

Example 6.10: Illustration of Indexed Indirect Addressing

Cycle	Address Bus	Data Bus	External Operation	Internal Operation
1	0100	OP CODE	Fetch OP CODE	Finish Previous Operation, 0101 → PC
2	0101	BAL	Fetch BAL	Interpret Instruction, 0102 → PC
3	00, BAL	DATA (Discarded)	Fetch Discarded DATA	Add BAL + X
4	00, BAL + X	ADL	Fetch ADL	Add 1 to BAL + X
5	00, BAL + X + 1	ADH	Fetch ADH	Hold ADL
6	ADH, ADL	DATA	Fetch DATA	
7	0102	Next OP	Fetch Next OP CODE	Finish Operation 0103 → PC

One of the advantages of this type of indexing is that a 16-bit address can be fetched with only two bytes of memory, the byte that contains the OP CODE and the byte that contains the indirect pointer. It does require, however, that there be a table of addresses kept in a read/write memory which is more expensive than having it in read-only memory. Therefore, this approach is normally reserved for applications where use of indexed indirect results in significant coding improvement or where the address being fetched is a variable computed address.

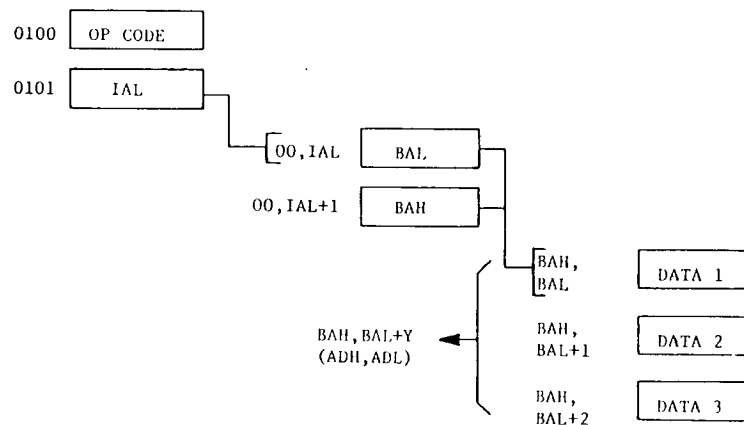
It is also obvious from the example that the user pays a minor time penalty for this form of addressing in that indexed indirect always takes six cycles to fetch a single operand, which is 25% more than an absolute address and 50% more than a Zero Page reference to an operand. As in the Zero Page indexed, the operation in cycles three and four are located in Zero Page and there is no ability to carry over into the next page. It is possible to develop a value of the index plus the base address where the result exceeds 255; in this case, the address put out is a wrap-around to the low part of the Page Zero.

Instructions which allow the use of indexed indirect are ADC, AND, CMP, EOR, LDA, ORA, SBC, and STA.

6.5 INDIRECT INDEXED ADDRESSING

The indirect indexed instruction combines a feature of indirect addressing and a capability of indexing. The usefulness of this instruction is primarily for those operations in which one of several values could be used as part of a subroutine. By having an indirect pointer to the base operation and by using the index register Y in the normal counter type form, one can have the advantages of an address that points anywhere in memory, combined with the advantages of the counter offset capability of the index register.

Figure 6.6 illustrates the indirect indexed concept in flow form while Example 6.11 indicates the internal operation of a non-page rollover of an indirect index.



Indirect Indexed Addressing

FIGURE 6.6

Example 6.11: Indirect Indexed Addressing (No Page Crossing)

<u>Cycle</u>	<u>Address Bus</u>	<u>Data Bus</u>	<u>External Operation</u>	<u>Internal Operation</u>
1	0100	OP CODE	Fetch OP CODE	Finish Previous Operation, 0101 → PC
2	0101	IAL	Fetch IAL	Interpret Instruction, 0102 → PC
3	00, IAL	BAL	Fetch BAL	Add 1 to IAL
4	00, IAL + 1	BAH	Fetch BAH	Add BAL + Y
5	BAH, BAL + Y	DATA	Fetch Operand	
6	0102	Next OP CODE	Fetch Next OP CODE	Finish Operation 0103 → PC

The indirect index still requires two bytes of program storage -- one for the OP CODE, one for the indirect pointer. Once beyond the indirect, the indexing of the indirect memory location is just the same as though it were an absolute indexed operation, in the sense that if there is no page crossing, pipelining occurs in the adding of the index register Y to address low while fetching address high, and, therefore, the non-page crossing solution is one cycle shorter than the indexed indirect. In Example 6.12 it is seen that the page crossing problem that occurs with absolute indexed page crossing also occurs with indirect indexed addressing.

Example 6.12: Indirect Indexed Addressing (With Page Crossing)

<u>Cycle</u>	<u>Address Bus</u>	<u>Data Bus</u>	<u>External Operation</u>	<u>Internal Operation</u>
1	0100	OP CODE	Load OP CODE	Finish Previous Operation, 0101 → PC
2	0101	IAL	Fetch IAL	Interpret Instruction, 0102 → PC
3	00, IAL	BAL	Fetch BAL	Add 1 to IAL
4	00, IAL + 1	BAH	Fetch BAH	Add BAL to Y
5	BAH, BAL + Y	DATA (Discarded)	Fetch DATA (Discarded)	Add 1 to BAH
6	BAH + 1	DATA BAL + Y	Fetch Data	
7	0102	Next OP CODE	Fetch Next OP CODE	Finish This Operation, 0103 → PC

When there is a page crossing, the base address high and base address low plus Y are pointing to an incorrect location within a referenced page. However, it should be noted that the programmer has control of this incorrect reference in the sense that it is always pointing to the page of the base address. In one more cycle the correct address is referenced. As was true in the case of absolute indexed, the data at the incorrect address is only read. STA and the various read, modify, and write memory commands all operate assuming that there will be a page crossing; take the extra cycle time to perform the add and carry; and perform a write only on the sixth cycle, rather than taking advantage of the five-cycle short-cut which is available to read operations. This added cycle guarantees that a memory location will never be written into with incorrect data.

Instructions which allow the use of indexed indirect are ADC, AND, CMP, EOR, LDA, ORA, SBC, and STA.

In the following two examples a comparison can be seen between the use of absolute modified by Y and indirect indexed addressing.

In both examples, the same function is performed. Values from two memory locations are added, and the result is stored in a third memory location, assuming that there are several values to be added. The first example deals with known field locations. The second example, such as might be traditionally used in subroutines, deals with field locations that vary between routines. A two-byte pointer for each routine using the subroutine is stored in Page Zero. The number of values to be added for each routine is also stored.

Example 6.13: Absolute Indexed Add -- Sample Program

<u>#Bytes</u>	<u>Cycles</u>	<u>Label</u>	<u>Instruction</u>	<u>Comments</u>
2	2	START	LDY #COUNT -1	Set Y = End of FIELD
3	4	LOOP	LDA FIELD 1,Y	Load Location 1
3	4		ADC FIELD 2,Y	Add Location 2
3	4		STA FIELD 3,Y	Store in Location 3
1	2		DEY	
2	3		BPL LOOP	Check for Less Than Zero
<u>14</u>	<u>19</u>		Time for 10 Bytes = 171 Cycles	

Example 6.14: Indirect Indexed Add -- Sample Program

<u>#Bytes</u>	<u>Cycles</u>	<u>Label</u>	<u>Instruction</u>	<u>Comments</u>
2	2	START	LDY #COUNT -1	Set Y = End of FIELD
2	5	LOOP	LDA (PNT1), Y	Load FIELD 1 Value
2	5		ADC (PNT2), Y	Add FIELD 2 Value
2	5		STA (PNT3), Y	Store FIELD 3 Value
1	2		DEY	
2	3		BPL LOOP	
<u>11</u>	<u>22</u>		Time for 10 Bytes = 201 Cycles	

+ 6 Bytes for Pointers

The "count" term in these examples represents the number of sets of values to be added and stored. Loading the index register with COUNT-1 will allow a fall through the BPL instruction when computation on all set of values has been completed.

There is a definite saving in program storage using indirect because it requires only two bytes for each indirect pointer, the OP CODE plus the pointer of the Page Zero location, whereas in the case of the absolute, it takes three bytes, the OP CODE, address low and address high.

It is to be noted that there are six bytes of Page Zero memory used for pointers, two bytes for each pointer. The number of memory locations allocated to the problem are 17 for the indirect and 14 for the problem where the values are known. The execution time is longer in the indirect loop. Even though the increase in time for a single pass through the loop is only three cycles, if many values are to be transferred, it adds up. It is important to note that loops require time for setup but this time is consumed only once. In the loop itself, however, additional time is multiplied by the number of times the program goes through the loop; therefore, on problems where execution time is important, the time reduction effort should be placed on the loop.

Even though the loop time is longer and the actual memory expended is greater for the indexed indirect add, it has the advantage of not requiring determination of the locations of FIELD 1, FIELD 2, and FIELD 3 at the time the program was written as is necessary with absolute.

An attempt to define problems to take advantage of this shorter memory and execution time by defining fields should be investigated first. However, in almost every program, the same operation must be performed several times. In those cases, it is sometimes more useful to define a subroutine, and to set the values that the subroutine will operate on as fields in memory. Pointers to these fields are placed in the Zero Page of memory, and then the indexed indirect operation is used to perform the function. This is the primary application of the indexed indirect operation.

6.6 INDIRECT ABSOLUTE

In the case of all of the indirect operations previously described, the indirect reference is always to a Page Zero location from which are picked up the effective address low and effective address high. There is an exception in the R6500-series microprocessor family for the jump instruction in which absolute indirect jumps are allowed. The use of the absolute indirect jump is best described in the discussion on interrupts where the addressing mode and its capabilities are explained.

6.7 APPLICATION OF INDEXES

As has been developed in many of the previous examples, an index register has primary values as a modifier and as a counter. As a modifier to a base address operation, it allows the accessing of contiguous groups of data by simple modification of the index. This is the primary application of indexes, and it is for this purpose that they were created. By virtue of the fact that all of the R6500 instructions have the base address in the instruction or, in the case of the indirect, in the pointer, a single index can usually service an entire loop, because each of the many instructions in the loop are normally referring to the same relative value in each of the lists. An example is adding the third byte of a number to the corresponding third byte of another number, then storing the result in the memory location representing the third byte of the result; therefore, the index register needs to contain only three to accomplish all three of these offset functions.

In some other microprocessors internal registers serve as indirect pointers. The single register requirement is a significant advantage of the type of indexing done in the R6500. Even though the R6500 has two indexes, more often than not, a single index will solve many of the problems because of the fact that the data is normally organized in corresponding fields.

The second feature of the R6500 type of indexing is that, if applied properly, the index register also contains the count of the operations to be performed.

The examples given in this manner are presented in an attempt to show how to take advantage of that feature. There are two approaches to counting: forward counting and reverse counting. In forward counting, the data in memory can be organized in such a manner that the index register starts at zero and is added to on each successive operation. The disadvantage of this type of approach is that the compare index instruction, as used in Example 6.13, must be inserted into the loop in order to determine that the correct number of operations is completed.

The reverse counting approach has been applied in the latter examples. The data must be organized for reverse counting operation. The first value to be operated on is at the end of the FIELD, the next value is one memory location in front of that, etc. The advantage of this type of operation is that it takes advantage of the combined decrement and test capability of the processor. There are two ways to use the test. First there is the case where the actual number of operations to be performed is loaded into the index register as was done in Example 6.13. In this case, the index contains the correct count but if added to the base directly, would be pointing to one value beyond the FIELD because the base address contains the first byte. Therefore, when using the actual count in the index register, one always references to the base address minus one. This is easily accomplished as shown in the examples. The assembler accepts symbolic references in the form of base address minus one, and the microprocessor very carefully performs the operation shown.

The advantage of putting the actual count in the register is that the branch if not equal instruction (BNE) can be used because the value of the register goes to zero on the last operation.

The second alternative is to load the counter with the count minus one as done in Example 6.14. In this case, the actual value of the base address is used in the offset. However, the branch back to loop now is a branch plus, remembering that the value in the index register will not go to minus (all ones) until one decrement is past zero.

Values of count minus one through zero will all take the branch. It is only when attempting to reference less than the base address that the loop will be completed.

Either approach gives minimum coding and requires only that the user develop a philosophy of always organizing his data with the first value at the end. In many cases, the operations such as MOVE can be performed even if the data is organized the other way. Experienced programmers find that this reverse counting form is actually more convenient to use and always results in minimum loop time and space.

Although for most applications, the 8-bit index register permits simple count in offset operations, there are a few operations where the 256 count that is available in the 8-bit register is not adequate to perform the indexed operations. There are two solutions to this problem. First to code the program with two sets of bases -- that is, to duplicate the coding for the loop with two different address highs, each a page apart. The second, more useful solution, is to go to indirect operations because the indirect pointer can be modified to allow an infinite indexed operation. An example of the move done under 256 and over 256 is shown in the following examples:

Example: 6.15: Move N Bytes (N < 256)

<u>Number of Cycles</u>	<u>Program Label</u>	<u>Instruction Mnemonics</u>	<u>OPERAND FIELD</u>	<u>Comments</u>
2		LDX	#BLOCK	Set Up 2 Cycles
4	LOOP	LDA	FROM-1,X	
4		STA	TO -1,X	LOOP Time:
2		DEX		13 Cycles
3		BNE	LOOP	

Memory Required:
11 Bytes

Example 6.16: Move N Bytes (N < 256)

<u>Number of Cycles</u>	<u>Program Label</u>	<u>Instruction Mnemonics</u>	<u>operand FIELD</u>	<u>Comments</u>
2	MOVE	LDA	#FROML	
3		STA	FRPOINT	
2		LDA	#FROMH	
3		STA	FRPOINT + 1	Move FROM Address to an Indirect Pointer
2		LDA	#TOL	
3		STA	TOPOINT	Move TO Address to an Indirect Pointer
2		LDA	#TOH	
3		STA	TOPOINT + 1	
2		LDX	#BLOCKS	Setup Number of 256
2		LDY	#0	Blocks to Move
5	LOOP	LDA	(FRPOINT),Y	Loop Time: 16 Cycles/
6		STA	(TOPOINT),Y	Byte. Move 256 Bytes
2		DEY		
3		BNE	LOOP	
5	SPECIAL	INC	FRPOINT + 1	Increase High
5		INC	TOPOINT + 1	Pointer
2		DEX		
2		BMI	OUT	Check for Last Move
3		BNE	LOOP	
2		LDY	#COUNT	
3		BNE	LOOP	Set Up Last Move
	OUT	---	---	

Memory Required:
40 Bytes

CHAPTER 7

INDEX REGISTER INSTRUCTIONS

The index registers can be treated as auxiliary general-purpose registers, having the added ability of being incremented and decremented because of the normal operations which they are required to perform.

7.0 LDX -- LOAD INDEX REGISTER X FROM MEMORY

Load the index register X from memory.

The symbolic notation is $M \rightarrow X$.

LDX does not affect the C or V flags; sets Z if the value loaded was zero, otherwise resets it; sets N if the value loaded in bit 7 is a 1, otherwise N is reset; and affects only the X register. The addressing modes for LDX are Immediate; Absolute; Zero Page; Absolute Indexed by Y; and Zero Page Indexed by Y.

7.1 LDY -- LOAD INDEX REGISTER Y FROM MEMORY

Load the index register Y from memory.

The symbolic notation is $M \rightarrow Y$.

LDY does not affect the C or V flags, sets the N flag if the value loaded in bit 7 is a 1, otherwise resets N; sets Z flag if the loaded value is zero, otherwise resets Z; and affects only the Y register. The addressing modes for load Y are Immediate; Absolute; Zero Page; Zero Indexed by X, Absolute Indexed by X.

7.2 STX -- STORE INDEX REGISTER X IN MEMORY

Transfers value of X register to addressed memory location.
The symbolic notation is $X \rightarrow M$.

No flags or registers in the microprocessor are affected by the store operation. The addressing modes for STX are Absolute, Zero Page, and Zero Page Indexed by Y.

7.3 STY -- STORE INDEX REGISTER Y IN MEMORY

Transfer the value of the Y register to the addressed memory location. The symbolic notation is $Y \rightarrow M$. STY does not affect any flags or registers in the microprocessor. The addressing modes for STY are Absolute; Zero Page; and Zero Page Indexed by X.

7.4 INX -- INCREMENT INDEX REGISTER X BY ONE

Increment X adds 1 to the current value of the X register. This is an 8-bit increment which does not affect the carry operation, therefore, if the value of X before the increment was FF, the resulting value is 00. The symbolic notation is $X + 1 \rightarrow X$. INX does not affect the carry or overflow flags; it sets the N flag if the result of the increment has a one in bit 7, otherwise resets N; sets the Z flag if the result of the increment is 0, otherwise it resets the Z flag. INX does not affect any other register other than the X register. INX is a single byte instruction and the only addressing mode is Implied.

7.5 INY -- INCREMENT INDEX REGISTER Y BY ONE

Increment Y increments or adds one to the current value in the Y register, storing the result in the Y register. As in the case of INX the primary application is to step through a set of values using the Y register. The symbolic notation is $Y + 1 \rightarrow Y$. The INY does not affect the carry or overflow flags; sets the N flag if the result of the increment has a one in bit 7, otherwise resets N; sets Z if

as a result of the increment the Y register is zero, otherwise resets the Z flag. Increment Y is a single-byte instruction and the only addressing mode is Implied.

7.6 DEX -- DECREMENT INDEX REGISTER X BY ONE

This instruction subtracts one from the current value of the index register X and stores the result in the index register X.

The symbolic notation is $X - 1 \rightarrow X$.

DEX does not affect the carry or overflow flag, it sets the N flag if it has bit 7 on as a result of the decrement, otherwise it resets the N flag; sets the Z flag if X is a 0 as a result of the decrement, otherwise it resets the Z flag.

DEX is a single-byte instruction, the addressing mode is Implied.

7.7 DEY -- DECREMENT INDEX REGISTER Y BY ONE

This instruction subtracts one from the current value in the index register Y and stores the result into the index register Y. The result does not affect or consider carry so that the value in the index register Y is decremented to 0 and then through 0 to FF.

Symbolic notation is $Y - 1 \rightarrow Y$.

Decrement Y does not affect the carry or overflow flags; if the Y register contains bit 7 on as a result of the decrement the N flag is set, otherwise the N flag is reset. If the Y register is 0 as a result of the decrement, the Z flag is set, otherwise the Z flag is reset. This instruction affects only the index register Y.

DEY is a single-byte instruction and the addressing mode is Implied.

NOTE: Decrement of the index registers is the most convenient method of using the index registers as a counter, in that the decrement involves setting the value N as a result of having passed through 0 and sets Z when the results of the decrement are 0.

7.8 CPX -- COMPARE INDEX REGISTER X TO MEMORY

This instruction subtracts the value of the addressed memory location from the content of index register X using the adder but does not store the result; therefore, its only use is to set the N, Z and C flags to allow for comparison between the index register X and the value in memory.

The symbolic notation is $X - M$.

The CPX instruction does not affect any register in the machine; it also does not affect the overflow flag. It causes the carry to be set on if the absolute value of the index register X is equal to or greater than the data from memory. If the value of the memory is greater than the content of the index register X, carry is reset. If the results of the subtraction contain a bit 7, then the N flag is set, if not, it is reset. If the value in memory is equal to the value in index register X, the Z flag is set, otherwise it is reset.

The addressing modes for CPX are Immediate, Absolute and Zero Page.

7.9 CPY -- COMPARE INDEX REGISTER Y TO MEMORY

This instruction performs a two's complement subtraction between the index register Y and the specified memory location. The results of the subtraction are not stored anywhere. The instruction is strictly used to set the flags.

The symbolic notation for CPY is $Y - M$.

CPY affects no registers in the microprocessor and also does not affect the overflow flag. If the value in the index register Y is equal to or greater than the value in the memory, the carry flag will be set, otherwise it will be cleared. If the results of the subtraction contain bit 7 on the N bit will be set, otherwise it will be cleared. If the value in the index register Y and the value in the memory are equal, the zero flag will be set, otherwise it will be reset.

The addressing modes for CPY are Immediate, Absolute and Zero Page.

7.10 TRANSFERS BETWEEN THE INDEX REGISTERS AND ACCUMULATOR

There are four instructions which allow the accumulator and index registers to be interchanged. They are TXA, TAX which transfer the contents of the index register X to the accumulator A and back, and TYA, TAY which transfer the contents of the index register Y to the accumulator A and back. The usefulness of these will be discussed after the instructions.

7.11 TAX -- TRANSFER ACCUMULATOR TO INDEX X

This instruction takes the value from accumulator A and transfers or loads it into the index register X without disturbing the contents of the accumulator A.

The symbolic notation for this is $A \rightarrow X$.

TAX only affects the index register X, does not affect the carry or overflow flags. The N flag is set if the resultant value in the index register X has bit 7 on, otherwise N is reset. The Z bit is set if the content of the register X is 0 as a result of the operation, otherwise it is reset. TAX is a single-byte instruction and its addressing mode is Implied.

7.12 TXA -- TRANSFER INDEX X TO ACCUMULATOR

This instruction moves the value that is in the index register X to the accumulator A without disturbing the contents of the index register X.

The symbolic notation is $X \rightarrow A$.

TXA does not affect any register other than the accumulator and does not affect the carry or overflow flag. If the result in A has bit 7 on, then the N flag is set, otherwise it is reset. If the resultant value in the accumulator is 0, then the Z flag is set, otherwise it is reset.

The addressing mode is Implied, it is a single-byte instruction.

7.13 TAY-- TRANSFER ACCUMULATOR TO INDEX Y

This instruction moves the value of the accumulator into index register Y without affecting the accumulator.

The symbolic notation is $A \rightarrow Y$.

TAY instruction only affects the Y register and does not affect either the carry or overflow flags. If the index register Y has bit then N is set, otherwise it is reset. If the content of the index register Y equals 0 as a result of the operation, Z is set on, otherwise it is

TAY is a single-byte instruction and the addressing mode is Im

7.14 TYA -- TRANSFER INDEX Y TO ACCUMULATOR

This instruction moves the value that is in the index register to accumulator A without disturbing the contents of the register Y.

The symbolic notation is $Y \rightarrow A$.

TYA does not affect any other register other than the accumulator and does not affect the carry or overflow flag. If the result in the accumulator A has bit 7 on, the N flag is set, otherwise it is reset. If the resultant value in the accumulator A is 0, then the Z flag is set, otherwise it is reset.

The addressing mode is Implied and it is a single-byte instruction.

Some of the applications of the transfer instructions between accumulator A and index registers X, Y are those when the user wishes to use the index register to access memory locations where there are multiple byte values between the addresses. In this application a count is loaded into the index register, the index register is transferred to the accumulator, a value such as 5, 7, 10, etc. is added immediate to the accumulator and results stored back into the index

register using the TAX or TAY instruction. The consequence of this type of operation is that it allows the microprocessor to address non-consecutive locations in memory. Another application is where the internal transfer instructions allow the index registers to hold intermediate values for the accumulator which allows rapid transfer to and from the accumulator to help solve high speed data shuffling problems.

7.15 SUMMARY OF INDEX REGISTER APPLICATIONS AND MANIPULATIONS

Primary use of index register X and Y is as offset and counters for data manipulation in which the index register is used to compute an address based on the value of the index register plus base address specified by the user, either in a fixed instruction format or in a variable pointer type format. In order to operate as both an offset and counter, index registers may be incremented or decremented by one or compared to values from memory. There are limitations on the applications of each of the index registers which have to do with formats which are unique to certain instruction addressing modes. Because of the ability of the index registers to be loaded, changed and stored, they are also useful as general purpose registers. They can be used as interim storages for moves between memory locations or for moves between memory and the accumulator.

One of the optimum uses of the indexing concept is the case when the index register is being used both as an offset and a counter. This type of operation makes use of the ability of the microprocessor to perform a decrement function on the index registers and set flags. Therefore, a single decrement instruction not only changes the value in the counter but can also perform a test on the count value.

CHAPTER 8

STACK PROCESSING

8.0 INTRODUCTION TO STACK AND TO PUSH DOWN STACK CONCEPT

In all of the discussions on addressing, it has been assumed that either the exact location or at least a relation to an exact location of a memory address was known.

Although this is true in most of the programming for control applications, there are certain types of programming and applications which require the basic program not to be working with known memory locations but only with a known order for accessing memory. This type of programming is called "re-entrant coding" and it is often used in servicing interrupts.

To implement this type of addressing, the microprocessor maintains a separate address generator which is used by the program to access memory. This address generator employs a push down stack concept.

Discussions of push-down stacks are usually best stated considering that if one were given three cards -- an ace, a king, and a ten -- and were told that the order of cards was important, and were then asked to lay them down on the table in the order in which they were given, ace first, the king on top of it and finally the ten, and then if they were retrieved one card at a time, the ten would be retrieved first even though it was put on last, the king would be retrieved second, and the ace would be retrieved last, even though it was put on first.

The only commands needed to implement this operation are "put next card on stack" and "pull next card from the stack." The stack could be processing clubs and then go to diamonds and back to clubs. However, we know that while we are processing clubs, we will always find the ten first, king second, etc.

The hardware implementation of the ordered card stack which was just described is a 8-bit counter, into which the address of a memory location is stored. This counter is called a "Stack Pointer." Every time data are to be pushed onto the stack, the stack pointer is put out on the address bus, data are written into the memory addressed by the stack pointer, and the stack pointer is decremented by 1 as may be seen in Example 8.1. Every time data are pulled from the stack, the stack pointer is incremented by 1. The stack pointer is put out on the address bus, and data are read from the memory location addressed by the stack pointer. This implementation using the stack pointer gives the effect of a push-down stack which is program-independent addressing.

Example 8.1: Basic Stack Map for 3-Deep JMP to Subroutine Sequence


<u>Stack Address</u>	<u>Data</u>
01FF	PCH 1
01FE	PCL 1
01FD	PCH 2
01FC	PCL 2
01FB	PCH 3
01FA	PCL 3
01F9	

In the above example, the stack pointer starts out at 01FF. The stack pointer is used to store the first state of the program counter by storing the content of program counter high at 01FF, and the content of program counter low at 01FE. The stack pointer would now be pointed at 01FD. The second time the store program count is performed, the program counter high number is stored on the stack at 01FD and the program counter low is stored at 01FC. The stack pointer would now be pointing at 01FB. The same procedure is used to store the third program counter.

When data are taken from the stack, the PCL 3 will come first and the PCH 3 will come second just by adding 1 to the stack pointer before each memory read. The example above contains the program count for 3 successive jump and store operations where the jump transfers control to a subroutine and stores the value of the program counter onto the stack in order to remember to which address the program should return after completion of the subroutine.

Following is an example of a program that would create the Example 8.1 stack operation.

Example 8.2: Basic Stack Operation



This is known as subroutine nesting and is often encountered in solving complex control equations.

To correctly use the stack for this type of operation requires a jump to subroutine and a return from subroutine instruction.

8.1 JSR -- JUMP TO SUBROUTINE

This instruction transfers control of the program counter to a subroutine location but leaves a return pointer on the stack to allow the user to return to perform the next instruction in the main program after the subroutine is complete. To accomplish this, the JSR instruction stores the program counter address, which points to the last byte of the jump instruction, onto the stack using the stack pointer. The stack byte contains the program count high first, followed by program count low. The JSR then transfers the addresses following the jump instruction to the program counter low and the program counter high, thereby directing the program to begin at that new address.

The symbolic notation for this is $PC + 24$, $(PC + 1) \rightarrow PCL$, $(PC + 2) \rightarrow PCH$.

The JSR instruction affects no flags, causes the stack pointer to be decremented by 2 and substitutes new values into the program counter low and the program counter high. The addressing mode for the JSR is always Absolute.

Example 8.3 gives the details of a JSR instruction.

Example 8.3: Illustration of JSR Instruction

Program Memory	
PC	Data
0100	JSR
0101	ADL
0102	ADH Subroutine

Stack Memory	
Stack Pointer	Stack
01FF	01
01FE	02
01FD	

Cycle	Address Bus	Data Bus	External Operations	Internal Operations
1	0100	OP CODE	Fetch Instruction	Finish Previous Operation; Increment PC to 0101
2	0101	New ADL	Fetch New ADL	Decode JSR; Increment PC to 0102
3	01FF			Store ADL
4	01FF	PCH	Store PCH	Hold ADL, Decrement S* to 01FE
5	01FE	PCL	Store PCL	Hold ADL, Decrement S* to 01FD
6	0102	ADH	Fetch ADH	Store Stack Pointer
7	ADH, ADL	New OP CODE	Fetch New OP CODE	ADL → PCL ADH → PCH

* S denotes "Stack Pointer."

In this example, it can be seen that during the first cycle the microprocessor fetches the JSR instruction. During the second cycle, address low for new program counter low is fetched. At the end of cycle 2, the microprocessor has decoded the JSR instruction and holds the address low in the microprocessor until the stack operations are complete. NOTE: The stack is always stored in Page 1 (Hex address 0100-01FF).

The operation of the stack in the R6500-series microprocessors is such that the stack pointer is always pointing at the next memory location into which data can be stored. In Example 8.3, the stack pointer is assumed to be at 01FF in the beginning and PC at location 0100. During the third cycle, the microprocessor puts the stack pointer onto the address lines and on the fourth writes the contents of the current value of the program counter high, 01, into the memory location indicated by the stack pointer address. During the time that the write is being accomplished, the stack pointer is being automatically decremented by 1 to 01FE. During the fifth cycle, the PCL is stored in the next memory location with the stack pointer being automatically decremented.

It should be noted that the program counter low, which is now stored in the stack, is pointing at the last address in the JSR sequence. This is not what would be expected as a result of a JSR instruction. It would be expected that the stack points at the next instruction. This apparent anomaly in the machine is corrected during the Return from Subroutine instruction.

NOTE: At the end of the JSR instruction, the values on the stack contain the program counter low and the program counter high which referenced the last address of the JSR instruction. Any subroutine calls which want to use the program counter as an intermediate pointer must consider this fact. It should be noted also that the Return from Subroutine instruction performs an automatic increment at the end of the RTS which means that any program counters which are substituted on the stack must be 1 byte or 1 pointer count less than the program count to which the programmer expects the RTS to return.

The advantage of delaying the accessing of the address high until after the current program counter can be written in the stack is that only the address low has to be stored in the microprocessor. This has the effect of shortening the JSR instruction by one byte and also minimizing internal storage requirements.

After both program counter low and high have been transferred to the stack, the program counter is used to access the next byte which is the address high for the JSR. During this operation, the sixth cycle, internally the microprocessor is storing the stack pointer which is now pointing at 01FD or the next location at which memory can be loaded.

During the seventh cycle, the address high from the data bus and the address low stored in the microprocessor are transferred to the new program counter and are used to access the next OP CODE, thus making JSR a 6-cycle instruction.

At the completion of the subroutine the programmer wants to return to the instruction following the Jump-to-Subroutine instruction. This is accomplished by transferring the last 2 stack bytes to the program counter which allows the microprocessor to resume operations at the instruction following the JSR, and it is done by means of the RTS instruction.

8.2 RTS -- RETURN FROM SUBROUTINE

This instruction loads the program count low and program count high from the stack into the program counter and increments the program counter so that it points to the instruction following the JSR. The stack pointer is adjusted by incrementing it twice.

The symbolic notation for the RTS is PC+, INC PC.

The RTS instruction does not affect any flags and affects only PCL and PCH. RTS is a single-byte instruction and its addressing mode is implied.

The following Example 8.4 gives the details of the RTS instruction. It is the complete reverse of the JSR shown in Example 8.3.

Example 8.4: Illustration of RTS Instruction

Program Memory

PC	Data
0300	RTS
0301	?

Stack Memory

Stack Pointer	Stack
01FF	01
01FE	02
01FD	?

Return from Subroutine (Example)

Cycle	Address Bus	Data Bus	External Operations	Internal Operations
1	0300	OP CODE	Fetch OP CODE	Finish Previous Operation, 0301 → PC
2	0301	Discarded Data	Fetch Discarded Data	Decode RTS
3	01FD	Discarded Data	Fetch Discarded Data	Increment Stack Pointer to 01FE
4	01FE	02	Fetch PCL	Increment Stack Pointer to 01FF
5	01FF	01	Fetch PCH	
6	0102	Discarded Data	Put Out PC	Increment PC by 1 to 0103
7	0103	Next OP CODE	Fetch Next OP CODE	

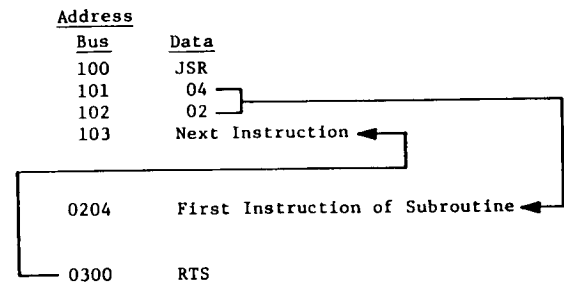
As we can see, the RTS instruction effectively unwinds what was done to the stack in the JSR instruction. Because RTS is a single-byte instruction it wastes the second memory access in doing a look-ahead operation. During the second cycle, the value located at the next program address after the RTS is read but not used in this operation. It should be noted that the stack is always left pointing at the next empty location, which means that to pull off the stack, the microprocessor has to wait 1 cycle while it adds 1 to the stack address. This is done to shorten the interrupt sequence which will be discussed below; therefore, cycle 3 is a dead cycle in which the microprocessor fetches but does not use the current value of the stack and, like the fetch of address low on Indexed and Zero Page Indexed operations, does nothing other than initialize the microprocessor to the proper state. It can be seen that the stack pointer decrements as data is pushed on to the stack and increments as data is pulled from the stack. In the fourth cycle of the RTS, the microprocessor puts out the 01FE address and reads the data stored there which is the program count low which was written in the second write cycle of the JSR. During the fifth cycle, the microprocessor puts out the incremented stack picking up the program count high which was written in the first write cycle of the JSR.

As is indicated during the discussions of JSR, the program counter stored on the stack really points to the last address of the JSR instruction itself; therefore, during the sixth cycle the RTS causes the program count from the stack to be incremented. That is the only purpose of the sixth cycle. Finally, in the seventh cycle, the incremented program counter is used to fetch the next instruction; therefore, RTS takes six cycles.

Because every subroutine requires one JSR followed by one RTS, the time to jump to and return from a subroutine is 12 cycles.

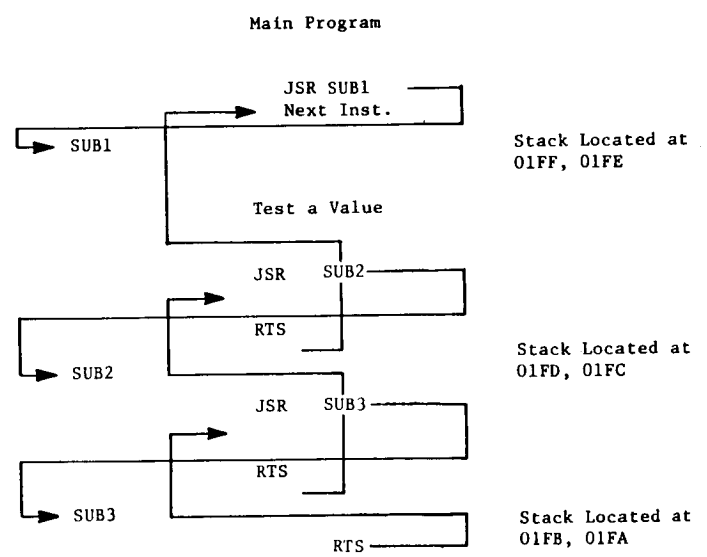
In the previous two examples, we have shown the operations of the JSR located in location 100 and the RTS located in location 300. The following pictorial diagram, Example 8.5, illustrates how the memory map for this operation might look:

Example 8.5: Memory Map for RTS Instruction



With this capability of subroutines, the microprocessor is allowed to go from the main program to 1 subroutine, to the second subroutine, to a third subroutine, and then finally to work its way back to the main program. Example 8.6 is an expansion of Example 8.2 with the returns included.

Example 8.6: Expansion of RTS Memory Map



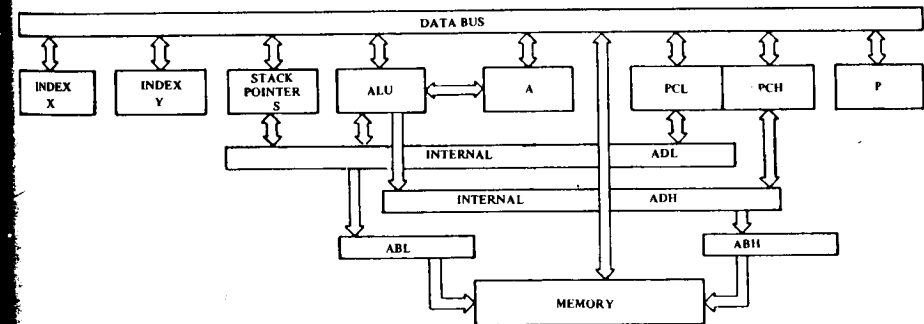
This concept is known as nesting of subroutines, and the number of subroutines which can be called in such a manner is limited only by the length of the stack.

8.3 IMPLEMENTATION OF STACK

As we have seen, the primary requirement for the stack is that irrespective of where or when a stack operation is called, the microprocessor must have an independent counter or register which contains the current memory location value of the stack address. This register is called the Stack Pointer, S. The stack becomes an auxiliary field in memory which is basically independent of programmer control. We will discuss later how the stack pointer becomes initialized, but once it is initialized, the primary requirement is that it be self-adjusted. In other words, operations which put data on the stack cause the pointer to be decremented automatically; operations which take data off the stack cause the pointer to be incremented automatically. Only under rare circumstances should the programmer find it necessary to move his stack from one location to another if he is using the stack as designed.

On this basis, there is no need for a stack to be longer than 256 bytes. To perform a single subroutine call takes only two bytes of stack memory. To perform an interrupt takes only three bytes of stack memory. Therefore, with 256 bytes, one can access 128 subroutines deep or interrupt 85 times. Therefore the length of the stack is extremely unlikely to be limiting. Microprocessors in the R6500 series have a 256-byte stack length.

Figure 8.1, which is now the complete block diagram, shows all of the microprocessor registers. The 8-bit stack pointer register, S, has been added. It is initialized by the programmer and thereafter automatically increments or decrements, depending on whether data is being put on the stack or taken off the stack by the microprocessor under control of the program or the interrupt lines.



Partial Block Diagram of a R6500 Microcomputer System Microprocessor Including Stack Pointer. S

FIGURE 8.1

The primary purpose of the stack is to furnish a block of memory locations in which the microprocessor can write data such as the program counter for use in later processing. In many control systems the requirements for read/write memory are very small and the stack just represents another demand on read/write memory. Therefore for these applications it is preferable for the stack to be in the Page Zero location in order that memory allocation for the stack, the Zero Page operations, and the indirect addresses can be performed. Therefore, one of the requirements of a stack is that it be easily locatable into Page Zero.

On the other hand, if more than one page of RAM is needed because of the amount of data that must be handled by the user programs, having the stack in Page Zero is an unnecessary waste of Page Zero memory in the sense that the stack can take no real advantage of being located in Zero Page, whereas other operations can.

In each of the examples, the stack has been located at high-order address 01 followed by a low-order address. In the same manner that the microprocessor forces locations 00 onto the high-order 8 bits of the address lines for Zero Page operations, the microprocessor automatically puts 01 Hex onto the high-order 8-bit address lines during each stack operation. This has the advantage to the user of locating the stack into Page One of memory which would be the next memory location added if the Zero Page operation requirements exceed Page Zero memory capacity. This has the advantage of the stack not requiring memory to be added specifically for the stack but only requiring the allocation of existing memory locations. It should be noted that the selected addressing concepts of the R6500-series microprocessor support devices would involve connecting the memories such that bit 8, which is the selection bit for the Page One versus Page Zero, is a "don't care" for operations in which the user does not need more than one page of Read/Write memory. This gives the user the effect of locating the stack in Page Zero for those applications.

The second feature that should be noted from the examples is that the stack was located at the end of Page One and decremented from that point towards the beginning of the page. This is the natural operation of the stack. RAM memory comes in discrete increments (64, 128, 256) bytes so the normal method of allocating stack addressing is for the user to calculate the number of bytes probably needed for stack access. This could be done by analyzing the number of subroutines which might be called and the amount of data which might be put onto the stack in order to communicate between subroutines, or the number of interrupts plus subroutines which might occur with the respective data that would be stored on the stack for each of them. By counting three bytes for each interrupt, two bytes for each jump to subroutine, plus one byte for each programmer-controlled stack operation, the microprocessor designer can estimate the amount of memory which must be allocated for the stack. This is part of his decision-making process in deciding how much memory is necessary for his whole program.

Once the allocation has been made, it is recommended that the user assign his working storage from the beginning of memory forward and always load his stack at the end of either Page Zero, Page One, or at the end of his physical memory which is located in one of those locations. This will

give the effect of having the highest bytes of memory allocated to the stack, with lower bytes of memory allocated to user working storage; hopefully, the two will never overlap.

It should be noted that the natural operation of the stack, which often is called by hardware not totally under program control, is such that it will continue to decrement throughout the page to which it is allocated irrespective of the user's desire to have it do so. A normal mistake in allocation of memory can result in the user's writing data into a memory location and later accessing it with another subroutine or another part of his program, only to find that the stack has very carefully written over that area as the result of its performing hardware control operations. This is one of the more difficult problems to diagnose. If this problem is suspected by the programmer, he should analyze memory locations higher than unexplained disturbed locations.

There is a distinctive pattern for stack operations which are unique to the user's program but which are quite predictable. An analysis of the value which has been destroyed will often indicate that it is part of an address which would normally be expected during the execution of the program between the time data were stored and the time they were fetched. This is a very strong indication of the fact that the stack somehow or other did get into the user's program area. This is almost always caused by improper control of interrupt lines or unexpected operations of interrupt or subroutine calls and has only two solutions: (1) If the operation is normal and predictable, the user must assign more memory to his program and particularly reassign his memory such that the stack has more room to operate; or (2) if the operation of the interrupt lines is not predictable, attention must be given to solving the hardware problem that causes this type of unpredictable operation.

8.3.1 Summary of Stack Implementation

The R6500 series microprocessors have a single 8-bit stack register. This register is automatically incremented and decremented under control of the microprocessor to perform stack manipulation operations, under direction of the user program or the interrupt lines. Once the programmer has initialized the stack pointer to the

end of whatever memory he wants the stack to operate in, the programmer can ignore stack addressing other than in those cases where there is an interference between stack operations and his normal program working space.

In the R6500 series, the stack is automatically located in Page One. The microprocessor always puts out the address 0100 plus stack register for every stack operation. By selected memory techniques, the user can either locate the stack in Page Zero or Page One, depending on whether or not Page One exists for his hardware.

8.4 USE OF THE STACK BY THE PROGRAMMER

Discussed in Section 8.1 was the use of the JSR to call a subroutine. However, not indicated was the technique by which the subroutine knew which data to operate on. There are three classical techniques for communicating data between subroutines. The first and most straightforward technique is that each subroutine has a defined set of working registers located in Page Zero in which the user has left values to be operated on by the subroutine. The registers can either contain the values directly or can contain indirect pointers to values which would be operated on. The following example shows the combination of these:

Example 8.7: Call-a-Move Subroutine Using Preassigned Memory Locations

		Location 10	= Count
		Location 11, 12	= Base from Address
		Location 13, 14	= Base to Address
	<u>Main Line Routine</u>		
<u>No. of</u>	<u>Instruction</u>	<u>Comment</u>	
<u>Bytes</u>			
2	LDA #Count -1	Load Fixed Value for the Move	
2	STA 10		
2	LDA #FRADH	Set up "FROM" Pointer	
2	STA 12		
2	LDA #FRADL		
2	STA 11		
2	LDA #TOADL		
2	STA 13		
2	LDA #TOADH	Set up "TO" Pointer	
2	STA 14		
3	JSR SUB1		
<u>23</u>			bytes

Subroutine Coding

<u>No. of</u>	<u>Label</u>	<u>Instruction</u>
<u>Bytes</u>		
2	SUB1	LDY 10
2	LOOP	LDA (11), Y
2		STA (13), Y
1		DEY
2		BNE LOOP
1		RTS
Total 33 Bytes		

As has been previously discussed, the loop time is the overriding consideration rather than setup time for a large number of executions.

It can be seen that we have used the techniques developed in previous sections of the indirect referencing, the jump to subroutine and the return from subroutine to perform this type of subroutine value communication. In this operation, there was no use of the stack except for the program counter value.

A second form of communication is the use of the stack itself as an intermediate storage for data which are going to be communicated to the subroutine. In order for the programmer to use the stack as an intermediate storage, he needs instructions which allow him to put data on the stack and to read from the stack. Such instructions are known as "push" and "pull" instructions.

8.5 PHA - PUSH ACCUMULATOR ON STACK

This instruction transfers the current value of the accumulator to the next location on the stack, automatically decrementing the stack to point to the next empty location.

The symbolic notation for this operation is $A \downarrow$. It should be remembered that the notation \downarrow means push to the stack, while \uparrow means pull from the stack.

The Push A instruction affects only the stack pointer register which is decremented by 1 as a result of the operation. It affects no flags.

PHA is a single-byte instruction and its addressing mode is Implied.

The following example shows the operations which occur during Push A instruction.

Example 8.3: Operation of PHA, Assuming Stack at 01FF

Cycles	Address Bus	Data Bus	External Operations	Internal Operations
1	0100	OP CODE	Fetch Instruction	Finish Previous Operation, Increment PC to 0101
2	0101	Next OP CODE	Fetch Next OP CODE and Discard	Interpret PHA Instruction, Hold P-Counter
3	01FF	(A)	Write A on Stack	Decrement Stack Pointer to 01FE
4	0101	Next OP CODE	Fetch Next OP CODE	

As can be seen, the PHA requires three cycles and takes advantage of the fact that the stack pointer is pointing to the correct location to write the value of A. As a result of this operation, the stack pointer will be setting at 01FE. The notation (A) implies contents of A. Now that the data are on the stack, later on in the program the programmer will call for the data to be retrieved from the stack with a PLA instruction.

8.6 PLA -- PULL ACCUMULATOR FROM STACK

This instruction adds 1 to the current value of the stack pointer and uses it to address the stack and loads the contents of the stack into the A register.

The symbolic notation for this is A+.

The PLA instruction does not affect the carry or overflow flags. It sets N if the bit 7 is on in accumulator A as a result of the instruction, otherwise it is reset. If accumulator A is zero as a result of the PLA, then the Z flag is set, otherwise it is reset. The PLA instruction changes the contents of accumulator A to the contents of the memory location at stack register plus 1 and also increments the stack register.

The PLA instruction is a single-byte instruction and the addressing mode is Implied.

In the following example, the data stored on the stack in Example 8.8 are transferred to the accumulator.

Example 8.9: Operation of PLA stack from Example 8.8

Cycles	Address Bus	Data Bus	External Operations	Internal Operations
1	0200	PLA	Fetch Instruction	Finish Previous Operation, Increment PC to 201
2	0201	Next OP CODE	Fetch Next OP CODE and Discard	Interpret Instruction, Hold P-Counter
3	01FE	Discarded Data	Read Stack	Increment Stack Pointer to 01FF
4	01FF	Data	Fetch Data	Save Stack
5	0201	Next OP CODE	Fetch Next OP CODE	Data → A

When data are being taken off the stack, there is one extra cycle during which time the current contents of the stack register are accessed but not used and the stack pointer is incremented by 1 to allow access to the value that was previously stored on the stack. The stack pointer is left pointing at this location because it is now considered to be an empty location to be used by the stack during a subsequent operation.

8.7 USE OF PUSHES AND PULLS TO COMMUNICATE VARIABLES BETWEEN SUBROUTINE OPERATIONS

In Example 8.10, we perform the same operation as we did in Example 8.7, except that here, instead of using fixed locations to pick up the pointers, we use the stack as a communications vehicle:

Example 8.10: Call-a-Move Subroutine Using the Stack to Communicate

<u>Main Line Routine</u>		
Bytes	Instruction	
2	LDA #Count -1	Location 11, 12 = Base "FROM" Address
1	PHA	Location 13, 14 = Base "TO" Address
2	LDA #FRADL	
1	PHA	
2	LDA #FRADH	
1	PHA	
2	LDA #TOADL	
1	PHA	
2	LDA #TOADH	
1	PHA	
3	JSR SUB1	
18		

<u>Bytes</u>	<u>Label</u>	<u>Instruction</u>	<u>Comments</u>
2	SUB1	LDX 6	
1	LOOP1	PLA	
2		STA 10,X	
1		DEX	Move Stack to Memory
2		BNE LOOP 1	
1		PLA	Set up Count
1		TAY	
2	LOOP2	LDA (11),Y	
2		STA (13),Y	Move Memory Location
1		DEY	
2		BNE LOOP 2	
2		LDA 15	
1		PHA	Restore PC to Stack
2		LDA 16	
1		PHA	
1		RTS	
Total <u>42</u> Bytes			

We can see from this example that using the stack as a communication vehicle actually increases the number of bytes in the subroutine and the total bytes overall. However, the only time one should be using subroutines in this case is when the subroutine is fairly long and is used fairly frequently. This technique does reduce the number of bytes in the calling sequence. The calling sequence is normally repeated once for every time the instruction is called; therefore the use of the stack to communicate should result in a net reduction in the number of bytes used in the total program.

Up until this time, we have been considering that the stack is at a fixed location and that all stack references use the stack pointer. It has not been explained how the stack pointer in the microprocessor gets loaded and accessed. This is done through communication between the stack pointer and index register X.

8.8 TXS -- TRANSFER INDEX X TO STACK POINTER

This instruction transfers the value in the index register X to the stack pointer.

Symbolic notation is $X \rightarrow S$.

TXS changes only the stack pointer, making it equal to the content of the index register X. It does not affect any of the flags.

TXS is a single-byte instruction and its addressing mode is Implied. Another application for TXS is the concept of passing parameters to the subroutine by storing them immediately after the jump to subroutine instruction.

In Example 8.11, the from and to address, plus the count of number of values would be written right after the JSR instruction and its address.

By locating the stack in Page Zero, the address of the last byte of the JSR can be incremented to point at the parameter bytes and then used as an indirect pointer to move the parameter to its memory location.

The key to this approach is transferring the stack pointer to X, which allows the program to operate directly on the address while it is in the stack.

It should be noted that this approach automatically leaves the address on the stack, positioned so that the RTS picks up the next OP CODE address.

Example 8.11: Jump to Subroutine (JSR) Followed by Parameters

<u>Address Bus</u>	<u>Data</u>
0100	JSR
0101	ADL
0102	ADH
0103	To High
0104	To Low
0105	From High
0106	From Low
0107	Count
0108	Next OP CODE

Before concluding this discussion on subroutines and parameter passing, it should again be noted that the use of subroutines should be limited to those cases where the user expects to duplicate code of significant length several times in the program. In these cases, and only in these cases, is subroutine call warranted rather than the normal mode of knowing the addresses and specifying them in an instruction. In all cases where timing is of significant interest, subroutines should also be avoided. Subroutines add significantly to the setup and execution time of problem solution. However, subroutines definitely have their place in microcomputer code, and three alternatives have been presented for use in application programs. The user will find a combination of the above techniques most valuable for solving his particular problem.

8.9 TSX -- TRANSFER STACK POINTER TO INDEX X

This instruction transfers the value in the stack pointer to the index register X.

Symbolic notation is $S \rightarrow X$.

TSX does not affect the carry or overflow flags. It sets N if bit 7 is on in index X as a result of the instruction, otherwise it is reset. If index X is zero as a result of the TSX, the Z flag is set, otherwise it is reset. TSX changes the value of index X, making it equal to the content of the stack pointer.

TSX is a single-byte instruction and the addressing mode is Implied.

8.10 SAVING OF THE PROCESSOR STATUS REGISTER

During the interrupt sequences, the current contents of the processor status register (P) are saved on the stack automatically. However, there are times in a program where the current contents of the P register must be saved for performing some type of other operation. A particular example of this would be the case of a subroutine which is called independently and which involves decimal arithmetic. It is important that the programmer keeps track of the arithmetic mode the program is in at all times. One way to do this is to establish the convention that the machine will always be in binary or decimal mode, with every subroutine changing its mode being responsible for restoring it back to the known state. This is a superior convention to the one that is about to be described.

A more general convention would be one in which the subroutine that wanted to change modes of operation would push P onto the stack, then set the decimal mode to perform the subroutine and then pull P back from the stack prior to returning from the subroutine.

Instructions which allow the user to accomplish this are described in Section 8.11 and 8.12.

8.11 PHP -- PUSH PROCESSOR STATUS ON STACK

This instruction transfers the contents of the processor status register unchanged to the stack, as governed by the stack pointer.

Symbolic notation for this is $P\uparrow$.

The PHP instruction affects no registers or flags in the microprocessor.

PHP is a single-byte instruction and the addressing mode is Implied.

8.12 PLP -- PULL PROCESSOR STATUS FROM STACK

This instruction transfers the next value on the stack to the Processor Status register, thereby changing all of the flags and setting the mode switches to the values from the stack.

Symbolic notation is $P\uparrow$.

The PLP instruction affects no registers in the processor other than the status register. This instruction could affect all flags in the status register.

PLP is a single-byte instruction and the addressing mode is Implied.

8.13 SUMMARY OF THE STACK

The stack in the R6500 family is a push-down stack implemented by a processor register called the stack pointer which the programmer initializes by means of a Load X immediately followed by a TXS instruction and thereafter is controlled by the microprocessor which loads data into memory based on an address constructed by adding the contents of the stack pointer to a fixed address, Hex address 0100. Every time the microprocessor loads data into memory using the stack pointer, it automatically decrements the stack pointer, thereby leaving the stack pointer pointing at the next open memory byte. Every time the microprocessor accesses data from the stack, it adds 1 to the current value of the stack pointer and reads the memory location by putting out the address 0100 plus the stack pointer. The status register is automatically pointing at the next memory location to which data can now be written. The stack makes an interesting place to store interim data without the programmer having to worry about the actual memory location in which data will be directly stored.

There are 8 instructions which affect the stack. They are: BRK, JSR, PHA, PHP, PLA, PLP, RTI, and RTS.

BRK and RTI involve the handling of the interrupts.

CHAPTER 9

RESET AND INTERRUPT CONSIDERATIONS

9.0 VECTORS

Before developing the concepts of how the R6500-series microprocessors handle interrupts and start-up, a brief definition of the concept of vector pointers should be developed.

In the sections on Jumps and Branches, it was always assumed that the program counter is changed by the microprocessor under control of the programmer while accessing addresses which were in program sequence. In order to get the microprocessor started and in order to properly handle external control or interrupt, there has been developed a different way of setting the program counter to point at a specific location. This concept is called vectored pointers. A vector pointer consists of a program counter high and program counter low value which, under control of the microprocessor, is loaded in the program counter when certain external events occur. The word vector is developed from the fact that the microprocessor directly controls the memory location from which a particular operation will fetch the program counter value and hence the concept of vector.

By allowing the programmer to specify the vector address and then by allowing the programmer to write coding that the address points to, the microprocessor makes available to the programmer all of the control necessary to develop a general purpose control program. The microprocessor has fixed addresses in memory from which it picks up the vectors. By this

implementation, minimum hardware in the microprocessor is required. Locations FFFA through FFFF are reserved for vector pointers for the microprocessor. Into these locations are stored the interrupt vectors or pointers for non-maskable interrupt, reset, and interrupt request.

9.1 RESET OR RESTART

In the microprocessor, there is a state counter which controls when the microprocessor is going to use the program counter to access memory to pick up an instruction; then, after the instruction is loaded, the microprocessor goes through a fixed sequence of interpreting instructions and develops a series of operations which are based on the OP CODE decoding.

Up to this point, it has been assumed that the program counter was set at some location, and that all program counter changes are directed by the program once the program counter had been initialized.

Instructions exist for the initialization and loading of all other registers in the microprocessor except for the initial setting of the program counter. It is for this initial setting of the program counter to a fixed location in the restart vector location specified by the microprocessor programmer that the reset line in the microprocessor is primarily used.

The reset line is controlled during power on initialization and is a common line which is connected to all devices in the microcomputer system which have to be initialized to a known state. The initialization of most I/O devices is such that they are brought up in a benign state so that with minimum coding in the microcomputer, the programmer can configure and control the I/O in an orderly fashion.

The concept has important implications in systems where damage can be done if peripheral devices came up in unknown states. Therefore, in the R6500, power on or reset control operates at two levels.

First, by holding of an external line to ground, and having this external line connected to all the devices during power-up transient conditions, the entire microcomputer system is initialized to a known disabled state. Second, the releases of the reset line from the ground or TTL zero condition to a TTL one condition causes the microprocessor to be automatically initialized, first by the internal hardware vector which causes it to be pointed to a known program location, and secondly through a software program which is written by the user to control the orderly start-up of the microcomputer system.

All of the R6500 family parts also obey a discipline that while the reset line is low, the system is in a stop or reset state. The microprocessor is guaranteed to be in a Read state and upon release of the reset line from ground to positive, the microprocessor will continue to hold the line in a Read state until it has addressed the specified vectored count location, at which time control of the microprocessor is available to the programmer.

9.2 START FUNCTION

While the reset line is in the low state, it can be assumed that internal registers may be initialized to any random condition; therefore, no conditions about the internal state of the microprocessor are assumed other than that the microprocessor will, one cycle after the reset line goes high, implement the following sequence:

Example 9.1: Illustration of Start Cycle

<u>Cycles</u>	<u>Address Bus</u>	<u>Data Bus</u>	<u>External Operation</u>	<u>Internal Operation</u>
1	?	?	Don't Care	Hold During Reset
2	? + 1	?	Don't Care	First Start State
3	0100 + SP	?	Don't Care	Second Start State
4	0100 + SP-1	?	Don't Care	Third Start State
5	0100 + SP-2	?	Don't Care	Fourth Start State
6	FFFC	Start PCL	Fetch First Vector	
7	FFFD	Start PCH	Fetch Second Vector	Hold PCL
8	PCH PCL	First OP CODE	Load First OP CODE	

The start cycle actually takes seven cycles from the time the reset line is let go to TTL plus. On the eighth cycle, the vector fetched from the memory location FFFC and FFFD is used to access the next instruction. The microprocessor is now in a normal program load sequence, the location where the vector points should be the first OP CODE which the programmer desires to perform.

The second point that should be noted is that the microprocessor actually accesses the stack three times during the start sequence in cycles 3, 4 and 5. This is because the start sequence is in effect a specialized form of interrupt with the exception that the read/write line is disabled so that no writes to stack are accomplished during any of the cycles.

9.3 PROGRAMMER CONSIDERATIONS FOR INITIALIZATION SEQUENCES

There are two major facts to remember about initialization. First, the only automatic operations of the microprocessor during reset are to turn on the interrupt disable bit and to force the program counter to the vector location specified in locations FFFC and FFFD and to load the first instruction from that location. Therefore, the first operation in any normal program will be to initialize the stack. This should be done by having previously decided what the stack value should be for initial operations and then doing a LDX immediate of this value followed by a TXS. By this simple operation, the microprocessor is ready for any interrupt or non-maskable interrupt operation which might occur during the rest of the start-up sequence.

Once this is accomplished, the two nonvariable operations of the machine are under control. The program counter is initialized and under program control, and the stack is initialized and under program control. The next operations during the initialization sequences will consist of configuring and setting up the various control functions necessary to perform the I/O desired for the microprocessor.

Specific discussion for considerations regarding the start-up are covered in Section 11.

The major things which have to be considered include the current state of the I/O device and the nondestructive operations that will allow the state to be changed to the active state.

The initialization programs mostly consist of loading accumulator A immediately with a bit pattern and storing it in the data control register of an I/O device.

NOTE : The interrupt disable is automatically set by the microprocessor during the start sequence. This is to minimize the possibility of a series of interrupts occurring during the start-up sequence because of uncontrolled external values, although it is usually possible to control interrupts as part of the configuration.

The programmer should consider two effects: 1) The non-maskable interrupt is not blockable by this technique since it would be possible to configure a device that was connected to a non-maskable interrupt and have to service the interrupt immediately; and 2) the mask must be cleared at the end of the start sequence unless the user has specific reason to inhibit interrupts after he has carried out the start-up sequence. Therefore, the next-to-last instruction of the start-up sequence should be CLI.

It should be noted that the start-up routine is a series of sequential operations which should occur only during power on initialization and is the first step in the programmed logic of the machine.

Because the execution of the routine during power on occurs very seldom in the normal operation of the machine, the coding for power on sequence should tend to minimize the use of memory space rather than speed.

The last instruction in the start-up sequence should initialize the decimal mode flag to the normal setting for the program.

The next instruction should be the beginning of the user's normal programming for his device, everything preceding that being known as "housekeeping."

9.4 RESTART

It should be noted that the basic microprocessor control philosophy allows for a single common reset line which initializes all devices. This line can be used to clear the microprocessor to a known state and to reset all peripherals to a known state; therefore, it can be used as a result of power interruption, during the power on sequence, or as an external clear by the user to re-initialize the system.

As discussed in the hardware manual, restart is often used as an aid to making sure the microprocessor has been properly interconnected and that programs have been loaded in the correct locations.

9.5 INTERRUPT CONSIDERATIONS

Up until this point, the microprocessor has to proceed under programmer control through a variety of sequences. The only way for the programmer to change the sequence of operations of the microprocessor was to change the program counter location to point at new operations. The microprocessor is in control of fetching the next instruction at the conclusion of the current instruction. The only way that external events could control the microprocessor, if it were not for interrupts, would be for the programmer to periodically interrupt or stop processing data and check to see whether or not an external event which might cause him to change his direction has occurred. The problem with this technique is that

I/O events are usually asynchronous, i.e., not timed with the microprocessor internal instructions; therefore, it would be possible for the event to occur shortly after the programmer has stopped to look at I/O events which would mean that the event would not be sampled until the programmer took the time to stop his coding and sample again.

Because the sampling of I/O devices normally takes several byte counts or cycles to accomplish, the frequent insertion of checking routines into straight line code results in significant delays to the entire program. In trying to use this technique, there has to be a tradeoff between the fact that the program wastes a significant amount of time checking events which have not yet occurred versus delaying checking of an event which has occurred and if not timely serviced the data may be lost.

In order to solve this dichotomy, the concept of interrupt is utilized to signal the microprocessor that an external event has occurred and the microprocessor should devote attention to it immediately. This technique accomplishes processing in which the microprocessor's program is interrupted and the event that caused the interrupt is serviced.

Transferring most of data and control to I/O devices in an interrupt driven environment will usually result in maximum program and/or programmer efficiency. Each event is serviced when it occurs, which means that there is a minimum amount of delay in servicing events, also a minimum amount of coding because of elimination of the need to determine occurrence of several events simultaneously; each interrupting event is handled as a unique combination. It is possible to interrupt an interrupt processing routine and, therefore, all the interrupt logic uses the stack which allows processing of successive interrupts without any penalty other than increasing the stack length.

A real-world example of an event which should interrupt is when the user is given a panic button indicating to the microcomputer some event has occurred that requires total immediate attention of the microprocessor to solving that problem.

The action and events are as follows: The microprocessor user pushes the panic button; the panic switch sensor causes an external device to indicate to the microprocessor an interrupt is desired; the microprocessor checks the status of the internal interrupt inhibit signal; if the internal inhibit is set, then the interrupt is ignored. However, if it is reset, or when it becomes reset through some program action, the following set of operations occur:

Example 9.2: Interrupt Sequence

<u>Cycles</u>	<u>Address Bus</u>	<u>Data Bus</u>	<u>External Operation</u>	<u>Internal Operation</u>
1	PC	OP CODE	Fetch OP CODE	Hold Program Counter, Finish Previous Operation
2	PC	OP CODE	Fetch OP CODE	Force a BRK Instruction, Hold P-Counter
3	01FF	PCH	Store PCH on Stack	Decrement Stack Pointer to 01FE
4	01FE	PCL	Store PCL on Stack	Decrement Stack Pointer to 01FD
5	01FD	P	Store P on Stack	Decrement Stack Pointer to 01FC
6	FFFE	New PCL	Fetch Vector Low	Put Away Stack Vector Low →
7	FFFF	New PCH	Fetch Vector High	PCL and Set I
8	Vector PCH PCL	OP CODE	Fetch Interrupt Program	Increment PC to PC + 1

As can be seen in Example 9.2, the microprocessor uses the stack to save the reentrant or recovery code and then uses the interrupt vectors FFFE and FFFF, (or FFFA and FFFB), depending on whether or not an interrupt request or a non-maskable interrupt request had occurred. It should be noted that the interrupt disable is turned on automatically at this point by the microprocessor.

Because the interrupt disable had to be off for an interrupt request to have been honored, the return from interrupt which loads the processor status from before the interrupt occurred has the effect of clearing the interrupt disable bit. After the interrupt has been acknowledged by the microprocessor by transferring to the proper vector location, there are a variety of operations which the user can perform to service the interrupt; however, all operations should end with a single instruction which reinitializes the microprocessor back to the point at which the interrupt occurred. This instruction is called the RTI instruction.

9.6 RTI - RETURN FROM INTERRUPT

This instruction transfers from the stack into the microprocessor the processor status and the program counter location for the instruction which was interrupted. By virtue of the interrupt having stored this data before executing the instruction, and due to the fact that the RTI reinitializes the microprocessor to the same state as when it was interrupted, the combination of interrupt plus RTI allows truly reentrant coding.

The symbolic notation for RTI is P↑ PC↑.

The RTI instruction reinitializes all flags to the position to the point at which they were when the interrupt was taken, and sets the program counter back to its pre-interrupt state. It affects no other registers in the microprocessor.

RTI is a single-byte instruction and its addressing mode is Implied.

In the following example, we can see the internal operation of the RTI which restores the microprocessor: