Style Guidelines for Assembly Language Programmers

1.0   Introduction

Most people consider assembly language programs difficult to read. While there are a multitude of reasons why people feel this way, the primary reason is that assembly language does not make it easy for programmers to write readable programs. This doesn't mean it's impossible to write readable programs, only that it takes an extra effort on the part of an assembly language programmer to produce readable code.

To demonstrate some common problems with assembly language programs, consider the following programs or program segments. These are actual programs written in assembly language taken from the internet. Each example demonstrates a separate problem. (By the way, the choice of these examples is not intended to embarass the original authors. These programs are typical of assembly language source code found on the Internet.)

1.1   ADDHEX.ASM

```
%TITLE "Sums TWO hex values"

        IDEAL
        DOSSEG
        MODEL   small
        STACK   256

        DATASEG

exitCode        db      0
prompt1         db      'Enter value 1: ', 0
prompt2         db      'Enter value 2: ', 0
string          db      20 DUP (?)

        CODESEG

        EXTRN   StrLength:proc
        EXTRN   StrWrite:proc, StrRead:proc, NewLine:proc
        EXTRN   AscToBin:proc, BinToAscHex:proc

Start:
        mov     ax,@data
        mov     ds,ax
        mov     es,ax
        mov     di, offset prompt1
        call    GetValue
        push    ax
        mov     di, offset prompt2
        call    GetValue
        pop     bx
        add     ax,bx
        mov     cx,4
        mov     di, offset string
```

```
        call    BinToAscHex
        call    StrWrite
Exit:
        mov     ah,04Ch
        mov     al,[exitCode]
        int     21h

PROC    GetValue
        call    StrWrite
        mov     di, offset string
        mov     cl,4
        call    StrRead
        call    NewLine
        call    StrLength
        mov     bx,cx
        mov     [word bx + di], 'h'
        call    AscToBin
        ret
ENDP    GetValue

        END     Start
```

Well, the biggest problem with this program should be fairly obvious - it has absolutely no comments other than the title of the program. Another problem is the fact that strings that prompt the user appear in one part of the program and the calls that print those strings appear in another. While this is typical assembly language programming, it still makes the program harder to read. Another, relatively minor, problem is that it uses TASM's "less-than" IDEAL syntax[1].

This program also uses the MASM/TASM "simplified" segment directives. How typically Microsoft to name a feature that adds complexity to a product "simplified." It turns out that programs that use the standard segmentation directives will be easier to read[2].

Before moving one, it is worthwhile to point out two good features about this program (with respect to readability). First, the programmer chose a reasonable set of names for the procedures and variables this program uses (I'll assume the author of this code segment is also the author of the library routines it calls). Another positive aspect to this program is that the mnemonic and operand fields are nicely aligned.

Okay, after complaining about how hard this code is to read, how about a more readable version? The following program is, arguably, more readable than the version above. Arguably, because this version uses the UCR Standard Library v2.0 and it assumes that the reader is familiar with features of that particular library.

```
;*************************************************
;
; AddHex-
;
; This simple program reads two integer values from
; the user, computes their sum, and prints the
; result to the display.
;
; This example uses the "UCR Standard Library for
```

---

1. Someone who uses TASM all the time may think this is fine, but consider those individuals who don't. They're not familiar with TASM's funny syntax so they may find several statements in this program to be confusing.
2. Simplified segment directives *do* make it easier to write assembly language programs that interface with HLLs. However, they only complicate matters in stand-alone assembly language programs.

```
; 80x86 Assembly Language Programmers v2.0"
;
; Randall Hyde
; 12/13/96

                title           AddHex
                .xlist
                include         ucrlib.a
                includelib      ucrlib.lib
                .list




cseg            segment para public 'code'
                assume  cs:cseg


; GetInt-
;
; This function reads an integer value from the keyboard and
; returns that value in the AX register.
;
; This routine traps illegal values (either too large or
; incorrect digits) and makes the user re-enter the value.

GetInt          textequ <call GetInt_p>
GetInt_p        proc
                push    dx              ;DX hold error code.

GetIntLoop:     mov     dx, false       ;Assume no error.
                try                     ;Trap any errors.

                FlushGetc               ;Force input from a new line.
                geti                    ;Read the integer.

                except  $Conversion     ;Trap if bad characters.
                print   "Illegal numeric conversion, please re-enter", nl
                mov     dx, true

                except  $Overflow       ;Trap if # too large.
                print   "Value out of range, please re-enter.",nl
                mov     dx, true



                endtry
                cmp     dx, true
                je      GetIntLoop
                pop     dx
                ret
GetInt_p        endp
```

Style Guidelines for Assembly Language Programmers

```
Main            proc

                InitExcept

                print   'Enter value 1: '
                GetInt
                mov     bx, ax

                print   'Enter value 2: '
                GetInt
                print   cr, lf, 'The sum of the two values is '
                add     ax, bx
                puti
                putcr

Quit:           CleanUpEx
                ExitPgm                 ;DOS macro to quit program.
Main            endp

cseg            ends

sseg            segment para stack 'stack'
stk             db      256 dup (?)
sseg            ends


zzzzzzseg       segment para public 'zzzzzz'
LastBytes       db      16 dup (?)
zzzzzzseg       ends
                end     Main
```

It is well worth pointing out that this code does quite a bit more than the original AddHex program. In particular, it validates the user's input; something the original program did not do. If one were to exactly simulate the original program, the program could be simplified to the following:

```
                print   nl, 'Enter value 1: '
                Geti
                mov     bx, ax

                print   nl, 'Enter value 2: '
                Geti
                add     ax, bx
                putcr
                puti
                putcr
```

In this example, the two sample solutions improved the readability of the program by adding comments, formatting the program a little bit better, and by using the high-level features of the UCR Standard Library to simplify the coding and keep output string literals with the statements that print them.

1.2   Graphics Example

The following program segment comes from a much larger program named "MODEX.ASM" on the net. It deals with setting up the color graphics display.

```
;==================================
;SET_POINT (Xpos%, Ypos%, ColorNum%)
;==================================
;
; Plots a single Pixel on the active display page
;
; ENTRY: Xpos     = X position to plot pixel at
;        Ypos     = Y position to plot pixel at
;        ColorNum = Color to plot pixel with
;
; EXIT:  No meaningful values returned
;


SP_STACK      STRUC
                DW  ?,? ; BP, DI
                DD  ?   ; Caller
    SETP_Color  DB  ?,? ; Color of Point to Plot
    SETP_Ypos   DW  ?   ; Y pos of Point to Plot
    SETP_Xpos   DW  ?   ; X pos of Point to Plot
SP_STACK      ENDS

        PUBLIC SET_POINT

SET_POINT    PROC     FAR

    PUSHx   BP, DI              ; Preserve Registers
    MOV     BP, SP              ; Set up Stack Frame

    LES     DI, d CURRENT_PAGE  ; Point to Active VGA Page

    MOV     AX, [BP].SETP_Ypos  ; Get Line # of Pixel
    MUL     SCREEN_WIDTH        ; Get Offset to Start of Line

    MOV     BX, [BP].SETP_Xpos  ; Get Xpos
    MOV     CX, BX              ; Copy to extract Plane # from
    SHR     BX, 2               ; X offset (Bytes) = Xpos/4
    ADD     BX, AX              ; Offset = Width*Ypos + Xpos/4

    MOV     AX, MAP_MASK_PLANE1 ; Map Mask & Plane Select Register
    AND     CL, PLANE_BITS      ; Get Plane Bits
    SHL     AH, CL              ; Get Plane Select Value
    OUT_16  SC_Index, AX        ; Select Plane

    MOV     AL,[BP].SETP_Color  ; Get Pixel Color
    MOV     ES:[DI+BX], AL      ; Draw Pixel

    POPx    DI, BP              ; Restore Saved Registers
```

```
     RET     6                       ; Exit and Clean up Stack


SET_POINT          ENDP
```

Unlike the previous example, this one has lots of comments. Indeed, the comments are not bad. However, this particular routine suffers from its own set of problems. First, most of the instructions, register names, and identifiers appear in upper case. Upper case characters are *much* harder to read than lower case letters. Considering the extra work involved in entering upper case letters into the computer, it's a real shame to see this type of mistake in a program[3]. Another big problem with this particular code segment is that the author didn't align the label field, the mnemonic field, and the operand field very well (it's not horrible, but it's bad enough to affect the readability of the program.

Here is an improved version of the program:

```
;==================================
;
;SetPoint (Xpos%, Ypos%, ColorNum%)
;
;
; Plots a single Pixel on the active display page
;
; ENTRY: Xpos      = X position to plot pixel at
;        Ypos      = Y position to plot pixel at
;        ColorNum = Color to plot pixel with
;
;        ES:DI    = Screen base address (??? I added this without really
;                                       knowing what is going on here [RLH]).
;
; EXIT:  No meaningful values returned
;
dp              textequ <dword ptr>

Color           textequ <[bp+6]>
YPos            textequ <[bp+8]>
XPos            textequ <[bp+10]>

                public  SetPoint
SetPoint        proc    far
                push    bp
                mov     bp, sp
                push    di
                les     di, dp CurrentPage      ;Point at active VGA Page

                mov     ax, YPos                ;Get line # of Pixel
                mul     ScreenWidth             ;Get offset to start of line

                mov     bx, XPos                ;Get offset into line
                mov     cx, bx                  ;Save for plane computations
```

_____

3. A lot of old-time programmers believe that assembly instructions should appear in upper case. A lot of this has to do with the fact that old IBM mainframes and certain personal computers like the original Apple II only supported upper case characters.

```
            shr     bx, 2                   ;X offset (bytes)= XPos/4
            add     bx, ax                  ;Offset=Width*YPos + XPos/4

            mov     ax, MapMaskPlane1       ;Map mask & plane select reg
            and     cl, PlaneBits           ;Get plane bits
            shl     ah, cl                  ;Get plane select value
            out_16  SCIndex, ax             ;Select plane

            mov     al, Color               ;Get pixel color
            mov     es:[di+bx], al          ;Draw pixel

            pop     di
            pop     bp
            ret     6
SetPoint        endp
```

Most of the changes here were purely mechanical: reducing the number of upper case letters in the program, spacing the program out better, adjusting some comments, etc. Nevertheless, these small, subtle, changes have a big impact on how easy the code is to read (at least, to an experienced assembly langage programmer).

1.3   S.COM Example

The following code sequence came from a program labelled "S.COM" that was also found in an archive on the internet.

```
;Get all file names matching filespec and set up tables
GetFileRecords:
    mov dx, OFFSET DTA              ;Set up DTA
    mov ah, 1Ah
    int 21h
    mov dx, FILESPEC                ;Get first file name
    mov cl, 37h
    mov ah, 4Eh
    int 21h
    jnc FileFound                   ;No files.  Try a different filespec.
    mov si, OFFSET NoFilesMsg
    call Error
    jmp NewFilespec
FileFound:
    mov di, OFFSET fileRecords      ;DI -> storage for file names
    mov bx, OFFSET files            ;BX -> array of files
    sub bx, 2
StoreFileName:
    add bx, 2                       ;For all files that will fit,
    cmp bx, (OFFSET files) + NFILES*2
    jb @@L1
    sub bx, 2
    mov [last], bx
    mov si, OFFSET tooManyMsg
    jmp DoError
@@L1:
    mov [bx], di                    ;Store pointer to status/filename in files[]
```

Style Guidelines for Assembly Language Programmers

```
        mov al, [DTA_ATTRIB]       ;Store status byte
        and al, 3Fh                ;Top bit is used to indicate file is marked
        stosb
        mov si, OFFSET DTA_NAME     ;Copy file name from DTA to filename storage
        call CopyString
        inc di
        mov si, OFFSET DTA_TIME     ;Copy time, date and size
        mov cx, 4
        rep movsw
        mov ah, 4Fh                ;Next filename
        int 21h
        jnc StoreFileName
        mov [last], bx             ;Save pointer to last file entry
        mov al, [keepSorted]       ;If returning from EXEC, need to resort files?
        or al, al
        jz DisplayFiles
        jmp Sort0
```

The primary problem with this program is the formatting. The label fields overlap the mnemonic fields (in almost every instance), the operand fields of the various instructions are not aligned, there are very few blank lines to organize the code, the programmer makes excessive use of "local" label names, and, although not prevalent, there are a few items that are all upper-case (remember, upper case characters are harder to read). This program also makes considerable use of "magic numbers," especially with respect to opcodes passed on to DOS.

Another subtle problem with this program is the way it organizes control flow. At a couple of points in the code it checks to see if an error condition exists (file not found and too many files processed). If an error exists, the code above branches around some error handling code that the author places in the middle of the routine. Unfortunately, this interrupts the flow of the program. Most readers will want to see a straight-line version of the program's typical operation without having to worry about details concerning error conditions. Unfortunately, the organization of this code is such that the user must skip over seldomly-executed code in order to follow what is happening with the common case[4].

Here is a slightly improved version of the above program:

```
;Get all file names matching filespec and set up tables


GetFileRecords  mov     dx, offset DTA          ;Set up DTA
                DOS     SetDTA


; Get the first file that matches the specified filename (that may
; contain wildcard characters).  If no such file exists, then
; we've got an error.


                mov     dx, FileSpec
                mov     cl, 37h
                DOS     FindFirstFile
                jc      FileNotFound


; As long as there are no more files matching our file spec (that contains
; wildcard characters), get the file information and place it in the
; "files" array.  Each time through the "StoreFileName" loop we've got
```

_____

4. Note, by the way, that I am not suggesting that this error checking/handling code should be absent from the program. I am only suggesting that it not interrupt the normal flow of the program while reading the code.

```
; a new file name via a call to DOS' FindNextFile function (FindFirstFile
; for the first iteration).  Store the info concerning the file away and
; move on to the next file.

                mov     di, offset fileRecords  ;DI -> storage for file names
                mov     bx, offset files        ;BX -> array of files
                sub     bx, 2                   ;Special case for 1st iteration
StoreFileName:  add     bx, 2
                cmp     bx, (offset files) + NFILES*2
                jae     TooManyFiles

; Store away the pointer to the status/filename in files[] array.
; Note that the H.O. bit of the status byte indicates that the file is
; is marked.

                mov     [bx], di                ;Store pointer in files[]
                mov     al, [DTAattrib]         ;Store status byte
                and     al, 3Fh                 ;Clear file is marked bit
                stosb

; Copy the filename from the DTA storage area to the space we've set aside.

                mov     si, offset DTAname
                call    CopyString
                inc     di                      ;Skip zero byte (???).

                mov     si, offset DTAtime      ;Copy time, date and size
                mov     cx, 4
        rep     movsw

; Move on to the next file and try again.

                DOS     FindNextFile
                jnc     StoreFileName

; After processing the last file entry, do some clean up.
; (1) Save pointer to last file entry.
; (2) If returning from EXEC, we may need to resort and display the files.

                mov     [last], bx
                mov     al, [keepSorted]
                or      al, al
                jz      DisplayFiles
                jmp     Sort0

; Jump down here if there were no files to process.

FileNotFound:   mov     si, offset NoFilesMsg
                call    Error
                jmp     NewFilespec
```

```
; Jump down here if there were too many files to process.

TooManyFiles:   sub     bx, 2
                mov     [last], bx
                mov     si, offset tooManyMsg
                jmp     DoError
```

This improved version dispenses with the local labels, formats the code better by aligning all the statement fields and inserting blank lines into the code. It also eliminates much of the uppercase characters appearing in the previous version. Another improvment is that this code moves the error handling code out of the main stream of this code segment, allowing the reader to follow the typical execution in a more linear fashion.

1.4   Intended Audience

Of course, an assembly language program is going to be nearly unreadable to someone who doesn't know assembly language. This is true for almost any programming language. In the examples above, it's doubtful that the "improved" versions are really any more readable than the original version if you don't know 80x86 assembly language. Perhaps the improved versions are more aesthetic in a generic sense, but if you don't know 80x86 assembly language it's doubtful you'd make any more sense of the second version than the first. Other than burying a tutorial on 80x86 assembly language in a program's comments, there is no way to address this problem[5].

In view of the above, it makes sense to define an "intended audience" that we intend to have read our assembly language programs. Such a person should:

·   Be a reasonably competent 80x86 assembly language programmer.
·   Be reasonably familiar with the problem the assembly language program is attempting to solve.
·   Fluently read English[6].
·   Have a good grasp of high level language concepts.
·   Possess appropriate knowledge for someone working in the field of Computer Science (e.g., understands standard algorithms and data structures, understands basic machine architecture, and understands basic discrete mathmatics).

1.5   Readability Metrics

One has to ask "What is it that makes one program more readable than another?" In other words, how do we measure the "readability" of a program? The usual metric, "I know a well-written program when I see one" is inappropriate; for most people, this translates to "If your programs look like my better programs then they are readable, otherwise they are not." Obviously, such a metric is of little value since it changes with every person.

To develop a metric for measuring the readability of an assembly language program, the first thing we must ask is "Why is readability important?" This question has a simple (though somewhat flippant) answer: Readability is important because programs are *read* (furthermore, a line of code is typically read ten times more often than it is written). To expand on this, consider the fact that most programs are read and maintained by other programmers (Steve McConnell claims that up to ten generations of maintenance programmers work on a typically real world program before it is rewritten; furthermore, they spend up to 60% of their effort on that code simply figuring out how it works). The more readable your programs are, the less time these other people will have to spend figuring out what your program does. Instead, they can concentrate on adding features or correcting defects in the code.

For the purposes of this document, we will define a "readable" program as one that has the following trait:

·   A "readable" program is one that a competent programmer (one who is familiar with the problem the program is

---

5. Doing so (inserting an 80x86 tutorial into your comments) would wind up making the program *less* readable to those who already know assembly language since, at the very least, they'd have to skip over this material; at the worst they'd have to read it (wasting their time).
6. Or whatever other natural language is in use at the site(s) where you develop, maintain, and use the software.

attempting to solve) can pick up, without ever having seen the program before, and fully comprehend the entire program in a minimal amount of time.

That's a tall order! This definition doesn't sound very difficult to achieve, but few non-trivial programs ever really achieve this status. This definition suggests that an appropriate programmer (i.e., one who is familiar with the problem the program is trying to solve) can pick up a program, read it at their normal reading pace (just once), and fully comprehend the program. Anything less is not a "readable" program.

Of course, in practice, this definition is unusable since very few programs reach this goal. Part of the problem is that programs tend to be quite long and few human beings are capable of managing a large number of details in their head at one time. Furthermore, no matter how well-written a program may be, "a competent programmer" does not suggest that the programmer's IQ is so high they can read a statement a fully comprehend its meaning without expending much thought. Therefore, we must define readabilty, not as a boolean entity, but as a scale. Although truly unreadable programs exist, there are many "readable" programs that are less readable than other programs. Therefore, perhaps the following definition is more realistic:

·   A readable program is one that consists of one or more *modules*. A competent program should be able to pick a given module in that program and achieve an 80% comprehension level by expending no more than an average of one minute for each statement in the program.

An 80% comprehension level means that the programmer can correct bugs in the program and add new features to the program without making mistakes due to a misunderstanding of the code at hand.

1.6   How to Achieve Readability

The "I'll know one when I see one" metric for readable programs provides a big hint concerning how one should write programs that are readable. As pointed out early, the "I'll know it when I see it" metric suggests that an individual will consider a program to be readable if it is very similar to (good) programs that this particular person has written. This suggests an important trait that readable programs must possess: consistency. If all programmers were to write programs using a consistent style, they'd find programs written by others to be similar to their own, and, therefore, easier to read. This single goal is the primary purpose of this paper - to suggest a consistent standard that everyone will follow.

Of course, consistency by itself is not good enough. Consistently bad programs are not particularly easy to read. Therefore, one must carefully consider the guidelines to use when defining an all-encompassing standard. The purpose of this paper is to create such a standard. However, don't get the impression that the material appearing in this document appears simply because it sounded good at the time or because of some personal preferences. The material in this paper comes from several software engineering texts on the subject (including Elements of Programming Style, Code Complete, and Writing Solid Code), nearly 20 years of personal assembly language programming experience, and a set of generic programming guidelines developed for Information Management Associates, Inc.

This document assumes consistent usage by its readers. Therefore, it concentrates on a lot of mechanical and psychological issues that affect the readability of a program. For example, uppercase letters are harder to read than lower case letters (this is a well-known result from psychology research). It takes longer for a human being to recognize uppercase characters, therefore, an average human being will take more time to read text written all in upper case. Hence, this document suggests that one should avoid the use of uppercase sequences in a program. Many of the other issues appearing in this document are in a similar vein; they suggest minor changes to the way you might write your programs that make it easier for someone to recognize some pattern in your code, thus aiding in comprehension.

1.7   How This Document is Organized

This document follows a top-down discussion of readability. It starts with the concept of a program. Then it discusses modules. From there it works its way down to procedures. Then it talks about individual statements. Beyond that, it talks about components that make up statements (e.g., instructions, names, and operators). Finally, this paper concludes by discussing some orthogonal issues.

Section Two discusses programs in general. It primarily discusses documentation that must accompany a program and the organization of source files. It also discusses, briefly, configuration management and source code control issues. Keep in mind

that figuring out how to build a program (make, assemble, link, test, debug, etc.) is important. If your reader fully understands the "heapsort" algorithm you are using, but cannot build an executable module to run, they still do not fully understand your program.

Section Three discusses how to organize modules in your program in a logical fashion. This makes it easier for others to locate sections of code and organizes related sections of code together so someone can easily find important code and ignore unimportant or unrelated code while attempting to understand what your program does.

Section Four discusses the use of procedures within a program. This is a continuation of the theme in Section Three, although at a lower, more detailed, level.

Section Five discusses the program at the level of the statement. This (large) section provides the meat of this proposal. Most of the rules this paper presents appear in this section.

Section Six discusses those items that make up a statement (labels, names, instructions, operands, operators, etc.) This is another large section that presents a large number of rules one should follow when writing readable programs. This section discusses naming conventions, appropriateness of operators, and so on.

Section Seven discusses data types and other related topics.

Section Eight covers miscellaneous topics that the previous sections did not cover.

1.8   Guidelines, Rules, Enforced Rules, and Exceptions

Not all rules are equally important. For example, a rule that you check the spelling of all the words in your comments is probably less important than suggesting that the comments all be in English[7]. Therefore, this paper uses three designations to keep things straight: Guidelines, Rules, and Enforced Rules.

A Guideline is a suggestion. It is a rule you should follow unless you can verbally defend why you should break the rule. As long as there is a good, defensible, reason, you should feel no apprehension violated a guideline. Guidelines exist in order to encourage consistency in areas where there are no good reasons for choosing one methodology over another. You shouldn't violate a Guideline just because you don't like it -- doing so will make your programs inconsistent with respect to other programs that do follow the Guidline (and, therefore, harder to read -- however, you shouldn't lose any sleep because you violated a Guideline.

Rules are much stronger than Guidelines. You should never break a rule unless there is some external reason for doing so (e.g., making a call to a library routine forces you to use a bad naming convention). Whenever you feel you must violate a rule, you should verify that it is reasonable to do so in a peer review with at least two peers. Furthermore, you should explain in the program's comments why it was necessary to violate the rule. Rules are just that -- rules to be followed. However, there are certain situations where it may be necessary to violate the rule in order to satisfy external requirements or even make the program more readable.

Enforced Rules are the toughest of the lot. You should *never* violate an enforced rule. If there is ever a true need to do this, then you should consider demoting the Enforced Rule to a simple Rule rather than treating the violation as a reasonable alternative.

An Exception is exactly that, a known example where one would commonly violate a Guideline, Rule, or (very rarely) Enforced Rule. Although exceptions are rare, the old adage "Every rule has its exceptions..." certainly applies to this document. The Exceptions point out some of the common violations one might expect.

Of course, the categorization of Guidelines, Rules, Enforced Rules, and Exceptions herein is one man's opinion. At some organizations, this categorization may require reworking depending on the needs of that organization.

1.9   Source Language Concerns

This document will assume that the entire program is written in 80x86 assembly language. Although this organization is rare in commercial applications, this assumption will, in no way, invalidate these guidelines. Other guidelines exist for various

---

7. You may substitute the local language in your area if it is not English.

high level languages (including a set written by this paper's author). You should adopt a reasonable set of guidelines for the other languages you use and apply these guidelines to the 80x86 assembly language modules in the program.

2.0   Program Organization

A source program generally consists of one or more source, object, and library files. As a project gets larger and the number of files increases, it becomes difficult to keep track of the files in a project. This is especially true if a number of different projects share a common set of source modules. This section will address these concerns.

2.1   Library Functions

A library, by its very nature, suggests stability. Ignoring the possibility of software defects, one would rarely expect the number or function of routines in a library to vary from project to project. A good example is the "UCR Standard Library for 80x86 Assembly Language Programmers." One would expect "printf" to behave identically in two different programs that use the Standard Library. Contrast this against two programs, each of which implement their own version of *printf*. One could not reasonably assume both programs have identical implementations[8]. This leads to the following rule:

Rule: Library functions are those routines intended for common reuse in many different assembly language programs. All assembly language (callable) libraries on a system should exist as ".lib" files and should appear in the "/lib" or "/asmlib" subdirectory.

Guideline: "/asmlib" is probably a better choice if you're using multiple languages since those other languages may need to put files in a "/lib" directory.

Exception: It's probably reasonable to leave the UCR Standard Library's "stdlib.lib" file in the "/stdlib/lib" directory since most people expect it there.

The rule above ensures that the library files are all in one location so they are easy to find, modify, and review. By putting all your library modules into a single directory, you avoid configuration management problems such as having outdated versions of a library linking with one program and up-to-date versions linking with other programs.

2.2   Common Object Modules

This document defines a *library* as a collection of object modules that have wide application in many different programs. The UCR Standard Library is a typical example of a library. Some object modules are not so general purpose, but still find application in two or more different programs. Two major configuration management problems exist in this situation: (1) making sure the ".obj" file is up-to-date when linking it with a program; (2) Knowing which modules use the module so one can verify that changes to the module won't break existing code.

The following rules takes care of case one:

Rule: If two different program share an object module, then the associated source, object, and makefiles for that module should appear in a subdirectory that is specific to that module (i.e., no other files in the subdirectory). The subdirectory name should be the same as the module name. If possible, you should create a set of link/alias/shortcuts to this subdirectory and place these links in the main directory of each of the projects that utilize the module. If links are not possible, you should place the module's subdirectory in the "/common" subdirectory.

Enforced Rule: Every subdirectory containing one or more modules should have a make file that will automatically generate the appropriate, up-to-date, ".obj" files. An individual, a batch file, or another make file should be able to automatically generate new object modules (if necessary) by simply executing the make program.

Guideline: Use Microsoft's nmake program. At the very least, use nmake acceptable syntax in your makefiles.

_____

8. In fact, just the opposite is true. One should get concerned if both implementations *are* identical. This would suggest poor planning on the part of the program's author(s) since the same routine must now be maintained in two different programs.

The other problem, noting which projects use a given module is much more difficult. The obvious solution, commenting the source code associated with the module to tell the reader which programs use the module, is impractical. Maintaining these comments is too error-prone and the comments will quickly get out of phase and be worse than useless -- they would be incorrect. A better solution is to create a dummy file using the module's name with a ".elw" (elsewhere) suffix and placing this file in the main subdirectory of each program that links the module. Now, using one of the venerable "whereis" programs, you can easily locate all projects that use the module.

Guideline: If a project uses a module that is not local to the project's subdirectory, create a dummy file (using "TOUCH" or a comparable program) that uses the module's main name with a ".elw" suffix. This will allow someone to easily search for all the projects that use a common object module by using a "whereis" program.

2.3   Local Modules

Local modules are those that a single program/project uses. Typically, the source and object code for each module appears in the same directory as the other files associated with the project. This is a reasonable arrangement until the number of files increases to the point that it is difficult to find a file in a directory listing. At that point, most programmers begin reorganizing their directory by creating subdirectories to hold many of these source modules. However, the placement, name, and contents of these new subdirectories can have a big impact on the overall readability of the program. This section will address these issues.

The first issue to consider is the contents of these new subdirectories. Since programmers rummaging through this project in the future will need to easily locate source files in a project, it is important that you organize these new subdirectories so that it is easy to find the source files you are moving into them. The best organization is to put each source module (or a small group of *strongly related* modules) into its own subdirectory. The subdirectory should bear the name of the source module minus its suffix (or the main module if there is more than one present in the subdirectory). If you place two or more source files in the same directory, ensure this set of source files forms a *cohesive* set (meaning the source files contain code that solve a single problem). A discussion of cohesiveness appears later in this document.

Rule: If a project directory contains too many files, try to move some of the modules to subdirectories within the project directory; give the subdirectory the same name as the source file without the suffix. This will nearly reduce the number of files in half. If this reduction is insufficient, try categorizing the source modules (e.g., FileIO, Graphics, Rendering, and Sound) and move these modules to a subdirectory bearing the name of the category.

Enforced Rule: Each new subdirectory you create should have its own make file that will automatically assemble all source modules within that subdirectory, as appropriate.

Enforced Rule: Any new subdirectories you create for these source modules should appear within the directory containing the project. The only excepts are those modules that are, or you anticipate, sharing with other projects. See "Common Object Modules" on page 13 for more details.

Stand-alone assembly language programs generally contain a "main" procedure – the first program unit that executes when the operating system loads the program into memory. For any programmer new to a project, this procedure is the *anchor* where one first begins reading the code and the point where the reader will continually refer. Therefore, the reader should be able to easily locate this source file. The following rule helps ensure this is the case:

Rule: The source module containing the *main program* should have the same name as the executable (obviously the suffix will be different). For example, if the "Simulate 886" program's executable name is "Sim886.exe" then you should find the main program in the "Sim886.asm" source file.

Finding the souce file that contains the main program is one thing. Finding the main program itself can be almost as hard. Assembly language lets you give the main program any name you want. However, to make the main procedure easy to find (both in the source code and at the O/S level), you should actually name this program "main". See "Module Organization" on page 15 for more details about the placement of the main program.

Rule: The name of the main procedure in an assembly language program should be "main".

2.4   Program Make Files

Every project, even if it contains only a single source module, should have an associated make file. If someone want to assemble your program, they should have to worry about what program (e.g., MASM) to use, what command line options to use, what library modules to use, etc. They should be able to type "nmake"[9] and wind up with an executable program. Even if assembling the program consists of nothing more than typing the name of the assembler and the source file, you should still have a make file. Someone else may not realize that's all that is necessary.

Enforced Rule: The main project directory should contain a make file that will automatically generate an executable (or other expected object module) in response to a simple make/nmake command.

Rule: If your project uses object modules that are not in the same subdirectory as the main program's module, you should test the ".obj" files for those modules and execute the corresponding make files in their directories if the object code is out of date. You can assume that library files are up to date.

Guideline: Avoid using fancy "make" features. Most programmers only learn the basics about make and will not be able to understand what your make file is doing if you fully exploit the make language. Especially avoid the use of default rules since this can create havoc if someone arbitrarily adds or removes files from the directory containing the make file.

3.0   Module Organization

A module is a collection of objects that are logically related. Those objects may include constants, data types, variables, and program units (e.g., functions, procedures, etc.). Note that objects in a module need not be *physically* related. For example, it is quite possible to construct a module using several different source files. Likewise, it is quite possible to have several different modules in the same source file. However, the best modules are physically related as well as logically related; that is, all the objects associated with a module exist in a single source file (or directory if the source file would be too large) and nothing else is present.

Modules contain several different objects including constants, types, variables, and program units (routines). Modules shares many of the attributes with routines (program units); this is not surprising since routines are the major component of a typical module. However, modules have some additional attributes of their own. The following sections describe the attributes of a well-written module.

Note: *Unit* and *package* are both synonyms for the term *module*.

3.1   Module Attributes

A module is a generic term that describes a set of program related objects (program units as well as data and type objects) that are somehow coupled. Good modules share many of the same attributes as good program units as well as the ability to hide certain details from code outside the module.

3.1.1   Module Cohesion

Modules exhibit the following different kinds of *cohesion* (listed from good to bad):

·   Functional or logical cohesion exists if the module accomplishes exactly one (simple) task.
·   Sequential or *pipelined* cohesion exists when a module does several sequential operations that must be performed in a certain order with the data from one operation being fed to the next in a "filter-like" fashion.
·   Global or *communicational* cohesion exists when a module performs a set of operations that make use of a common set of data, but are otherwise unrelated.
·   Temporal cohesion exists when a module performs a set of operations that need to be done at the same time (though not necessarily in the same order). A typical initialization module is an example of such code.
·   Procedural cohesion exists when a module performs a sequence of operations in a specific order, but the only thing

---

9. Or whatever make program you normally use.

that binds them together is the order in which they must be done. Unlike sequential cohesion, the operations do not share data.

·   State cohesion occurs when several different (unrelated) operations appear in the same module and a state variable (e.g., a parameter) selects the operation to execute. Typically such modules contain a case (switch) or **if..elseif..elseif...** statement.

·   No cohesion exists if the operations in a module have no apparent relationship with one another.

The first three forms of cohesion above are generally acceptable in a program. The fourth (temporal) is probably okay, but you should rarely use it. The last three forms should almost never appear in a program. For some reasonable examples of module cohesion, you should consult "Code Complete".

Guideline: Design good modules! *Good modules exhibit strong cohesion*. That is, a module should offer a (small) group of services that are logically related. For example, a "printer" module might provide all the services one would expect from a printer. The individual routines within the module would provide the individual services.

3.1.2   Module Coupling

Coupling refers to the way that two modules communicate with one another. There are several criteria that define the level of coupling between two modules:

·   Cardinality- the number of objects communicated between two modules. The fewer objects the better (i.e., fewer parameters).

·   Intimacy- how "private" is the communication? Parameter lists are the most private form; private data fields in a class or object are next level; public data fields in a class or object are next, global variables are even less intimate, and passing data in a file or database is the least intimate connection. Well-written modules exhibit a high degree of intimacy.

·   Visibility- this is somewhat related to intimacy above. This refers to how visible the data is to the entire system that you pass between two modules. For example, passing data in a parameter list is direct and very visible (you always see the data the caller is passing in the call to the routine); passing data in global variables makes the transfer less visible (you could have set up the global variable long before the call to the routine). Another example is passing simple (scalar) variables rather than loading up a bunch of values into a structure/record and passing that structure/record to the callee.

·   Flexibility- This refers to how easy it is to make the connection between two routines that may not have been originally intended to call one another. For example, suppose you pass a structure containing three fields into a function. If you want to call that function but you only have three data objects, not the structure, you would have to create a dummy structure, copy the three values into the field of that structure, and then call the function. On the other hand, had you simply passed the three values as separate parameters, you could still pass in structures (by specifying each field) as well as call the function with separate values. The module containing this later function is more flexible.

A module is *loosely coupled* if its functions exhibit low cardinality, high intimacy, high visibility, and high flexibility. Often, these features are in conflict with one another (e.g., increasing the flexibility by breaking out the fields from a structures [a good thing] will also increase the cardinality [a bad thing]). It is the traditional goal of any engineer to choose the appropriate compromises for each individual circumstance; therefore, you will need to carefully balance each of the four attributes above.

A module that uses loose coupling generally contains fewer errors per KLOC (thousands of lines of code). Furthermore, modules that exhibit loose coupling are easier to reuse (both in the current and future projects). For more information on coupling, see the appropriate chapter in "Code Complete".

Guideline: Design good modules! *Good modules exhibit loose coupling*. That is, there are only a few, well-defined (visible) interfaces between the module and the outside world. Most data is private, accessible only through accessor functions (see information hiding below). Furthermore, the interface should be flexible.

Guideline: Design good modules! *Good modules exhibit information hiding*. Code outside the module should only have access to the module through a small set of public routines. All data should be private to that mod-

ule. A module should implement an *abstract data type*. All interface to the module should be through a well-defined set of operations.

3.1.3   Physical Organization of Modules

Many languages provide direct support for modules (e.g., packages in Ada, modules in Modula-2, and units in Delphi/Pascal). Some languages provide only indirect support for modules (e.g., a source file in C/C++). Others, like BASIC, don't really support modules, so you would have to simulate them by physically grouping objects together and exercising some discipline. Assembly language falls into the middle ground. The primary mechanism for hiding names from other modules is to implement a module as an individual source file and publish only those names that are part of the module's interface to the outside world.

Rule: Each module should completely reside in a single source file. If size considerations prevent this, then all the source files for a given module should reside in a subdirectory specifically designated for that module.

Some people have the crazy idea that modularization means putting each function in a separate source file. Such *physical modularization* generally impairs the readability of a program more than it helps. Strive instead for *logical modularization*, that is, defining a module by its actions rather than by source code syntax (e.g., separating out functions).

This document does not address the decomposition of a problem into its modular components. Presumably, you can already handle that part of the task. There are a wide variety of texts on this subject if you feel weak in this area.

3.1.4   Module Interface

In any language system that supports modules, there are two primary components of a module: the interface component that publicizes the module visible names and the implementation component that contains the actual code, data, and private objects. MASM (and most assemblers) uses a scheme that is very similar to the one C/C++ uses. There are directives that let you import and export names. Like C/C++, you could place these directives directly in the related source modules. However, such code is difficult to maintain (since you need to change the directives in every file whenever you modify a public name). The solution, as adopted in the C/C++ programming languages, is to use *header files*. Header files contain all the public definitions and exports (as well as common data type definitions and constant definitions). The header file provides the *interface* to the other modules that want to use the code present in the implementation module.

The MASM 6.x *externdef* directive is perfect for creating interface files. When you use externdef within a source module that defines a symbol, externdef behaves like the *public* directive, exporting the name to other modules. When you use externdef within a source modules that refers to an external name, externdef behaves like the *extern* (or *extrn*) directive. This lets you place an externdef directive in a single file and include this file into both the modules that import and export the public names.

If you are using an assembler that does not support externdef, you should probably consider switching to MASM 6.x. If switching to a better assembler (that supports externdef) is not feasible, the last thing you want to do is have to maintain the interface information in several separate files. Instead, use the assembler's *ifdef conditional assembly* directives to assemble a set of public statements in the header file if a symbol with the module's name is defined prior to including the header file. It should assemble a set of extrn statements otherwise. Although you still have to maintain the public and external information in two places (in the ifdef true and false sections), they are in the same file and located near one another.

Rule: Keep all module interface directives (public, extrn, extern, and externdef) in a single header file for a given module. Place any other common data type definitions and constant definitions in this header file as well.

Guideline: There should only be a single header file associated with any one module (even if the module has multiple source files associated with it). If, for some reason, you feel it is necessary to have multiple header files associated with a module, you should create a single file that includes all of the other interface files. That way a program that wants to use all the header files need only include the single file.

When designing header files, make sure you can include a file more than once without ill effects (e.g., duplicate symbol errors). The traditional way to do this is to put an IFDEF statement like the following around all the statements in a header file:

```
; Module: MyHeader.a

            ifndef   MyHeader_A
MyHeader_A      =       0
            .
            .       ;Statements in this header file.
            .
            endif
```

The first time a source file includes "MyHeader.a" the symbol "MyHeader_A" is undefined. Therefore, the assembler will process all the statements in the header file. In successive include operations (during the same assembly) the symbol "MyHeader_A" is already defined, so the assembler ignores the body of the include file.

My would you ever include a file twice? Easy. Some header files may include other header files. By including the file "YourHeader.a" a module might also be including "MyHeader.a" (assuming "YourHeader.a" contains the appropriate include directive). Your main program, that includes "YourHeader.a" might also need "MyHeader.a" so it explicitly includes this file not realizing "YourHeader.a" has already processed "MyHeader.a" thereby causing symbol redefinitions.

Rule:	Always put an appropriate IFNDEF statement around all the definitions in a header file to allow multiple inclusion of the header file without ill effect.
Guideline:	Use the ".a" suffix for assembly language header/interface files.
Rule:	Include files for library functions on a system should exist as ".a" files and should appear in the "/include" or "/asminc" subdirectory.
Guideline:	"/asminc" is probably a better choice if you're using multiple languages since those other languages may need to put files in a "/include" directory.
Exception:	It's probably reasonable to leave the UCR Standard Library's "stdlib.a" file in the "/stdlib/include" directory since most people expect it there.

4.0   Program Unit Organization

A program unit is any procedure, function, coroutine, iterator, subroutine, subprogram, routine, or other term that describes a section of code that abstracts a set of common operations on the computer. This text will simply use the term *procedure* or *routine* to describe these concepts.

Routines are closely related to modules, since they tend to be the major component of a module (along with data, constants, and types). Hence, many of the attributes that apply to a module also apply to routines. The following paragraphs, at the expense of being redundant, repeat the earlier definitions so you don't have to flip back to the previous sections.

4.1   Routine Cohesion

Routines exhibit the following kinds of *cohesion* (listed from good to bad):

· Functional or logical cohesion exists if the routine accomplishes exactly one (simple) task.
· Sequential or *pipelined* cohesion exists when a routine does several sequential operations that must be performed in a certain order with the data from one operation being fed to the next in a "filter-like" fashion.
· Global or *communicational* cohesion exists when a routine performs a set of operations that make use of a common set of data, but are otherwise unrelated.
· Temporal cohesion exists when a routine performs a set of operations that need to be done at the same time (though not necessarily in the same order). A typical initialization routine is an example of such code.
· Procedural cohesion exists when a routine performs a sequence of operations in a specific order, but the only thing that binds them together is the order in which they must be done. Unlike sequential cohesion, the operations do not share data.
· State cohesion occurs when several different (unrelated) operations appear in the same module and a state variable

(e.g., a parameter) selects the operation to execute. Typically such routines contain a case (switch) or **if..elseif..elseif...** statement.

· No cohesion exists if the operations in a routine have no apparent relationship with one another.

The first three forms of cohesion above are generally acceptable in a program. The fourth (temporal) is probably okay, but you should rarely use it. The last three forms should almost never appear in a program. For some reasonable examples of routine cohesion, you should consult "Code Complete".

Guideline: All routines should exhibit good cohesiveness. Functional cohesiveness is best, followed by sequential and global cohesiveness. Temporal cohesiveness is okay on occasion. You should avoid the other forms.

4.1.1 Routine Coupling

Coupling refers to the way that two routines communicate with one another. There are several criteria that define the level of coupling between two routines:

· Cardinality- the number of objects communicated between two routines. The fewer objects the better (i.e., fewer parameters).

· Intimacy- how "private" is the communication? Parameter lists are the most private form; private data fields in a class or object are next level; public data fields in a class or object are next, global variables are even less intimate, and passing data in a file or database is the least intimate connection. Well-written routines exhibit a high degree of intimacy.

· Visibility- this is somewhat related to intimacy above. This refers to how visible the data is to the entire system that you pass between two routines. For example, passing data in a parameter list is direct and very visible (you always see the data the caller is passing in the call to the routine); passing data in global variables makes the transfer less visible (you could have set up the global variable long before the call to the routine). Another example is passing simple (scalar) variables rather than loading up a bunch of values into a structure/record and passing that structure/record to the callee.

· Flexibility- This refers to how easy it is to make the connection between two routines that may not have been originally intended to call one another. For example, suppose you pass a structure containing three fields into a function. If you want to call that function but you only have three data objects, not the structure, you would have to create a dummy structure, copy the three values into the field of that structure, and then call the routine. On the other hand, had you simply passed the three values as separate parameters, you could still pass in structures (by specifying each field) as well as call the routine with separate values.

A function is *loosely coupled* if it exhibits low cardinality, high intimacy, high visibility, and high flexibility. Often, these features are in conflict with one another (e.g., increasing the flexibility by breaking out the fields from a structures [a good thing] will also increase the cardinality [a bad thing]). It is the traditional goal of any engineer to choose the appropriate compromises for each individual circumstance; therefore, you will need to carefully balance each of the four attributes above.

A program that uses loose coupling generally contains fewer errors per KLOC (thousands of lines of code). Furthermore, routines that exhibit loose coupling are easier to reuse (both in the current and future projects). For more information on coupling, see the appropriate chapter in "Code Complete".

Guideline: Coupling between routines in source code should be loose.

4.1.2 Routine Size

Sometime in the 1960's, someone decided that programmers could only look at one page in a listing at a time, therefore routines should be a maximum of one page long (66 lines, at the time). In the 1970's, when interactive computing became popular, this was adjusted to 24 lines -- the size of a terminal screen. In fact, there is very little empirical evidence to suggest that small routine size is a good attribute. In fact, several studies on code containing artificial constraints on routine size indicate just the opposite -- shorter routines often contain more bugs per KLOC[10].

10.This happens because shorter function invariable have stronger coupling, leading to integration errors.

Style Guidelines for Assembly Language Programmers

A routine that exhibits functional cohesiveness is the right size, almost regardless of the number of lines of code it contains. You shouldn't artificially break up a routine into two or more subroutines (e.g., sub_partI and sub_partII) just because you feel a routine is getting to be too long. First, verify that your routine exhibits strong cohesion and loose coupling. If this is the case, the routine is not too long. Do keep in mind, however, that a long routine is probably a good indication that it is performing several actions and, therefore, does not exhibit strong cohesion.

Of course, you can take this too far. Most studies on the subject indicate that routines in excess of 150-200 lines of code tend to contain more bugs and are more costly to fix than shorter routines. Note, by the way, that you do not count blank lines or lines containing only comments when counting the lines of code in a program.

Also note that most studies involving routine size deal with HLLs. A comparable assembly language routine will contain more lines of code than the corresponding HLL routine. Therefore, you can expect your routines in assembly language to be a little longer.

Guideline: Do not let artificial constraints affect the size of your routines. If a routine exceeds about 200-250 lines of code, make sure the routine exhibits functional or sequential cohesion. Also look to see if there aren't some generic subsequences in your code that you can turn into stand alone routines.

Rule: Never shorten a routine by dividing it into *n* parts that you would always call in the appropriate sequence as a way of shortening the original routine.

4.2   Placement of the Main Procedure and Data

As noted earlier, you should name the main procedure *main* and place it in the source file bearing the same name as the executable file. If this module is rather long, it can still be difficult to locate the main program. A good solution is to always place the main procedure at the same point in the source file. By convention (meaning everyone expects it this way), most programmers make their main program the first or last procedure in an assembly language program. Either position is fine. Putting the main program anywhere else makes it hard to find.

Rule: Always make the main procedure the first or last procedure in a source file.

MASM, because it is a *multiphase* assembler, does not require that you define a symbol before you use it. This is necessary because many instructions (like JMP) need to refer to symbols found later in the program. In a similar manner, MASM doesn't really care where you define your data – before or after its use[11]. However, most programmers "grew up" with high level languages that require the definition of a symbol before its first use. As a result, they expect to be able to find a variable declaration by looking backwards in the source file. Since everyone expects this, it is a good tradition to continue in an assembly language program.

Rule: You should declare all variables, constants, and macros prior to their use in an assembly language program.

Rule: You should define all static variables (those you declare in a segment) at the beginning of the source module.

5.0   Statement Organization

In an assembly language program, the author must work extra hard to make a program readable. By following a large number of rules, you can produce a program that is readable. However, by breaking a single rule *no matter how many other rules you've followed,* you can render a program unreadable. Nowhere is this more true than how you organize the statements within your program. Consider the following example taken from "The Art of Assembly Language Programming":

The Microsoft Macro Assembler is a free form assembler. The various fields of an assembly language statement may appear in any column (as long as they appear in the proper order). Any number of spaces or tabs can separate the various fields in the statement. To the assembler, the following two code sequences are identical:

---

11. Technically, this is incorrect. In some very special cases MASM will generate better machine code if you define your variables before you use them in a program.

```
                        mov     ax, 0
                        mov     bx, ax
                        add     ax, dx
                        mov     cx, ax
```

```
mov             ax,                     0
        mov bx,                     ax
    add             ax, dx
                        mov             cx, ax
```

The first code sequence is much easier to read than the second (if you don't think so, perhaps you should go see a doctor!). With respect to readability, the judicial use of spacing within your program can make all the difference in the world.

While this is an extreme example, do note that it only takes a few mistakes to have a large impact on the readability of a program. Consider (a short section from) an example presented earlier:

```
GetFileRecords:
    mov dx, OFFSET DTA              ;Set up DTA
    mov ah, 1Ah
    int 21h
    mov dx, FILESPEC               ;Get first file name
    mov cl, 37h
    mov ah, 4Eh
    int 21h
    jnc FileFound                  ;No files.  Try a different filespec.
    mov si, OFFSET NoFilesMsg
    call Error
    jmp NewFilespec
FileFound:
    mov di, OFFSET fileRecords     ;DI -> storage for file names
    mov bx, OFFSET files           ;BX -> array of files
    sub bx, 2
```

Improved version:

```
GetFileRecords: mov     dx, offset DTA              ;Set up DTA
                DOS     SetDTA

                mov     dx, FileSpec
                mov     cl, 37h
                DOS     FindFirstFile
                jc      FileNotFound

                mov     di, offset fileRecords  ;DI -> storage for file names
                mov     bx, offset files        ;BX -> array of files
                sub     bx, 2                   ;Special case for 1st iteration
```

Style Guidelines for Assembly Language Programmers

An assembly language statement consists of four possible *fields*: a label field, a mnemonic field, an operand field, and a comment file. The mnemonic and comment fields are always optional. The label field is generally optional although certain instructions (mnemonics) do not allow labels while others require labels. The operand field's presence is tied to the mnemonic field. For most instructions the actual mnemonic determines whether an operand field must be present.

MASM is a free-form assembler insofar as it does not require these fields to appear in any particular column[12]. However, the freedom to arrange these columns in any manner is one of the primary contributors to hard to read assembly language programs. Although MASM lets you enter your programs in free-form, there is absolutely no reason you cannot adopt a fixed field format, always starting each field in the same column. Doing so generally helps make an assembly language program much easier to read. Here are the rules you should use:

Rule: If an identifier is present in the label field, always start that identifier in column one of the source line.

Rule: All mnemonics should start in the same column. Generally, this should be column 17 (the second tab stop) or some other convenient position.

Rule: All operands should start in the same column. Generally, this should be column 25 (the third tab stop) or some other convenient position.

Exception: If a mnemonic (typically a macro) is longer than seven characters and requires an operand, you have no choice but to start the operand field beyond column 25 (this is an exception assuming you've chosen columns 17 and 25 for your mnemonic and operand fields, respectively).

Guideline: Try to always start the comment fields on adjacent source lines in the same column (note that it is impractical to always start the comment field in the same column throughout a program).

Most people learn a high level language prior to learning assembly language. They have been firmly taught that readable (HLL) programs have their control structures properly indented to show the structure of the program. Indentation works great when you have a *block structured* language. Assembly language, however, is the original unstructured language and indentation rules for a structured programming language simply do not apply. While it *is* important to be able to denote that a certain sequence of instructions is special (e.g., belong to the "then" portion of an if..then..else..endif statement), indentation is not the proper way to achieve this in an assembly language program.

If you need to set off a sequence of statements from surrounding code, the best thing you can do is use blank lines in your source code. For a small amount of detachment, to separate one computation from another for example, a single blank line is sufficient. To really show that one section of code is special, use two, three, or even four blank lines to separate one block of statements from the surrounding code. To separate two totally unrelated sections of code, you might use several blank lines and a row of dashes or asterisks to separate the statements. E.g.,

```
        mov     dx, FileSpec
        mov     cl, 37h
        DOS     FindFirstFile
        jc      FileNotFound


;       **********************************************


        mov     di, offset fileRecords  ;DI -> storage for file names
        mov     bx, offset files        ;BX -> array of files
        sub     bx, 2                   ;Special case for 1st iteration
```

Guideline: Use blank lines to separate special blocks of code from the surrounding code. Use an aesthetic looking row of asterisks or dashes if you need a stronger separation between two blocks of code (do not overdo this, however).

---

12. Older assemblers on other machines have required the labels to begin in column one, the mnemonic to appear in a specific column, the operand to appear in a specific column, etc. These were examples of *fixed-formant source line* translators.

If two sequences of assembly language statements correspond to roughly two HLL statements, it's generally a good idea to put a blank line between the two sequences. This helps clarify the two segments of code in the reader's mind. Of course, it is easy to get carried away and insert too much white space in a program, so use some common sense here.

Guideline: If two sequences of code in assembly language correspond to two adjacent statements in a HLL, then use a blank line to separate those two assembly sequences (assuming the sequences are real short).

A common problem in any language (not just assembly language) is a line containing a comment that is adjacent to one or two lines containing code. Such a program is very difficult read because it is hard to determine where the code ends and the comment begins (or vice-versa). This is especially true when the comments contain sample code. It is often quite difficult to determine if what you're looking at is code or comments; hence the following enforced rule:

Enforced Rule: Always put at least one blank line between code and comments (assuming, of course, the comment is sitting only a line by itself; that is, it is not an endline comment[13]).

6.0   Comments

Comments in an assembly language program generally come in two forms: *endline* comments and *standalone* comments[14]. As their names suggest, endline lines comments always occur at the end of a source statement and standalone comments sit on a line by themselves[15]. These two types of comments have distinct purposes, this section will explore their use and describe the attributes of a well-commented program.

6.1   What is a Bad Comment?

It is amazing how many programmers claim their code is well-commented. Were you to count characters between (or after) the comment delimiters, they might have a point. Consider, however, the following comment:

```
mov   ax, 0      ;Set AX to zero.
```

Quite frankly, this comment is worse than no comment at all. It doesn't tell the reader anything the instruction itself doesn't tell and it requires the reader to take some of his or her precious time to figure out that the comment is worthless. If someone cannot tell that this instruction is setting AX to zero, they have no business reading an assembly language program. This brings up the first guideline of this section:

Guideline: Choose an intended audience for your source code and write the comments to that audience. For assembly language source code, you can usually assume the target audience are those who know a reasonable amount of assembly language.

Don't explain the actions of an assembly language instruction in your code unless that instruction is doing something that isn't obvious (and most of the time you should consider changing the code sequence if it isn't obvious what is going on). Instead, explain how that instruction is helping to solve the problem at hand. The following is a much better comment for the instruction above:

```
mov   ax, 0      ;AX is the resulting sum.  Initialize it.
```

Note that the comment does not say "Initialize it to zero." Although there would be nothing intrinsically wrong with saying this, the phrase "Initialize it" remains true no matter what value you assign to AX. This makes maintaining the code (and comment) much easier since you don't have to change the comment whenever you change the constant associated with the instruction.

Guideline: Write your comments in such a way that minor changes to the instruction do not require that you change the corresponding comment.

---

13. See the next section concerning comments for more information.
14. This document will simply use the term *comments* when refering to standalone comments.
15. Since the label, mnemonic, and operand fields are all optional, it is legal to have a comment on a line by itself.

Style Guidelines for Assembly Language Programmers

Note: Although a trivial comment is bad (indeed, worse than no comment at all), the worst comment a program can have is one that is wrong. Consider the following statement:

```
mov    ax, 1      ;Set AX to zero.
```

It is amazing how long a typical person will look at this code trying to figure out how on earth the program sets AX to zero when it's obvious it does not do this. *People will always believe comments over code.* If there is some ambiguity between the comments and the code, they will assume that the code is tricky and that the comments are correct. Only after exhausting all possible options is the average person likely to concede that the comment must be incorrect.

Enforced Rule: Never allow incorrect comments in your program.

This is another reason not to put trivial comments like "Set AX to zero" in your code. As you modify the program, these are the comments most likely to become incorrect as you change the code and fail to keep the comments in sync. However, even some non-trivial comments can become incorrect via changes to the code. Therefore, always follow this rule:

Enforced Rule: Always update *all* comments affected by a code change immediately after making the code change.

Undoubtedly you've heard the phrase "make sure you comment your code as though someone else wrote it for you; otherwise in six months you'll wish you had." This statement encompasses two concepts. First, don't ever think that your understanding of the current code will last. While working on a given section of a program you're probably investing considerable thought and study to figure out what's going on. Six months down the road, however, you will have forgotten much of what you figured out and the comments can go a long way to getting you back up to speed quickly. The second point this code makes is the implication that others read and write code too. You will have to read someone else's code, they will have to read yours. If you write the comments the way you would expect others to write it for you, chances are pretty good that your comments will work for them as well.

Rule: Never use racist, sexist, obscene, or other exceptionally politically incorrect language in your comments. Undoubtedly such language in your comments will come back to embarass you in the future. Furthermore, it's doubtful that such language would help someone better understand the program.

It's much easier to give examples of bad comments than it is to discuss good comments. The following list describes some of the worst possible comments you can put in a program (from worst up to barely tolerable):

·   The absolute worst comment you can put into a program is an incorrect comment. Consider the following assembly statement:

```
    mov ax, 10;  { Set AX to 11 }
```

It is amazing how many programmers will automatically assume the comment is correct and try to figure out how this code manages to set the variable "A" to the value 11 when the code so obviously sets it to 10.

·   The second worst comment you can place in a program is a comment that explains what a statement is doing. The typical example is something like "mov ax, 10; { Set 'A' to 10 }". Unlike the previous example, this comment is correct. But it is still worse than no comment at all because it is redundant and forces the reader to spend additional time reading the code (reading time is directly proportional to reading difficulty). This also makes it harder to maintain since slight changes to the code (e.g., "mov ax, 9") requires modifications to the comment that would not otherwise be required.

·   The third worst comment in a program is an irrelevant one. Telling a joke, for example, may seem cute, but it does little to improve the readability of a program; indeed, it offers a distraction that breaks concentration.

·   The fourth worst comment is no comment at all.

·   The fifth worst comment is a comment that is obsolete or out of date (though not incorrect). For example, comments at the beginning of the file may describe the current version of a module and who last worked on it. If the last programmer to modify the file did not update the comments, the comments are now out of date.

6.2   What is a Good Comment?

Steve McConnell provides a long list of suggestions for high-quality code. These suggestions include:

·   **Use commenting styles that don't break down or discourage modification.** Essentially, he's saying pick a com-

menting style that isn't so much work people refuse to use it. He gives an example of a block of comments sur-rounded by asterisks as being hard to maintain. This is a poor example since modern text editors will automatically "outline" the comments for you. Nevertheless, the basic idea is sound.

· **Comment as you go along**. If you put commenting off until the last moment, then it seems like another task in the software development process always comes along and management is likely to discourage the completion of the commenting task in hopes of meeting new deadlines.

· **Avoid self-indulgent comments.** Also, you should avoid sexist, profane, or other insulting remarks in your com-ments. Always remember, someone else will eventually read your code.

· **Avoid putting comments on the same physical line as the statement they describe.** Such comments are very hard to maintain since there is very little room. McConnell suggests that endline comments are okay for variable declara-tions. For some this might be true but many variable declarations may require considerable explanation that simply won't fit at the end of a line. One exception to this rule is "maintenance notes." Comments that refer to a defect tracking entry in the defect database are okay (note that the CodeWright text editor provides a much better solution for this -- buttons that can bring up an external file). Of course, endline comments are marginally more useful in assembly language than in the HLLs that McConnell addresses, but the basic idea is sound.

· **Write comments that describe blocks of statements rather than individual statements.** Comments covering sin-gle statements tend to discuss the mechanics of that statement rather than discussing what the program is doing.

· **Focus paragraph comments on the *why* rather than the *how*.** Code should explain what the program is doing and why the programmer chose to do it that way rather than explain what each individual statement is doing.

· **Use comments to prepare the reader for what is to follow.** Someone reading the comments should be able to have a good idea of what the following code does without actually looking at the code. Note that this rule also suggests that comments should always precede the code to which they apply.

· **Make every comment count.** If the reader wastes time reading a comment of little value, the program is harder to read; period.

· **Document surprises and tricky code.** Of course, the best solution is not to have any tricky code. In practice, you can't always achieve this goal. When you do need to restore to some tricky code, make sure you fully document what you've done.

· **Avoid abbreviations.** While there may be an argument for abbreviating identifiers that appear in a program, no way does this apply to comments.

· **Keep comments close to the code they describe.** The prologue to a program unit should give its name, describe the parameters, and provide a short description of the program. It should not go into details about the operation of the module itself. Internal comments should to that.

· **Comments should explain the parameters to a function**, assertions about these parameters, whether they are input, output, or in/out parameters.

· **Comments should describe a routine's limitations, assumptions, and any side effects**.

Rule: All comments will be high-quality comments that describe the actions of the surrounding code in a con-cise manner

6.3   Endline vs. Standalone Comments

Guideline: Whenever a comment appears on a line by itself, always put the semicolon in column one. You may indent the text if this is appropriate or aesthetic.

Guideline: Adjacent lines of comments should not have any interspersed blank lines. A blank comment line should, at least, contain a semicolon in column one.

The guidline above suggests that your code should look like this:

```
; This is a comment with a blank line between it and the next comment.
;
; This is another line with a comment on it.
```

Rather than like this:

```
; This is a comment with a blank line between it and the next comment.


; This is another line with a comment on it.
```

The semicolon appearing between the two statements suggest continuity that is not present when you remove the semicolon. If two blocks of comments are truly separate and whitespace between them is appropriate, you should consider separating them by a large number of blank lines to completely eliminate any possible association between the two.

Standalone comments are great for describing the actions of the code that immediately follows. So what are endline comments useful for? Endline comments can explain how a sequence of instructions are implimenting the algorithm described in a previous set of standalone comments. Consider the following code:

```
; Compute the transpose of a matrix using the algorithm:
;
;        for i := 0 to 3 do
;               for j := 0 to 3 do
;                      swap( a[i][j], b[j][i] );

               forlp   i, 0, 3

               forlp   j, 0, 3

               mov     bx, i              ;Compute address of a[i][j] using
               shl     bx, 2              ; row major ordering (i*4 + j)*2.
               add     bx, j
               add     bx, bx
               lea     bx, a[bx]
               push    bx                 ;Push address of a[i][j] onto stack.

               mov     bx, j              ;Compute address of b[j][i] using
               shl     bx, 2              ;row major ordering (j*4 + i)*2.
               add     bx, i
               add     bx, bx
               lea     bx, b[bx]
               push    bx                 ;Push address of b[j][i] onto stack.

               call    swap               ;Swap a[i][j] with b[j][i].

               next
               next
```

Note that the block comments before this sequence explain, in high level terms, what the code is doing. The endline comments explain how the statement sequence implements the general algorithm. Note, however, that the endline comments do not explain what each statement is doing (at least at the machine level). Rather than claiming "add bx, bx" is multiplying the quantity in BX by two, this code assumes the reader can figure that out for themselves (any reasonable assembly programmer would know this). Once again, keep in mind your audience and write your comments for them.

6.4   Unfinished Code

Often it is the case that a programmer will write a section of code that (partially) accomplishes some task but needs further work to complete a feature set, make it more robust, or remove some known defect in the code. It is common for such pro-

grammers to place comments into the code like "This needs more work," "Kludge ahead," etc. The problem with these comments is that they are often forgotten. It isn't until the code fails in the field that the section of code associated with these comments is found and their problems corrected.

Ideally, one should never have to put such code into a program. Of course, ideally, programs never have any defects in them, either. Since such code inevitably finds its way into a program, it's best to have a policy in place to deal with it, hence this section.

Unfinished code comes in five general categories: non-functional code, partially functioning code, suspect code, code in need of enhancement, and code documentation. Non-functional code might be a stub or driver that needs to be replaced in the future with actual code or some code that has severe enough defects that it is useless except for some small special cases. This code is really bad, fortunately its severity prevents you from ignoring it. It is unlikely anyone would miss such a poorly constructed piece of code in early testing prior to release.

Partially functioning code is, perhaps, the biggest problem. This code works well enough to pass some simple tests yet contains serious defects that should be corrected. Moreover, these defects are known. Software often contains a large number of unknown defects; it's a shame to let some (prior) known defects ship with the product simply because a programmer forgot about a defect or couldn't find the defect later.

Suspect code is exactly that- code that is suspicious. The programmer may not be aware of a quantifiable problem but may suspect that a problem exists. Such code will need a later review in order to verify whether it is correct.

The fourth category, code in need of enhancement, is the least serious. For example, to expedite a release, a programmer might choose to use a simple algorithm rather than a complex, faster algorithm. S/he could make a comment in the code like "This linear search should be replaced by a hash table lookup in a future version of the software." Although it might not be absolutely necessary to correct such a problem, it would be nice to know about such problems so they can be dealt with in the future.

The fifth category, documentation, refers to changes made to software that will affect the corresponding documentation (user guide, design document, etc.). The documentation department can search for these defects to bring existing documentation in line with the current code.

This standard defines a mechanism for dealing with these five classes of problems. Any occurrence of unfinished code will be preceded by a comment that takes one of the following forms (where "_" denotes a single space):

```
;_#defect#severe_;
;_#defect#functional_;
;_#defect#suspect_;
;_#defect#enhancement_;
;_#defect#documentation_;
```

It is important to use all lower case and verify the correct spelling so it is easy to find these comments using a text editor search or a tool like grep. Obviously, a separate comment explaining the situation must follow these comments in the source code.

Examples:

```
; #defect#suspect ;
; #defect#enhancement ;
; #defect#documentation ;
```

Notice the use of comment delimiters (the semicolon) on both sides even though assembly language, doesn't require them.

Enforced Rule: If a module contains some defects that cannot be immediately removed because of time or other constraints, the program will insert a standardized comment before the code so that it is easy to locate such problems in the future. The five standardized comments are ";_#defect#severe_;,

";_#defect#functional_;", ";_#defect#suspect_;", ";_#defect#enhancement_;", and ";_#defect#documentation_;" where "_" denotes a single space. The spelling and spacing should be exact so it is easy to search for these strings in the source tree.

6.5   Cross References in Code to Other Documents

In many instances a section of code might be intrinsically tied to some other document. For example, you might refer the reader to the user document or the design document within your comments in a program. This document proposes a standard way to do this so that it is relatively easy to locate cross references appearing in source code. The technique is similar to that for defect reporting, except the comments take the form:

```
                  ;    text #link#location text ;
```

"*Text*" is optional and represents arbitrary text (although it is really intended for embedding html commands to provide hyperlinks to the specified document). "*Location*" describes the document and section where the associated information can be found.

`Examples:`

```
; #link#User's Guide Section 3.1 ;
; #link#Program Design Document, Page 5 ;
; #link#Funcs.pas module, "xyz" function ;
; <A HREF="DesignDoc.html#xyzfunc"> #link#xyzfunc </a> ;
```

Guideline: If a module contains some cross references to other documents, there should be a comment that takes the form "; *text #link#location text* ;" that provides the reference to that other document. In this comment "*text*" represents some optional text (typically reserved for html tags) and "location" is some descriptive text that describes the document (and a position in that document) related to the current section of code in the program.

7.0   Names, Instructions, Operators, and Operands

Although program features like good comments, proper spacing of statements, and good modularization can help yield programs that are more readable; ultimately, a programmer must read the instructions in a program to understand what it does. Therefore, do not underestimate the importance of making your statements as readable as possible. This section deals with this issue.

7.1   Names

According to studies done at IBM, the use of high-quality identifiers in a program contributes more to the readability of that program than any other single factor, including high-quality comments. The quality of your identifiers can make or break your program; program with high-quality identifiers can be very easy to read, programs with poor quality identifiers will be very difficult to read. There are very few "tricks" to developing high-quality names; most of the rules are nothing more than plain old-fashion common sense. Unfortunately, programmers (especially C/C++ programmers) have developed many arcane naming conventions that ignore common sense. The biggest obstacle most programmers have to learning how to create good names is an unwillingness to abandon existing conventions. Yet their only defense when quizzed on why they adhere to (existing) bad conventions seems to be "because that's the way I've always done it and that's the way everybody else does it."

The aforementioned researchers at IBM developed several programs with the following set of attributes:

·   Bad comments, bad names
·   Bad comments, good names
·   Good comments, bad names
·   Good comments, good names

As should be obvious, the programs that had bad comments and names were the hardest to read; likewise, those programs with good comments and names were the easiest to read. The surprising results concerned the other two cases. Most people assume good comments are more important than good names in a program. Not only did IBM find this to be false, they found it to be *really* false.

As it turns out, good names are even more important that good comments in a program. This is not to say that comments are unimportant, they are extremely important; however, it is worth pointing out that if you spend the time to write good comments and then choose poor names for your program's identifiers, you've damaged the readability of your program despite the work you've put into your comments. Quickly read over the following code:

```
mov     ax, SignedValue
cwd
add     ax, -1
rcl     dx, 1
mov     AbsoluteValue, dx
```

Question: What does this code compute and store in the AbsoluteValue variable?

· The sign extension of SignedValue.
· The negation of SignedValue.
· The absolute value of SignedValue.
· A boolean value indicating that the result is positive or negative.
· Signum(SignedValue) (-1, 0, +1 if neg, zero, pos).
· Ceil(SignedValue)
· Floor(SignedValue)

The obvious answer is the absolute value of SignedValue. This is also incorrect. The correct answer is signum:

```
mov     ax, SignedValue ;Get value to check.
cwd                     ;DX = FFFF if neg, 0000 otherwise.
add     ax, 0ffffh      ;Carry=0 if ax is zero, one otherwise.
rcl     dx, 1           ;DX = FFFF if AX is neg, 0 if ax=0,
mov     Signum, dx      ; 1 if ax>0.
```

Granted, this is a tricky piece of code[16]. Nonetheless, even without the comments you can probably figure out what the code sequence does even if you can't figure out how it does it:

```
mov     ax, SignedValue
cwd
add     ax, 0ffffh
rcl     dx, 1
mov     Signum, dx
```

Based on the names alone you can probably figure out that this code computes the signum function. This is the "understanding 80% of the code" referred to earlier. Note that you don't need misleading names to make this code unphathomable. Consider the following code that doesn't trick you by using misleading names:

```
mov     ax, x
cwd
```

16.It could be worse, you should see what the "superoptimizer" outputs for the signum function. It's even shorter and harder to understand than this code.

Style Guidelines for Assembly Language Programmers

```
add     ax, 0ffffh
rcl     dx, 1
mov     y, dx
```

This is a very simple example. Now imagine a large program that has many names. As the number of names increase in a program, it becomes harder to keep track of them all. If the names themselves do not provide a good clue to the meaning of the name, understanding the program becomes very difficult.

Enforced Rule: All identifiers appearing in an assembly language program must be descriptive names whose meaning and use are clear.

Since labels (i.e., identifiers) are the target of jump and call instructions, a typical assembly language program will have a large number of identifiers. Therefore, it is tempting to begin using names like "label1, label2, label3, ..." Avoid this temptation! There is always a reason you are jumping to some spot in your code. Try to describe that reason and use that description for your label name.

Rule: Never use names like "Lbl0, Lbl1, Lbl2, ..." in your program.

7.1.1   Naming Conventions

Naming conventions represent one area in Computer Science where there are far too many divergent views (program layout is the other principle area). The primary purpose of an object's name in a programming language is to describe the use and/or contents of that object. A secondary consideration may be to describe the type of the object. Programmers use different mechanisms to handle these objectives. Unfortunately, there are far too many "conventions" in place, it would be asking too much to expect any one programmer to follow several different standards. Therefore, this standard will apply across all languages as much as possible.

The vast majority of programmers know only one language - English. Some programmers know English as a second language and may not be familiar with a common non-English phrase that is not in their own language (e.g., rendezvous). Since English is the common language of most programmers, all identifiers should use easily recognizable English words and phrases.

Rule: All identifiers that represent words or phrases must be English words or phrases.

7.1.2   Alphabetic Case Considerations

A case-neutral identifier will work properly whether you compile it with a compiler that has case sensitive identifiers or case insensitive identifiers. In practice, this means that all uses of the identifiers must be spelled exactly the same way (including case) *and* that no other identifier exists whose only difference is the case of the letters in the identifier. For example, if you declare an identifier "ProfitsThisYear" in Pascal (a case-insensitive language), you could legally refer to this variable as "profitsThisYear" and "PROFITSTHISYEAR". However, this is not a case-neutral usage since a case sensitive language would treat these three identifiers as different names. Conversely, in case-sensitive languages like C/C++, it is possible to create two different identifiers with names like "PROFITS" and "profits" in the program. This is not case-neutral since attempting to use these two identifiers in a case insensitive language (like Pascal) would produce an error since the case-insensitive language would think they were the same name.

Enforced Rule: All identifiers must be "case-neutral."

Different programmers (especially in different languages) use alphabetic case to denote different objects. For example, a common C/C++ coding convention is to use all upper case to denote a constant, macro, or type definition and to use all lower case to denote variable names or reserved words. Prolog programmers use an initial lower case alphabetic to denote a variable. Other comparable coding conventions exist. Unfortunately, there are so many different conventions that make use of alphabetic case, they are nearly worthless, hence the following rule:

Rule: You should never use alphabetic case to denote the type, classification, or any other program-related attribute of an identifier (unless the language's syntax specifically requires this).

There are going to be some obvious exceptions to the above rule, this document will cover those exceptions a little later. Alphabetic case does have one very useful purpose in identifiers - it is useful for separating words in a multi-word identifier; more on that subject in a moment.

To produce readable identifiers often requires a multi-word phrase. Natural languages typically use spaces to separate words; we can not, however, use this technique in identifiers. Unfortunatelywritingmultiwordidentifiers makesthemalmostimpossibletoreadifyoudonotdosomethingtodistiguishtheindividualwords (Unfortunately writing multiword identifiers makes them almost impossible to read if you do not do something to distinguish the individual words). There are a couple of good conventions in place to solve this problem. This standard's convention is to capitalize the first alphabetic character of each word in the middle of an identifier.

Rule: Capitalize the first letter of interior words in all multi-word identifiers.

Note that the rule above does not specify whether the first letter of an identifier is upper or lower case. Subject to the other rules governing case, you can elect to use upper or lower case for the first symbol, although you should be consistent throughout your program.

Lower case characters are easier to read than upper case. Identifiers written completely in upper case take almost twice as long to recognize and, therefore, impair the readability of a program. Yes, all upper case does make an identifier stand out. Such emphasis is rarely necessary in real programs. Yes, common C/C++ coding conventions dictate the use of all upper case identifiers. Forget them. They not only make your programs harder to read, they also violate the first rule above.

Rule: Avoid using all upper case characters in an identifier.

7.1.3   Abbreviations

The primary purpose of an identifier is to describe the use of, or value associated with, that identifier. The best way to create an identifier for an object is to describe that object in English and then create a variable name from that description. Variable names should be meaningful, concise, and non-ambiguous to an average programmer fluent in the English language. Avoid short names. Some research has shown that programs using identifiers whose average length is 10-20 characters are generally easier to debug than programs with substantially shorter or longer identifiers.

Avoid abbreviations as much as possible. What may seem like a perfectly reasonable abbreviation to you may totally confound someone else. Consider the following variable names that have actually appeared in commercial software:

NoEmployees, NoAccounts, pend

The "NoEmployees" and "NoAccounts" variables seem to be boolean variables indicating the presence or absence of employees and accounts. In fact, this particular programmer was using the (perfectly reasonable in the real world) abbreviation of "number" to indicate the number of employees and the number of accounts. The "pend" name referred to a procedure's end rather than any pending operation.

Programmers often use abbreviations in two situations: they're poor typists and they want to reduce the typing effort, or a good descriptive name for an object is simply too long. The former case is an unacceptable reason for using abbreviations. The second case, especially if care is taken, may warrant the occasional use of an abbreviation.

Guideline: Avoid all identifier abbreviations in your programs. When necessary, use standardized abbreviations or ask someone to review your abbreviations. Whenever you use abbreviations in your programs, create a "data dictionary" in the comments near the names' definition that provides a full name and description for your abbreviation.

The variable names you create should be pronounceable. "NumFiles" is a much better identifier than "NmFls". The first can be spoken, the second you must generally spell out. Avoid homonyms and long names that are identical except for a few syllables. If you choose good names for your identifiers, you should be able to read a program listing over the telephone to a

peer without overly confusing that person.

Rule: All identifiers should be pronounceable (in English) without having to spell out more than one letter.

7.1.4   The Position of Components Within an Identifier

When scanning through a listing, most programmers only read the first few characters of an identifier. It is important, therefore, to place the most important information (that defines and makes this identifier unique) in the first few characters of the identifier. So, you should avoid creating several identifiers that all begin with the same phrase or sequence of characters since this will force the programmer to mentally process additional characters in the identifier while reading the listing. Since this slows the reader down, it makes the program harder to read.

Guideline: Try to make most identifiers unique in the first few character positions of the identifier. This makes the program easier to read.

Corollary: Never use a numeric suffix to differentiate two names.

Many C/C++ Programmers, especially Microsoft Windows programmers, have adopted a formal naming convention known as "Hungarian Notation." To quote Steve McConnell from Code Complete: "The term 'Hungarian' refers both to the fact that names that follow the convention look like words in a foreign language and to the fact that the creator of the convention, Charles Simonyi, is originally from Hungary." One of the first rules given concerning identifiers stated that all identifiers are to be English names. Do we really want to create "artificially foreign" identifiers? Hungarian notation actually violates another rule as well: names using the Hungarian notation generally have very common prefixes, thus making them harder to read.

Hungarian notation does have a few minor advantages, but the disadvantages far outweigh the advantages. The following list from Code Complete and other sources describes what's *wrong* with Hungarian notation:

· Hungarian notation generally defines objects in terms of basic machine types rather than in terms of abstract data types.
· Hungarian notation combines *meaning* with *representation*. One of the primary purposes of high level language is to abstract representation away. For example, if you declare a variable to be of type *integer*, you shouldn't have to change the variable's name just because you changed its type to *real*.
· Hungarian notation encourages lazy, uninformative variable names. Indeed, it is common to find variable names in Windows programs that contain *only* type prefix characters, without an descriptive name attached.
· Hungarian notation prefixes the descriptive name with some type information, thus making it harder for the programming to find the descriptive portion of the name.

Guideline: Avoid using Hungarian notation and any other formal naming convention that attaches low-level type information to the identifier.

Although attaching *machine* type information to an identifier is generally a bad idea, a well thought-out name can successfully associate some high-level type information with the identifier, especially if the name implies the type or the type information appears as a suffix. For example, names like "PencilCount" and "BytesAvailable" suggest integer values. Likewise, names like "IsReady" and "Busy" indicate boolean values. "KeyCode" and "MiddleInitial" suggest character variables. A name like "StopWatchTime" probably indicates a real value. Likewise, "CustomerName" is probably a string variable. Unfortunately, it isn't always possible to choose a great name that describes both the content and type of an object; this is particularly true when the object is an instance (or definition of) some abstract data type. In such instances, some additional text can improve the identifier. Hungarian notation is a raw attempt at this that, unfortunately, fails for a variety of reasons.

A better solution is to use a *suffix phrase* to denote the type or class of an identifier. A common UNIX/C convention, for example, is to apply a "_t" suffix to denote a type name (e.g., size_t, key_t, etc.). This convention succeeds over Hungarian notation for several reasons including (1) the "type phrase" is a suffix and doesn't interfere with reading the name, (2) this particular convention specifies the *class* of the object (const, var, type, function, etc.) rather than a low level *type*, and (3) It cer-

tainly makes sense to change the identifier if it's classification changes.

Guideline:　　　　　If you *want* to differentiate identifiers that are constants, type definitions, and variable names, use the suffixes "_c", "_t", and "_v", respectively.

Rule:　　　　　The classification suffix should not be the only component that differentiates two identifiers.

Can we apply this suffix idea to variables and avoid the pitfalls? Sometimes. Consider a high level data type "button" corresponding to a button on a Visual BASIC or Delphi form. A variable name like "CancelButton" makes perfect sense. Likewise, labels appearing on a form could use names like "ETWWLabel" and "EditPageLabel". Note that these suffixes still suffer from the fact that a change in type will require that you change the variable's name. However, changes in high level types are far less common than changes in low-level types, so this shouldn't present a big problem.

7.1.5　　　Names to Avoid

Avoid using symbols in an identifier that are easily mistaken for other symbols. This includes the sets {"1" (one), "I" (upper case "I"), and "l" (lower case "L")}, {"0" (zero) and "O" (upper case "O")}, {"2" (two) and "Z" (upper case "Z")}, {"5" (five) and "S" (upper case "S")}, and ("6" (six) and "G" (upper case "G")).

Guideline:　　　　　Avoid using symbols in identifiers that are easily mistaken for other symbols (see the list above).

Avoid misleading abbreviations and names. For example, FALSE shouldn't be an identifier that stands for "Failed As a Legitimate Software Engineer." Likewise, you shouldn't compute the amount of free memory available to a program and stuff it into the variable "Profits".

Rule:　　　　　Avoid misleading abbreviations and names.

You should avoid names with similar meanings. For example, if you have two variables "InputLine" and "InputLn" that you use for two separate purposes, you will undoubtedly confuse the two when writing or reading the code. If you can swap the names of the two objects and the program still makes sense, you should rename those identifiers. Note that the names do not have to be similar, only their meanings. "InputLine" and "LineBuffer" are obviously different but you can still easily confuse them in a program.

Rule:　　　　　Do not use names with similar meanings for different objects in your programs.

In a similar vein, you should avoid using two or more variables that have different meanings but similar names. For example, if you are writing a teacher's grading program you probably wouldn't want to use the name "NumStudents" to indicate the number of students in the class along with the variable "StudentNum" to hold an individual student's ID number. "NumStudents" and "StudentNum" are too similar.

Rule:　　　　　Do not use similar names that have different meanings.

Avoid names that sound similar when read aloud, especially out of context. This would include names like "hard" and "heart", "Knew" and "new", etc. Remember the discussion in the section above on abbreviations, you should be able to discuss your problem listing over the telephone with a peer. Names that sound alike make such discussions difficult.

Guideline:　　　　　Avoid homonyms in identifiers.

Avoid misspelled words in names and avoid names that are commonly misspelled. Most programmers are notoriously bad spellers (look at some of the comments in our own code!). Spelling words correctly is hard enough, remembering how to spell an identifier *incorrectly* is even more difficult. Likewise, if a word is often spelled incorrectly, requiring a programer to spell it correctly on each use is probably asking too much.

Style Guidelines for Assembly Language Programmers

Guideline: Avoid misspelled words and names that are often misspelled in identifiers.

If you redefine the name of some library routine in your code, another program will surely confuse your name with the library's version. This is especially true when dealing with standard library routines and APIs.

Enforced Rule: Do not reuse existing standard library routine names in your program unless you are specifically replacing that routine with one that has similar semantics (i.e., don't reuse the name for a different purpose).

7.2   Instructions, Directives, and Pseudo-Opcodes

Your choice of assembly language sequences, the instructions themselves, and your choice of directives and pseudo-opcodes can have a big impact on the readability of your programs. The following subsections discuss these problems.

7.2.1   Choosing the Best Instruction Sequence

Like any language, you can solve a given problem using a wide variety of solutions involving different instruction sequences. As a continuing example, consider (again) the following code sequence:

```
        mov     ax, SignedValue ;Get value to check.
        cwd                     ;DX = FFFF if neg, 0000 otherwise.
        add     ax, 0ffffh      ;Carry=0 if ax is zero.
        rcl     dx, 1           ;DX = FFFF if AX is neg, 0 if AX=0,
        mov     Signum, dx      ; 1 if AX>0.
```

Now consider the following code sequence that also computes the signum function:

```
        mov     ax, SignedValue ;Get value to check.
        cmp     ax, 0           ;Check the sign.
        je      GotSignum       ;We're done if it's zero.
        mov     ax, 1           ;Assume it was positive.
        jns     GotSignum       ;Branch if it was positive.
        neg     ax              ;Else return –1 for negative values.
GotSignum:      mov     Signum, ax
```

Yes, the second version is longer and slower. However, an average person can read the instruction sequence and figure out what it's doing; hence the second version is much easier to read than the first. Which sequence is best? Unless speed or space is an extremely critical factor and you can show that this routine is in the critical execution path, then the second version is obviously better. There is a time and a place for tricky assembly code; however, it's rare that you would need to pull tricks like this throughout your code.

So how does one choose appropriate instruction sequences when there are many possible ways to accomplish the same task? The best way is to ensure that you have a choice. Although there are many different ways to accomplish an operation, few people bother to consider any instruction sequence other than the first one that comes to mind. Unfortunatley, the "best" instruction sequence is rarely the first instruction sequence that comes to most people's minds[17]. In order to make a choice, you have to have a choice to make. That means you should create at least two different code sequences for a given operation if there is ever a question concerning the readability of your code. Once you have at least two versions, you can choose between them based on your needs at hand. While it is impractical to "write your program twice" so that you'll have a choice for every sequence of instructions in the program, you should apply this technique to particularly bothersome code sequences.

Guideline: For particularly difficult to understand sections of code, try solving the problem several different ways.

_____

17.This is true regardless of what metric you use to determine the "best" code sequence.

Then choose the most easily understood solution for actual incorporation into your program.

One problem with the above suggestion is that you're often too close to your own work to make decisions like "this code isn't too hard to understand, I don't have to worry about it." It is often a good idea to have someone else review your code and point out those sections they find hard to understand[18].

Guideline: Take advantage of reviews to determine those sections of code in your program that may need to be rewritten to make them easier to understand.

7.2.2 Control Structures

Ralph Griswold[19] once said (roughly) the following about C, Pascal, and Icon: "C makes it easy to write hard to read programs[20], Pascal makes it hard to write hard to read programs, and Icon makes it easy to write easy to read programs." Assembly language can be summed up like this: "Assembly language makes it hard to write easy to read programs and easy to write hard to read programs." It takes considerable discipline to write readable assembly language programs; *but it can be done*. Sadly, most assembly code you find today is extremely poorly written. Indeed, that state of affairs is the whole reason for this document. Once you get past issues like comments and naming conventions, issues like program control flow and data structure design have among the largest impacts on program readability. Since most assembly languages lack structured control flow constructs, this is one area where undisciplined programmers can really show how poorly they can write their code. One need look no farther than the public domain code on the Internet, or at Microsoft's sample code for that matter[21], to see abundant examples of poorly written assembly language code.

Fortunately, with a little discipline it is possible to write readable assembly language programs. How you design your control structures can have a big impact on the readability of your programs. The best way to do this can be summed up in two words: avoid spaghetti.

Spaghetti code is the name given to a program that has a large number of intertwined branches and branch targets within a code sequence. Consider the following example:

```
            jmp     L1
L1:         mov     ax, 0
            jmp     L2
L3:         mov     ax, 1
            jmp     L2
L4:         mov     ax, -1
            jmp     L2
L0:         mov     ax, x
            cmp     ax, 0
            je      L1
            jns     L3
            jmp     L4
L2:         mov     y, ax
```

This code sequence, by the way, is our good friend the Signum function. It takes a few moments to figure this out because as you manually trace through the code you find yourself spending more time following jumps around than you do looking at code that computes useful results. Now this is a rather extreme example, but it is also fairly short. A longer code sequence code become just as obfuscated with even fewer branches all over the place.

Spaghetti code is given this name because it resembles a bowl of spaghetti. That is, if we consider a control path in the program a spaghetti noodle, spaghetti code contains lots of intertwined branches into and out of different sections of the pro-

---

18. Of course, if the program is a *class assignment*, you may want to check your instructor's cheating policy before showing your work to your classmates!
19. The designer of the SNOBOL4 and Icon programming languages.
20. Note that this does not infer that it is hard to write easy to read C programs. Only that if one is sloppy, one can easily write something that is near impossible to understand.
21. Okay, this is a cheap shot. In fact, most of the assembly code on this planet is poorly written.

Style Guidelines for Assembly Language Programmers

gram. Needless to say, most spaghetti programs are difficult to understand, generally contain lots of bugs, and are often ineffi-
cient (don't forget that branches are among the slowest executing instructions on most modern processors).

So how to we resolve this? Easy by physically adopting structured programming techniques in assembly language code.
Of course, 80x86 assembly language doesn't provide if..then..else..endif, while..endwhile, repeat..until, and other such state-
ments[22], but we can certainly simulate them. Consider the following high level language code sequence:

```
if(expression) then

        << statements to execute if expression is true >>

else

        << statements to execute if expression is false >>

endif
```

Almost any high level language program can figure out what this type of statement will do. Assembly langue program-
mers should leverage this knowledge by attempting to organize their code so it takes this same form. Specifically, the assem-
bly language version should look something like the following:

```
        << Assembly code to compute value of expression >>

        JNxx    ElsePart ;xx is the opposite condition we want to check.

        << Assembly code corresponding to the then portion >>

        jmp     AroundElsePart

ElsePart:
        << Assembly code corresponding to the else portion >>

AroundElsePart:
```

For an concrete example, consider the following:

```
if ( x=y ) then

        write( 'x = y' );

else

        write( 'x <> y' );

endif;

; Corresponding Assembly Code:

        mov     ax, x
        cmp     ax, y
```

22.Actually, MASM 6.x *does*, but we'll ignore that fact here.

```
            jne     ElsePart

            print   "x=y",nl
            jmp     IfDone

ElsePart:   print   "x<>y",nl
IfDone:
```

While this may seem like the obvious way to organize an if..then.else..endif statement, it is suprising how many people would naturally assume they've got to place the else part somewhere else in the program as follows:

```
            mov     ax, x
            cmp     ax, y
            jne     ElsePart

            print   "x=y",nl
IfDone:
                .
                .
                .
ElsePart:   print   "x<>y",nl
            jmp     IfDone
```

This code organization makes the program more difficult to follow. Most programmers have a HLL background and despite a current assignment, they still work mostly in HLLs. Assembly language programs will be more readable if they mimic the HLL control constructs[23].

For similar reasons, you should attempt to organize your assembly code that simulates while loops, repeat..until loops, for loops, etc., so that the code resembles the HLL code (for example, a while loop should physically test the condition at the beginning of the loop with a jump at the bottom of the loop).

Rule: Attempt to design your programs using HLL control structures. The organization of the assembly code that you write should physically resemble the organization of some corresponding HLL program.

Assembly language offers you the flexibility to design arbitrary control structures. This flexibility is one of the reasons good assembly language programmers can write better code than that produced by a compiler (that can only work with high level control structures). However, keep in mind that a fast program doesn't have to contain the tightest possible code in every sequence. Execution speed is nearly irrelevant in most parts of the program. Sacrificing readability for speed isn't a big win in most of the program.

Guideline: Avoid control structures that don't easily map to well-known high level language control structures in your assembly language programs. Deviant control structures should only appear in small sections of code when efficiency demands their use.

7.2.3   Instruction Synonyms

MASM defines several synonyms for common instructions. This is especially true for the conditional jump and "set on condition code" instructions. For example, JA and JNBE are synonyms for one another. Logically, one could use either instruction in the same context. However, the choice of synonym can have an impact on the readability of a code sequence. To see why, consider the following:

```
            if( x <= y ) then
```

_____

23.Sometimes, for performance reasons, the code sequence above is justified since straight-line code executes faster than code with jumps. If the program rarely executes the ELSE portion of an if statement, always having to jump over it could be a waste of time. But if you're optimizing for speed, you will often need to sacrafice readability.

```
                << true statements>>
        else
            << false statements>>
        endif

; Assembly code:

        mov     ax, x
        cmp     ax, y
        ja      ElsePart
        << true code >>
        jmp     IfDone

ElsePart:       << false code >>
IfDone:
```

When someone reads this program, the "JA" statement skips over the true portion. Unfortunately, the "JA" instruction gives the illusion we're checking to see if something is greater than something else; in actuality, we're testing to see if some condition is less than or equal, not greater than. As such, this code sequence hides some of the original intent of high level algorithm. One solution is to swap the false and true portions of the code:

```
        mov     ax, x
        cmp     ax, y
        jbe     ThenPart
        << false code >>
        jmp     IfDone

ThenPart:       << true code >>
IfDone:
```

This code sequence uses the conditional jump that matches the high level algorithm's test (less than or equal). However, this code is now organized in a non-standard fashion (it's an if..else..then..endif statement). This hurts the readability more than using the proper jump helped it. Now consider the following solution:

```
        mov     ax, x
        cmp     ax, y
        jnbe    ElsePart
        << true code >>
        jmp     IfDone

ElsePart:       << false code >>
IfDone:
```

This code is organized in the traditional if..then..else..endif fashion. Instead of using JA to skip over the then portion, it uses JNBE to do so. This helps indicate, in a more readable fashion, that the code falls through on below or equal and branches if it is not below or equal. Since the instruction (JNBE) is easier to relate to the original test (<=) than JA, this makes this section of code a little more readable.

Rule: When skipping over some code because some condition has failed (e.g., you fall into the code because the condition is successful), always use a conditional jump of the form "JN*xx*" to skip over the code section. For example, to fall through to a section of code if one value is less than another, use the JNL or JNB instruction to skip over the code. Of course, if you are testing a negative condition (e.g., testing for equality) then use an instruction of the form J*x* to skip over the code.

8.0   Data Types

Prior to the arrival of MASM, most assemblers provided very little capability for declaring and allocated complex data types. Generally, you could allocate bytes, words, and other primitive machine structures. You could also set aside a block of bytes. As high level languages improved their ability to declare and use abstract data types, assembly language fell farther and farther behind. Then MASM came along and changed all that[24]. Unfortunately, many long time assembly language programmers haven't bothered to learn the new MASM syntax for things like arrays, structures, and other data types. Likewise, many new assembly language programmers don't bother learning and using these data typing facilities because they're already overwhelmed by assembly language and want to minimize the number of things they've got to learn. This is really a shame because MASM data typing is one of the biggest improvements to assembly language since using mnemonics rather than binary opcodes for machine level programming.

Note that MASM is a "high-level" assembler. It does things assemblers for other chips won't do like checking the types of operands and reporting errors if there are mismatches. Some people, who are used to assemblers on other machines find this annoying. However, it's a great idea in assembly language for the same reason it's a great idea in HLLs[25]. These features have one other beneficial side-effect: they help other understand what you're trying to do in your programs. It should come as no surprise, then, that this style guide will encourage the use of these features in your assembly language programs.

8.1   Defining New Data Types with TYPEDEF

MASM provides a small number of *primitive* data types. For typical applications, bytes, sbytes, words, swords, dwords, sdwords, and various floating point formats are the most commonly used scalar data types available. You may construct more abstract data types by using these built-in types. For example, if you want a character, you'd normally declare a byte variable. If you wanted a 16-bit integer, you'd typically use the sword (or word) declaration. Of course, when you encounter a variable declaration like "answer byte ?" it's a little difficult to figure out what the real type is. Do we have a character, a boolean, a small integer, or something else here? Ultimately it doesn't matter to the machine; a byte is a byte is a byte. It's interpretation as a character, boolean, or integer value is defined by the machine instructions that operate on it, not by the way you define it. Nevertheless, this distinction is important to someone who is reading your program (perhaps they are verifying that you've supplied the *correct* instruction sequence for a given data object). MASM's *typedef* directive can help make this distinction clear.

In its simplest form, the *typedef* directive behaves like a textequ. It let's you replace one string in your program with another. For example, you can create the following definitions with MASM:

```
char           typedef byte
integer        typedef sword
boolean        typedef byte
float          typedef real4
IntPtr         typedef far ptr integer
```

Once you have declared these names, you can define char, integer, boolean, and float variables as follows:

```
MyChar         char    ?
I              integer ?
Ptr2I          IntPtr  I
IsPresent      boolean ?
ProfitsThisYear float   ?
```

Rule: Use the existing MASM data types as type building blocks. For most data types you create in your program, you should declare explicit type names using the typedef directive. There is really no excuse for using the built-in primitive types[26].

---

24. Okay, MASM wasn't the first, but such techniques were not popularized until MASM appeared.
25. Of course, MASM gives you the ability to override this behavoir when necessary. Therefore, the complaints from "old-hand" assembly language programmers that this is insane are groundless.
26. Okay, using some assembler that doesn't support typedef would probably be a good excuse!

8.2   Creating Array Types

MASM provides an interesting facility for reserving blocks of storage – the DUP operator. This operator is unusual (among assembly languages) because its definition is *recursive*. The basic definition is (using HyGram notation):

```
DupOperator = expression ws* 'DUP' ws* '(' ws*  operand ws* ') %%
```

Note that "expression" expands to a valid numeric value (or numeric expression), "ws*" means "zero or more whitespace characters" and "operand" expands to anything that is legal in the operand field of a MASM word/dw, byte/db, etc., directive[27]. One would typically use this operator to reserve a block of memory locations as follows:

```
ArrayName        integer 16 dup (?)       ;Declare array of 16 words.
```

This declaration would set aside 16 contiguous words in memory.

The interesting thing about the DUP operator is that any legal operand field for a directive like byte or word may appear inside the parentheses, *including additional DUP expressions*. The DUP operator simply says "duplicate this object the specified number of times." For example, "16 dup (1,2)" says "give me 16 copies of the value pair one and two. If this operand appeared in the operand field of a byte directive, it would reserve 32 bytes, containing the alternating values one and two.

So what happens if we apply this technique recursively? Well, "4 dup (3 dup (0))" when read recursively says "give me four copies of whatever is inside the (outermost) parentheses. This turns out to be the expression "3 dup (0)" that says "give me three zeros." Since the original operand says to give four copies of three copies of a zero, the end result is that this expression produces 12 zeros. Now consider the following two declarations:

```
Array1               integer  4 dup ( 3 dup (0))
Array2               integer  12 dup (0)
```

Both definitions set aside 12 integers in memory (initializing each to zero). To the assembler these are nearly identical; to the 80x86 they are absolutely identical. To the reader, however, they are obviously different. Were you to declare two identical one-dimensional arrays of integers, using two different declarations makes your program inconsistent and, therefore, harder to read.

However, we can exploit this difference to declare multidimensional arrays. The first example above suggests that we have four copies of an array containing three integers each. This corresponds to the popular row-major array access function. The second example above suggests that we have a single dimensional array containing 12 integers.

Guideline: Take advantage of the recursive nature of the DUP operator to declare multidimensional arrays in your programs.

8.3   Declaring Structures in Assembly Language

MASM provides an excellent facility for declaring and using structures, unions, and records[28]; for some reason, many assembly language programmers ignore them and manually compute offsets to fields within structures in their code. Not only does this produce hard to read code, the result is nearly unmaintainable as well.

Rule: When a structure data type is appropriate in an assembly language program, declare the corresponding structure in the program and use it. Do not compute the offsets to fields in the structure manually, use the standard structure "dot-notation" to access fields of the structure.

One problem with using structures occurs when you access structure fields indirectly (i.e., through a pointer). Indirect access always occurs through a register (for near pointers) or a segment/register pair (for far pointers). Once you load a pointer value into a register or register pair, the program doesn't readily indicate what pointer you are using. This is especially true if you use the indirect access several times in a section of code without reloading the register(s). One solution is to use a textequ to create a special symbol that expands as appropriate. Consider the following code:

27.For brevity, the productions for these objects do not appear here.
28.MASM records are equivalent to bit fields in C/C++. They are not equivalent to records in Pascal.

```
s               struct
a               Integer ?
b               integer ?
s               ends
                 .
                 .
                 .
r               s       {}
ptr2r           dword   r
                 .
                 .
                 .
                les     di, ptr2r
                mov     ax, es:[di].s.a         ;No indication this is ptr2r!
                 .
                 .
                 .
                mov     es:[di].b, bx           ;Really no indication!
```

Now consider the following:

```
s               struct
a               Integer ?
b               integer ?
s               ends

sPtr            typedef far ptr s
                 .
                 .
                 .
q               s       {}
r               sPtr    q
r@              textequ <es:[di].s>
                 .
                 .
                 .
                les     di, ptr2r
                mov     ax, r@.a        ;Now it's clear this is using r
                 .
                 .
                 .
                mov     r@.b, bx        ;Ditto.
```

Note that the "@" symbol is a legal identifier character to MASM, hence "r@" is just another symbol. As a general rule you should avoid using symbols like "@" in identifiers, but it serves a good purpose here – it indicates we've got an indirect pointer. Of course, you must always make sure to load the pointer into ES:DI when using the textequ above. If you use several different segment/register pairs to access the data that "r" points at, this trick may not make the code anymore readable since you will need several text equates that all mean the same thing.

8.4   Data Types and the UCR Standard Library

The UCR Standard Library for 80x86 Assembly Language Programmers (version 2.0 and later) provide a set of macros that let you declare arrays and pointers using a C-like syntax. The following example demonstrates this capability:

```
var
        integer  i, j, array[10], array2[10][3], *ptr2Int
        char     *FirstName, LastName[32]
endvar
```

These declarations emit the following assembly code:

```
i               integer ?
j               integer 25
array           integer 10 dup (?)
array2          integer 10 dup ( 3 dup (?))
ptr2Int         dword   ?

LastName        char    32 dup (?)
Name            dword   LastName
```

For those comfortable with C/C++ (and other HLLs) the UCR Standard Library declarations should look very familiar. For that reason, their use is a good idea when writing assembly code that uses the UCR Standard Library.