
Low-Level Control Structures

Chapter Two

2.1 Chapter Overview

This chapter discusses “pure” assembly language control statements. The last section of this chapter discusses *hybrid* control structures that combine the features of HLA’s high level control statements with the 80x86 control instructions.

2.2 Low Level Control Structures

Until now, most of the control structures you’ve seen and have used in your programs have been very similar to the control structures found in high level languages like Pascal, C++, and Ada. While these control structures make learning assembly language easy they are not true assembly language statements. Instead, the HLA compiler translates these control structures into a sequence of “pure” machine instructions that achieve the same result as the high level control structures. This text uses the high level control structures to avoid your having to learn too much all at once. Now, however, it’s time to put aside these high level language control structures and learn how to write your programs in *real* assembly language, using low-level control structures.

2.3 Statement Labels

HLA low level control structures make extensive use of *labels* within your code. A low level control structure usually transfers control from one point in your program to another point in your program. You typically specify the destination of such a transfer using a statement label. A statement label consists of a valid (unique) HLA identifier and a colon, e.g.,

```
aLabel:
```

Of course, like procedure, variable, and constant identifiers, you should attempt to choose descriptive and meaningful names for your labels. The identifier “aLabel” is hardly descriptive or meaningful.

Statement labels have one important attribute that differentiates them from most other identifiers in HLA: you don’t have to declare a label before you use it. This is important, because low-level control structures must often transfer control to a label at some point later in the code, therefore the label may not be defined at the point you reference it.

You can do three things with labels: transfer control to a label via a jump (*goto*) instruction, call a label via the *CALL* instruction, and you can take the address of a label. There is very little else you can directly do with a label (of course, there is very little else you would want to do with a label, so this is hardly a restriction). The following program demonstrates two ways to take the address of a label in your program and print out the address (using the *LEA* instruction and using the “&” address-of operator):

```
program labelDemo;
#include( "stdlib.hhf" );

begin labelDemo;

    lbl1:

        lea( ebx, lbl1 );
        lea( eax, lbl2 );
        stdout.put( "&lbl1=$", ebx, " &lbl2=", eax, nl );
```

```

    lbl2:

end labelDemo;

```

Program 2.1 Displaying the Address of Statement Labels in a Program

HLA also allows you to initialize dword variables with the addresses of statement labels. However, there are some restrictions on labels that appear in the initialization portions of variable declarations. The most important restriction is that you must define the statement label at the same lex level as the variable declaration. That is, if you reference a statement label in the initialization section of a variable declaration appearing in the main program, the statement label must also be in the main program. Conversely, if you take the address of a statement label in a local variable declaration, that symbol must appear in the same procedure as the local variable. The following program demonstrates the use of statement labels in variable initialization:

```

program labelArrays;
#include( "stdlib.hhf" );

static
    labels:dword[2] := [ &lbl1, &lbl2 ];

    procedure hasLabels;
    static
        stmtLbls: dword[2] := [ &label1, &label2 ];

    begin hasLabels;

        label1:

            stdout.put
            (
                "stmtLbls[0]= $", stmtLbls[0], nl,
                "stmtLbls[1]= $", stmtLbls[4], nl
            );

        label2:

    end hasLabels;

begin labelArrays;

    hasLabels();
    lbl1:

        stdout.put( "labels[0]= $", labels[0], " labels[1]=", labels[4], nl );

    lbl2:

end labelArrays;

```

Program 2.2 Initializing DWORD Variables with the Address of Statement Labels

Once in a really great while, you'll need to refer to a label that is not within the current procedure. The need for this is sufficiently rare that this text will not describe all the details. However, you can look up the details on HLA's LABEL declaration section in the HLA documentation should the need to do this ever arise.

2.4 Unconditional Transfer of Control (JMP)

The JMP (jump) instruction unconditionally transfers control to another point in the program. There are three forms of this instruction: a direct jump, and two indirect jumps. These instructions take one of the following three forms:

```

jmp label;
jmp( reg32 );
jmp( mem32 );

```

For the first (direct) jump above, you normally specify the target address using a statement label (see the previous section for a discussion of statement labels). The statement label is usually on the same line as an executable machine instruction or appears by itself on a line preceding an executable machine instruction. The direct jump instruction is the most commonly used of these three forms. It is completely equivalent to a GOTO statement in a high level language¹. Example:

```

<< statements >>
  jmp laterInPgm;
  .
  .
  .
laterInPgm:
  << statements >>

```

The second form of the JMP instruction above, "jmp(reg₃₂);", is a *register indirect* jump instruction. This instruction transfers control to the instruction whose address appears in the specified 32-bit general purpose register. To use this form of the JMP instruction you must load the specified register with the address of some machine instruction prior to the execution of the JMP. You could use this instruction to implement a state machine (see "State Machines and Indirect Jumps" on page 784) by loading a register with the address of some label at various points throughout your program; then, arriving along different paths, a point in the program can determine what path it arrived upon by executing the indirect jump. The following short sample program demonstrates how you could use the JMP in this manner:

```

program regIndJmp;
#include( "stdlib.hhf" );

static
  i:int32;

begin regIndJmp;

  // Read an integer from the user and set EBX to
  // denote the success or failure of the input.

  try

    stdout.put( "Enter an integer value between 1 and 10: " );
    stdin.get( i );
    mov( i, eax );

```

1. Unlike high level languages, where your instructors usually forbid you to use GOTO statements, you will find that the use of the JMP instruction in assembly language is absolutely essential.

```

    if( eax in 1..10 ) then
        mov( &GoodInput, ebx );
    else
        mov( &valRange, ebx );
    endif;

exception( ex.ConversionError )

    mov( &convError, ebx );

exception( ex.ValueOutOfRange )

    mov( &valRange, ebx );

endtry;

// Okay, transfer control to the appropriate
// section of the program that deals with
// the input.

jmp( ebx );

valRange:
    stdout.put( "You entered a value outside the range 1..10" nl );
    jmp Done;

convError:
    stdout.put( "Your input contained illegal characters" nl );
    jmp Done;

GoodInput:
    stdout.put( "You entered the value ", i, nl );

Done:

end regIndJump;

```

Program 2.3 Using Register Indirect JMP Instructions

The third form of the JMP instruction is a memory indirect JMP. This form of the JMP instruction fetches a dword value from the specified memory location and transfers control to the instruction at the address specified by the contents of the memory location. This is similar to the register indirect JMP except the address appears in a memory location rather than in a register. The following program demonstrates a rather trivial use of this form of the JMP instruction:

```

program memIndJump;
#include( "stdlib.hhf" );

static
    LabelPtr:dword := &stmtLabel;

```

```

begin memIndJump;

    stdout.put( "Before the JMP instruction" nl );
    jmp( LabelPtr );

    stdout.put( "This should not execute" nl );

    stmtLabel:

        stdout.put( "After the LabelPtr label in the program" nl );

end memIndJump;

```

Program 2.4 Using Memory Indirect JMP Instructions

Warning: unlike the HLA high level control structures, the low-level JMP instructions can get you into a lot of trouble. In particular, if you do not initialize a register with the address of a valid instruction and you jump indirect through that register, the results are undefined (though this will usually cause a general protection fault). Similarly, if you do not initialize a dword variable with the address of a legal instruction, jumping indirect through that memory location will probably crash your program.

2.5 The Conditional Jump Instructions

Although the JMP instruction provides transfer of control, it does not allow you to make any serious decisions. The 80x86's conditional jump instructions handle this task. The conditional jump instructions are the basic tool for creating loops and other conditionally executable statements like the IF..ENDIF statement.

The conditional jumps test one or more flags in the flags register to see if they match some particular pattern (just like the SETcc instructions). If the flag settings match the instruction control transfers to the target location. If the match fails, the CPU ignores the conditional jump and execution continues with the next instruction. Some conditional jump instructions simply test the setting of the sign, carry, overflow, and zero flags. For example, after the execution of a SHL instruction, you could test the carry flag to determine if the SHL shifted a one out of the H.O. bit of its operand. Likewise, you could test the zero flag after a TEST instruction to see if any specified bits were one. Most of the time, however, you will probably execute a conditional jump after a CMP instruction. The CMP instruction sets the flags so that you can test for less than, greater than, equality, etc.

The conditional JMP instructions take the following form:

```
Jcc label;
```

The "cc" in Jcc indicates that you must substitute some character sequence that specifies the type of condition to test. These are the same characters the SETcc instruction uses. For example, "JS" stands for jump if the sign flag is set." A typical JS instruction looks like this

```
js ValueIsNegative;
```

In this example, the JS instruction transfers control to the *ValueIsNegative* statement label if the sign flag is currently set; control falls through to the next instruction following the JS instruction if the sign flag is clear.

Unlike the unconditional JMP instruction, the conditional jump instructions do not provide an indirect form. The only form they allow is a branch to a statement label in your program. Conditional jump instructions have a restriction that the target label must be within 32,768 bytes of the jump instruction. However, since this generally corresponds to somewhere between 8,000 and 32,000 machine instructions, it is unlikely you will ever encounter this restriction.

Note: Intel's documentation defines various synonyms or instruction aliases for many conditional jump instructions. The following tables list all the aliases for a particular instruction. These tables also list out the opposite branches. You'll soon see the purpose of the opposite branches.

Table 1: Jcc Instructions That Test Flags

Instruction	Description	Condition	Aliases	Opposite
JC	Jump if carry	Carry = 1	JB, JNAE	JNC
JNC	Jump if no carry	Carry = 0	JNB, JAE	JC
JZ	Jump if zero	Zero = 1	JE	JNZ
JNZ	Jump if not zero	Zero = 0	JNE	JZ
JS	Jump if sign	Sign = 1		JNS
JNS	Jump if no sign	Sign = 0		JS
JO	Jump if overflow	Ovrflw=1		JNO
JNO	Jump if no Ovrflw	Ovrflw=0		JO
JP	Jump if parity	Parity = 1	JPE	JNP
JPE	Jump if parity even	Parity = 1	JP	JPO
JNP	Jump if no parity	Parity = 0	JPO	JP
JPO	Jump if parity odd	Parity = 0	JNP	JPE

Table 2: Jcc Instructions for Unsigned Comparisons

Instruction	Description	Condition	Aliases	Opposites
JA	Jump if above (>)	Carry=0, Zero=0	JNBE	JNA
JNBE	Jump if not below or equal (not <=)	Carry=0, Zero=0	JA	JBE
JAE	Jump if above or equal (>=)	Carry = 0	JNC, JNB	JNAE
JNB	Jump if not below (not <)	Carry = 0	JNC, JAE	JB
JB	Jump if below (<)	Carry = 1	JC, JNAE	JNB
JNAE	Jump if not above or equal (not >=)	Carry = 1	JC, JB	JAE
JBE	Jump if below or equal (<=)	Carry = 1 or Zero = 1	JNA	JNBE
JNA	Jump if not above (not >)	Carry = 1 or Zero = 1	JBE	JA
JE	Jump if equal (=)	Zero = 1	JZ	JNE
JNE	Jump if not equal (≠)	Zero = 0	JNZ	JE

Table 3: Jcc Instructions for Signed Comparisons

Instruction	Description	Condition	Aliases	Opposite
JG	Jump if greater (>)	Sign = Ovrflw or Zero=0	JNLE	JNG
JNLE	Jump if not less than or equal (not <=)	Sign = Ovrflw or Zero=0	JG	JLE
JGE	Jump if greater than or equal (>=)	Sign = Ovrflw	JNL	JNGE
JNL	Jump if not less than (not <)	Sign = Ovrflw	JGE	JL
JL	Jump if less than (<)	Sign ≠ Ovrflw	JNGE	JNL
JNGE	Jump if not greater or equal (not >=)	Sign ≠ Ovrflw	JL	JGE
JLE	Jump if less than or equal (<=)	Sign ≠ Ovrflw or Zero = 1	JNG	JNLE
JNG	Jump if not greater than (not >)	Sign ≠ Ovrflw or Zero = 1	JLE	JG
JE	Jump if equal (=)	Zero = 1	JZ	JNE
JNE	Jump if not equal (≠)	Zero = 0	JNZ	JE

One brief comment about the “opposites” column is in order. In many instances you will need to be able to generate the opposite of a specific branch instructions (lots of examples of this appear throughout the remainder of this chapter). With only two exceptions, a very simple rule completely describes how to generate an opposite branch:

- If the second letter of the *Jcc* instruction is *not* an “n”, insert an “n” after the “j”. E.g., JE becomes JNE and JL becomes JNL.
- If the second letter of the *Jcc* instruction *is* an “n”, then remove that “n” from the instruction. E.g., JNG becomes JG and JNE becomes JE.

The two exceptions to this rule are JPE (jump if parity is even) and JPO (jump if parity is odd). These exceptions cause few problems because (a) you’ll hardly ever need to test the parity flag, and (b) you can use the aliases JP and JNP synonyms for JPE and JPO. The “N/No N” rule applies to JP and JNP.

Though you *know* that JGE is the opposite of JL, get in the habit of using JNL rather than JGE as the opposite jump instruction for JL. It’s too easy in an important situation to start thinking “greater is the opposite of less” and substitute JG instead. You can avoid this confusion by always using the “N/No N” rule.

The 80x86 conditional jump instructions give you the ability to split program flow into one of two paths depending upon some logical condition. Suppose you want to increment the AX register if BX is equal to CX. You can accomplish this with the following code:

```
cmp( bx, cx );
jne SkipStmts;
```



```

    inc( ax );
SkipStmts:

```

The trick is to use the *opposite* branch to skip over the instructions you want to execute if the condition is true. Always use the “opposite branch (N/no N)” rule given earlier to select the opposite branch.

You can also use the conditional jump instructions to synthesize loops. For example, the following code sequence reads a sequence of characters from the user and stores each character in successive elements of an array until the user presses the Enter key (carriage return):

```

    mov( 0, edi );
RdLnLoop:
    stdin.getc();           // Read a character into the AL register.
    mov( al, Input[ edi ] ); // Store away the character
    inc( edi );            // Move on to the next character
    cmp( al, stdio.cr );   // See if the user pressed Enter
    jne RdLnLoop;

```

For more information concerning the use of the conditional jumps to synthesize IF statements, loops, and other control structures, see “Implementing Common Control Structures in Assembly Language” on page 759.

Like the SETcc instructions, the conditional jump instructions come in two basic categories – those that test specific processor flags (e.g., JZ, JC, JNO) and those that test some condition (less than, greater than, etc.). When testing a condition, the conditional jump instructions almost always follow a CMP instruction. The CMP instruction sets the flags so you can use a JA, JAE, JB, JBE, JE, or JNE instruction to test for unsigned less than, less than or equal, equality, inequality, greater than, or greater than or equal. Simultaneously, the CMP instruction sets the flags so you can also do a signed comparison using the JL, JLE, JE, JNE, JG, and JGE instructions.

The conditional jump instructions only test flags, they do not affect any of the 80x86 flags.

2.6 “Medium-Level” Control Structures: JT and JF

HLA provides two special conditional jump instructions: JT (jump if true) and JF (jump if false). These instructions take the following syntax:

```

    jt( boolean_expression ) target_label;
    jf( boolean_expression ) target_label;

```

The *boolean_expression* is the standard HLA boolean expression allowed by IF..ENDIF and other HLA high level language statements. These instructions evaluate the boolean expression and jump to the specified label if the expression evaluates true (JT) or false (JF).

These are not real 80x86 instructions. HLA compiles them into a sequence of one or more 80x86 machine instructions that achieve the same result. In general, you should not use these two instructions in your main code; they offer few benefits over using an IF..ENDIF statement and they are no more readable than the pure assembly language sequences they compile into. HLA provides these “medium-level” instructions so that you may create your own high level control structures using macros (see the chapters on Macros, the HLA Run-Time Language, and Domain Specific Languages for more details).

2.7 Implementing Common Control Structures in Assembly Language

Since a primary goal of this chapter is to teach you how to use the low-level machine instructions to implement decisions, loops, and other control constructs, it would be wise to show you how to simulate

these high level statements using “pure” assembly language. The following sections provide this information.

2.8 Introduction to Decisions

In its most basic form, a decision is some sort of branch within the code that switches between two possible execution paths based on some condition. Normally (though not always), conditional instruction sequences are implemented with the conditional jump instructions. Conditional instructions correspond to the IF..THEN..ENDIF statement in HLA:

```
if( expression ) then
    << statements >>
endif;
```

Assembly language, as usual, offers much more flexibility when dealing with conditional statements. Consider the following C/C++ statement:

```
if( ( ( x < y ) && ( z > t ) ) || ( a != b ) )
    stmt1;
```

A “brute force” approach to converting this statement into assembly language might produce:

```
mov( x, eax );
cmp( eax, y );
setl( bl );          // Store X<Y in bl.
mov( z, eax );
cmp( eax, t );
setg( bh );          // Store Z > T in bh.
and( bh, bl );       // Put (X<Y) && (Z>T) into bl.
mov( a, eax );
cmp( eax, b );
setne( bh );         // Store A != B into bh.
or( bh, bl );        // Put (X<Y) && (Z>T) || (A!=B) into bl
je SkipStmt1;        // Branch if result is false (OR sets Z-Flag if false).
```

<Code for stmt1 goes here>

SkipStmt1:

As you can see, it takes a considerable number of conditional statements just to process the expression in the example above. This roughly corresponds to the (equivalent) C/C++ statements:

```
bl = x < y;
bh = z > t;
bl = bl && bh;
bh = a != b;
bl = bl || bh;
if( bl )
    stmt1;
```

Now compare this with the following “improved” code:

```
mov( a, eax );
cmp( eax, b );
jne DoStmt;
mov( x, eax );
cmp( eax, y );
jnl SkipStmt;
mov( z, eax );
cmp( eax, t );
jng SkipStmt;
```

```

DoStmt:
    << Place code for Stmt1 here >>
SkipStmt:

```

Two things should be apparent from the code sequences above: first, a single conditional statement in C/C++ (or some other HLL) may require several conditional jumps in assembly language; second, organization of complex expressions in a conditional sequence can affect the efficiency of the code. Therefore, care should be exercised when dealing with conditional sequences in assembly language.

Conditional statements may be broken down into three basic categories: IF statements, SWITCH/CASE statements, and indirect jumps. The following sections will describe these program structures, how to use them, and how to write them in assembly language.

2.8.1 IF..THEN..ELSE Sequences

The most common conditional statement is the IF..THEN or IF..THEN..ELSE statement. These two statements take the form shown in Figure 2.1:

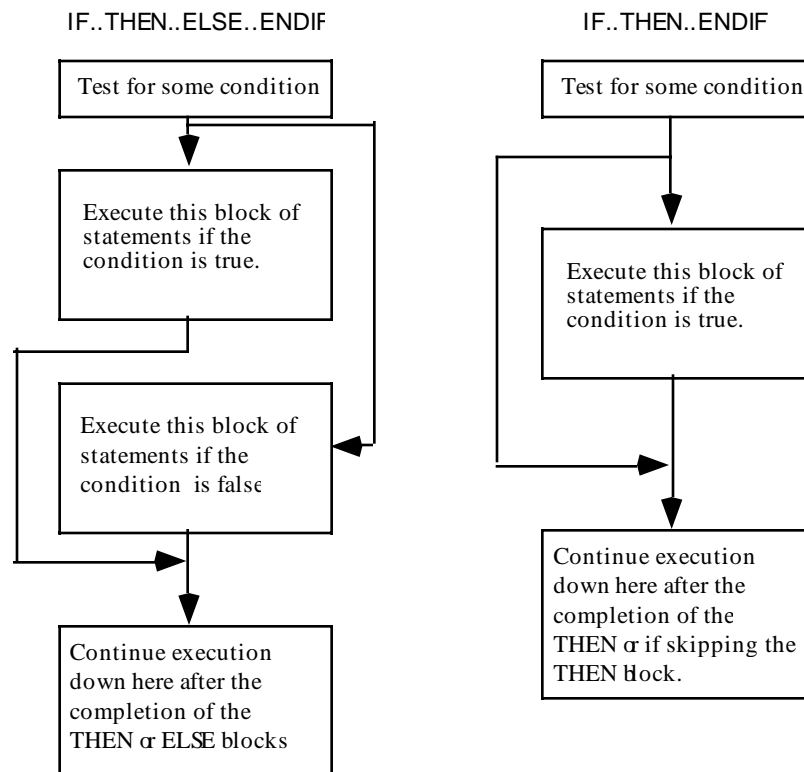


Figure 2.1 IF..THEN..ELSE..ENDIF and IF..ENDIF Statement Flow

The IF..ENDIF statement is just a special case of the IF..ELSE..ENDIF statement (with an empty ELSE block). Therefore, we'll only consider the more general IF..ELSE..ENDIF form. The basic implementation of an IF..THEN..ELSE statement in 80x86 assembly language looks something like this:

```

{Sequence of statements to test some condition}
    Jcc ElseCode
{Sequence of statements corresponding to the THEN block}

    jmp EndOfIF

```

```

ElseCode:
    {Sequence of statements corresponding to the ELSE block}
EndOfIF:

```

Note: *Jcc* represents some conditional jump instruction.

For example, to convert the C/C++ statement:

```

if( a == b )
    c = d;
else
    b = b + 1;

```

to assembly language, you could use the following 80x86 code:

```

    mov( a, eax );
    cmp( eax, b );
    jne ElsePart;
    mov( d, c );
    jmp EndOfIf;

ElseBlk:
    inc( b );

EndOfIf:

```

For simple expressions like “(a == b)” generating the proper code for an IF..ELSE..ENDIF statement is almost trivial. Should the expression become more complex, the associated assembly language code complexity increases as well. Consider the following C/C++ IF statement presented earlier:

```

if( ( ( x > y ) && ( z < t ) ) || ( a != b ) )
    c = d;

```

When processing complex IF statements such as this one, you’ll find the conversion task easier if you break this IF statement into a sequence of three different IF statements as follows:

```

if( a != b ) C = D;
else if( x > y )
    if( z < t )
        C = D;

```

This conversion comes from the following C/C++ equivalences:

```

if( expr1 && expr2 ) stmt;

```

is equivalent to

```

if( expr1 ) if( expr2 ) stmt;

```

and

```

if( expr1 || expr2 ) stmt;

```

is equivalent to

```

if( expr1 ) stmt;
else if( expr2 ) stmt;

```

In assembly language, the former IF statement becomes:

```

// if( ( ( x > y ) && ( z < t ) ) || ( a != b ) )
//     c = d;

```

```

        mov( a, eax );
        cmp( eax, b );
        jne DoIF;
        mov( x, eax );
        cmp( eax, y );
        jng EndOfIF;
        mov( z, eax );
        cmp( eax, t );
        jnl EndOfIf;
DoIf:
        mov( d, c );
EndOfIF:

```

As you can probably tell, the code necessary to test a condition can easily become more complex than the statements appearing in the ELSE and THEN blocks. Although it seems somewhat paradoxical that it may take more effort to test a condition than to act upon the results of that condition, it happens all the time. Therefore, you should be prepared for this situation.

Probably the biggest problem with the implementation of complex conditional statements in assembly language is trying to figure out what you've done after you've written the code. Probably the biggest advantage high level languages offer over assembly language is that expressions are much easier to read and comprehend in a high level language. This is one of the primary reasons HLA supports high level language control structures. The high level language version is self-documenting whereas assembly language tends to hide the true nature of the code. Therefore, well-written comments are an essential ingredient to assembly language implementations of if..then..else statements. An elegant implementation of the example above is:

```

// IF ((X > Y) && (Z < T)) OR (A != B) C = D;
// Implemented as:
// IF (A != B) THEN GOTO DoIF;

        mov( a, eax );
        cmp( eax, b );
        jne DoIF;

// if NOT (X > Y) THEN GOTO EndOfIF;

        mov( x, eax );
        cmp( eax, y );
        jng EndOfIF;

// IF NOT (Z < T) THEN GOTO EndOfIF ;

        mov( z, eax );
        cmp( eax, t );
        jnl EndOfIf;

// THEN Block:
DoIf:
        mov( d, c );

// End of IF statement
EndOfIF:

// if( (( x > y ) && ( z < t )) || ( a != b ) ) c = d;
// Test the boolean expression:

        mov( a, eax );
        cmp( eax, b );
        jne DoIF;
        mov( x, eax );

```

```

    cmp( eax, y );
    jng EndOfIF;
    mov( z, eax );
    cmp( eax, t );
    jnl EndOfIf;

; THEN Block:

DoIf:
    mov( d, c );

; End of IF statement

EndOfIF:

```

However, as your IF statements become complex, the density (and quality) of your comments become more and more important.

2.8.2 Translating HLA IF Statements into Pure Assembly Language

Translating HLA IF statements into pure assembly language is very easy. The boolean expressions that the HLA IF supports were specifically chosen to expand into a few simple machine instructions. The following paragraphs discuss the conversion of each supported boolean expression into pure machine code.

if(*flag_specification*) then <<stmts>> endif;

This form is, perhaps, the easiest HLA IF statement to convert. To execute the code immediately following the THEN keyword if a particular flag is set (or clear), all you need do is skip over the code if the flag is clear (set). This requires only a single conditional jump instruction for implementation as the following examples demonstrate:

```

// if( @c ) then inc( eax ); endif;

    jnc SkipTheInc;

    inc( eax );

SkipTheInc:

// if( @ns ) then neg( eax ); endif;

    js SkipTheNeg;

    neg( eax );

SkipTheNeg:

```

if(*register*) then <<stmts>> endif;

This form of the IF statement uses the TEST instruction to check the specified register for zero. If the register contains zero (false), then the program jumps around the statements after the THEN clause with a JZ instruction. Converting this statement to assembly language requires a TEST instruction and a JZ instruction as the following examples demonstrate:

```

// if( eax ) then mov( false, eax ); endif;

    test( eax, eax );
    jz DontSetFalse;

    mov( false, eax );

```

```

DontSetFalse:

// if( al ) then mov( bl, cl ); endif;

test( al, al );
jz noMove;

    mov( bl, cl );

noMove:

```

if(!register) then <<stmts>> endif;

This form of the IF statement uses the TEST instruction to check the specified register to see if it is zero. If the register is not zero (true), then the program jumps around the statements after the THEN clause with a JNZ instruction. Converting this statement to assembly language requires a TEST instruction and a JNZ instruction in a manner identical to the previous examples.

if(boolean_variable) then <<stmts>> endif;

This form of the IF statement compares the boolean variable against zero (false) and branches around the statements if the variable does contain false. HLA implements this statement by using the CMP instruction to compare the boolean variable to zero and then it uses a JZ (JE) instruction to jump around the statements if the variable is false. The following example demonstrates the conversion:

```

// if( bool ) then mov( 0, al ); endif;

cmp( bool, false );
je SkipZeroAL;

    mov( 0, al );

SkipZeroAL:

```

if(!boolean_variable) then <<stmts>> endif;

This form of the IF statement compares the boolean variable against zero (false) and branches around the statements if the variable contains true (i.e., the opposite condition of the previous example). HLA implements this statement by using the CMP instruction to compare the boolean variable to zero and then it uses a JNZ (JNE) instruction to jump around the statements if the variable contains true. The following example demonstrates the conversion:

```

// if( !bool ) then mov( 0, al ); endif;

cmp( bool, false );
jne SkipZeroAL;

    mov( 0, al );

SkipZeroAL:

```

if(mem_reg relop mem_reg_const) then <<stmts>> endif;

HLA translates this form of the IF statement into a CMP instruction and a conditional jump that skips over the statements on the opposite condition specified by the *relop* operator. The following table lists the correspondence between operators and conditional jump instructions:

Table 4: IF Statement Conditional Jump Instructions

Relop	Conditional jump instruction if both operands are unsigned	Conditional jump instruction if either operand is signed
= or ==	JNE	JNE
<> or !=	JE	JE
<	JNB	JNL
<=	JNBE	JNLE
>	JNA	JNG
>=	JNAE	JNGE

Here are a few examples of IF statements translated into pure assembly language that use expressions involving relational operators:

```
// if( al == ch ) then inc( cl ); endif;

    cmp( al, ch );
    jne SkipIncCL;

    inc( cl );

SkipIncCL:

// if( ch >= 'a' ) then and( $5f, ch ); endif;

    cmp( ch, 'a' );
    jnae NotLowerCase

    and( $5f, ch );

NotLowerCase:

// if( (type int32 eax) < -5 ) then mov( -5, eax ); endif;

    cmp( eax, -5 );
    jnl DontClipEAX;

    mov( -5, eax );

DontClipEAX:

// if( si <> di ) then inc( si ); endif;

    cmp( si, di );
    je DontIncSI;

    inc( si );

DontIncSI:
```


if(*reg/mem* in *LowConst..HiConst*) then <<stmts>> endif;

HLA translates this IF statement into a pair of CMP instructions and a pair of conditional jump instructions. It compares the register or memory location against the lower valued constant and jumps if less than (below) past the statements after the THEN clause. If the register or memory location's value is greater than or equal to *LowConst*, the code falls through to the second CMP/conditional jump pair that compares the register or memory location against the higher constant. If the value is greater than (above) this constant, a conditional jump instruction skips the statements in the THEN clause. Example:

```
// if( eax in 1000..125_000 ) then sub( 1000, eax ); endif;

    cmp( eax, 1000 );
    jb DontSub1000;
    cmp( eax, 125_000 );
    ja DontSub1000;

    sub( 1000, eax );

DontSub1000:

// if( i32 in -5..5 ) then add( 5, i32 ); endif;

    cmp( i32, -5 );
    jl NoAdd5;
    cmp( i32, 5 );
    jg NoAdd5;

    add( 5, i32 );

NoAdd5:
```

if(*reg/mem* not in *LowConst..HiConst*) then <<stmts>> endif;

This form of the HLA IF statement tests a register or memory location to see if its value is outside a specified range. The implementation is very similar to the code above exception you branch to the THEN clause if the value is less than the *LowConst* value or greater than the *HiConst* value and you branch over the code in the THEN clause if the value is within the range specified by the two constants. The following examples demonstrate how to do this conversion:

```
// if( eax not in 1000..125_000 ) then add( 1000, eax ); endif;

    cmp( eax, 1000 );
    jb Add1000;
    cmp( eax, 125_000 );
    jbe SkipAdd1000;

    Add1000:
    add( 1000, eax );

SkipAdd1000:

// if( i32 not in -5..5 ) then mov( 0, i32 ); endif;

    cmp( i32, -5 );
    jl Zeroi32;
    cmp( i32, 5 );
    jle SkipZero;
```

```

Zeroi32:
mov( 0, i32 );

SkipZero:

```

if(*reg₈* in *CSetVar/CSetConst*) then <<stmts>> endif;

This statement checks to see if the character in the specified eight-bit register is a member of the specified character set. HLA emits code that is similar to the following for instructions of this form:

```

movzx( reg8, eax );
bt( eax, CsetVar/CsetConst );
jnc SkipPastStmts;

<< stmts >>

SkipPastStmts:

```

This example modifies the EAX register (the code HLA generates does not, because it pushes and pops the register it uses). You can easily swap another register for EAX if you've got a value in EAX you need to preserve. In the worst case, if no registers are available, you can push EAX, execute the MOVZX and BT instructions, and then pop EAX's value from the stack. The following are some actual examples:

```

// if( al in { 'a'..'z' } ) then or( $20, al ); endif;

movzx( al, eax );
bt( eax, { 'a'..'z' } ); // See if we've got a lower case char.
jnc DontConvertCase;

or( $20, al ); // Convert to uppercase.

DontConvertCase:

// if( ch in { '0'..'9' } ) then and( $f, ch ); endif;

push( eax );
movzx( ch, eax );
bt( eax, { 'a'..'z' } ); // See if we've got a lower case char.
pop( eax );
jnc DontConvertNum;

and( $f, ch ); // Convert to binary form.

DontConvertNum:

```

2.8.3 Implementing Complex IF Statements Using Complete Boolean Evaluation

The previous section did not discuss how to translate boolean expressions involving conjunction (AND) or disjunction (OR) into assembly language. This section will begin that discussion. There are two different ways to convert complex boolean expressions involving conjunction and disjunction into assembly language: using complete boolean evaluation or short circuit evaluation. This section discusses complete boolean evaluation. The next section discusses short circuit boolean evaluation, which is the scheme that HLA uses when converting complex boolean expressions to assembly language.

Using complete boolean evaluation to evaluate a boolean expression for an IF statement is almost identical to converting arithmetic expressions into assembly language. Indeed, the previous volume covers this conversion process (see "Logical (Boolean) Expressions" on page 604). About the only thing worth noting

about that process is that you do not need to store the ultimate boolean result in some variable; once the evaluation of the expression is complete you check to see if you have a false (zero) or true (one, or non-zero) result to determine whether to branch around the THEN portion of the IF statement. As you can see in the examples in the preceding sections, you can often use the fact that the last boolean instruction (AND/OR) sets the zero flag if the result is false and clears the zero flag if the result is true. This lets you avoid explicitly testing the result. Consider the following IF statement and its conversion to assembly language using complete boolean evaluation:

```

    if( (( x < y ) && ( z > t )) || ( a != b ) )
        Stmt1;

        mov( x, eax );
        cmp( eax, y );
        setl( bl );           // Store x<y in bl.
        mov( z, eax );
        cmp( eax, t );
        setg( bh );         // Store z > t in bh.
        and( bh, bl );      // Put (x<y) && (z>t) into bl.
        mov( a, eax );
        cmp( eax, b );
        setne( bh );        // Store a != b into bh.
        or( bh, bl );       // Put (x<y) && (z>t) || (a != b) into bl
        je SkipStmt1;       // Branch if result is false (OR sets Z-Flag if false).

    << Code for Stmt1 goes here >>

SkipStmt1:

```

This code computes a boolean value in the BL register and then, at the end of the computation, tests this resulting value to see if it contains true or false. If the result is false, this sequence skips over the code associated with *Stmt1*. The important thing to note in this example is that the program will execute each and every instruction that computes this boolean result (up to the JE instruction).

For more details on complete boolean evaluation, see “Logical (Boolean) Expressions” on page 604.

2.8.4 Short Circuit Boolean Evaluation

If you are willing to spend a little more effort studying a complex boolean expression, you can usually convert it to a much shorter and faster sequence of assembly language instructions using *short-circuit boolean evaluation*. Short-circuit boolean evaluation attempts to determine whether an expression is true or false by executing only a portion of the instructions that compute the complete expression. By executing only a portion of the instructions, the evaluation is often much faster. For this reason, plus the fact that short circuit boolean evaluation doesn’t require the use of any temporary registers, HLA uses short circuit evaluation when translating complex boolean expressions into assembly language.

To understand how short-circuit boolean evaluation works, consider the expression “A && B”. Once we determine that A is false, there is no need to evaluate B since there is no way the expression can be true. If A and B represent sub-expressions rather than simple variables, you can begin to see the savings that are possible with short-circuit boolean evaluation. As a concrete example, consider the sub-expression “((x<y) && (z>t))” from the previous section. Once you determine that x is not less than y, there is no need to check to see if z is greater than t since the expression will be false regardless of z and t’s values. The following code fragment shows how you can implement short-circuit boolean evaluation for this expression:

```

// if( (x<y) && (z>t) ) then ...

    mov( x, eax );
    cmp( eax, y );
    jnl TestFails;
    mov( z, eax );

```

```

cmp( eax, t );
jng TestFails;

    << Code for THEN clause of IF statement >>

TestFails:

```

Notice how the code skips any further testing once it determines that x is not less than y . Of course, if x is less than y , then the program has to test z to see if it is greater than t ; if not, the program skips over the THEN clause. Only if the program satisfies both conditions does the code fall through to the THEN clause.

For the logical OR operation the technique is similar. If the first sub-expression evaluates to true, then there is no need to test the second operand. Whatever the second operand's value is at that point, the full expression still evaluates to true. The following example demonstrates the use of short-circuit evaluation with disjunction (OR):

```

// if( ch < 'A' || ch > 'Z' ) then stdout.put( "Not an upper case char" ); endif;

cmp( ch, 'A' );
jb ItsNotUC
cmp( ch, 'Z' );
jna ItWasUC;

    ItsNotUC:
        stdout.put( "Not an upper case char" );

    ItWasUC:

```

Since the conjunction and disjunction operators are commutative, you can evaluate the left or right operand first if it is more convenient to do so. As one last example in this section, consider the full boolean expression from the previous section:

```

// if( ( ( x < y ) && ( z > t ) ) || ( a != b ) ) Stmt1;

mov( a, eax );
cmp( eax, b );
jne DoStmt1;
mov( x, eax );
cmp( eax, y );
jnl SkipStmt1;
mov( z, eax );
cmp( eax, t );
jng SkipStmt1;

    DoStmt1:
        << Code for Stmt1 goes here >>

    SkipStmt1:

```

Notice how the code in this example chose to evaluate “ $a \neq b$ ” first and the remaining sub-expression last. This is a common technique assembly language programmers use to write better code.

2.8.5 Short Circuit vs. Complete Boolean Evaluation

One fact about complete boolean evaluation is that every statement in the sequence will execute when evaluating the expression. Short-circuit boolean evaluation may not require the execution of every statement associated with the boolean expression. As you've seen in the previous two sections above, code based on short-circuit evaluation is usually shorter and faster². So it would seem that short-circuit evaluation is the technique of choice when converting complex boolean expressions to assembly language.

Sometimes, unfortunately, short-circuit boolean evaluation may not produce the correct result. In the presence of *side-effects* in an expression, short-circuit boolean evaluation will produce a different result than complete boolean evaluation. Consider the following C/C++ example:

```
if( ( x == y ) && ( ++z != 0 ) ) stmt;
```

Using complete boolean evaluation, you might generate the following code:

```
mov( x, eax );      // See if x == y
cmp( eax, y );
sete( bl );
inc( z );           // ++z
cmp( z, 0 );        // See if incremented z is zero.
setne( bh );
and( bh, bl );      // Test x == y && ++z != 0
jz SkipStmt;

<< code for stmt goes here >>
```

SkipStmt:

Using short-circuit boolean evaluation, you might generate the following code:

```
mov( x, eax );      // See if x == y
cmp( eax, y );
jne SkipStmt;
inc( z );           // ++z
cmp( z, 0 );        // See if incremented z is zero.
je SkipStmt;

<< code for stmt goes here >>
```

SkipStmt:

Notice a very subtle, but important difference between these two conversions: if it turns out that x is equal to y , then the first version above *still increments* z and compares it to zero before it executes the code associated with *stmt*; the short-circuit version, on the other hand skips over the code that increments z if it turns out that x is equal to y . Therefore, the behavior of these two code fragments is different with respect to what happens to z if x is equal to y . Neither implementation is particularly wrong, depending on the circumstances you may or may not want the code to increment z if x is equal to y . However, it is important that you realize that these two schemes produce different results so you can choose an appropriate implementation if the effect of this code on z matters to your program.

Many programs take advantage of short-circuit boolean evaluation and rely upon the fact that the program may not evaluate certain components of the expression. The following C/C++ code fragment demonstrates what is probably the most common example that requires short-circuit boolean evaluation:

```
if( Ptr != NULL && *Ptr == 'a' ) stmt;
```

If it turns out that Ptr is NULL in this IF statement, then the expression is false and there is no need to evaluate the remainder of the expression (and, therefore, code that uses short-circuit boolean evaluation will not evaluate the remainder of this expression). This statement relies upon the semantics of short-circuit boolean evaluation for correct operation. Were C/C++ to use complete boolean evaluation, and the variable Ptr contained NULL, then the second half of the expression would attempt to dereference a NULL pointer (which tends to crash most programs). Consider the translation of this statement using complete and short-circuit boolean evaluation:

```
// Complete boolean evaluation:
```

2. Note that this does not always mean that the program will run faster. Jumps (conditional or otherwise) are often very slow executing instructions. Sometimes it's faster to execute several instructions in a row rather than execute a few instructions that include a conditional jump.

```

mov( Ptr, eax );
test( eax, eax );    // Check to see if EAX is zero (NULL is zero).
setne( bl );
mov( [eax], al );    // Get *Ptr into AL.
cmp( al, 'a' );
sete( bh );
and( bh, bl );
jz SkipStmt;

<< code for stmt goes here >>

SkipStmt:

```

Notice in this example that if *Ptr* contains NULL (zero), then this program will attempt to access the data at location zero in memory via the “mov([eax], al);” instruction. Under most operating systems this will cause a memory access fault (general protection fault). Now consider the short-circuit boolean conversion:

```

// Short-circuit boolean evaluation

mov( Ptr, eax );    // See if Ptr contains NULL (zero) and
test( eax, eax );    // immediately skip past Stmt if this
jz SkipStmt;        // is the case

mov( [eax], al );    // If we get to this point, Ptr contains
cmp( al, 'a' );      // a non-NULL value, so see if it points
jne SkipStmt;        // at the character 'a'.

<< code for stmt goes here >>

SkipStmt:

```

As you can see in this example, the problem with dereferencing the NULL pointer doesn't exist. If *Ptr* contains NULL, this code skips over the statements that attempt to access the memory address *Ptr* contains.

2.8.6 Efficient Implementation of IF Statements in Assembly Language

Encoding IF statements efficiently in assembly language takes a bit more thought than simply choosing short-circuit evaluation over complete boolean evaluation. To write code that executes as quickly as possible in assembly language you must carefully analyze the situation and generate the code appropriately. The following paragraphs provide some suggestions you can apply to your programs to improve their performance.

Know your data!

A mistake programmers often make is the assumption that data is random. In reality, data is rarely random and if you know the types of values that your program commonly uses, you can use this knowledge to write better code. To see how, consider the following C/C++ statement:

```
if(( a == b ) && ( c < d )) ++i;
```

Since C/C++ uses short-circuit evaluation, this code will test to see if *a* is equal to *b*. If so, then it will test to see if *c* is less than *d*. If you expect *a* to be equal to *b* most of the time but don't expect *c* to be less than *d* most of the time, this statement will execute slower than it should. Consider the following HLA implementation of this code:

```

mov( a, eax );
cmp( eax, b );
jne DontIncI;

mov( c, eax );

```

```

cmp( eax, d );
jnl DontIncI;

    inc( i );

DontIncI:

```

As you can see in this code, if a is equal to b most of the time and c is not less than d most of the time, you will have to execute the first three instructions nearly every time in order to determine that the expression is false. Now consider the following implementation of the above C/C++ statement that takes advantage of this knowledge and the fact that the “&&” operator is commutative:

```

mov( c, eax );
cmp( eax, d );
jnl DontIncI;

mov( a, eax );
cmp( eax, b );
jne DontIncI;

    inc( i );

DontIncI:

```

In this example the code first checks to see if c is less than d . If most of the time c is less than d , then this code determines that it has to skip to the label *DontIncI* after executing only three instructions in the typical case (compared with six instructions in the previous example). This fact is much more obvious in assembly language than in a high level language; this is one of the main reasons that assembly programs are often faster than their high level language counterparts: optimizations are more obvious in assembly language than in a high level language. Of course, the key here is to understand the behavior of your data so you can make intelligent decisions such as the one above.

Rearranging Expressions

Even if your data is random (or you can't determine how the input values will affect your decisions), there may still be some benefit to rearranging the terms in your expressions. Some calculations take far longer to compute than others. For example, the DIV instruction is much slower than a simple CMP instruction. Therefore, if you have a statement like the following you may want to rearrange the expression so that the CMP comes first:

```
if( (x % 10 = 0 ) && (x != y ) ++x;
```

Converted to assembly code, this IF statement becomes:

```

mov( x, eax );           // Compute X % 10
cdq();                  // Must sign extend EAX -> EDX:EAX
imod( 10, edx:eax );    // Remember, remainder goes into EDX
test( edx, edx );       // See if EDX is zero.
jnz SkipIF

mov( x, eax );
cmp( eax, y );
je SkipIF

    inc( x );

SkipIF:

```

The IMOD instruction is very expensive (often 50-100 times slower than most of the other instructions in this example). Unless it is 50-100 times more likely than the remainder is zero rather than x is equal to y, it would be better to do the comparison first and the remainder calculation afterwards:

```

mov( x, eax );
cmp( eax, y );
je SkipIF

mov( x, eax );           // Compute X % 10
cdq();                 // Must sign extend EAX -> EDX:EAX
imod( 10, edx:eax );    // Remember, remainder goes into EDX
test( edx, edx );       // See if EDX is zero.
jnz SkipIF

    inc( x );

SkipIF:

```

Of course, in order to rearrange the expression in this manner, the code must not assume the use of short-circuit evaluation semantics (since the && and || operators are not commutative if the code must compute one subexpression before another).

Destructuring Your Code

Although there is a lot of good things to be said about structured programming techniques, there are some drawbacks to writing structured code. Specifically, structured code is sometimes less efficient than unstructured code. Most of the time this is tolerable because unstructured code is difficult to read and maintain; it is often acceptable to sacrifice some performance in exchange for maintainable code. In certain instances, however, you may need all the performance you can get. In those rare instances you might choose to compromise the readability of your code in order to gain some additional performance.

One classic way to do this is to use code movement to move code your program rarely uses out of the way of code that executes most of the time. For example, consider the following pseudo C/C++ statement:

```

if( See_If_an_Error_Has_Ocurred )
{
    << Statements to execute if no error >>
}
else
{
    << Error handling statements >>
}

```

In normal code, one does not expect errors to be frequent. Therefore, you would normally expect the THEN section of the above IF to execute far more often than the ELSE clause. The code above could translate into the following assembly code:

```

cmp( See_If_an_Error_Has_Ocurred, true );
je HandleTheError

    << Statements to execute if no error >>
    jmp EndOfIF;

HandleTheError:
    << Error handling statements >>

EndOfIf:

```

Notice that if the expression is false this code falls through to the normal statements and then jumps over the error handling statements. Instructions that transfer control from one point in your program to

another (e.g., JMP instructions) tend to be slow. It is much faster to execute a sequential set of instructions rather than jump all over the place in your program. Unfortunately, the code above doesn't allow this. One way to rectify this problem is to move the ELSE clause of the code somewhere else in your program. That is, you could rewrite the code as follows:

```

cmp( See_If_an_Error_Has_Occurred, true );
je HandleTheError

    << Statements to execute if no error >>

EndOfIf:

```

At some other point in your program (typically after a JMP instruction) you would insert the following code:

```

HandleTheError:
    << Error handling statements >>
    jmp EndOfIf;

```

Note that the program isn't any shorter. The JMP you removed from the original sequence winds up at the end of the ELSE clause. However, since the ELSE clause rarely executes, moving the JMP instruction from the THEN clause (which executes frequently) to the ELSE clause is a big performance win because the THEN clause executes using only straight-line code. This technique is surprisingly effective in many time-critical code segments.

There is a difference between writing *destructured* code and writing *unstructured* code. Unstructured code is written in an unstructured way to begin with. It is generally hard to read, difficult to maintain, and it often contains defects. Destructured code, on the other hand, starts out as structured code and you make a conscious decision to eliminate the structure in order to gain a small performance boost. Generally, you've already tested the code in its structured form before destructuring it. Therefore, destructured code is often easier to work with than unstructured code.

Calculation Rather than Branching

On many processors in the 80x86 family, branches are very expensive compared to many other instructions (perhaps not as bad as IMOD or IDIV, but typically an order of magnitude worse than instructions like ADD and SUB). For this reason it is sometimes better to execute more instructions in a sequence rather than fewer instructions that involve branching. For example, consider the simple assignment "EAX = abs(EAX);" Unfortunately, there is no 80x86 instruction that computes the absolute value of an integer value. The obvious way to handle this is with an instruction sequence like the following:

```

test( eax, eax );
jns ItsPositive;

    neg( eax );

ItsPositive:

```

However, as you can plainly see in this example, it uses a conditional jump to skip over the NEG instruction (that creates a positive value in EAX if EAX was negative). Now consider the following sequence that will also do the job:

```

// Set EDX to $FFFF_FFFF if EAX is negative, $0000_0000 if EAX is
// zero or positive:

    cdq();

// If EAX was negative, the following code inverts all the bits in EAX,
// otherwise it has no effect on EAX.

```

```

xor( edx, edx);

// If EAX was negative, the following code adds one to EAX, otherwise
// it doesn't modify EAX's value.

and( 1, edx );      // EDX = 0 or 1 (1 if EAX was negative).
add( edx, eax );

```

This code will invert all the bits in EAX and then add one to EAX if EAX was negative prior to the sequence (i.e., it takes the two's complement [negates] the value in EAX). If EAX was zero or positive, then this code does not change the value in EAX.

Note that this sequence takes four instructions rather than the three the previous example requires. However, since there are no transfer of control instructions in this sequence, it may execute faster on many CPUs in the 80x86 family.

2.8.7 SWITCH/CASE Statements

The HLA (Standard Library) SWITCH statement takes the following form :

```

switch( reg32 )
  case( const1 )
    <<stmts1>>

  case( const2 )
    <<stmts2>>
  .
  .
  .
  case( constn )
    <<stmtsn >>

  default      // Note that the default section is optional
    <<stmtsdefault >>

endswitch;

```

When this statement executes, it checks the value of register against the constants $const_1 \dots const_n$. If a match is found then the corresponding statements execute. HLA places a few restrictions on the SWITCH statement. First, the HLA SWITCH statement only allows a 32-bit register as the SWITCH expression. Second, all the constants appearing as CASE clauses must be unique. The reason for these restrictions will become clear in a moment.

Most introductory programming texts introduce the SWITCH/CASE statement by explaining it as a sequence of IF..THEN..ELSEIF statements. They might claim that the following two pieces of HLA code are equivalent:

```

switch( eax )
  case(0) stdout.put("I=0");
  case(1) stdout.put("I=1");
  case(2) stdout.put("I=2");
endswitch;

if( eax = 0 ) then
  stdout.put("I=0")
elseif( eax = 1 ) then
  stdout.put("I=1")
elseif( eax = 2 ) then
  stdout.put("I=2");
endif;

```

While semantically these two code segments may be the same, their implementation is usually different. Whereas the IF.THEN.ELSEIF chain does a comparison for each conditional statement in the sequence, the SWITCH statement normally uses an indirect jump to transfer control to any one of several statements with a single computation. Consider the two examples presented above, they could be written in assembly language with the following code:

```
// IF..THEN..ELSE form:

    mov( i, eax );
    test( eax, eax ); // Check for zero.
    jnz Not0;
        stdout.put( "I=0" );
        jmp EndCase;

Not0:
    cmp( eax, 1 );
    jne Not1;
        stdou.put( "I=1" );
        jmp EndCase;

Not1:
    cmp( eax, 2 );
    jne EndCase;
        stdout.put( "I=2" );
EndCase:

// Indirect Jump Version

readonly
    jmpTbl:dword[3] := [ &Stmt0, &Stmt1, &Stmt2 ];
    .
    .
    .
    mov( i, eax );
    jmp( jmpTbl[ eax*4 ] );

    Stmt0:
        stdout.put( "I=0" );
        jmp EndCase;

    Stmt1:
        stdout.put( "I=1" );
        jmp EndCase;

    Stmt2:
        stdout.put( "I=2" );

EndCase:
```

The implementation of the IF.THEN.ELSEIF version is fairly obvious and doesn't need much in the way of explanation. The indirect jump version, however, is probably quite mysterious to you; so let's consider how this particular implementation of the SWITCH statement works.

Remember that there are three common forms of the JMP instruction (see "Unconditional Transfer of Control (JMP)" on page 753). The standard unconditional JMP instruction, like the "jmp EndCase;" instructions in the previous examples, transfer control directly to the statement label specified as the JMP operand. The second form of the JMP instruction (i.e., "jmp Reg₃₂;") transfers control to the memory location specified by the address found in a 32-bit register. The third form of the JMP instruction, the one the example above uses, transfers control to the instruction specified by the contents of a dword memory location. As this example clearly illustrates, that memory location can use any addressing mode. You are not limited to the

displacement-only addressing mode. Now let's consider exactly how this second implementation of the SWITCH statement works.

To begin with, a SWITCH statement requires that you create an array of pointers with each element containing the address of a statement label in your code (those labels must be attached to the sequence of instructions to execute for each case in the SWITCH statement). In the example above, the *JmpTbl* array serves this purpose. Note that this code initializes *JmpTbl* with the address of the statement labels *Stmt0*, *Stmt1*, and *Stmt2*. The program places this array in the READONLY section because the program should never change these values during execution.

Warning: whenever you initialize an array with a set of address of statement labels as in this example, the declaration section in which you declare the array (e.g., READONLY in this case) must be in the same procedure that contains the statement labels³.

During the execution of this code sequence, the program loads the EAX register with *I*'s value. Then the program uses this value as an index into the *JmpTbl* array and transfers control to the four-byte address found at the specified location. For example, if EAX contains zero, the "jmp(JmpTbl[eax*4]);" instruction will fetch the dword at address *JmpTbl+0* (*eax*4=0*). Since the first double word in the table contains the address of *Stmt0*, the JMP instruction will transfer control to the first instruction following the *Stmt0* label. Likewise, if *I* (and therefore, EAX) contains one, then the indirect JMP instruction fetches the double word at offset four from the table and transfers control to the first instruction following the *Stmt1* label (since the address of *Stmt1* appears at offset four in the table). Finally, if *I/EAX* contains two, then this code fragment transfers control to the statements following the *Stmt2* label since it appears at offset eight in the *JmpTbl* table.

Two things should become readily apparent: the more (consecutive) cases you have, the more efficient the jump table implementation becomes (both in terms of space and speed) over the IF/ELSEIF form. Except for trivial cases, the SWITCH statement is almost always faster and usually by a large margin. As long as the CASE values are consecutive, the SWITCH statement version is usually smaller as well.

What happens if you need to include non-consecutive CASE labels or you cannot be sure that the SWITCH value doesn't go out of range? With the HLA SWITCH statement, such an occurrence will transfer control to the first statement after the ENDSWITCH clause. However, this doesn't happen in the example above. If variable *I* does not contain zero, one, or two, the result of executing the code above is undefined. For example, if *I* contains five when you execute the code in the previous example, the indirect JMP instruction will fetch the dword at offset 20 ($5*4$) in *JmpTbl* and transfer control to that address. Unfortunately, *JmpTbl* doesn't have six entries; so the program will wind up fetching the value of the third double word following *JmpTbl* and using that as the target address. This will often crash your program or transfer control to an unexpected location. Clearly this code does not behave like the HLA SWITCH statement, nor does it have desirable behavior.

The solution is to place a few instructions before the indirect JMP to verify that the SWITCH selection value is within some reasonable range. In the previous example, we'd probably want to verify that *I*'s value is in the range 0..2 before executing the JMP instruction. If *I*'s value is outside this range, the program should simply jump to the *EndCase* label (this corresponds to dropping down to the first statement after the ENDSWITCH clause). The following code provides this modification:

```
readonly
  JmpTbl:dword[3] := [ &Stmt0, &Stmt1, &Stmt2 ];
  .
  .
  .
  mov( i, eax );
  cmp( eax, 2 );           // Verify that I is in the range
  ja EndCase;             // 0..2 before the indirect JMP.
  jmp( JmpTbl[ eax*4 ] );
```

3. If the SWITCH statement appears in your main program, you must declare the array in the declaration section of your main program.

```

Stmt0:
    stdout.put( "I=0" );
    jmp EndCase;

Stmt1:
    stdout.put( "I=1" );
    jmp EndCase;

Stmt2:
    stdout.put( "I=2" );

EndCase:

```

Although the example above handles the problem of selection values being outside the range zero through two, it still suffers from a couple of severe restrictions:

- The cases must start with the value zero. That is, the minimum CASE constant has to be zero in this example.
- The case values must be contiguous; there cannot be any gaps between any two case values.

Solving the first problem is easy and you deal with it in two steps. First, you must compare the case selection value against a lower and upper bounds before determining if the case value is legal, e.g.,

```

// SWITCH statement specifying cases 5, 6, and 7:
// WARNING: this code does *NOT* work. Keep reading to find out why.

```

```

mov( i, eax );
cmp( eax, 5 );
jb EndCase
cmp( eax, 7 );           // Verify that I is in the range
ja EndCase;             // 5..7 before the indirect JMP.
jmp( JumpTbl[ eax*4 ] );

Stmt5:
    stdout.put( "I=5" );
    jmp EndCase;

Stmt6:
    stdout.put( "I=6" );
    jmp EndCase;

Stmt7:
    stdout.put( "I=7" );

EndCase:

```

As you can see, this code adds a pair of extra instructions, `CMP` and `JB`, to test the selection value to ensure it is in the range five through seven. If not, control drops down to the `EndCase` label, otherwise control transfers via the indirect `JMP` instruction. Unfortunately, as the comments point out, this code is broken. Consider what happens if variable `i` contains the value five: The code will verify that five is in the range five through seven and then it will fetch the dword at offset 20 (`5*@size(dword)`) and jump to that address. As before, however, this loads four bytes outside the bounds of the table and does not transfer control to a defined location. One solution is to subtract the smallest case selection value from `EAX` before executing the `JMP` instruction. E.g.,

```

// SWITCH statement specifying cases 5, 6, and 7:
// WARNING: there is a better way to do this. Keep reading.

```

```

readonly
JumpTbl:dword[3] := [ &Stmt5, &Stmt6, &Stmt7 ];

```

```

    .
    .
    .
mov( i, eax );
cmp( eax, 5 );
jb EndCase
cmp( eax, 7 );           // Verify that I is in the range
ja EndCase;             // 5..7 before the indirect JMP.
sub( 5, eax );           // 5->0, 6->1, 7->2.
jmp( JumpTbl[ eax*4 ] );

    Stmt5:
        stdout.put( "I=5" );
        jmp EndCase;

    Stmt6:
        stdout.put( "I=6" );
        jmp EndCase;

    Stmt7:
        stdout.put( "I=7" );

EndCase:

```

By subtracting five from the value in EAX this code forces EAX to take on the values zero, one, or two prior to the JMP instruction. Therefore, case selection value five jumps to *Stmt5*, case selection value six transfers control to *Stmt6*, and case selection value seven jumps to *Stmt7*.

There is a sneaky way to slightly improve the code above. You can eliminate the SUB instruction by merging this subtraction into the JMP instruction's address expression. Consider the following code that does this:

```

// SWITCH statement specifying cases 5, 6, and 7:

readonly
    JumpTbl:dword[3] := [ &Stmt5, &Stmt6, &Stmt7 ];
    .
    .
    .
mov( i, eax );
cmp( eax, 5 );
jb EndCase
cmp( eax, 7 );           // Verify that I is in the range
ja EndCase;             // 5..7 before the indirect JMP.
jmp( JumpTbl[ eax*4 - 5*@size(dword)] );

    Stmt5:
        stdout.put( "I=5" );
        jmp EndCase;

    Stmt6:
        stdout.put( "I=6" );
        jmp EndCase;

    Stmt7:
        stdout.put( "I=7" );

EndCase:

```

The HLA SWITCH statement provides a DEFAULT clause that executes if the case selection value doesn't match any of the case values. E.g.,

```
switch( ebx )

    case( 5 ) stdout.put( "ebx=5" );
    case( 6 ) stdout.put( "ebx=6" );
    case( 7 ) stdout.put( "ebx=7" );
    default
        stdout.put( "ebx does not equal 5, 6, or 7" );

endswitch;
```

Implementing the equivalent of the DEFAULT clause in pure assembly language is very easy. Just use a different target label in the JB and JA instructions at the beginning of the code. The following example implements an HLA SWITCH statement similar to the one immediately above:

// SWITCH statement specifying cases 5, 6, and 7 with a DEFAULT clause:

```
readonly
    jmpTbl:dword[3] := [ &stmt5, &stmt6, &stmt7 ];
    .
    .
    .
    mov( i, eax );
    cmp( eax, 5 );
    jb DefaultCase;
    cmp( eax, 7 );           // Verify that I is in the range
    ja DefaultCase;         // 5..7 before the indirect JMP.
    jmp( jmpTbl[ eax*4 - 5*@size(dword)] );

    stmt5:
        stdout.put( "I=5" );
        jmp EndCase;

    stmt6:
        stdout.put( "I=6" );
        jmp EndCase;

    stmt7:
        stdout.put( "I=7" );
        jmp EndCase;

    DefaultCase:
        stdout.put( "I does not equal 5, 6, or 7" );
    EndCase:
```

The second restriction noted earlier, that the case values need to be contiguous, is easy to handle by inserting extra entries into the jump table. Consider the following HLA SWITCH statement:

```
switch( ebx )

    case( 1 ) stdout.put( "ebx = 1" );
    case( 2 ) stdout.put( "ebx = 2" );
    case( 4 ) stdout.put( "ebx = 4" );
    case( 8 ) stdout.put( "ebx = 8" );
    default
        stdout.put( "ebx is not 1, 2, 4, or 8" );

endswitch;
```

The minimum switch value is one and the maximum value is eight. Therefore, the code before the indirect JMP instruction needs to compare the value in EBX against one and eight. If the value is between one and eight, it's still possible that EBX might not contain a legal case selection value. However, since the JMP instruction indexes into a table of dwords using the case selection table, the table must have eight dword entries. To handle the values between one and eight that are not case selection values, simply put the statement label of the default clause (or the label specifying the first instruction after the ENDSWITCH if there is no DEFAULT clause) in each of the jump table entries that don't have a corresponding CASE clause. The following code demonstrates this technique:

```
readonly
  JumpTbl2: dword :=
    [
      &Case1, &Case2, &dfltCase, &Case4,
      &dfltCase, &dfltCase, &dfltCase, &Case8
    ];
    .
    .
    .
  cmp( ebx, 1 );
  jb dfltCase;
  cmp( ebx, 8 );
  ja dfltCase;
  jmp( JumpTbl2[ ebx*4 - 1*@size(dword)] );

  Case1:
    stdout.put( "ebx = 1" );
    jmp EndOfSwitch;

  Case2:
    stdout.put( "ebx = 2" );
    jmp EndOfSwitch;

  Case4:
    stdout.put( "ebx = 4" );
    jmp EndOfSwitch;

  Case8:
    stdout.put( "ebx = 8" );
    jmp EndOfSwitch;

  dfltCase:
    stdout.put( "ebx is not 1, 2, 4, or 8" );

  EndOfSwitch:
```

There is a problem with this implementation of the SWITCH statement. If the CASE values contain non-consecutive entries that are widely spaced the jump table could become exceedingly large. The following SWITCH statement would generate an extremely large code file:

```
switch( ebx )

  case( 1      ) stmt1;
  case( 100    ) stmt2;
  case( 1_000  ) stmt3;
  case( 10_000 ) stmt4;
  default stmt5;

endswitch;
```

In this situation, your program will be much smaller if you implement the SWITCH statement with a sequence of IF statements rather than using an indirect jump statement. However, keep one thing in mind-

the size of the jump table does not normally affect the execution speed of the program. If the jump table contains two entries or two thousand, the SWITCH statement will execute the multi-way branch in a constant amount of time. The IF statement implementation requires a linearly increasing amount of time for each case label appearing in the case statement.

Probably the biggest advantage to using assembly language over a HLL like Pascal or C/C++ is that you get to choose the actual implementation. In some instances you can implement a SWITCH statement as a sequence of IF..THEN..ELSEIF statements, or you can implement it as a jump table, or you can use a hybrid of the two:

```
switch( eax )

    case( 0 ) stmt0;
    case( 1 ) stmt1;
    case( 2 ) stmt2;
    case( 100 ) stmt3;
    default stmt4;

endswitch;
```

could become:

```
cmp( eax, 100 );
je DoStmt3;
cmp( eax, 2 );
ja TheDefaultCase;
jmp( JumpTbl[ eax*4 ]);
etc.
```

Of course, you could do this in HLA using the following code high-level control structures:

```
if( ebx = 100 ) then stmt3
else
    switch( eax )
        case(0) stmt0;
        case(1) stmt1;
        case(2) stmt2;
        Otherwise stmt4
    endswitch;
endif;
```

But this tends to destroy the readability of the program. On the other hand, the extra code to test for 100 in the assembly language code doesn't adversely affect the readability of the program (perhaps because it's so hard to read already). Therefore, most people will add the extra code to make their program more efficient.

The C/C++ SWITCH statement is very similar to the HLA SWITCH statement⁴. There is only one major semantic difference: the programmer must explicitly place a BREAK statement in each CASE clause to transfer control to the first statement beyond the SWITCH. This BREAK corresponds to the JMP instruction at the end of each CASE sequence in the assembly code above. If the corresponding BREAK is not present, C/C++ transfers control into the code of the following CASE. This is equivalent to leaving off the JMP at the end of the CASE'S sequence:

```
switch (i)
{
    case 0: stmt1;
    case 1: stmt2;
    case 2: stmt3;
        break;
    case 3: stmt4;
```

4. The HLA Standard Library SWITCH statement actually provides an option to support C semantics. See the HLA Standard Library documentation for details.

```

        break;
    default: stmt5;
}

```

This translates into the following 80x86 code:

```

readonly
    JumpTbl: dword[4] := [ &case0, &case1, &case2, &case3 ];
    .
    .
    .
    mov( i, ebx );
    cmp( ebx, 3 );
    ja DefaultCase;
    jmp( JumpTbl[ ebx*4 ]);

    case0:
        stmt1;

    case1:
        stmt2;

    case2:
        stmt3;
        jmp EndCase;    // Emitted for the break stmt.

    case3:
        stmt4;
        jmp EndCase;    // Emitted for the break stmt.

    DefaultCase:
        stmt5;

EndCase:

```

2.9 State Machines and Indirect Jumps

Another control structure commonly found in assembly language programs is the *state machine*. A state machine uses a *state variable* to control program flow. The FORTRAN programming language provides this capability with the assigned GOTO statement. Certain variants of C (e.g., GNU's GCC from the Free Software Foundation) provide similar features. In assembly language, the indirect jump provides a mechanism to easily implement state machines.

So what is a state machine? In very basic terms, it is a piece of code⁵ that keeps track of its execution history by entering and leaving certain "states". For the purposes of this chapter, we'll not use a very formal definition of a state machine. We'll just assume that a state machine is a piece of code which (somehow) remembers the history of its execution (its *state*) and executes sections of code based upon that history.

In a very real sense, all programs are state machines. The CPU registers and values in memory constitute the "state" of that machine. However, we'll use a much more constrained view. Indeed, for most purposes only a single variable (or the value in the EIP register) will denote the current state.

Now let's consider a concrete example. Suppose you have a procedure which you want to perform one operation the first time you call it, a different operation the second time you call it, yet something else the third time you call it, and then something new again on the fourth call. After the fourth call it repeats these four different operations in order. For example, suppose you want the procedure to ADD EAX and EBX the

5. Note that state machines need not be software based. Many state machines' implementation are hardware based.

first time, subtract them on the second call, multiply them on the third, and divide them on the fourth. You could implement this procedure as follows:

```

procedure StateMachine;
static
  State:byte := 0;
begin StateMachine;

  cmp( State, 0 );
  jne TryState1;

  // State 0: Add EBX to EAX and switch to state 1:

  add( ebx, eax );
  inc( State );
  exit StateMachine;

TryState1:
  cmp( State, 1 );
  jne TryState2;

  // State 1: subtract ebx from EAX and switch to state 2:

  sub( ebx, eax );
  inc( State );      // State 1 becomes State 2.
  exit StateMachine;

TryState2:
  cmp( State, 2 );
  jne MustBeState3;

  // If this is state 2, multiply EAX by EAX and switch to state 3:

  intmul( ebx, eax );
  inc( State );      // State 2 becomes State 3.
  exit StateMachine;

  // If it isn't one of the above states, we must be in state 3
  // So divide eax by ebx and switch back to state zero.

MustBeState3:
  push( edx );      // Preserve this 'cause it gets whacked by DIV.
  xor( edx, edx );  // Zero extend EAX into EDX.
  div( ebx, edx:eax);
  pop( edx );       // Restore EDX's value preserved above.
  mov( 0, State ); // Reset the state back to zero.

end StateMachine;

```

Technically, this procedure is not the state machine. Instead, it is the variable *State* and the CMP/JNE instructions which constitute the state machine.

There is nothing particularly special about this code. It's little more than a SWITCH statement implemented via the IF.THEN..ELSEIF construct. The only thing special about this procedure is that it remembers how many times it has been called⁶ and behaves differently depending upon the number of calls. While this is a *correct* implementation of the desired state machine, it is not particularly efficient. The astute reader, of course, would recognize that this code could be made a little faster using an actual SWITCH statement rather than the IF.THEN..ELSEIF implementation. However, there is a better way...

6. Actually, it remembers how many times, *MOD 4*, that it has been called.

The more common implementation of a state machine in assembly language is to use an *indirect jump*. Rather than having a state variable which contains a value like zero, one, two, or three, we could load the state variable with the *address* of the code to execute upon entry into the procedure. By simply jumping to that address, the state machine could save the tests above needed to execute the proper code fragment. Consider the following implementation using the indirect jump:

```

procedure StateMachine;
static
    State:dword := &State0;
begin StateMachine;

    jmp( State );

    // State 0: Add EBX to EAX and switch to state 1:

State0:
    add( ebx, eax );
    mov( &State1, State );
    exit StateMachine;

State1:

    // State 1: subtract ebx from EAX and switch to state 2:

    sub( ebx, eax );
    mov( &State2, State );    // State 1 becomes State 2.
    exit StateMachine;

State2:

    // If this is state 2, multiply EAX by EAX and switch to state 3:

    intmul( ebx, eax );
    mov( &State3, State );    // State 2 becomes State 3.
    exit StateMachine;

    // State 3: divide eax by ebx and switch back to state zero.

State3:
    push( edx );                // Preserve this 'cause it gets whacked by DIV.
    xor( edx, edx );            // Zero extend EAX into EDX.
    div( ebx, edx:eax);
    pop( edx );                 // Restore EDX's value preserved above.
    mov( &State0, State );    // Reset the state back to zero.

end StateMachine;

```

The JMP instruction at the beginning of the *StateMachine* procedure transfers control to the location pointed at by the *State* variable. The first time you call *StateMachine* it points at the *State0* label. Thereafter, each subsection of code sets the *State* variable to point at the appropriate successor code.

2.10 Spaghetti Code

One major problem with assembly language is that it takes several statements to realize a simple idea encapsulated by a single HLL statement. All too often an assembly language programmer will notice that s/he can save a few bytes or cycles by jumping into the middle of some program structure. After a few such observations (and corresponding modifications) the code contains a whole sequence of jumps in and out of portions of the code. If you were to draw a line from each jump to its destination, the resulting listing would end up looking like someone dumped a bowl of spaghetti on your code, hence the term “spaghetti code”.

Spaghetti code suffers from one major drawback- it's difficult (at best) to read such a program and figure out what it does. Most programs start out in a "structured" form only to become spaghetti code at the altar of efficiency. Alas, spaghetti code is rarely efficient. Since it's difficult to figure out exactly what's going on, it's very difficult to determine if you can use a better algorithm to improve the system. Hence, spaghetti code may wind up less efficient than structured code.

While it's true that producing some spaghetti code in your programs may improve its efficiency (e.g., destructuring your code, see "Efficient Implementation of IF Statements in Assembly Language" on page 772), doing so should always be a last resort after you've tried everything else and you still haven't achieved what you need. Always start out writing your programs with straight-forward IFs and SWITCH statements. Start combining sections of code (via JMP instructions) once everything is working and well understood. Of course, you should never obliterate the structure of your code unless the gains are worth it.

A famous saying in structured programming circles is "After GOTOs, pointers are the next most dangerous element in a programming language." A similar saying is "Pointers are to data structures what GOTOs are to control structures." In other words, avoid excessive use of pointers. If pointers and gotos are bad, then the indirect jump must be the worst construct of all since it involves both GOTOs and pointers! Seriously though, the indirect jump instructions should be avoided for casual use. They tend to make a program harder to read. After all, an indirect jump can (theoretically) transfer control to any label within a program. Imagine how hard it would be to follow the flow through a program if you have no idea what a pointer contains and you come across an indirect jump using that pointer. Therefore, you should always exercise care when using jump indirect instructions.

2.11 Loops

Loops represent the final basic control structure (sequences, decisions, and loops) that make up a typical program. Like so many other structures in assembly language, you'll find yourself using loops in places you've never dreamed of using loops. Most HLLs have implied loop structures hidden away. For example, consider the BASIC statement "IF A\$ = B\$ THEN 100". This IF statement compares two strings and jumps to statement 100 if they are equal. In assembly language, you would need to write a loop to compare each character in A\$ to the corresponding character in B\$ and then jump to statement 100 if and only if all the characters matched. In BASIC, there is no loop to be seen in the program. In assembly language, this very simple IF statement requires a loop to compare the individual characters in the string⁷. This is but a small example which shows how loops seem to pop up everywhere.

Program loops consist of three components: an optional initialization component, a loop termination test, and the body of the loop. The order with which these components are assembled can dramatically change the way the loop operates. Three permutations of these components appear over and over again. Because of their frequency, these loop structures are given special names in high-level languages: WHILE loops, REPEAT..UNTIL loops (do..while in C/C++), and infinite loops (e.g., FOREVER..ENDFOR in HLA).

2.11.1 While Loops

The most general loop is the WHILE loop. In HLA it takes the following form:

```
while( expression ) do <<statements>> endwhile;
```

There are two important points to note about the WHILE loop. First, the test for termination appears at the beginning of the loop. Second as a direct consequence of the position of the termination test, the body of the loop may never execute. If the termination condition is always true, the loop body will never execute.

7. Of course, the HLA Standard Library provides the *str.eq* routine that compares the strings for you, effectively hiding the loop even in an assembly language program.

Consider the following HLA WHILE loop:

```
mov( 0, I );
while( I < 100 ) do

    inc( I );

endwhile;
```

The “mov(0, I);” instruction is the initialization code for this loop. *I* is a loop control variable, because it controls the execution of the body of the loop. “I<100” is the loop termination condition. That is, the loop will not terminate as long as *I* is less than 100. The single instruction “inc(I);” is the loop body. This is the code that executes on each pass of the loop.

Note that an HLA WHILE loop can be easily synthesized using IF and JMP statements. For example, the HLA WHILE loop presented above can be replaced by:

```
mov( 0, I );
WhileLp:
if( I < 100 ) then

    inc( i );
    jmp WhileLp;

endif;
```

More generally, any WHILE loop can be built up from the following:

```
<< optional initialization code >>

UniqueLabel:
if( not_termination_condition ) then

    <<loop body>>
    jmp UniqueLabel;

endif;
```

Therefore, you can use the techniques from earlier in this chapter to convert IF statements to assembly language along with a single JMP instruction to produce a WHILE loop. The example we’ve been looking at in this section translates to the following “pure” 80x86 assembly code⁸:

```
mov( 0, i );
WhileLp:
    cmp( i, 100 );
    jnl WhileDone;
    inc( i );
    jmp WhileLp;

WhileDone:
```

2.11.2 Repeat..Until Loops

The REPEAT..UNTIL (do..while) loop tests for the termination condition at the end of the loop rather than at the beginning. In HLA, the REPEAT..UNTIL loop takes the following form:

```
<< optional initialization code >>
repeat
```

8. Note that HLA will actually convert most WHILE statements to different 80x86 code than this section presents. The reason for the difference appears a little later in this text when we explore how to write more efficient loop code.

```
<<loop body>>
```

```
until( termination_condition );
```

This sequence executes the initialization code, the loop body, then tests some condition to see if the loop should repeat. If the boolean expression evaluates to false, the loop repeats; otherwise the loop terminates. The two things to note about the REPEAT..UNTIL loop are that the termination test appears at the end of the loop and, as a direct consequence of this, the loop body always executes at least once.

Like the WHILE loop, the REPEAT..UNTIL loop can be synthesized with an IF statement and a JMP . You could use the following:

```
<< initialization code >>
```

```
SomeUniqueLabel:
```

```
<< loop body >>
```

```
if( not_the_termination_condition ) then jmp SomeUniqueLabel; endif;
```

Based on the material presented in the previous sections, you can easily synthesize REPEAT..UNTIL loops in assembly language. The following is a simple example:

```
repeat
```

```
    stdout.put( "Enter a number greater than 100: " );
```

```
    stdin.get( i );
```

```
until( i > 100 );
```

```
// This translates to the following IF/JMP code:
```

```
RepeatLbl:
```

```
    stdout.put( "Enter a number greater than 100: " );
```

```
    stdin.get( i );
```

```
if( i <= 100 ) then jmp RepeatLbl; endif;
```

```
// It also translates into the following "pure" assembly code:
```

```
RepeatLabel:
```

```
    stdout.put( "Enter a number greater than 100: " );
```

```
    stdin.get( i );
```

```
    cmp( i, 100 );
```

```
    jng RepeatLbl;
```

2.11.3 FOREVER..ENDFOR Loops

If WHILE loops test for termination at the beginning of the loop and REPEAT..UNTIL loops check for termination at the end of the loop, the only place left to test for termination is in the middle of the loop. The HLA FOREVER..ENDFOR loop, combined with the BREAK and BREAKIF statements, provide this capability. The FOREVER..ENDFOR loop takes the following form:

```

forever

    << loop body >>

endfor;

```

Note that there is no explicit termination condition. Unless otherwise provided for, the FOREVER..ENDFOR construct simply forms an infinite loop. Loop termination is typically handled by a BREAKIF statement. Consider the following HLA code that employs a FOREVER..ENDFOR construct:

```

forever

    stdin.get( character );
    breakif( character = '.' );
    stdout.put( character );

endfor;

```

Converting a FOREVER loop to pure assembly language is trivial. All you need is a label and a JMP instruction. The BREAKIF statement in this example is really nothing more than an IF and a JMP instruction. The “pure” assembly language version of the code above looks something like the following:

```

foreverLabel:

    stdin.get( character );
    cmp( character, '.' );
    je ForIsDone;
    stdout.put( character );
    jmp foreverLabel;

ForIsDone:

```

2.11.4 FOR Loops

The FOR loop is a special form of the WHILE loop that repeats the loop body a specific number of times. In HLA, the FOR loop takes the following form:

```

for( <<Initialization Stmt>>; <<Termination Expression>>; <<inc_Stmt>> ) do

    << statements >>

endfor;

```

This is completely equivalent to the following:

```

<< Initialization Stmt>>;
while( <<Termination Expression>> ) do

    << statements >>

    <<inc_Stmt>>

endwhile;

```

Traditionally, the FOR loop has been used to process arrays and other objects accessed in sequential numeric order. One normally initializes a loop control variable with the initialization statement and then uses the loop control variable as an index into the array (or other data type), e.g.,

```

for( mov(0, esi); esi < 7; inc( esi ) ) do

```



```

stdout.put( "Array Element = ", SomeArray[ esi*4], nl );
endfor;

```

To convert this to “pure” assembly language, begin by translating the FOR loop into an equivalent WHILE loop:

```

mov( 0, esi );
while( esi < 7 ) do

    stdout.put( "Array Element = ", SomeArray[ esi*4], nl );

    inc( esi );
endwhile;

```

Now, using the techniques from the section on WHILE loops (see “While Loops” on page 787), translate the code into pure assembly language:

```

mov( 0, esi );
WhileLp:
cmp( esi, 7 );
jnl EndWhileLp;

    stdout.put( "Array Element = ", SomeArray[ esi*4], nl );

    inc( esi );
    jmp WhileLp;

EndWhileLp:

```

2.11.5 The BREAK and CONTINUE Statements

The HLA BREAK and CONTINUE statements both translate into a single JMP instruction. The BREAK instruction exits the loop that immediately contains the BREAK statement; the CONTINUE statement restarts the loop that immediately contains the CONTINUE statement.

Converting a BREAK statement to “pure” assembly language is very easy. Just emit a JMP instruction that transfers control to the first statement following the ENDxxx clause of the loop to exit. This is easily accomplished by placing a label after the associated END clause and jumping to that label. The following code fragments demonstrate this technique for the various loops:

```

// Breaking out of a forever loop:

forever
    <<stmts>>
    //break;
    jmp BreakFromForever;
    <<stmts>>
endfor;
BreakFromForever:

// Breaking out of a FOR loop;
for( <<initStmt>>; <<expr>>; <<incStmt>> ) do
    <<stmts>>
    //break;
    jmp BrkFromFor;
    <<stmts>>
endfor;

```

```
BrkFromFor:

// Breaking out of a WHILE loop:

while( <<expr>> ) do
  <<stmts>>
  //break;
  jmp BrkFromWhile;
<<stmts>>
endwhile;
BrkFromWhile:

// Breaking out of a REPEAT..UNTIL loop:

repeat
  <<stmts>>
  //break;
  jmp BrkFromRpt;
<<stmts>>
until( <<expr>> );
BrkFromRpt:
```

The CONTINUE statement is slightly more difficult to implement than the BREAK statement. The implementation still consists of a single JMP instruction, however the target label doesn't wind up going in the same spot for each of the different loops. The following figures show where the CONTINUE statement transfers control for each of the HLA loops:

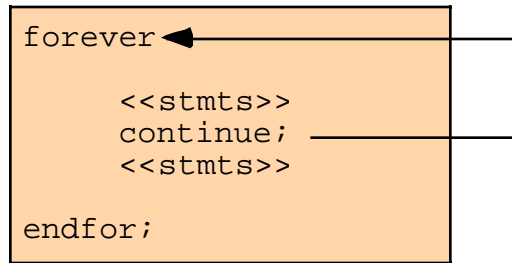


Figure 2.2 CONTINUE Destination for the FOREVER Loop

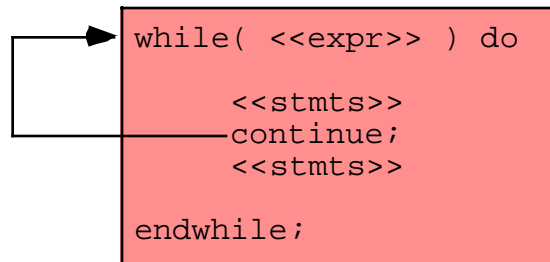



Figure 2.3 CONTINUE Destination and the WHILE Loop



Note: CONTINUE forces the execution of the <<incStmt>> clause and then transfers control to the test for loop termination.

Figure 2.4 CONTINUE Destination and the FOR Loop

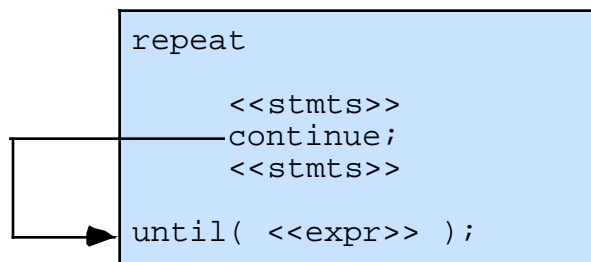


Figure 2.5 CONTINUE Destination and the REPEAT..UNTIL Loop

The following code fragments demonstrate how to convert the CONTINUE statement into an appropriate JMP instruction for each of these loop types.

forever..continue..endfor

```
// Conversion of forever loop w/continue
// to pure assembly:
forever
    <<stmts>>
    continue;
    <<stmts>>
endfor;
```

```
// Converted code:
```

```
foreverLbl:
    <<stmts>>
    jmp foreverLbl;
    <<stmts>>
```

while..continue..endwhile

```
// Conversion of while loop w/continue
// into pure assembly:
```

```
while( <<expr>> ) do
    <<stmts>>
    continue;
    <<stmts>>
endwhile;
```

```
// Converted code:
```

```
whlLabel:
<<Code to evaluate expr>>
Jcc EndOfWhile; // Skip loop on expr failure.
    <<stmts>>
    jmp whlLabel; // Jump to start of loop on continue.
    <<stmts>>
    jmp whlLabel; // Repeat the code.
EndOfwhile:
```

for..continue..endfor

```
// Conversion for a for loop w/continue
// into pure assembly:
```

```
for( <<initStmt>>; <<expr>>; <<incStmt>> ) do
    <<stmts>>
    continue;
    <<stmts>>
endfor;
```

```
// Converted code
```

```
<<initStmt>>
ForLpLbl:
```

```

<<Code to evaluate expr>>
Jcc EndOfFor; // Branch if expression fails.
  <<stmts>>
  jmp ContFor; // Branch to <<incStmt>> on continue.
  <<stmts>>

  ContFor:
  <<incStmt>>
  jmp ForLpLbl;
EndOfFor:

```

repeat..continue..until

```

repeat
  <<stmts>>
  continue;
  <<stmts>>
until( <<expr>> );

// Converted Code:

RptLpLbl:
  <<stmts>>
  jmp ContRpt; // Continue branches to loop termination test.
  <<stmts>>
ContRpt:
  <<code to test expr>>
  Jcc RptLpLbl; // Jumps if expression evaluates false.

```

2.11.6 Register Usage and Loops

Given that the 80x86 accesses registers much faster than memory locations, registers are the ideal spot to place loop control variables (especially for small loops). However, there are some problems associated with using registers within a loop. The primary problem with using registers as loop control variables is that registers are a limited resource. Therefore, the following will not work properly:

```

mov( 8, cx );
loop1:
  mov( 4, cx );
  loop2:
    <<stmts>>
    dec( cx );
    jnz loop2;

  dec( cx );
  jnz loop1;

```

The intent here, of course, was to create a set of nested loops, that is, one loop inside another. The inner loop (*Loop2*) should repeat four times for each of the eight executions of the outer loop (*Loop1*). Unfortunately, both loops use the same register as a loop control variable.. Therefore, this will form an infinite loop since CX will be set to zero at the end of the first loop. Since CX is always zero upon encountering the second DEC instruction, control will always transfer to the LOOP1 label (since decrementing zero produces a non-zero result). The solution here is to save and restore the CX register or to use a different register in place of CX for the outer loop:

```

mov( 8, cx );
loop1:
  push( cx );
  mov( 4, cx );
  loop2:

```

```

<<stmts>>
dec( cx );
jnz loop2;

```

```

pop( cx );
dec( cx );
jnz loop1;

```

or:

```

mov( 8, dx );
loop1:
mov( 4, cx );
loop2:
<<stmts>>
dec( cx );
jnz loop2;

dec( dx );
jnz loop1;

```

Register corruption is one of the primary sources of bugs in loops in assembly language programs, always keep an eye out for this problem.

2.12 Performance Improvements

The 80x86 microprocessors execute sequences of instructions at blinding speeds. Therefore, you'll rarely encounter a program that is slow which doesn't contain any loops. Since loops are the primary source of performance problems within a program, they are the place to look when attempting to speed up your software. While a treatise on how to write efficient programs is beyond the scope of this chapter, there are some things you should be aware of when designing loops in your programs. They're all aimed at removing unnecessary instructions from your loops in order to reduce the time it takes to execute one iteration of the loop.

2.12.1 Moving the Termination Condition to the End of a Loop

Consider the following flow graphs for the three types of loops presented earlier:

REPEAT..UNTIL loop:

```

Initialization code
  Loop body
  Test for termination
  Code following the loop

```

WHILE loop:

```

Initialization code
  Loop termination test
  Loop body
  Jump back to test
  Code following the loop

```

FOREVER..ENDFOR loop:

```

Initialization code
  Loop body, part one
  Loop termination test

```

```

    Loop body, part two
    Jump back to loop body part 1
Code following the loop

```

As you can see, the REPEAT..UNTIL loop is the simplest of the bunch. This is reflected in the assembly language code required to implement these loops. Consider the following REPEAT..UNTIL and WHILE loops that are identical:

```

// Example involving a WHILE loop:

mov( edi, esi );
sub( 20, esi );
while( esi <= edi ) do

    <<stmts>>
    inc( esi );

endwhile;

// Conversion of the code above into pure assembly language:

mov( edi, esi );
sub( 20, esi );
whlLbl:
cmp( esi, edi );
jnl EndOfWhile;

    <<stmts>>
    inc( esi );
    <<stmts>>
    jmp whlLbl;

EndOfWhile:

//Example involving a REPEAT..UNTIL loop:

mov( edi, esi );
sub( 20, esi );
repeat

    <<stmts>>
    inc( esi );

until( esi > edi );

// Conversion of the REPEAT..UNTIL loop into pure assembly:

rptLabel:
    <<stmts>>
    inc( esi );
    cmp( esi, edi );
    jng rptLabel;

```

As you can see by carefully studying the conversion to pure assembly language, testing for the termination condition at the end of the loop allowed us to remove a JMP instruction from the loop. This can be significant if this loop is nested inside other loops. In the preceding example there wasn't a problem with executing the body at least once. Given the definition of the loop, you can easily see that the loop will be executed exactly 20 times. This suggests that the conversion to a REPEAT..UNTIL loop is trivial and always possible. Unfortunately, it's not always quite this easy. Consider the following HLA code:

```

while( esi <= edi ) do
  <<stmts>>
  inc( esi );
endwhile;

```

In this particular example, we haven't the slightest idea what ESI contains upon entry into the loop. Therefore, we cannot assume that the loop body will execute at least once. So we must test for loop termination before executing the body of the loop. The test can be placed at the end of the loop with the inclusion of a single JMP instruction:

```

jmp WhlTest;
TopOfLoop:
  <<stmts>>
  inc( esi );
  WhlTest:
  cmp( esi, edi );
  jle TopOfLoop;

```

Although the code is as long as the original WHILE loop, the JMP instruction executes only once rather than on each repetition of the loop. Note that this slight gain in efficiency is obtained via a slight loss in readability. The second code sequence above is closer to spaghetti code than the original implementation. Such is often the price of a small performance gain. Therefore, you should carefully analyze your code to ensure that the performance boost is worth the loss of clarity. More often than not, assembly language programmers sacrifice clarity for dubious gains in performance, producing impossible to understand programs.

Note, by the way, that HLA translates its WHILE statement into a sequence of instructions that test the loop termination condition at the bottom of the loop using exactly the technique this section describes. Therefore, you do not have to worry about the HLA WHILE statement introducing slower code into your programs.

2.12.2 Executing the Loop Backwards

Because of the nature of the flags on the 80x86, loops which range from some number down to (or up to) zero are more efficient than any other. Compare the following HLA FOR loop and the code it generates:

```

for( mov( 1, J ); J <= 8; inc(J) ) do
  <<stmts>>
endfor;

```

// Conversion to pure assembly (as well as using a repeat..until form):

```

mov( 1, J );
ForLp:
  <<stmts>>
  inc( J );
  cmp( J, 8 );
  jnge ForLp;

```

Now consider another loop that also has eight iterations, but runs its loop control variable from eight down to one rather than one up to eight:

```

mov( 8, J );
LoopLbl:
  <<stmts>>
  dec( J );
  jnz LoopLbl;

```

Note that by running the loop from eight down to one we saved a comparison on each repetition of the loop.

Unfortunately, you cannot force all loops to run backwards. However, with a little effort and some coercion you should be able to write many FOR loops so they operate backwards. By saving the execution time of the CMP instruction on each iteration of the loop the code may run faster.

The example above worked out well because the loop ran from eight down to one. The loop terminated when the loop control variable became zero. What happens if you need to execute the loop when the loop control variable goes to zero? For example, suppose that the loop above needed to range from seven down to zero. As long as the upper bound is positive, you can substitute the JNS instruction in place of the JNZ instruction above to repeat the loop some specific number of times:

```
mov( 7, J );
LoopLbl:
  <<stmts>>
  dec( J );
  jns LoopLbl;
```

This loop will repeat eight times with *J* taking on the values seven down to zero on each execution of the loop. When it decrements zero to minus one, it sets the sign flag and the loop terminates.

Keep in mind that some values may look positive but they are negative. If the loop control variable is a byte, then values in the range 128..255 are negative. Likewise, 16-bit values in the range 32768..65535 are negative. Therefore, initializing the loop control variable with any value in the range 129..255 or 32769..65535 (or, of course, zero) will cause the loop to terminate after a single execution. This can get you into a lot of trouble if you're not careful.

2.12.3 Loop Invariant Computations

A loop invariant computation is some calculation that appears within a loop that always yields the same result. You needn't do such computations inside the loop. You can compute them outside the loop and reference the value of the computation inside the loop. The following HLA code demonstrates a loop which contains an invariant computation:

```
for( mov( 0, eax ); eax < n; inc( eax ) ) do

  mov( eax, edx );
  add( j, edx );
  sub( 2, edx );
  add( edx, k );

endfor;
```

Since *j* never changes throughout the execution of this loop, the sub-expression “*j*-2” can be computed outside the loop and its value used in the expression inside the loop:

```
mov( j, ecx );
sub( 2, ecx );
for( mov( 0, eax ); eax < n; inc( eax ) ) do

  mov( eax, edx );
  add( ecx, edx );
  add( edx, k );

endfor;
```

Still, the value in ECX never changes inside this loop, so although we've eliminated a single instruction by computing the subexpression “*j*-2” outside the loop, there is still an invariant component to this calculation. Since we note that this invariant component executes *n* times in the loop, we can translate the code above to the following:

```
mov( j, ecx );
sub( 2, ecx );
```

```

intmul( n, ecx ); // Compute n*(j-2) and add this into k outside
add( ecx, k ); // the loop.
for( mov( 0, eax ); eax < n; inc( eax ) ) do

    add( eax, k );

endfor;

```

As you can see, we've shrunk the loop body from four instructions down to one. Of course, if you're really interested in improving the efficiency of this particular loop, you'd be much better off (most of the time) computing k using the formula:

$$k = k + ((n + 1) \times temp) + \frac{(n + 2) \times (n + 2)}{2}$$

This computation for k is based on the formula:

$$\sum_{i=0}^n i = \frac{(n + 1) \times (n)}{2}$$

However, simple computations such as this one aren't always possible. Still, this demonstrates that a better algorithm is almost always better than the trickiest code you can come up with.

Removing invariant computations and unnecessary memory accesses from a loop (particularly an inner loop in a set of nested loops) can produce dramatic performance improvements in a program.

2.12.4 Unraveling Loops

For small loops, that is, those whose body is only a few statements, the overhead required to process a loop may constitute a significant percentage of the total processing time. For example, look at the following Pascal code and its associated 80x86 assembly language code:

```

FOR I := 3 DOWNTO 0 DO A [I] := 0;

mov( 3, I );
LoopLbl:
    mov( I, ebx );
    mov( 0, A[ebx*4] );
    dec( I );
    jns LoopLbl;

```

Each iteration of the loop requires four instructions. Only one instruction is performing the desired operation (moving a zero into an element of A). The remaining three instructions control the repetition of the loop. Therefore, it takes 16 instructions to do the operation logically required by four.

While there are many improvements we could make to this loop based on the information presented thus far, consider carefully exactly what it is that this loop is doing-- it's simply storing four zeros into $A[0]$ through $A[3]$. A more efficient approach is to use four MOV instructions to accomplish the same task. For example, if A is an array of dwords, then the following code initializes A much faster than the code above:

```

mov( 0, A[0] );
mov( 0, A[4] );
mov( 0, A[8] );
mov( 0, A[12] );

```

Although this is a trivial example, it shows the benefit of loop unraveling. If this simple loop appeared buried inside a set of nested loops, the 4:1 instruction reduction could possibly double the performance of that section of your program.

Of course, you cannot unravel all loops. Loops that execute a variable number of times cannot be unraveled because there is rarely a way to determine (at assembly time) the number of times the loop will execute.

Therefore, unraveling a loop is a process best applied to loops that execute a known number of times (and the number of times is known at assembly time).

Even if you repeat a loop some fixed number of iterations, it may not be a good candidate for loop unraveling. Loop unraveling produces impressive performance improvements when the number of instructions required to control the loop (and handle other overhead operations) represent a significant percentage of the total number of instructions in the loop. Had the loop above contained 36 instructions in the body of the loop (exclusive of the four overhead instructions), then the performance improvement would be, at best, only 10% (compared with the 300-400% it now enjoys). Therefore, the costs of unraveling a loop, i.e., all the extra code which must be inserted into your program, quickly reaches a point of diminishing returns as the body of the loop grows larger or as the number of iterations increases. Furthermore, entering that code into your program can become quite a chore. Therefore, loop unraveling is a technique best applied to small loops.

Note that the superscalar x86 chips (Pentium and later) have *branch prediction hardware* and use other techniques to improve performance. Loop unrolling on such systems many actually *slow down* the code since these processors are optimized to execute short loops.

2.12.5 Induction Variables

Consider the following modification of the loop presented in the previous section:

```
FOR I := 0 TO 255 DO csetVar[I] := {};
```

Here the program is initializing each element of an array of character sets to the empty set. The straight-forward code to achieve this is the following:

```
mov( 0, i );
FLp:

// Compute the index into the array (note that each element
// of a CSET array contains 16 bytes).

mov( i, ebx );
shl( 4, ebx );

// Set this element to the empty set (all zero bits).

mov( 0, csetVar[ ebx ] );
mov( 0, csetVar[ ebx+4 ] );
mov( 0, csetVar[ ebx+8 ] );
mov( 0, csetVar[ ebx+12 ] );

inc( i );
cmp( i, 256 );
jb FLp;
```

Although unraveling this code will still produce a performance improvement, it will take 1024 instructions to accomplish this task, too many for all but the most time-critical applications. However, you can reduce the execution time of the body of the loop using *induction variables*. An induction variable is one whose value depends entirely on the value of some other variable. In the example above, the index into the array *csetVar* tracks the loop control variable (it's always equal to the value of the loop control variable times 16). Since *i* doesn't appear anywhere else in the loop, there is no sense in performing all the computations on *i*. Why not operate directly on the array index value? The following code demonstrates this technique:

```
mov( 0, ebx );
FLp:
mov( 0, csetVar[ ebx ] );
mov( 0, csetVar[ ebx+4 ] );
mov( 0, csetVar[ ebx+8 ] );
```

```

mov( 0, csetVar[ ebx+12 ] );

add( 16, ebx );
cmp( ebx, 256*16 );
jb FLp;

```

The induction that takes place in this example occurs when the code increments the loop control variable (moved into EBX for efficiency reasons) by 16 on each iteration of the loop rather than by one. Multiplying the loop control variable by 16 allows the code to eliminate multiplying the loop control variable by 16 on each iteration of the loop (i.e., this allows us to remove the SHL instruction from the previous code). Further, since this code no longer refers to the original loop control variable (*i*), the code can maintain the loop control variable strictly in the EBX register.

2.13 Hybrid Control Structures in HLA

The HLA high level language control structures have a few drawbacks: (1) they're not true assembly language instructions, (2) complex boolean expressions only support short circuit evaluation, and (3) they often introduce inefficient coding practices into a language that most people only use when they need to write high-performance code. On the other hand, while the 80x86 low level control structures let you write efficient code, the resulting code is very difficult to read and maintain. HLA provides a set of hybrid control structures that allow you to use pure assembly language statements to evaluate boolean expressions while using the high level control structures to delineate the statements controlled by the boolean expressions. The result is code that is much more readable than pure assembly language without being a whole lot less efficient.

HLA provides hybrid forms of the IF.ELSEIF.ELSE..ENDIF, WHILE..ENDWHILE, REPEAT..UNTIL, BREAKIF, EXITIF, and CONTINUEIF statements (i.e., those that involve a boolean expression). For example, a hybrid IF statement takes the following form:

```

if( #{ <<statements>> }# ) then <<statements>> endif;

```

Note the use of #{ and }# operators to surround a sequence of instructions within this statement. This is what differentiates the hybrid control structures from the standard high level language control structures. The remaining hybrid structures take the following forms:

```

while( #{ <<statements>> }# ) <<statements>> endwhile;
repeat <<statements>> until( #{ <<statements>> }# );
breakif( #{ <<statements>> }# );
exitif( #{ <<statements>> }# );
continueif( #{ <<statements>> }# );

```

The statements within the curly braces replace the normal boolean expression in an HLA high level control structure. These particular statements are special insofar as HLA defines two labels, *true* and *false*, within their context. HLA associates the label *true* with the code that would normally execute if a boolean expression were present and that expression's result was true. Similarly, HLA associates the label *false* with the code that would execute if a boolean expression in one of these statements evaluated false. As a simple example, consider the following two (equivalent) IF statements:

```

if( eax < ebx ) then inc( eax ); endif;

if
( #{
  cmp( eax, ebx );
  jnl false;
}# ) then
  inc( eax );

```

```
endif;
```

The JNL that transfers control to the *false* label in this latter example will skip over the INC instruction if EAX is not less than EBX. Note that if EAX is less than EBX then control falls through to the INC instruction. This is roughly equivalent to the following pure assembly code:

```
cmp( eax, ebx );
jnl falseLabel;
    inc( eax );
falseLabel:
```

As a slightly more complex example, consider the statement

```
if( eax >= J && eax <= K ) then sub( J, eax ); endif;
```

The following hybrid IF statement accomplishes the above:

```
if
( #{
    cmp( eax, J );
    jnge false;
    cmp( eax, K );
    jnle false;
}# ) then
    sub( J, eax );
endif;
```

As one final example of the hybrid IF statement, consider the following:

```
// if( ((eax > ebx) && (eax < ecx)) || (eax = edx) ) then mov( ebx, eax ); endif;

if
( #{
    cmp( eax, edx );
    je true;
    cmp( eax, ebx );
    jng false;
    cmp( eax, ecx );
    jnl false;
}# ) then
    mov( ebx, eax );
endif;
```

Since these examples are rather trivial, they don't really demonstrate how much more readable the code can be when using hybrid statements rather than pure assembly code. However, one thing you notice is that the use of hybrid statements eliminate the need to insert labels throughout your code. This is what makes your programs easier to read and understand.

For the IF statement, the *true* label corresponds to the THEN clause of the statement; the *false* label corresponds to the ELSEIF, ELSE, or ENDIF clause (whichever follows the THEN clause). For the WHILE loop, the *true* label corresponds to the body of the loop while the *false* label is attached to the first statement following the corresponding ENDWHILE. For the REPEAT.UNTIL statement, the *true* label is attached to the code following the UNTIL clause while the *false* label is attached to the first statement of the body of the loop. The BREAKIF, EXITIF, and CONTINUEIF statements associate the *false* label with the statement immediately following one of these statements, they associate the *true* label with the code normally associated with a BREAK, EXIT, or CONTINUE statement.

2.14 Putting It All Together

In this chapter we've taken a look at the low-level, or "pure" assembly language, implementation of several common control structures. Although HLA's high level control structures are easy to use and quite a bit more readable than their low-level equivalents, sometimes efficiency demands a low-level implementation. This chapter presents the blueprints for such transformations.

While this chapter covers the principle high level control structures and their translation to assembly language, there are some additional control structures that this chapter does not consider. Iterators and the TRY..ENDTRY statement are two good examples. Fear not, however, the volume on Advanced Assembly Language Programming will tidy up those loose ends.