in front of the operator.
*	Multiplication. Both expressions must be absolute immediate values, and the result is an absolute immediate value.

Table 24.7: Definitions of built-in assembler expression operators (continued)

/	Integer division. Both expressions must be absolute immediate values, and the result is an absolute immediate value.
MOD	Remainder after integer division. Both expressions must be absolute immediate values, and the result is an absolute immediate value.
SHL	Logical shift left. Both expressions must be absolute immediate values, and the result is an absolute immediate value.
SHR	Logical shift right. Both expressions must be absolute immediate values, and the result is an absolute immediate value.
+	Addition. The expressions can be immediate values or memory references, but only one of the expressions can be a relocatable value. If one of the expressions is a relocatable value, the result is also a relocatable value. If either of the expressions are memory references, the result is also a memory reference.
-	Subtraction. The first expression can have any class, but the second expression must be an absolute immediate value. The result has the same class as the first expression.
NOT	Bitwise negation. The expression must be an absolute immediate value, and the result is an absolute immediate value.
AND	Bitwise AND. Both expressions must be absolute immediate values, and the result is an absolute immediate value.
OR	Bitwise OR. Both expressions must be absolute immediate values, and the result is an absolute immediate value.
XOR	Bitwise exclusive OR. Both expressions must be absolute immediate values, and the result is an absolute immediate value.

Assembler procedures and functions

So far, every **asm...end** construct you've seen has been a statement within a normal **begin...end** statement part. Borland Pascal's assembler directive allows you to write complete procedures and functions in built-in assembler, without the need for a **begin...end** statement part. Here's an example of an assembler function:

```
function LongMul(X, Y: Integer): Longint; assembler;
asm
    MOV     AX, X
    IMUL   Y
end;
```

The **assembler** directive causes Borland Pascal to perform a number of code generation optimizations:

- The compiler doesn't generate code to copy value parameters into local variables. This affects all string-type value parameters, and other value parameters whose size isn't 1, 2, or 4 bytes. Within the procedure or function, such parameters must be treated as if they were **var** parameters.
- The compiler doesn't allocate a function result variable, and a reference to the *@Result* symbol is an error. String functions, however, are an exception to this rule—they always have a *@Result* pointer that is allocated by the caller.
- The compiler generates no stack frame for procedures and functions that aren't nested and have no parameters and no local variables.
- The automatically generated entry and exit code for an assembler procedure or function looks like this:

```

PUSH  BP                ;Present if Locals <> 0 or Params <> 0
MOV   BP,SP            ;Present if Locals <> 0 or Params <> 0
SUB   SP,Locals        ;Present if Locals <> 0
:
MOV   SP,BP            ;Present if Locals <> 0
POP   BP                ;Present if Locals <> 0 or Params <> 0
RET   Params            ;Always present

```

- *Locals* is the size of the local variables, and *Params* is the size of the parameters. If both *Locals* and *Params* are zero, there is no entry code, and the exit code consists simply of a RET instruction.

Functions using the **assembler** directive must return their results as follows:

- Ordinal-type function results (integer, boolean, enumerated types, and *Char*) are returned in AL (8-bit values), AX (16-bit values), or DX:AX (32-bit values).
- Real-type function results (type *Real*) are returned in DX:BX:AX.
- 8087-type function results (type *Single*, *Double*, *Extended*, and *Comp*) are returned in ST(0) on the 8087 coprocessor's register stack.
- Pointer-type function results are returned in DX:AX.
- String-type function results are returned in the temporary location pointed to by the *@Result* function result symbol.

The **assembler** directive is comparable to the **external** directive, and **assembler** procedures and functions must obey the same rules as **external** procedures and functions. The following

examples demonstrate some of the differences between **asm** statements in Pascal functions and assembler functions. The first example uses an **asm** statement in a Pascal function to convert a string to upper case. Notice that the value parameter *Str* in this case refers to a local variable, because the compiler automatically generates entry code that copies the actual parameter into local storage.

```

function UpperCase(Str: String): String;
begin
  asm
    CLD
    LEA    SI, Str
    LES    DI, @Result
    SEGSS LODSB
    STOSB
    XOR    AH, AH
    XCHG  AX, CX
    JCXZ  @3
  @1:
    SEGSS LODSB
    CMP   AL, 'a'
    JB   @2
    CMP   AL, 'z'
    JA   @2
    SUB   AL, 20H
  @2:
    STOSB
    LOOP @1
  @3:
  end;
end;

```

The second example is an assembler version of the *UpperCase* function. In this case, *Str* isn't copied into local storage, and the function must treat *Str* as a **var** parameter.

```

function UpperCase(Str: String): String; assembler;
asm
  PUSH   DS
  CLD
  LDS    SI, Str
  LES    DI, @Result
  LODSB
  STOSB
  XOR    AH, AH
  XCHG  AX, CX
  JCXZ  @3

```



```
@1:
  LODSB
  CMP    AL, 'a'
  JB     @2
  CMP    AL, 'z'
  JA     @2
  SUB    AL, 20H
@2:
  STOSB
  LOOP   @1
@3:
  POP    DS
end;
```


Linking assembler code

Procedures and functions written in assembly language can be linked with Borland Pascal programs or units using the **\$L** compiler directive. The assembly language source file must be assembled into an object file (extension .OBJ) using an assembler like Turbo Assembler. Multiple object files can be linked with a program or unit through multiple **\$L** directives.

Procedures and functions written in assembly language must be declared as **external** in the Pascal program or unit. For example,

```
function LoCase(Ch: Char): Char; external;
```

In the corresponding assembly language source file, all procedures and functions must be placed in a segment named CODE or CSEG, or in a segment whose name ends in **_TEXT**. The names of the external procedures and functions must appear in **PUBLIC** directives.

You must ensure that an assembly language procedure or function matches its Pascal definition with respect to call model (near or far), number of parameters, types of parameters, and result type.

An assembly language source file can declare initialized variables in a segment named **CONST** or in a segment whose name ends in **_DATA**. It can declare uninitialized variables in a segment named **DATA** or **DSEG**, or in a segment whose name ends in **_BSS**. Such variables are private to the assembly language source file and can't be referenced from the Pascal program or unit. However,

they reside in the same segment as the Pascal globals, and can be accessed through the DS segment register.

All procedures, functions, and variables declared in the Pascal program or unit, and the ones declared in the **interface** section of the used units, can be referenced from the assembly language source file through EXTRN directives. Again, it's up to you to supply the correct type in the EXTRN definition.

When an object file appears in a **\$L** directive, Borland Pascal converts the file from the Intel relocatable object module format (.OBJ) to its own internal relocatable format. This conversion is possible only if certain rules are observed:

- All procedures and functions must be placed in a segment named CODE or CSEG, or in a segment with a name that ends in _TEXT. All initialized private variables must be placed in a segment named CONST, or in a segment with a name that ends in _DATA. All uninitialized private variables must be placed in a segment named DATA or DSEG, or in a segment with a name that ends in _BSS. All other segments are ignored, and so are GROUP directives. The segment definitions can specify BYTE or WORD alignment, but when linked, code segments are always byte aligned, and data segments are always word aligned. The segment definitions can optionally specify PUBLIC and a class name, both of which are ignored.
- Borland Pascal ignores any data for segments other than the code segment (CODE, CSEG, or xxxx_TEXT) and the initialized data segment (CONST or xxxx_DATA). So, when declaring variables in the uninitialized data segment (DATA, DSEG, or xxxx_BSS), always use a question mark (?) to specify the value, for instance:

```
Count    DW    ?  
Buffer   DB   128 DUP(?)
```

- Byte-sized references to EXTRN symbols aren't allowed. For example, this means that the assembly language HIGH and LOW operators can't be used with EXTRN symbols.

Turbo Assembler and Borland Pascal

Turbo Assembler (TASM) makes it easy to program routines in assembly language and interface them into your Borland Pascal

programs. Turbo Assembler provides simplified segmentation and language support for Pascal programmers.

The `.MODEL` directive specifies the memory model for an assembler module that uses simplified segmentation. For linking with Pascal programs, the `.MODEL` syntax looks like this:

```
.MODEL xxxx, PASCAL
```

`xxxx` is the memory model (usually this is large).

Specifying the language `PASCAL` in the `.MODEL` directive tells Turbo Assembler that the arguments were pushed onto the stack from left to right, in the order they were encountered in the source statement that called the procedure.

The `PROC` directive lets you define your parameters in the same order as they are defined in your Pascal program. If you are defining a function that returns a string, notice that the `PROC` directive has a `RETURNS` option that lets you access the temporary string pointer on the stack without affecting the number of parameter bytes added to the `RET` statement.

Here's an example coded to use the `.MODEL` and `PROC` directives:

```
.MODEL LARGE, PASCAL
.CODE
MyProc PROC FAR I : BYTE, J : BYTE RETURNS Result : DWORD
PUBLIC MyProc
LES DI, Result ;get address of temporary string
MOV AL, I ;get first parameter I
MOV BL, J ;get second parameter J
:
RET
```

The Pascal function definition would look like this:

```
function MyProc(I, J: Char): string; external;
```

For more information about interfacing Turbo Assembler with Borland Pascal, refer to the *Turbo Assembler User's Guide*.

Examples of assembly language routines

The following code is an example of a unit that implements two assembly language string-handling routines. The *UpperCase* function converts all characters in a string to uppercase, and the

StringOf function returns a string of characters of a specified length.

```

unit Stringer;
interface
function UpperCase(S: String): String;
function StringOf(Ch: Char; Count: Byte): String;
implementation
{$L STRS}
function UpperCase; external;
function StringOf; external;
end.

```

The assembly language file that implements the *UpperCase* and *StringOf* routines is shown next. It must be assembled into a file called STRS.OBJ before the *Stringer* unit can be compiled. Note that the routines use the far call model because they are declared in the **interface** section of the unit. This example uses standard segmentation:

```

CODE    SEGMENT BYTE PUBLIC

        ASSUME    CS:CODE
        PUBLIC    UpperCase, StringOf      ;Make them known

; function UpperCase(S: String): String

UpperRes    EQU    DWORD PTR [BP + 10]
UpperStr    EQU    DWORD PTR [BP + 6]

UpperCase    PROC FAR

        PUSH    BP                ;Save BP
        MOV    BP, SP            ;Set up stack frame
        PUSH    DS                ;Save DS
        LDS    SI, Upperstr      ;Load string address
        LES    DI, Upperres      ;Load result address
        CLD                        ;Forward string-ops
        LODSB                       ;Load string length
        STOSB                       ;Copy to result
        MOV    CL, AL            ;String length to CX
        XOR    CH, CH
        JCXZ   U3                ;Skip if empty string
U1:      LODSB                       ;Load character
        CMP    AL, 'a'           ;Skip if not 'a'..'z'
        JB    U2
        CMP    AL, 'z'
        JA    U2
        SUB    AL, 'a'-'A'        ;Convert to uppercase
U2:      STOSB                       ;Store in result
        LOOP   U1                ;Loop for all characters

```

```

U3:    POP     DS                ;Restore DS
        POP     BP                ;Restore BP
        RET     4                ;Remove parameter and return

UpperCase    ENDP

; procedure StringOf(var S: String; Ch: Char; Count: Byte)

StrOfS      EQU     DWORD PTR [BP + 10]
StrOfChar   EQU     BYTE PTR [BP + 8]
StrOfCount  EQU     BYTE PTR [BP + 6]

StringOf    PROC FAR

        PUSH   BP                ;Save BP
        MOV   BP, SP            ;Set up stack frame
        LES   DI, StrOfRes      ;Load result address
        MOV   AL, StrOfCount    ;Load count
        CLD                     ;Forward string-ops
        STOSB                   ;Store length
        MOV   CL, AL            ;Count to CX
        XOR   CH, CH
        MOV   AL, StrOfChar     ;Load character
        REP   STOSB             ;Store string of characters
        POP   BP                ;Restore BP
        RET   8                ;Remove parameters and return

StringOf    ENDP

CODE    ENDS

        END

```

To assemble the example and compile the unit, use the following commands:

```

TASM STR5
TPCW stringer

```

Assembly language methods

Method implementations written in assembly language can be linked with Borland Pascal programs using the **\$L** compiler directive and the **external** reserved word. The declaration of an external method in an object type is no different than that of a normal method; however, the implementation of the method lists only the method header followed by the reserved word **external**. In an assembly language source text, an **@** is used instead of a period (.) to write qualified identifiers (the period already has a different meaning in assembly language and can't be part of an

identifier). For example, the Pascal identifier *Rect.Init* is written as *Rect@Init* in assembly language. The @ syntax can be used to declare both PUBLIC and EXTRN identifiers.

Inline machine code


For very short assembly language subroutines, Borland Pascal's **inline** statements and directives are very convenient. They let you insert machine code instructions directly into the program or unit text instead of through an object file.

Inline statements


An **inline** statement consists of the reserved word **inline** followed by one or more inline elements, separated by slashes and enclosed in parentheses:

```
inline(10/$2345/Count + 1/Data - Offset);
```

Here's the syntax of an **inline** statement:



Each inline element consists of an optional size specifier, < or >, and a constant or a variable identifier, followed by zero or more offset specifiers (see the syntax that follows). An offset specifier consists of a + or a - followed by a constant.



Each inline element generates 1 byte or 1 word of code. The value is computed from the value of the first constant or the offset of the variable identifier, to which is added or subtracted the value of each of the constants that follow it.

An inline element generates 1 byte of code if it consists of constants only and if its value is within the 8-bit range (0..255). If the

value is outside the 8-bit range or if the inline element refers to a variable, 1 word of code is generated (least-significant byte first).

The `<` and `>` operators can be used to override the automatic size selection we described earlier. If an inline element starts with a `<` operator, only the least-significant byte of the value is coded, even if it's a 16-bit value. If an inline element starts with a `>` operator, a word is always coded, even though the most-significant byte is 0. For example, the statement

```
inline(<$1234/>$44);
```

generates 3 bytes of code: `$34, $44, $00`.

*Registers BP, SP, SS, and DS must be preserved by **inline** statements; all other registers can be modified.*

The value of a variable identifier in an inline element is the offset address of the variable within its base segment. The base segment of global variables—variables declared at the outermost level in a program or a unit—and typed constants is the data segment, which is accessible through the DS register. The base segment of local variables—variables declared within the current subprogram—is the stack segment. In this case the variable offset is relative to the BP register, which automatically causes the stack segment to be selected.

The following example of an **inline** statement generates machine code for storing a specified number of words of data in a specified variable. When called, procedure *FillWord* stores *Count* words of the value *Data* in memory, starting at the first byte occupied by *Dest*.

```
procedure FillWord(var Dest; Count, Data: Word);
begin
  inline(
    $C4/$BE/Dest/           { LES DI, Dest[BP] }
    $8B/$8E/Count/         { MOV CX, Count[BP] }
    $8B/$86/Data/          { MOV AX, Data[BP] }
    $FC/                    { CLD }
    $F3/$AB);              { REP STOSW }
end;
```

Inline statements can be freely mixed with other statements throughout the statement part of a block.

Inline directives

With **inline** directives, you can write procedures and functions that expand into a given sequence of machine code instructions whenever they are called. These are comparable to macros in

assembly language. The syntax for an **inline** directive is the same as that of an **inline** statement:

inline directive → inline statement →

When a normal procedure or function is called (including one that contains **inline** statements), the compiler generates code that pushes the parameters (if any) onto the stack, and then generates a **CALL** instruction to call the procedure or function. However, when you call an inline procedure or function, the compiler generates code from the inline directive instead of the **CALL**. Here's a short example of two inline procedures:

```
procedure DisableInterrupts; inline($FA);      { CLI }
procedure EnableInterrupts; inline($FB);      { STI }
```

When *DisableInterrupts* is called, it generates 1 byte of code—a **CLI** instruction.

Procedures and functions declared with inline directives can have parameters; however, the parameters can't be referred to symbolically in the inline directive (other variables can, though). Also, because such procedures and functions are in fact macros, there is no automatic entry and exit code, nor should there be any return instruction.

The following function multiplies two *Integer* values, producing a *Longint* result:

```
function LongMul(X, Y: Integer): Longint;
inline(
    $5A/                          { POP AX ;Pop X }
    $58/                          { POP DX ;Pop Y }
    $F7/$EA);                     { IMUL DX ;DX : AX = X * Y }
```

Note the lack of entry and exit code and the missing return instruction. These aren't required, because the 4 bytes are inserted into the instruction stream when *LongMul* is called.

Inline directives are intended for very short procedures and functions only (less than 10 bytes).

Because of the macro-like nature of **inline** procedures and functions, they can't be used as arguments to the **@** operator and the *Addr*, *Ofs*, and *Seg* functions.

80486 processor *175*
 87 environment variable *182*
 @@ operator *79*
 ^ (pointer) symbol *43, 56*
 # (pound) character *19*
 @ operator *75*
 with a variable *75*
 with procedures and functions *75*
 80x87
 emulation *29, 176*
 floating-point model *28*
 numeric coprocessor *175-183*
 software emulation, selecting *28*

A

\$A compiler directive *309*
 AbortPrn procedure *171, 172*
 Abs function *148, 305*
 absolute
 clause *198*
 clause syntax *53*
 expressions, built-in assembler *330*
 variables *53, 198*
 AccessResource function *208*
 actual parameters *82*
 Addr function *150*
 address
 factor *66*
 functions *150*
 logical *193*
 physical *193*
 address-of (@) operator *43, 57, 75, 79*
 addressing
 protected mode *194*
 real-mode *194*
 aliased selectors *203*
 alignment, data *309*
 allocating
 protected-mode memory *201*
 Windows memory *268, 274*
 AllocDStoCSAlias function *209*
 AllocSelector function *209*
 ancestor of an object type *34*
 ancestors *34*
 and operator *70, 228*
 AnyFile constant *188*
 API, Windows *145*
 API, Windows 3.1 *146*
 apostrophes in character strings *19*
 Append procedure *155, 156*
 applications, spawned *212*
 Arc procedure *233*
 Archive constant *188*
 ArcTan function *148*
 arithmetic
 functions *148*
 operations precision rules *25*
 operators *68*
 array
 types *30*
 variables *55*
 array-type constant syntax *60*
 arrays *30, 55*
 accessing elements in *31*
 indexing multidimensional *55*
 number of elements in *31*
 of arrays *31*
 types *280*
 valid index types in *31*
 zero-based character *32, 61, 217, 219*
 defined *32*
 .ASM files *183*
 asm statement *316*
 assembler
 code
 in Borland Pascal *315*

- linked with Borland Pascal 339
 - declaration syntax 102
- assembly language
 - 80x87 emulation and 183
 - call model 339
 - inline
 - directives 345
 - statements 344
 - interfacing programs with 340
 - linking with Borland Pascal 339-346
 - overlays and 251
 - statements
 - multiple 316
 - syntax 316-321
- Assign procedure 155, 156, 172
- AssignCrt procedure 164, 168, 172
- AssignDefPrn procedure 172
- Assigned function 150
- assignment
 - compatibility 40, 48
 - object type 82
 - statement syntax 82
- AssignPrn procedure 172
- automatic
 - call model selection, overriding 294
 - data segment 271
 - jump sizing, built-in assembler 318
 - word alignment 309
- AutoTracking variable 167, 169
- AX register 293, 346

B

- \$B compiler directive 71, 306
- bar constants 237
- Bar3D procedure 223, 233
- Bar procedure 233
- base type 43
- .BGI files 223
- binary
 - arithmetic operators 68
 - operands 65
 - operators 25
- BIOS 162
- bit images 227
- bit-mapped fonts 226
- BitBlt
 - operations 228
 - operators 237
- bitwise operators 69
- blanks, defined 15
- block
 - defined 93
 - scope 95
 - subroutine 97
 - syntax 93
- BlockRead procedure 155, 159
- BlockWrite procedure 155, 159
- Boolean
 - data type 25, 276
 - expression evaluation 306
 - complete 70
 - short-circuit 70
 - operators 70
- boolean
 - data types 25
 - operators 26
 - variables 26
- Borland Graphics Interface 223-238
- Borland Pascal language overview 5-14
- BP register 251, 301, 303
- brackets, in expressions 76, 77
- Break procedure 148
- BufEnd variable 288
- buffer
 - overlay 241
 - loading and freeing up 242
 - optimization algorithm 243
 - probationary area 243
 - text, size 288
- BufPtr pointer 288
- BufSize variable 288
- built-in assembler
 - directives 315
 - expressions 322-334
 - classes 329-330
 - operators 332-334
 - Pascal expressions versus 322
 - types 330-332
 - instruction sizing 318-319
 - opcodes 317-319
 - operands 321
 - procedures and functions 334
 - registers, using 316
 - reserved words 321

- BX register *293, 303*
- Byte data type *25*
- ByteBool data type *276*

- C**
- \$C compiler directive *199*
- call model *339*
- calling conventions *291*
 - constructors and destructors *298*
 - methods *296*
- calls, near and far *294*
- case
 - sensitivity of Turbo Pascal *16*
 - statement syntax *85*
- CGA *223*
- ChangeSelector function *209*
- changing fonts *170*
- Char data type *26, 276*
- character
 - arrays *219*
 - pair special symbols *16*
 - pointer operators *71*
 - pointers
 - character arrays and *219*
 - indexing *219*
 - string literals and *217*
 - strings *20*
- ChDir procedure *155*
- CheckBreak variable *165, 170*
- CheckEOF variable *165, 168, 169*
- CheckSnow variable *165*
- .CHR files *223*
- Chr function *26, 148, 305*
- Circle procedure *233*
- circular unit references *123*
- ClearDevice procedure *233*
- ClearViewPort procedure *233*
- clipping constants *237*
- Close procedure *155, 157, 172, 174*
- CloseGraph procedure *224, 233*
- ClrEol procedure *164, 168*
- ClrScr procedure *164, 168*
- CmdLine variable *152*
- CmdShow variable *137, 152*
- code segment *340*
 - attributes *266, 270*
 - changing *268, 271*
 - controlling attributes of *199*
 - defined *195*
 - maximum size of *266, 270*
 - procedures and functions in *339*
 - writing to a *199*
- color constants *236*
 - text *165*
- colors, maximum number of *237*
- COM devices *161*
- command-line parameters *151*
- comments *20*
 - built-in assembler *316*
- common types of integer types *25*
- communication devices (COM1 and COM2) *161*
- Comp data type *28, 177, 278*
- comparing
 - character pointers *74*
 - packed strings *74*
 - pointers *74*
 - sets *74*
 - simple types *73*
 - strings *74*
 - values of real types *179*
- compatibility
 - assignment *40*
 - parameter type *111*
- compiler
 - directives
 - \$C *199*
 - \$N *29*
 - \$P *113*
 - \$A *309*
 - \$B *71, 306*
 - defined *20*
 - \$F *46, 99, 172, 246, 294*
 - \$G *317*
 - \$I *157*
 - \$L *339, 340, 343*
 - \$L filename *101, 183*
 - \$M *52, 259, 272, 273*
 - \$N *28, 69, 148, 176, 181, 317*
 - \$O *245*
 - nonoverlay units and *250*
 - \$R *39, 282*
 - \$S *52, 136*
 - \$T *48, 75*

- \$X 20, 32, 44, 54, 71
- optimization of code 305-311
- complete Boolean evaluation 70
- compound statement syntax 84
- CON device 160
- Concat function 149
- concatenation 71
- conditional statement syntax 84
- CONST segment 339
- constant
 - address expressions 59
 - declaration part syntax 94
 - declarations 21
 - defined 11
 - expressions 21
 - parameters 110
 - with an initial value 11
- constants 21, 191
 - array-type 60
 - Dos unit 188
 - folding 305
 - Graph unit 236
 - merging 306
 - numeric, built-in assembler 323
 - object-type 63
 - pointer-type 63
 - procedural-type 64
 - record-type 62
 - set-type 63
 - simple-type 21, 59
 - string, built-in assembler 324
 - string-type 60
 - structured-type 60
 - typed 59
- constructor syntax 106
- constructors 37, 38, 39, 283
 - calling conventions 298
 - declaring 105
 - defined 106
 - error recovery 108
 - virtual methods and 106
- Continue procedure 148
- control
 - characters 19, 168
 - defined 15
 - embedding in strings 19
 - in CRT window 163
 - in WinPrn unit 171
 - string syntax diagram 20
- Copy function 149
- Cos function 148
- CreateDir procedure 191
- creating objects 38
- Crt
 - mode constants 165
 - unit 144, 162
 - variables in 165
- CRT window 167
 - closing 167
 - control characters in 163
 - scrolling 167
- CS register 303
- CSeg function 150
- CSEG segment 340
- current
 - file position 157
 - file size 157
 - pointer 226
- Cursor variable 169
- CursorTo procedure 169
- CX register 303

D

- data
 - alignment 309
 - ports 288
 - segment 339, 340
 - defined 195
 - DLL 137
 - maximum size of 52
- DATA segment 339
- date and time procedures
 - Dos unit 186
 - WinDos unit 189
- DateTime type 189
- dead code eliminated 310
- debugging overlays 251
- Dec procedure 149
- decimal notation 18
- declaration part, defined 11
- declaring
 - an object type 36
 - methods 37
- Delay procedure 164

- Delete procedure *149*
- DelLine procedure *164*
- DEMANDLOAD code segment attribute *266, 270*
- dereferencing nil pointers *199*
- descendants *34*
- descriptor table *194*
- descriptors *195*
- designators
 - field *56*
 - method *56*
- destructor syntax *107*
- destructors *106*
 - calling conventions *298*
 - declaring *105*
 - defined *106*
- DetectGraph procedure *233*
- devices *160*
 - communication (COM1 and COM2) *161*
 - console (CON) *160*
 - DOS *160*
 - drivers *172*
 - handlers *303, 304*
 - line printer *161*
 - NUL *161*
 - text file *161*
- DI register *303*
- diagrams, syntax *14*
- digit syntax diagram *15*
- digits, defined *15*
- direct memory *288*
- directives
 - assembler, defined *319*
 - built-in assembler *315, 334, 335*
 - compiler, defined *20*
 - external *101*
 - far *98*
 - forward *100*
 - inline *102*
 - interrupt *100*
 - list of Borland Pascal *17*
 - near *98*
 - private *17*
 - public *17*
 - standard *16*
- Directory constant *188*
- directory-handling procedures and functions
 - WinDos unit *191*
- DirectVideo variable *165*
- DISCARDABLE code segment attribute *267, 271*
- disk status functions
 - Dos unit *187*
 - WinDos unit *190*
- DiskFree function *187, 190*
- DiskSize function *187, 190*
- dispatcher, RTL (run-time library) *297*
- Dispose procedure *150, 260, 261, 263*
 - extended syntax *282, 298*
 - constructor passed as parameter *106, 116*
- disposing of dynamic variables *260*
- div operator *69*
- DLL *127-139*
 - binary-compatible *139*
 - CmdShow variable *137*
 - contrasted with a unit *127*
 - data segment *137*
 - exit code *135*
 - files in a *137*
 - global memory in a *137*
 - global variables in *137*
 - gmem_DDEShare attribute *137*
 - HeapLimit variable *138*
 - HPrevInst variable *137*
 - initialization code *135*
 - KERNEL *128*
 - multi-language programming *128*
 - PrefixSeg variable *137*
 - run-time errors in a *138*
 - run-time manager and *197*
 - shared *139*
 - structure of *132*
 - syntax *132*
 - unloading a *136*
 - using a *128*
 - writing a *132-136*
- DMT (dynamic method table)
 - cache *286, 297*
 - entry count *286*
- domain of object type *34*
- DoneWinCrt procedure *167, 169*
- DOS
 - device handling *304*

- devices 160
- environment 257
- error level 302
- exit code 301
- operating system routines 185
- DOS3Call function 209
- DOS Protected Mode Interface (DPMI) 196
- Dos unit 144, 185
 - constants 188
 - date and time procedures 186
 - disk status functions 187
 - environment-handling functions 187
 - file-handling procedures and functions 187
 - interrupt support procedures 186
 - miscellaneous procedures and functions 188
 - types 189
- DosError variable
 - Dos unit 189
 - WinDos unit 192
- DosVersion function 188, 191
- double address-of (@@) operator 79
- Double data type 28, 177, 278
- DPMI 196
 - server 196
 - accessing directly 210
 - non-Borland 211
- DPMI16BLOVL file 196, 211
- DrawPoly procedure 233
- driver constants 236
- drivers
 - graphics 223-225
 - text-file device 172
- DS register 301, 303
- DSEg function 150
- DSEG segment 339
- DX register 293, 303
- dynamic
 - allocation procedures and functions 149
 - importing 131
 - linking 127
 - method calls 297
 - method index 38
 - method table 283
 - cache 286
 - entry count 286
 - methods 38, 284
 - how differ from virtual methods 38

- overriding 38
- object instances
 - allocation and disposal 106, 298
 - variables 43, 52, 56, 259
- dynamic-link libraries 127, *See also* DLL

E

- \$E compiler directive 28
- editing keys
 - in the CRT window 163
 - in WinCrt unit 168
- eliminate dead code 310
- Ellipse procedure 233
- embedding control characters in strings 19
- empty set 42
- EMS memory, overlay files and 240, 247
- emulating the 80x87 29, 176
- end-of-file character 161
- end-of-line character 15
- entry code 298, 335
- enumerated
 - constant ordinality 27
 - types 26, 276
- EnvCount function 187
- environment-handling functions
 - Dos unit 187
 - WinDos unit 191
- EnvStr function 187
- Eof function 155
- Eoln function 155
- Erase procedure 155
- error checking
 - dynamic object allocation 108
 - virtual method calls 282
- ErrorAddr variable 151, 152, 153, 302
- errors
 - fatal, in OvrInit 248
 - handling 228
 - reporting 301
- ES register 303
- examples
 - array type 31
 - avoiding ambiguity using subrange types 28
 - character strings 19
 - constant expressions 21, 22
 - constructor 39
 - control characters in strings 19

- enumerated type 27
- expressions 9
- function 6
- initializing virtual methods 39
- object-type declaration 35
- record type 32
- simple statements 8
- subrange type 27
- syntax diagram 14
- tokens 10
- variables 11
- variant part of a record 33
- Exclude procedure 151
- .EXE files 239
 - building 310
- Exec procedure 187
- exit
 - code 335
 - in DLL 135
 - functions 298
 - procedures 135, 136, 298, 301
- Exit procedure 148
- ExitCode variable 135, 151, 152, 153, 302
- exiting a program 301
- ExitProc variable 135, 136, 301
- Exp function 148
- exponents 277
- export directive 132, 133
- exporting procedures and functions 132-134
- exports clause 132
 - syntax 133
- expression syntax 66-68
- expressions 65-79
 - absolute, built-in assembler 330
 - built-in assembler 322-334
 - classes 329-330
 - elements of 323-328
 - versus Pascal 322
 - constant 21
 - standard functions permitted in 22
 - constant address 59
 - defined 9
 - elements of, built-in assembler 323
 - order of evaluation 306
 - relocation, built-in assembler 330
 - types, built-in assembler 330
- Extended data type 28, 177, 178, 278

- range arithmetic 178
- range of 176
- extended syntax 20, 32, 44
- external
 - (reserved word) 343
 - declaration 101, 128
 - directive 101, 335
 - in imported procedure and functions 128
 - procedures and functions 183, 339
- EXTRN directive 340

F

- \$F compiler directive 46, 99, 172, 246, 294
- faAnyFile constant 192
- faArchive constant 192
- factor syntax 66
- faDirectory constant 192
- faHidden constant 192
- Fail procedure 108
- False predefined constant identifier 26
- far
 - call 294
 - model 246
 - in imported procedures and function 128
 - model, forcing use of 301
 - requirement 241
 - directive 98
- faReadOnly constant 192
- faSysFile constant 192
- FatalExit function 209
- FAuxiliary constant 188
- fAuxiliary constant 192
- faVolumeID constant 192
- FCarry constant 188
- fCarry constant 192
- fcDirectory constant 192
- fcExtension constant 192
- fcFileName constant 192
- fcWildcards constant 192
- FExpand function 187
- Fibonacci numbers 180
- field
 - designators syntax 56
 - list (of records) 32
 - record 56

- fields
 - in record types 32
 - object 34
 - scope 105
- figures, graphics 227
- file *See also* files
 - buffer 288
 - format of protected-mode executable 211
 - handles 287
 - modes 287
 - types 42
- file-handling
 - functions
 - Dos unit 187
 - WinDos unit 190
 - procedures 187
 - WinDos unit 190
- FileExpand function 190
- FileMode variable 151, 152, 153, 159
- FilePos function 155, 157
- FileRec type 189
- files
 - access, read-only 159
 - .ASM 183
 - .BGI 223
 - .CHR 223
 - .EXE 239
 - building 310
 - functions for 155
 - I/O 162
 - .OBJ 339
 - .OVR 239
 - procedures for 155
 - text 157
 - typed 286
 - types of 286
 - untyped 159, 286
- FileSearch function 190
- FileSize function 155, 157
- FileSplit function 190
 - source code of 221
- fill pattern constants 237
- FillChar procedure 151
- FillEllipse procedure 234
- FillPoly procedure 227, 234
- FindFirst procedure 187, 190
- finding the size of a given string 30
- FindNext procedure 187, 190
- FindResource function 208
- FIXED code segment attribute 266, 270
- fixed part of records 32
- floating-point
 - calculations, type Real and 178
 - code generation, switching 176
 - numbers 28, 175
 - numeric coprocessor (80x87) 29
 - parameters 292
 - software 28
 - types 177
- FloodFill procedure 227, 234
- flow-control procedures 147
- Flush procedure 155
- fmClosed constant 188, 192
- fmInOut constant 188, 192
- fmInput constant 188, 192
- fmOutput constant 188, 192
- font constants 236
- fonts
 - changing 170
 - files 231
 - stroked 223, 226
- for statement syntax 88
- Force Far Calls option 246
- formal
 - parameter list syntax 109
 - parameters 76, 82, 109
- format of protected-mode executable file 211
- forward
 - declarations 100
 - directive 100
- FOverflow constant 188
- fOverflow constant 192
- FParity constant 188
- fParity constant 192
- Frac function 148
- free list 263
- FreeLibrary procedure 207
- FreeList variable 151
- FreeMem procedure 150, 260, 261, 263
- FreeResource function 208
- FreeSelector function 209
- FreeZero variable 151
- fsDirectory constant 192
- FSearch function 187

- fsExtension constant 192
- fsFileName constant 192
- FSign constant 188
- fSign constant 192
- fsPathName constant 192
- FSplit function 187
- function
 - calls 76
 - extended syntax and 76
 - syntax 76
 - declarations 103-104
 - assembler 101
 - external 101
 - headings 103
 - results 293
 - returns, built-in assembler 335
 - syntax 103
- functions 6, 97, *See also* procedures and functions
 - address 150
 - arithmetic 148
 - calls 291
 - directory-handling 191
 - disk status 187, 190
 - entry/exit code, built-in assembler 335
 - environment-handling 187, 191
 - far 294
 - file-handling 187, 190
 - graphics 233
 - heap error 108
 - miscellaneous 188, 191
 - near 294
 - nested 294
 - ordinal 149
 - OvrGetRetry 243
 - parameters, built-in assembler 334
 - pointer 150
 - private 121
 - program example 6
 - SizeOf 114
 - stack frame for, built-in assembler 335
 - standard 147
 - and constant expressions 22
 - string 149
 - transfer 148
- FZero constant 188
- fZero constant 192

G

- \$G compiler directive 317
- general protection fault 195, 197
 - practices that result in a 198
- GetArcCoords procedure 234
- GetArgCount function 191
- GetArgStr function 191
- GetAspectRatio procedure 234
- GetBkColor function 234
- GetCBreak procedure 188, 191
- GetColor function 234
- GetCurDir function 191
- GetDate procedure 186, 189
- GetDefaultPalette function 234, 237
- GetDir procedure 155
- GetDOSEnvironment function 209
- GetDriverName function 234
- GetEnv function 187
- GetEnvVar function 191
- GetFAttr procedure 187, 190
- GetFillPattern procedure 234
- GetFillSettings procedure 234
- GetFreeSpace function 204
- GetFTime procedure 186, 189
- GetGraphMode function 234
- GetImage procedure 223, 234
- GetIntVec procedure 186, 190
- GetLineSettings procedure 234
- GetMaxColor function 234
- GetMaxMode function 234
- GetMaxX function 234
- GetMaxY function 234
- GetMem function, in protected mode 233
- GetMem procedure 56, 150, 265
- GetModeName function 234
- GetModeRange procedure 234
- GetModuleFileName function 207
- GetModuleHandle function 207
- GetModuleUsage function 208
- GetPalette procedure 234, 237
- GetPaletteSize function 234
- GetPixel function 228, 234
- GetProcAddress function 208
- GetSelectorBase function 209
- GetSelectorLimit function 209
- GetTextSettings procedure 227, 234
- GetTime procedure 186, 189

GetVerify procedure *188, 191*
 GetVersion function *210*
 GetViewSettings procedure *234*
 GetWinFlags function *210*
 GetX function *234*
 GetY function *234*
 global
 DOS protected-mode heap *268*
 heap in Windows *273*
 memory in a DLL *137*
 variables in a DLL *137*
 GlobalAlloc function *204, 205*
 GlobalAllocPtr function *204, 207*
 GlobalCompact function *204*
 GlobalDiscard function *204*
 GlobalDosAlloc function *204*
 GlobalDosFree function *204*
 GlobalFix function *204*
 GlobalFlags function *204*
 GlobalFree function *204*
 GlobalFreePtr function *205, 207*
 GlobalHandle function *205*
 GlobalLock function *205*
 GlobalLockPtr function *205*
 GlobalLRUOldest function *205*
 GlobalNotify procedure *205*
 GlobalPageLock function *205*
 GlobalPageUnlock function *205*
 GlobalPtrHandle function *205*
 GlobalReAlloc function *205, 207*
 GlobalReAllocPtr function *205*
 GlobalRUNewest function *205*
 GlobalSize function *205*
 GlobalUnfix function *205*
 GlobalUnlock function *205*
 GlobalUnlockPtr function *205*
 gmem_DDEShare attribute *137, 138*
 gmem_Discardable attribute *206*
 gmem_Fixed attribute *206*
 gmem_Moveable attribute *138, 206*
 goto statement syntax *83*
 GotoXY procedure *164, 169*
 GP fault *195, 197*
 practices that result in a *198*
 Graph3 unit *145*
 GRAPH.TPP file *223*
 GRAPH.TPU file *223*

Graph unit *145, 223, 248*
 bit images in *227*
 colors *228*
 constants *236*
 error handling *228*
 figures and styles in *227*
 heap management routines *231*
 paging *228*
 procedures *233*
 sample program *229, 230*
 text in *226*
 types *237*
 variables *238*
 viewports in *227*
 GraphDefaults procedure *234*
 GraphDriver variable, IBM 8514 and *224*
 GraphErrorMsg function *234*
 GraphFreeMem procedure *231*
 GraphFreeMemPtr variable *238*
 GraphGetMem procedure *231*
 GraphGetMemPtr variable *238*
 graphics
 CloseGraph *224*
 current pointer in *226*
 drivers *223*
 figures and styles *227*
 in protected mode *223, 233*
 InitGraph in *224*
 mode constants *236*
 sample program *229, 230*
 using *223-238*
 GraphResult errors *236*
 GraphResult function *228, 235*
 grXXXX constants *236*

H

Halt procedure *148, 301*
 handles *206*
 file *287*
 hardware, interrupts *303*
 heading, program *5*
 heap
 error function *108*
 global
 DOS protected-mode *268*
 management *259*
 allocating *259, 260, 263, 265*

- deallocating 260
- fragmenting 259
- free list 263
- map 258
- pointers 258
- protected-mode 268
- routines 231
- Windows 273
 - fragmenting 273
- manager
 - DOS protected-mode 268
 - allocating memory blocks 269
 - protected-mode 200
 - Windows 273
 - allocating memory blocks 274
- Windows global 273
- Windows local 273
 - changing size of 273
 - size of 273
- HeapAllocFlags variable 138, 153
- HeapBlock variable 153, 268, 274
- HeapEnd variable 151
- HeapError variable 151, 153, 265, 269, 275
- HeapLimit variable 153, 233, 268, 274
- HeapList variable 153
- HeapOrg variable 151, 259, 260
- HeapPtr variable 151, 259
- hex digits 15
- hexadecimal
 - constants 18
 - numbers 19
- Hi function 151, 305
- Hidden constant 188
- high
 - bounds of index type of an array, finding 31
 - resolution graphics 224
- High function 24, 31, 114, 148
- highest value in a range, finding 24
- HightVideo procedure 164
- HInstance variable 137, 153, 208
- host type 27
- HPrevInst variable 153

I

- \$I compiler directive 157
- I/O 155
 - devices 172
 - error-checking 157
 - files 162
 - redirection 162
- IBM 8514 224
 - driver support 224-225
 - GraphDriver variable and 224
 - InitGraph procedure and 224
 - modes 224
 - SetRGBPalette and 225
- identifiers
 - as labels 19
 - defined 17
 - examples 18
 - how identified in manuals 18
 - length of 17
 - qualified 17
 - restrictions on naming 17
 - scope of 23
- if statement syntax 84
- ImageSize function 235
- immediate values, built-in assembler 329
- implementation part
 - of a unit 121, 294
 - syntax 121
- implementing methods 37
- import units 129
- importing procedures and functions 128-132
 - dynamically 131
 - statically 131
 - with import unit 129
- in operator 73, 75
- InactiveTitle variable 169
- Inc procedure 149
- Include procedure 151
- index
 - clause 128, 134
 - dynamic method 38
 - syntax 55
 - types valid in arrays 31
- indexes in arrays 31
- indexing character pointers 219
- indirect unit references 122
- infinite loop *See* loop, infinite
- inheritance, rules of 34
- inherited (reserved word) 41
- InitGraph procedure 224, 235

- initialization
 - code in DLL 135
 - part of a unit 122
- initialized variables 59
 - in assembler 339
- initializing virtual methods 37, 38
- InitWinCrt procedure 167, 169
- inline
 - directives 102, 345
 - statements 344
- InOutRes variable 151, 153
- Input text-file variable 158
 - and WinCrt unit 166
 - in WinCrt unit 158
- Input variable 151, 153
- Insert procedure 149
- InsLine procedure 164
- InstallUserDriver function 235
- InstallUserFont function 235
- instances
 - dynamic object 39
 - of an object type 38
- instantiating objects 38
- instruction opcodes, built-in assembler 317
- Int function 148
- Integer data type 25, 276
- integer types 25
- interface part of a unit 121, 294
- internal data formats 276
- interrupt
 - directives 100, 303
 - handlers 303
 - units and 250
 - handling routines 303
 - procedures, writing 303
 - service routines (ISRs) 303
 - support procedures
 - Dos unit 186
 - WinDos unit 190
- Intr procedure 186, 190
- IOResult function 155, 157
- IP flag 303
- ISRs (interrupt service routines) 303

J

- jump sizing, automatic, built-in assembler 318
- justify text constants 237

K

- Keep procedure 187
- KERNEL dynamic-link library 128
- keyboard status, testing 164
- KeyPressed function 164, 168, 169

L

- \$L compiler directive 339, 340, 343
- \$L filename compiler directive 101, 183
- label
 - declaration part syntax 93
 - syntax 19
- labels
 - built-in assembler 317
 - defined 19
- language overview 5
- LastMode variable 166
- late binding 37
- left
 - brace special symbol 16
 - bracket special symbol 16
- Length function 149, 305
- length of
 - a string-type value, finding 29
 - character strings 20, 279
 - identifiers 17
 - program lines 20
 - record 288
- letters, defined 15
- libraries
 - dynamic-link 127
 - run-time 143
- library header 132
- limit of a segment 199
- line
 - input editing keys 163, 168
 - style constants 236
- Line procedure 235
- LineRel procedure 235
- lines, maximum length of 20
- LineTo procedure 235
- linking
 - Borland Pascal with assembler code 339-346
 - dynamically 127
 - resources 208
 - smart 310

- statically 127
- Ln function 148
- Lo function 151, 305
- loading invalid values into segment registers 198
- LoadLibrary function 208
- LoadResource function 208
- LoadString function 208
- local
 - heap, Windows 272, 273
 - labels 317
- LockResource function 208
- LockSegment function 205
- logical
 - address 193
 - operators 69
- LongBool data type 25, 276
- Longint data type 25
- loop, infinite *See* infinite loop
- low bounds of index type of an array, finding 31
- Low function 24, 31, 114, 148
- lowest value in a range, finding 24
- LowVideo procedure 164
- LPT devices 161

M

- \$M compiler directive 52, 259, 272, 273
- machine code in program 344
- Mark procedure 260, 268, 273
- MaxAvail function 150
- Mem array 288
 - Ptr function and 198
- MemAvail function 150
- MemL array 288
 - Ptr function and 198
- memory
 - allocation 248
 - blocks
 - huge 201
 - types of 205
 - Borland Pascal and 257
 - manager, protected-mode 204-207, 268
 - map 258
 - model 341
 - references, built-in assembler 329
- MemW array 288

- Ptr function and 198
- MessageBox function 210
- method
 - declarations 104-109
 - designator 56
 - syntax of 40
- methods 34-42, 104-109
 - activating 40
 - assembly language 343
 - calling
 - conventions 296
 - dynamic 297
 - declaring 37, 104
 - defined 34
 - designators 56
 - dynamic 38, 284, 297
 - how differ from virtual methods 38
 - overriding 38
 - external 343
 - forward declaration 37
 - identifiers, qualified 37
 - implementation 37, 104
 - making them virtual 37
 - overriding inherited 38
 - parameters
 - Self 105
 - defined 296
 - type compatibility 111
 - qualifying method identifiers 37
 - static 37
 - virtual 37
 - calling 296
 - error checking 282
 - initializing 38
- miscellaneous procedures and functions
 - Dos unit 188
 - WinDos unit 191
- MkDir procedure 155
- mod operator 69
- .MODEL directive 341
- modular programming 120
- module management, protected-mode 207
- Move procedure 151
- MOVEABLE code segment attribute 266, 270
- MoveRel procedure 235
- MoveTo procedure 235
- MsDos procedure 186, 190

multi-language programming and DLLs 128

N

`$N` compiler directive 28, 29, 69, 148, 176, 181, 317

name clause 128, 134

near

call 294

directive 98

nested procedures and functions 46, 294

network file access, read-only 159

New procedure 43, 56, 150, 259, 265

extended syntax 282

constructor passed as parameter 106, 116, 298

used as function 117

nil (reserved word) 43, 57

nil pointers, dereferencing 199

NormVideo procedure 165

NoSound procedure 165

not operator 70, 228

NUL device 161

NULL character 215

null strings 19, 29

null-terminated strings 32, 145, 215-222

defined 215

NULL character 215

pointers and 217

standard procedures and 221

number

constants 18

syntax 18

numbers

counting 18, 276

hexadecimal 19

integer 19

real 18

numeric

constants, built-in assembler 323

coprocessor

detecting 182

emulating, assembly language and 183

evaluation stack 180

using 175-183

O

`$O` compiler directive 245

nonoverlay units and 250

.OBJ files 339

object

ancestor 34

component designators 56

descendant 34

files 339

scope 96

object-type

assignments 82

constants 63

object types 34-42. *See also* objects

components 34

declaring 36

domain 34

fields 34

instances 38

methods 34

rules of inheritance 34

scope

of identifier in 36

in private sections 36

in public sections 36

objects

ancestor 34

constructors 283

declaring 105

defined 106

error recovery 108

virtual methods and 106

creating 38

destructors 106

declaring 105

defined 106

domain of 34

dynamic instances 39

allocation and disposal of 106, 298

dynamic method table 283

fields

designators 56

scope 36, 105

files in `$L` directive 340

instantiating 38

internal data format 280

methods, scope 36

- pointers to 40
- virtual method table 282, 283
 - pointer 280
- virtual methods
 - call error checking 282
 - calling 296
- Odd function 149, 305
- Ofs function 150
- open parameters 110, 113
 - array 32, 110, 114
 - how passed 293
 - string 30, 113
- OpenString identifier 30
- operands 65
 - built-in assembler 321
- operators 65-75
 - @@ (double address-of) 79
 - @ (address-of) 43, 57, 75, 79
 - and 70, 228
 - arithmetic 68
 - binary arithmetic 68
 - bitwise 69
 - Boolean 70
 - built-in assembler, defined 333
 - character-pointer 71
 - div 69
 - logical 69
 - mod 69
 - not 70, 228
 - or 70, 228
 - precedence of 65, 69
 - built-in assembler 332
 - relational 73
 - set 72
 - shl 70
 - shr 70
 - string 71
 - structure member selector 328
 - types of 68
 - unary arithmetic 69
 - xor 70, 228
- optimization of code 305-311
- or operator 70, 228
- Ord function 24, 26, 27, 148, 305
- order of evaluation 308
- ordering between two string-type values 29
- ordinal
 - procedures and functions 149
 - types 24-28
 - predefined 25
 - user-defined 25
- ordinality
 - defined 24
 - enumerated constant 27
 - finding enumerated type's value 27
 - returning 24
 - Char values 26
- Origin variable 169
- Output text-file
 - in WinCrt unit 158
- Output text-file variable 158
 - and WinCrt unit 166
- Output variable 151, 153
- OutText procedure 227, 235
- OutTextXY procedure 227, 235
- overlaid
 - code, storing 259
 - initialization code 249
 - programs
 - designing 245-252
 - writing 240
 - routines, calling via procedure pointers 250
- overlay manager
 - initializing 246
 - protected-mode 197
- Overlay unit 144, 240
 - procedures and functions 243
- overlays 239-252
 - assembly language routines and 251
 - BP register and 251
 - buffer 241
 - loading and freeing up 242
 - optimization algorithm 243
 - probationary area 243
 - size
 - default 259
 - increasing with OvrSetBuf 259
 - cautions 250
 - debugging 251
 - defined 239
 - in .EXE files 254
 - installing a read function 252

- loading
 - into expanded memory 247
 - into memory 240
- manager, initializing 249
- using 239-254
- overriding
 - dynamic methods 38
 - inherited methods 38
- overview of Borland Pascal language 5
- .OVR files 239
- OvrClearBuf procedure 244
- OvrCodeList variable 152
- OvrDebugPtr variable 152
- OvrDosHandle variable 152
- OvrEmsHandle variable 152
- OvrFileMode variable 244
- OvrGetBuf function 244
- OvrGetRetry function 243, 244
- OvrHeapEnd variable 152
- OvrHeapOrg variable 152
- OvrHeapPtr variable 152
- OvrHeapSize variable 152
- OvrInit procedure 244
- OvrInitEMS procedure 244, 247
- OvrLoadCount variable 244
- OvrLoadList variable 152
- OvrReadBuf variable 244, 252
- OvrResult variable 244
- OvrSeg variable 252
- OvrSetBuf procedure 244, 247
 - increasing size of overlay buffer with 259
- OvrSetRetry procedure 243, 244
- OvrTrapCount variable 244

P

- \$P compiler directive 113
- Pack procedure 148
- packed
 - reserved word 30
 - string type 31
 - strings, comparing 74
- PackTime procedure 186, 189
- palette manipulation routines 225
- ParamCount function 151
- parameters 109-116
 - actual 82
 - command-line 151

- constant 110
- floating-point 292
- formal 82, 109
- open 113
 - array 114
 - string 113
- passing 83, 291-293
- Self 105
 - defined 296
- type compatibility 111
- types 109
- untyped 111
- value 110, 292
- variable 111
- virtual method 298
- ParamStr function 151
- Pascal strings 216
- passing parameters 291-293
 - by reference 291
 - by value 291
- passing string variables of varying sizes 30
- PChar data type 44
- PERMANENT code segment attribute 267, 271
- physical address 193
- Pi function 148
- PieSlice procedure 235
- pointer (^) symbol 43, 56
- pointer and address functions 150
- Pointer data type 43
- pointer-type constants 63
- pointers
 - assignment compatibility of 40
 - comparing 74
 - to objects 40
 - types 43, 279
 - values 56
 - variables 56
- Port array 288
- ports, accessing 288
- PortW array 288
- Pos function 149
- pound (#) character 19
- precedence of operators 65, 69
- precision
 - of real-type values 28
 - rules of arithmetic 25
- Pred function 24, 149, 305

- predecessor of a value, returning 24
- PrefixSeg variable 137, 152, 153, 257
- PRELOAD code segment attribute 266, 270
- PrestoChangoSelector function 209
- printer devices 161
- Printer unit 144, 161
- printing
 - from a DOS program 161
 - from a Windows program 170
- private
 - component sections 36
 - directive 17
 - procedures and functions 121
- PRN 161
- probationary area, overlay buffer 243
- PROC directive 341
- procedural
 - types 44-46
 - in expressions 78-79
 - type compatibility of 46
 - variable typecasts and 58
 - values 45
- procedural-type constants 64
- procedure
 - call models 98
 - declaration syntax 97
 - declarations 97-103
 - assembler 101
 - external 101
 - forward 100
 - inline 102
 - near and far 98
 - headings 97
 - statements 82
- procedure and function declaration part 94
- procedures 6, 97, *See also* procedures and functions
 - date and time 189
 - procedures 186
 - directory-handling 191
 - Dispose, extended syntax 282, 298
 - constructor passed as parameter 106, 116
 - entry/exit code, built-in assembler 335
 - exit 136
 - external 183
 - far 294
 - file-handling 187, 190
 - flow control 147
 - graphics 233
 - interrupt 100
 - support 186, 190
 - miscellaneous 188, 191
 - near 294
 - nested 294
 - New
 - extended syntax 282
 - constructor passed as parameter 106, 116, 298
 - used as function 117
 - ordinal 149
 - OvrSetRetry 243
 - parameters, built-in assembler 334
 - pointers, calling overlaid routines 250
 - process-handling 187
 - stack frame, built-in assembler 335
 - standard 147
 - string 149
- procedures and functions *See also* procedures; functions
 - exporting 132-134
 - importing 128, 128-132
 - nested 46
 - written in assembler 339
 - call model 339
- process-handling procedures
 - Dos unit 187
- program
 - block 5
 - comments 20
 - defined 5
 - heading 5, 119
 - lines, maximum length of 20
 - parameters 119
 - syntax 119
 - termination 301
- Program Segment Prefix (PSP) 257
- protected mode 193-213
 - addressing 194
 - application, writing a 197-210
 - benefits of 193
 - Borland DOS extensions 196
 - different than real mode 193
 - heap management in 200
 - memory manager 197, 204-207

- module management 207
- overlay manager 197
- resource management 208
- safe programming practices in 197
- Ptr function 43, 56, 150, 305
- public
 - component sections 36
 - directive 17
 - procedures and functions 121
- PUBLIC directives 339
- PutImage procedure 223, 228, 235
- PutPixel procedure 228, 235

Q

- qualified
 - identifiers 17
 - method
 - activating a 41
 - designator 41
 - identifiers 37, 56, 75, 104
- qualifier syntax 54

R

- \$R compiler directive 39, 282
 - virtual method checking 282
- Random function 151
- Randomize procedure 151
- RandSeed variable 152, 153
- range
 - checking 220
 - compile-time 309
 - finding highest value in 24
 - finding lowest value in 24
 - of real-type values 28
- read-only file access 159
- Read procedure
 - for values not of type Char 157
 - text files 156, 157
- ReadBuf function 169
- reading syntax diagrams 14
- ReadKey function 164, 165, 168, 169
- Readln procedure 156
- ReadOnly constant 188
- real
 - data types 28
 - mode addressing 194

- numbers 28, 175, 277
- Real data type 28
- real-type operations
 - 80x87 floating type 28
 - software floating point 28
- RealModeRegs variable 153
- record
 - scope 95
 - types 32
- record-type constant syntax 62
- records 32, 56, 62, 280
 - fields 56
 - variant part 32
- Rectangle procedure 235
- recursive loop *See* recursive loop
- redeclaration of variables 51
- redirection 162
- reentrant code 303, 304
- register-saving conventions 301
- RegisterBGIdriver function 224, 231, 235, 250
- RegisterBGIfont function 231, 235, 250
- registers
 - and inline statements 345
 - AX 293, 346
 - BP 301, 303
 - overlays and 251
 - built-in assembler 325, 329
 - BX 293, 303
 - CS 303
 - CX 303
 - DI 303
 - DS 301, 303
 - DX 293, 303
 - ES 303
 - segment
 - loading invalid values into 198
 - using to store temporary variables 199
 - SI 303
 - SP 301
 - SS 301
 - use, built-in assembler 316
 - using 293, 301, 303
- Registers type 189
- relational operators 73-75
- Release procedure 260, 268, 273
- relocation expressions, built-in assembler 330
- RemoveDir procedure 191

- Rename procedure 156
- repeat statement syntax 87
- repetitive statement syntax 86
- reserved words 16
 - built-in assembler 321
 - defined 16
 - external 343
 - how identified in manuals 16
 - list of 16
- Reset procedure 156
- resident option in exports clause 134
- resource management, protected-mode 208
- resources, linking into an application 208
- RestoreCrtMode procedure 224, 235
- RET instruction, built-in assembler 318
- RETF instruction, built-in assembler 318
- RETN instruction, built-in assembler 318
- return character, defined 15
- returning
 - Char values 26
 - the ordinality of a value 24
 - the predecessor of a value 24
 - the successor of a value 24
- Rewrite procedure 156
- right
 - brace special symbol 16
 - bracket special symbol 16
- Rmdir procedure 156
- Round function 148, 305
- round-off errors, minimizing 178
- RTM environment variable 212
- RTM.EXE file 196, 211
- rules
 - governing boolean variables 26
 - of inheritance 34
 - of scope 95-96
- run-time
 - errors 301, *See also the Programmer's Reference*
 - in a DLL 138
 - libraries 143-146
 - manager 196
 - controlling the memory used by 212
- RunError procedure 148
- running a DOS protected-mode application 211

S

- \$S compiler directive 52, 136

- SaveIntXX variables 152, 153
- scale factor syntax diagram 18
- scope
 - block 95
 - in object types 36
 - object 96
 - record 95
 - rules of 95, 95-96
 - type identifiers 23
 - unit 96
- screen output operations 162
- ScreenSize variable 167, 169
- ScrollTo procedure 169
- SearchRec type 189
- Sector procedure 235
- Seek procedure 156, 157
- SeekEof function 156
- SeekEoln function 156
- segment
 - arithmetic 198
 - attributes 266, 270
 - code 195
 - data 195
 - limit of 199
 - sub-allocator 268, 274
- segments 339
- SegXXXX variables 152, 154, 200
- SelectorInc variable 152, 153
 - using to deal with huge memory blocks 201
- selectors
 - aliased 203
 - defined 193
 - managing 209
 - predefined 200
- Self parameter 40, 41, 105, 298
 - defined 296
- separating tokens 15
- separators, defined 15
- Seq function 150
- set *See also* sets
 - constructors 66
 - syntax 76
 - membership testing 75
 - operators 72
 - types 42, 279
- set-type constants 63

- SetActivePage procedure 235
- SetAllPalette procedure 235, 237
- SetAspectRatio procedure 235
- SetBkColor procedure 235
- SetCBreak procedure 188, 191
- SetColor procedure 235
- SetCurDir procedure 191
- SetDate procedure 186, 189
- SetFAttr procedure 187, 190
- SetFillPattern procedure 227, 235
- SetFillStyle procedure 227, 235
- SetFTime procedure 186, 189
- SetGraphBufSize procedure 231, 235
- SetGraphMode procedure 224, 235
- SetIntVec procedure 186, 190
- SetLineStyle procedure 227, 235
- SetPalette procedure 235
- SetPrnFont function 170, 172
- SetRGBPalette procedure 225, 235
 - IBM 8514 and 225
- sets *See also set*
 - comparing 74
 - small 308
- SetSelectorBase function 209
- SetSelectorLimit function 209
- SetTextBuf procedure 156
- SetTextJustify procedure 227, 235
- SetTextStyle procedure 227, 236
- SetTime procedure 186, 190
- SetUserCharSize procedure 227, 236
- SetVerify procedure 188, 191
- SetViewPort procedure 236
- SetVisualPage procedure 236
- SetWriteMode procedure 236
- Shift instructions faster than multiply or divide 309
- shl operator 70
- short-circuit Boolean evaluation 70, 306
- Shortint data type 25
- shr operator 70
- SI register 303
- signed number syntax diagram 18
- significand 277
- simple
 - expression syntax 67
 - statement syntax 81
 - types 23-29
 - comparing 73
- simple-type constants 59
- Sin function 149
- single character special symbols 16
- Single data type 28, 177, 277
- size
 - of a given string, finding 30
 - of structured types, maximum 30
- SizeOf function 114, 151
- SizeOfResource function 208
- small sets 308
- smart linking 310
- software
 - floating-point
 - model 28
 - restrictions 28
 - interrupts 303
- sound operations
 - NoSound 165
 - Sound 165
- Sound procedure 165
- SP register 301
- space characters 15
- spawning applications 212
- special symbols
 - built-in assembler 325
 - character pairs listed 16
 - single characters listed 16
- SPtr function 150
- Sqr function 149
- Sqrt function 149
- SS register 301
- SSeg function 150
- stack 272
 - 80x87 180
 - changing size of 272
 - checking 136
 - frame, built-in assembler use of 335
 - overflow 52
 - passing parameters and 291
 - segment 52
 - DLLs and 138
- StackLimit variable 152
- standard
 - directives 17
 - functions 147

- procedure or function used as a procedural value 46
 - procedures 147
 - units, list of 143
- statement part syntax 94
- statements 81, 81-92
 - assignment 82
 - case 85
 - compound 84
 - conditional 84
 - for 88
 - goto 83
 - if 84
 - procedure 82
 - repeat 87
 - repetitive 86
 - simple 81
 - structured 83
 - while 87
 - with 90
- static
 - data area 272
 - importing 131
 - linking 127
 - methods 37
- stopping a Windows program print job 171
- storing
 - null-terminated strings 32
 - overlaid code 259
- Str procedure 149
- StrCat function 216
- StrComp function 216
- StrCopy function 216
- StrDispose function 216
- StrECopy function 216
- StrEnd function 216
- StrIComp function 216
- string *See also* strings
 - constants, built-in assembler 324
 - functions 149
 - literals, assigning to PChar 217
 - operator 71
 - procedures 149
 - type
 - default size 29
 - ordering between two values 29
 - packed 31
 - typed constants 60
 - types defined 29
 - variables 55
 - passing 30
- strings *See also* string
 - character 19
 - length of 20
 - comparing 74
 - concatenating 71
 - converting 216
 - embedding control characters in 19
 - length byte 279
 - maximum length 279
 - null 19, 29
 - null-terminated 32, 145, 215-222
 - Pascal 216
 - types 279
- Strings unit 145, 215
 - functions in 215
- StrLCat function 216
- StrLComp function 216
- StrLCopy function 216
- StrLen function 216
- StrLComp function 216
- StrLower function 216
- StrMove function 216
- StrNew function 216
- stroked fonts 223, 226
- StrPas function 216
- StrPCopy function 216
- StrPos function 216
- StrRScan function 216
- StrScan function 217
- structure member selector operator 328
- structured
 - statement syntax 83
 - types 30, 30-43
- structured-type constants 60
- StrUpper function 217
- stub 211
- styles, graphics 227
- subrange type 27
- subroutine block 97
- Succ function 24, 149, 305
- successor of a value, returning 24
- Swap function 151, 305
- SwapVectors procedure 187

- symbols 15
 - built-in assembler 325-328
 - invalid, built-in assembler 326
 - list of special 16
 - reserved, built-in assembler 325
 - scope access, built-in assembler 328
 - special, built-in assembler 325
- syntax diagrams, reading 14
- SysFile constant 188
- System unit 119, 144, 181
 - floating-point routines 177

T

- \$T compiler directive 48, 75
- tag field
 - identifier 33
- tag field (records) 33
- task header 272
- TDateTime type 192
- term syntax 67
- terminating a program 301
- Test8086 variable 152, 153, 154
- Test8087 variable 152, 154, 183
- testing
 - keyboard status 164
 - set membership 75
- text 226
 - files 157
 - devices 161
 - drivers 172
- text-color constants 165
- Text type 157
- TextAttr variable 166
- TextBackground procedure 165
- TextColor procedure 165
- TextHeight function 236
- TextMode procedure 165
- TextRec type 189
- TextWidth function 236
- TFileRec 286
 - type 192
- TitlePrn procedure 172
- tokens 15
 - categories of 15
 - defined 10, 15
 - examples of 10
 - separating 15

- TPP.TPL (DOS protected-mode run-time library) 143
- TPW.TPL (Windows run-time library) 143
- TrackCursor procedure 169
- transfer functions 148
- trapping interrupts 303
- TRegisters type 192
- True predefined constant identifier 26
- Trunc function 148, 305
- Truncate procedure 156
- TSearchRec type 192
- TTextRec records 172, 287
- Turbo3 unit 145
- Turbo Assembler 340
 - 80x87 emulation and 183
- TURBO.TPL (DOS real-mode run-time library) 143
- type *See also* types
 - declaration 23
 - declaration part syntax 94
 - defined 10
 - identifier 23
- type checking, built in assembler 330
- typecasting integer-type values 25
- typecasts, value 77
- typed
 - constant
 - defined 11
 - syntax 59
 - files 286
- TypeOf function 151
- types 23-50
 - array 30, 280
 - Boolean 276
 - boolean 25
 - Byte 25
 - ByteBool 276
 - Char 26, 276
 - Comp 24, 177
 - compatibility 47, 48
 - declaration part 49
 - Dos unit 189
 - Double 24, 177
 - enumerated 26, 276
 - Extended 24, 177
 - file 42
 - floating-point 28, 177, 277

- Comp 278
- comparing values of 179
- Double 278
- Extended 278
- Single 277
- Graph unit 237
- host 27
- identical 47
- identity 47
- Integer 25, 276
- integer
 - converting through typecasting 25
 - format of 25
 - range of 25
- LongBool 25, 276
- Longint 25
- major classes 23
- object 34-42
 - declaring 35
- ordinal 24-28
 - characteristics of 24
 - predefined 25
 - user-defined 25
- packed string 31
- PChar 44
- Pointer 43, 279
- procedural 44, 44-46, 78
- Real 24
- real 28
- real numbers 277
- record 32, 280
- set 42, 279
- Shortint 25
- simple 23-29
- Single 24, 177
- string 29, 279
- structured 30-43
- subrange 27
- Text 157
- WinDos 192
- Word 25
- WordBool 25, 276

U

unary

- arithmetic operators 69
- operands 65

- unit syntax 120
- units 120-126
 - 80x87 coprocessor and 181
 - circular references 123
 - Crt 144, 162
 - defined 13
 - Dos 144, 185
 - Graph 145, 223
 - Graph3 145
 - heading 120
 - identifiers 17
 - implementation part 121
 - import 129
 - indirect references 122
 - initialization code 249
 - initialization part 122
 - interface part 121
 - nonoverlay 250
 - Overlay 144, 240
 - overlays and 241
 - Printer 144, 161
 - reasons to use 13
 - scope of 96
 - standard, list of 143
 - Strings 145, 215
 - System 144
 - Turbo3 145
 - uses clause 119
 - version number 123
 - Win31 145
 - WinAPI 146, 197
 - using the 203-210
 - WinCrt 144, 166
 - WinDos 144
 - Windows 3.1 API 146
 - WinPrn 144
 - WinProcs 145
 - WinTypes 145
- UnlockResource function 208
- UnlockSegment function 205
- Unpack procedure 148
- UnpackTime procedure 186, 190
- unsigned
 - constant syntax 66
 - integer syntax diagram 18
 - number syntax diagram 18
 - real syntax diagram 18

- untyped
 - files 159, 286
 - parameters 111
- UpCase function 151
- uses clause 119
 - purpose of 13

V

- Val procedure 149
- value
 - parameters 110, 292
 - typecast syntax 77
- var
 - declaration section 311
 - parameters 111, 292
 - built-in assembler and 327
- variable *See also* variables
 - declaration part syntax 94
 - declaration syntax 51
 - defined 10
 - parameters 111
 - reference
 - qualifiers 54
 - syntax 54
 - typecast syntax 57
 - typecasts 57
 - and procedural types 58
- variables 51-64
 - absolute 53
 - array 55
 - declarations 51
 - dynamic 43, 56, 259
 - FileMode 159
 - global 52
 - Graph unit 238
 - in System unit 151
 - initialized in assembler 339
 - initializing 59
 - local 52
 - parameters 292
 - pointer 56
 - record 56
 - references 54
 - string 55
- variant part of records 32
- VGA emulated modes 224
- video memory 162

- viewports 227
- virtual
 - directive 37
 - method
 - parameter 298
 - table 280, 281
 - field 280
 - methods 37
 - calling 296
 - error checking 282
 - initializing 38
- VolumeID constant 188

W

- WEP exported function 136
- wep_Free_DLL value for ExitCode 135
- wep_System_Exit value for ExitCode 135
- WhereX function 165, 169
- WhereY function 165, 169
- while statement syntax 87
- Win31 unit 145
- WIN87EM.DLL 176
- WinAPI unit 146, 197
 - using the 203-210
- WinCrt unit 144, 158, 166
 - editing keys in 168
 - using the 166
 - variables in 169
- WindMax variable 166
- WindMin variable 166
- WinDos unit 144
 - constants 191
 - date and time procedures 189
 - directory-handling procedures and functions 191
 - disk status functions 190
 - environment-handling functions 191
 - file-handling procedures and functions 190
 - interrupt support procedures 190
 - miscellaneous procedures and functions 191
- Window procedure 163, 165
- Windows 3.1 API units 146
- WindowSize variable 169
- WindowTitle variable 170
- WinOrg variable 169
- WinPrn unit 144, 172
- WinProcs unit 145

- WinTypes unit *145*
- with statement syntax *90*
- word alignment, automatic *309*
- Word data type *25*
- WordBool data type *25, 276*
- Write procedure *156*
 - for values not of type Char *157*
- WriteBuf procedure *169*
- WriteChar procedure *169*
- Writeln procedure *156*
 - 80x87 coprocessor and *181*

- writing
 - DLLs *132-136*
 - overlaid programs *240, 245*
 - a protected-mode application *197-210*

X

- \$X compiler directive *20, 32, 44, 54, 71*
- xor operator *70, 228*

Z

- zero-based character arrays *61, 217, 219*

BORLAND PASCAL WITH OBJECTS

B O R L A N D

Corporate Headquarters: 1800 Green Hills Road, P.O. Box 660001, Scotts Valley, CA 95067-0001, (408) 438-8400. Offices in: Australia, Belgium, Canada, Denmark, France, Germany, Hong Kong, Italy, Japan, Korea, Malaysia, Netherlands, New Zealand, Singapore, Spain, Sweden, Taiwan, and United Kingdom ■ Part #11MN-BPL03-70 ■ BOR 4684