

# The Practical Side Of Assembly Language

## Part II: Loops And Arrays

Bruce D. Corbrey, Raleigh, NC

In the last installment (**COMPUTE!** #14) I presented some ideas on how to represent and use flags in 6502 assembly language programs. This time I will discuss methods for programming loop control structures for the manipulation of arrays of data. Let's start by writing a loop which simply initializes all elements of an array to zero. In BASIC, you might write:

```
100 DIM AR(50)
110 FOR I=1 TO 50
120 AR(I)=0
130 NEXT I
```

If you are a neophyte assembly language programmer and try to translate this program segment on a line-by-line basis, you might wind up with something like this:

```
AR    *=*+ 50      ;SPACE FOR ARRAY AR (50 BYTES)
I     *=*+ 1       ;LOOP COUNTER VARIABLE I
...
LDA   #1          ;INITIALIZE LOOP COUNTER
STA   I
LOOP  LDA   #0
      LDX   I      ;RECALL CURRENT LOOP COUNTER
      STA   AR,X   ;SET ELEMENT OF ARRAY TO 0
      INX     ;ADVANCE TO NEXT ELEMENT
      STX   I      ;STORE INDEX REGISTER
      CPX   #50    ;CHECK AGAINST LIMIT
      BNE   LOOP   ;REPEAT UNTIL DONE
```

If you run this program you'll be dismayed to find out that it only sets the last 49 elements of the array to 0 and skips the first element, because the first element of the array should be indexed with a zero, not a one.

**Rule #7.** To access the first element of an assembly language array, you should use an index of 0, not 1. The last element of an array of size N is indexed by N-1.

You may also recognize that it is not necessary to allocate space or save the variable I for the loop. Since it is only needed to control the loop, it can be simply kept in the index register (I chose the X register; the Y register will serve equally well). We can correct and improve the loop as follows:

```
AR    *=*+ 50      ;SPACE FOR ARRAY AR (50 BYTES)
...
LDA   #0          ;CONSTANT TO FILL ARRAY WITH
TAX                     ;X=0=INITIAL INDEX TO ARRAY AR
LOOP  STA   AR,X   ;ZERO OUT ONE ARRAY ELEMENT
      INX     ;ADVANCE TO NEXT ELEMENT
      CPX   #50    ;CHECK INDEX AGAINST LIMIT
      BNE   LOOP   ;REPEAT UNTIL DONE
```

You may notice some other subtle improvements in this program segment. The A register is only loaded with 0 once, outside the loop, since it does not change inside the loop. This will make the program run faster by eliminating 49 unneeded repetitions of the LDA #0 instruction.

**Rule #8.** Move code which does not need to be repeated out of the loop if possible.

Also note that one byte of code was saved by using TAX to initialize the X register instead of LDX #0. Naturally if we were filling our array with something other than 0, this trick won't work. We now have a correctly-functioning loop which is equivalent of our BASIC program (strictly speaking, is not exactly equivalent because the BASIC interpreter uses floating point arithmetic which uses four or five bytes for each array element instead of one).

Can our loop be further improved in terms of efficiency? Consider this alternative:

```
AR    *=*+ 50      ;SPACE FOR ARRAY (50 BYTES)
...
LDA   #0          ;CONSTANT TO FILL THE ARRAY WITH
LDX   #49         ;INDEX TO LAST ELEMENT OF THE ARRAY
LOOP  STA   AR,X   ;SET AN ELEMENT OF THE ARRAY TO 0
      DEX         ;BACKUP TO PREVIOUS ELEMENT
      BPL   LOOP   ;REPEAT UNTIL DONE
```

This code segment fills the loop backwards, filling the last element first and the first element last. Once the 0th element has been filled, the index register is decremented to -1 (\$FF) and the BPL LOOP instruction will exit the loop. Notice that we have eliminated the CMP instruction from the loop, saving two cycles.

**Rule #9.** Moving backwards through an array will usually be more efficient.

If you try to make the array bigger than 128 elements you will be in trouble! Suppose you increase the dimension of AR to 200. In this case your loop will be executed only once because on the first pass, the DEX instruction will change the index from 199 to 198. But 198 has the hexadecimal representation \$C6, which has bit 7 (the sign bit) set to 1. Therefore the 6502 will consider this a negative number (-58 decimal) and the BPL instruction will let control "fall through." Therefore, our BPL instruction will only work right up to +127 decimal, which is the largest signed 8-bit number. We can remedy this problem for index values up to 255 with a slightly more "tricky," but equivalent, method:

```
AR    *=*+ 200     ;SPACE FOR ARRAY (200 BYTES)
...
LDA   #0          ;CONSTANT TO FILL ARRAY WITH
LDX   #200        ;INDEX TO LAST ELEMENT OF THE ARRAY + 1
```

```

LOOP STA AR-1,X ;SET AN ELEMENT OF THE
          ;ARRAY TO 0
      DEX        ;BACKUP TO PREVIOUS
          ;ELEMENT
      BNE LOOP   ;REPEAT UNTIL DONE

```

Here we have replaced the BPL instruction with a BNE instruction so that the loop will terminate one pass earlier, but will not be stymied by index values greater than 127. Since our last pass through the array will now have an index of 1 instead of 0, we must compensate by changing the destination for our indexed STA instruction to AR-1. Therefore the last element set will be AR-1 + 1 = AR + 0. Finally, the starting index must be bumped from 199 to 200 for the same reason. Note that a starting index of 0 will clear a full 256 byte array.

The same technique can be used to move a block of data of up to 256 bytes from one known location to another:

```

ARA  *="+ 200    ;ARRAY A CONTAINING 200
                ;ELEMENTS
ARB  *="+ 200    ;ARRAY B CONTAINING 200
                ;ELEMENTS
...
      LDX #200    ;NUMBER OF ELEMENTS TO
                ;MOVE
LOOP LDA ARA-1,X ;FETCH ELEMENT OF ARRAY A
      STA ARB-1,X ;INSTALL IN ARRAY B
      DEX        ;DECREMENT TO PREVIOUS
                ;ELEMENT
      BNE LOOP   ;REPEAT UNTIL DONE
...

```

What happens if you have more than 256 bytes? Throw away your 6502 and get a processor with a 16-bit index register? Nope. The indirect,X and indirect,Y addressing modes will solve this problem.

**Rule #10.** To use arrays of more than 256 bytes or arrays whose location is not known at assembly time, plan on using indirect,X or indirect,Y addressing.

Unlike the absolute, indexed addressing modes, indirect,X and indirect,Y are not equivalent. You may remember that indirect,X addressing uses pre-indexing and indirect,Y uses post-indexing. As a practical matter, indirect,X addressing will almost always be used with a permanent index of 0, simulating simple indirect addressing. This mode lends itself to manipulating large data arrays in non-time-critical portions of a program. For example, the following loop initializes a 1000-element array to 0:

```

ARRAY *="+ 1000 ;ROOM ARRAY OF 1000
                ;ELEMENTS
PTR   *="+ 2     ;POINTER TO AN ARRAY
                ;ELEMENT
...
CLR1K LDA #ARRAY&SFF ;LOW 8 BITS OF ADDRESS
                ;OF START OF ARRAY
      STA PTR        ;INITIALIZE POINTER
      LDA #ARRAY/256 ;HIGH 8 BITS OF AD-
                ;DRESS OF START OF
                ;ARRAY
      STA PTR+1      ;INITIALIZE HIGH BYTE
                ;OF POINTER

```

```

LDX #0          ;PERMANENTLY LOAD
                ;X WITH 0
LOOP LDA #0     ;ZERO BYTE POINTED
      STA (PTR),X ;TO BY PTR
      INC PTR    ;BUMP POINTER UP TO
                ;NEXT ELEMENT
      BNE CHECK  ;BRANCH IF NOT
                ;CROSSING PAGE
                ;BOUNDARY
      INC PTR+1  ;ELSE BUMP HI-ORDER
                ;BYTE OF POINTER
CHECK LDA PTR
      CMP #ARRAY+1000&SFF ;CHECK POINTER
                ;AGAINST LIMIT
      BNE LOOP   ;REPEAT IF NOT DONE
      LDA PTR+1
      CMP #ARRAY+1000/256 ;ELSE CHECK HI BYTE
                ;OF POINTER
      BNE LOOP   ;REPEAT IF NOT DONE

```

Some assemblers use the notation #<ARRAY to mean the low byte of the address of ARRAY and #>ARRAY for the high byte instead of #ARRAY&SFF and #ARRAY/256. Clearly this program segment is quite a bit "messier" than the one for arrays of less than 256 bytes. When planning sizes for arrays, you should remember this and try to limit arrays to 256 bytes or less whenever practical.

Luckily, the indirect,Y addressing mode is considerably more powerful than indirect,X. For our final problem, let's use the indirect,Y mode to build a subroutine to move a large block of data from one place to another in memory as fast as possible. The source address, destination address, and number of bytes to be moved are to be specified as input to the routine as three 16-bit variables in page 0:

```

FROM *="+ 2 ;POINTER TO STARTING ADDRESS OF
                ;ARRAY TO MOVE
TO    *="+ 2 ;POINTER TO STARTING ADDRESS
                ;OF DESTINATION
COUNT *="+ 2 ;NUMBER OF BYTES TO COPY

```

In an earlier example we saved execution time by removing the need for a compare inside the loop. We can apply the same principle to speed up our block move by sub-dividing the routine into two loops, one which moves entire pages (1 page = 256 bytes), and one which moves the final fractional page. This allows us to avoid any compares in the part which moves entire pages (which is part of the routine executed the most when copying large blocks). This will also let us use both 8-bit index registers to maximum effectiveness by allocating one for counting pages and index registers to maximum effectiveness by allocating one for counting pages and the other for indexing bytes within the page. The resulting routine (shown in Program 4) can easily be converted into a block-fill routine instead by removing FROM and all lines that refer to it, and presetting A to 0 (or the value with which to fill the array).

**Rule #11.** To deal with large arrays, split your program into two loops, one to operate on entire pages and one to operate on the "leftover"

fractional page.

The routine in Program 4 moves data at about 16.1 machine cycles per byte for large blocks, which means a 16K byte array can be moved in 0.26 seconds using a 6502 with a 1MHz clock. In certain

applications where speed is of paramount importance, you may wish to improve even this super-fast copy routine. Can it be done? Yes, if you are willing to trade some increased program size for increased execution speed. Again, we employ the same gen-

**Program 2: Keyboard Driver with Alpha Lock Flag**  
Using 0 = False and Non-0 = True

*Editor's Note: Part of this program was not printed in*  
**COMPUTE! #14**, we reprint it entirely here. — RTM

```

;
;      SUBROUTINE INCH: KEYBOARD DRIVER FOR ASCII-ENCODED
;      KEYBOARD WITH PARALLEL INTERFACE.
;
;      ADDRESSES SHOWN ARE FOR 6530 ON KIM-1 COMPUTER.
;      KEYBOARD DATA LINES TO PORT A BITS 0 TO 6,
;      NEGATIVE-GOING STROBE TO BIT 7.
;
;      ON ENTRY: IF ALFALK IS NON-0, THEN ALL LOWERCASE LETTERS WILL
;      BE RETURNED AS THE EQUIVALENT UPPERCASE ALPHA.
;      ON RETURN: REGISTER A = ASCII CODE FOR KEY PRESSED;
;      X AND Y PRESERVED.
;
1700      PAD      =      $1700      ;KIM PORT A DATA REGISTER ON 6530
1701      PADD     =      $1701      ;KIM PORT A DATA DIRECTION REGISTER
;
0000      ;      *=      $1780      ;PROGRAM ORIGIN
;
1780 A900      INCH   LDA      #$00
1782 8D0117     STA      PADD      ;SET PORT DIRECTION = INPUTS
1785 AD0017     INCH1  LDA      PAD      ;TEST PORT
1788 30FB       BMI      INCH1      ;WAIT FOR STROBE PULSE
178A 2C0017     INCH2  BIT      PAD
178D 10FB       BPL      INCH2      ;WAIT FOR END OF STROBE
;
;      IF ALPHA-LOCK FLAG IS SET, FOLD ANY LOWERCASE LETTERS TO
;      EQUIVALENT UPPERCASE LETTERS.
;
178F 48         FOLD   PHA          ;SAVE CHARACTER TEMPORARILY
1790 ADA317     LDA      ALFALK     ;RECALL "ALPHA LOCK" FLAG
1793 F00C       BEQ     FOLD2       ;BRANCH IF NO FOLDING DESIRED
1795 68         PLA          ;ELSE RECALL CHARACTER
1796 C97B       CMP      #$7B      ;LOWER CASE "Z" + 1
1798 B006       BCS     FOLD1       ;BRANCH IF PUNCTUATION
179A C961       CMP      #$61      ;LOWER CASE "A"
179C 9002       BCC     FOLD1       ;BRANCH IF NOT LOWER CASE ALPHA
179E E920       SBC      #$20      ;ELSE FOLD TO EQUIVALENT UPPERCASE
17A0 60         FOLD1  RTS
;
17A1 68         FOLD2  PLA          ;RECALL CHARACTER
17A2 60         RTS
;
;      ALPHA LOCK FLAG (DEFAULT = ALLOW LOWER CASE)...
;
17A3      ALFALK  .BYTE    0        ;"ALPHA LOCK" FLAG; NON-0=UPPERCASE ONLY.
;
0000      .END

```

NO ERROR LINES

**Program 4: Block-Move Memory Routine**

GENERAL BLOCK-MOVE SUBROUTINE

MTU 6502 ASSEMBLER 1.0

```

0002 0000          .PAGE  'GENERAL BLOCK-MOVE SUBROUTINE'
0003 0000          *=      0      ;ZERO PAGE ORIGIN
0004 0000      FROM  *=*+    2      ;STARTING ADDRESS OF BLOCK TO BE COPIED
0005 0002          TO    *=*+    2      ;STARTING ADDRESS OF DESTINATION
0006 0004      COUNT *=*+    2      ;NUMBER OF BYTES TO BE MOVED
0007 0006          ;
0008 0006          *=      $2000 ;ORIGIN FOR PROGRAM
0009 2000          ;
0010 2000          ;      THIS ROUTINE COPIES A BLOCK OF ANY SIZE FROM ONE
0011 2000          ;      LOCATION TO ANOTHER.
0012 2000          ;
0013 2000          ;      ON ENTRY: FROM (2 BYTES) IS THE STARTING ADDRESS OF
0014 2000          ;      THE BLOCK TO BE COPIED; TO (2 BYTES) IS THE DESIRED
0015 2000          ;      STARTING DESTINATION ADDRESS FOR THE COPY; COUNT
0016 2000          ;      (2 BYTES) IS THE NUMBER OF BYTES TO COPY.
0017 2000          ;
0018 2000          ;      ON RETURN: NO REGISTERS PRESERVED; FROM, TO AND COUNT
0019 2000          ;      ARE ClobberED.
0020 2000          ;
0021 2000          ;      NOTE: THE DESTINATION BLOCK MAY OVERLAP THE SOURCE
0022 2000          ;      BLOCK ONLY IF "TO" IS AT A LOWER ADDRESS THAN "FROM".
0023 2000          ;
0024 2000 A000      BLKMOV LDY    #0      ;INITIAL INDEX WITHIN A PAGE
0025 2002 A605      LDX    COUNT+1 ;NUMBER OF PAGES TO BE MOVED
0026 2004 F00E      BEQ    FRCMOV ;BRANCH IF ONLY A FRACTIONAL PAGE
0027 2006          ;
0028 2006          ;      THIS LOOP COPIES ENTIRE PAGES...
0029 2006          ;
0030 2006 B100      PAGMOV LDA    (FROM),Y ;FETCH A BYTE FROM SOURCE
0031 2008 9102      STA    (TO),Y ;COPY TO DESTINATION
0032 200A C8        INY          ;BUMP POINTER
0033 200B D0F9      BNE    PAGMOV ;REPEAT TILL ENTIRE PAGE MOVED
0034 200D E601      INC    FROM+1 ;BUMP HI BYTE OF POINTERS
0035 200F E603      INC    TO+1
0036 2011 CA        DEX          ;DECREMENT COUNT OF PAGES TO COPY
0037 2012 D0F2      BNE    PAGMOV ;REPEAT TILL ALL WHOLE PAGES COPIED
0038 2014          ;
0039 2014          ;      THIS LOOP COPIES THE FINAL FRACTION OF A PAGE...
0040 2014          ;
0041 2014 A604      FRCMOV LDX    COUNT ;RECALL NUMBER OF BYTES LEFT TO COPY
0042 2016 F008      BEQ    DONEMV ;BRANCH IF COUNT IS EXACT PAGE MULTIPLE
0043 2018 B100      FRLOOP LDA    (FROM),Y ;FETCH A BYTE FROM SOURCE
0044 201A 9102      STA    (TO),Y ;COPY TO DESTINATION
0045 201C C8        INY          ;BUMP INDEX
0046 201D CA        DEX          ;DECREMENT COUNT OF BYTES LEFT
0047 201E D0F8      BNE    FRLOOP ;REPEAT UNTIL DONE
0048 2020          ;
0049 2020 60        DONEMV RTS
0050 2021          ;
0051 2021          .END

```

0 ERRORS IN PASS 2

eral principle of loop optimization:

**Rule #12.** To optimize loop execution speed, try to remove unnecessary compares from within the loop.

About the only way we can remove more compares from Program 4 is to "unwind" part of the loop and, instead, write some of the loop code "inline." Since we know the first loop will always move exactly 256 bytes, we can move two bytes at a time instead of one before checking for a page crossing:

```

...
PAGMOV LDA (FROM),Y ;FETCH A BYTE FROM SOURCE
      STA (TO),Y    ;COPY TO DESTINATION
      INY           ;BUMP POINTER
      LDA (FROM),Y  ;FETCH A BYTE FROM SOURCE
      STA (TO),Y    ;COPY TO DESTINATION
      INY           ;BUMP POINTER
      BNE PAGMOV    ;REPEAT TILL ENTIRE PAGE
                      MOVED
...

```

This loop now takes 14.5 cycles per byte moved versus 16 cycles for the equivalent loop of Program 4, because the three cycle BNE instruction is only executed for every other byte moved. The loop can be unwound still further to move four, eight or more bytes per pass, but the speed improvement gained drops off rapidly as more code is written inline.

In the next installment I will explore some techniques for optimizing jumps and subroutine calls.



