

Fig. 27. Am2910 block diagram.

### Architecture of the Am2910

The Am2910 is a bipolar microprogram controller intended for use in high-speed microprocessor applications. It allows addressing of up to 4096 words of microprogram. A block diagram of the Am2910 is shown in Fig. 27, and its application in a microcomputer is depicted in Fig. 28.

The controller contains a four-input multiplexer that is used to select either the register/counter, direct input, microprogram counter, or stack as the source of the next microinstruction address.

The register/counter consists of 12 *D*-type, edge-triggered flip-flops, with a common clock enable. When its load control,  $\overline{RLD}$ , is LOW, new data is loaded on a positive clock transition. A few instructions include load; in most systems, these instructions will be sufficient, simplifying the microcode. The output of the register/counter is available to the multiplexer as a source for the next microinstruction address. The direct input furnishes a source of data for loading the register/counter.

The Am2910 contains a microprogram counter ( $\mu PC$ ) that is

composed of a 12-bit incrementer followed by a 12-bit register. The  $\mu PC$  can be used in either of two ways: When the carry-in to the incrementer is HIGH, the microprogram register is loaded on the next clock cycle with the current *Y* output word plus one ( $Y + 1 \rightarrow \mu PC$ ). Sequential microinstructions are thus executed. When the carry-in is LOW, the incrementer passes the *Y* output word unmodified so that  $\mu PC$  is reloaded with the same *Y* word on the next clock cycle ( $Y \rightarrow \mu PC$ ). The same microinstruction is thus executed any number of times.

The third source for the multiplexer is the direct (*D*) input. This source is used for branching.

The fourth source available at the multiplexer input is a 5-word by 12-bit stack (file). The stack is used to provide return address linkage when executing microsubroutines or loops. The stack contains a built-in stack pointer (*SP*) which always points to the last file word written. This allows stack reference operations (looping) to be performed without a pop.

The stack pointer operates as an up/down counter. During microinstructions 1, 4, and 5, the PUSH operation is performed. This causes the stack pointer to increment and the file to be written with the required return linkage. On the cycle following the PUSH, the return data is at the new location pointed to by the stack pointer.

During five microinstructions, a POP operation may occur. The stack pointer decrements at the next rising clock edge following a POP, effectively removing old information from the top of the stack.

The stack pointer linkage is such that any sequence of pushes, pops, or stack references can be achieved. At RESET (instruction 0), the depth of nesting becomes 0. For each PUSH, the nesting depth increases by 1; for each POP, the depth increases by 1. The depth can grow to 5. After a depth of 5 is reached,  $\overline{FULL}$  goes LOW. Any further PUSHes onto a full stack overwrite information at the top of the stack but leave the stack pointer unchanged. This operation will usually destroy useful information and is normally avoided. A POP from an empty stack may place non-meaningful data on the *Y* outputs but is otherwise safe. The stack pointer remains at 0 whenever a POP is attempted from a stack already empty.

The register/counter is operated during three microinstructions (8, 9, and 15) as a 12-bit down-counter, with result = zero available as a microinstruction branch test criterion. This provides efficient iteration of microinstructions. The register/counter is arranged so that if it is preloaded with a number *n* and then used as a loop termination counter, the sequence will be executed exactly  $n + 1$  times. During instruction 15, a three-way branch under combined control of the loop counter and the condition code is available.

The device provides three-state *Y* outputs. These can be particularly useful in designs requiring automatic checkout of the



processor. The microprogram controller outputs can be forced into the high-impedance state, and pre-programmed sequences of microinstructions can be executed via external access to the address lines.

## Operation

Table 6 shows the result of each instruction in controlling the multiplexer which determines the Y outputs, and in controlling the three enable signals  $\overline{PL}$ ,  $\overline{MAP}$ , and  $\overline{VECT}$ . The effect on the register/counter and the stack after the next positive-going clock edge is also shown. The multiplexer determines which internal source drives the Y outputs. The value loaded into  $\mu PC$  is either identical to the Y output or else 1 greater, as determined by CI. For each instruction, one and only one of the three outputs  $\overline{PL}$ ,  $\overline{MAP}$ , and  $\overline{VECT}$  is LOW. If these outputs control three-state enables for the primary source of microprogram jumps (usually part of a pipeline register), a PROM which maps the instruction to a microinstruction starting location, and an optional third source (often a vector from a DMA or interrupt source), respectively, the three-state sources can drive the D inputs without further logic.

Several inputs, as shown in Table 7, can modify instruction execution. The combination  $\overline{CC}$  HIGH and  $\overline{CCEN}$  LOW is used as a test in 10 of the 16 instructions.  $\overline{RLD}$ , when LOW, causes the D input to be loaded into the register/counter, overriding any HOLD or DEC operation specified in the instruction. OE, normally LOW, may be forced HIGH to remove the Am2910 Y outputs from a three-state bus.

## The Am2910 Instruction Set

The Am2910 provides 16 instructions which select the address of the next microinstruction to be executed. Four of the instructions are unconditional—their effect depends only on the instruction. Ten of the instructions have an effect which is partially controlled by an external, data-dependent condition. Three of the instructions have an effect which is partially controlled by the contents of the internal register/counter. The instruction set is shown in Table 6. In this discussion it is assumed that  $C_n$  is tied HIGH.

In the 10 conditional instructions, the result of the data-dependent test is applied to  $\overline{CC}$ . If the  $\overline{CC}$  input is LOW, the test is considered to have been passed, and the action specified in the name occurs; otherwise, the test has failed and an alternate (often simply the execution of the next sequential microinstruction) occurs. Testing of  $\overline{CC}$  may be disabled for a specific microinstruction by setting  $\overline{CCEN}$  HIGH, which unconditionally forces the action specified in the name; that is, it forces a pass. Other ways of using  $\overline{CCEN}$  include (1) tying it HIGH, which is useful if no microinstruction is data-dependent; (2) tying it LOW if data-

dependent instructions are never forced unconditionally; or (3) tying it to the source of Am2910 instruction bit  $I_0$ , which leaves instructions 4, 6, and 10 as data-dependent but makes others unconditional. All of these tricks save one bit of microcode width.

The effect of three instructions depends on the contents of the register/counter. Unless the counter holds a value of zero, it is decremented; if it does hold zero, it is held and a different microprogram next address is selected. These instructions are useful for executing a microinstruction loop a known number of times. Instruction 15 is affected both by the external condition code and the internal register/counter.

Perhaps the best technique for understanding the Am2910 is to simply take each instruction and review its operation. In order to provide some feel for the actual execution of these instructions, Fig. 29 is included and depicts examples of all 16 instructions.

The examples given in Fig. 29 should be interpreted in the following manner: The intent is to show microprogram flow as various microprogram memory words are executed. For example, the CONTINUE instruction, instruction 14, as shown in Fig. 29, simply means that the contents of microprogram memory word 50 are executed and then the contents of word 51 are executed. This is followed by the contents of microprogram memory word 52 and the contents of microprogram memory word 53. The microprogram addresses used in the examples were arbitrarily chosen and have no meaning other than to show instruction flow. The exception to this is the first example, JUMP ZERO, which forces the microprogram location counter to address ZERO. Each dot refers to the time that the contents of the microprogram memory word is in the pipeline register. While no special symbology is used for the conditional instructions, the test to follow will explain what the conditional choices are in each example.

It might be appropriate at this time to mention that AMD has a microprogram assembler called AMDASM, which has the capability of using the Am2910 instructions in symbolic representation. AMDASM's Am2910 instruction symbolics (or mnemonics) are given in Fig. 29 for each instruction and are also shown in Table 6.

*Instruction 0.* JZ (JUMP and ZERO, or RESET) unconditional specifies that the address of the next microinstruction is zero. Many designs use this feature for power-up sequences and provide the power-up firmware beginning at microprogram memory word location 0.

*Instruction 1* is a CONDITIONAL JUMP-TO-SUBROUTINE via the address provided in the pipeline register. As shown in Fig. 29, the machine might have executed words at addresses 50, 51, and 52. When the contents of address 52 are in the pipeline register, the next address control function is the CONDITIONAL JUMP-TO-SUBROUTINE. Here, if the test is passed, the next instruction executed will be the contents of microprogram memory location 90. If the test has failed, the JUMP-TO-

Table 6 Instructions

Hex I <sub>5</sub> -I <sub>0</sub>	Mnemonic	Name	Reg/ cntr con- tents	Fail CCEN ≠ LOW and CC ≠ HIGH		PASS CCEN = High or CC = low		Reg/ cntr	Enable
				Y	STACK	Y	STACK		
0	JZ	JUMP ZERO	X	0	CLEAR	0	CLEAR	HOLD	PL
1	CJS	COND JSH PL	X	PC	HOLD	D	PUSH	HOLD	PL
2	JMAP	JUMP MAP	X	D	HOLD	D	HOLD	HOLD	MAP
3	CJP	COND JUMP PL	X	PC	HOLD	D	HOLD	HOLD	PL
4	PUSH	PUSH COND LD CNTR	X	PC	PUSH	PC	PUSH	†	PL
5	JSRP	COND JSB R/PL	X	R	PUSH	D	PUSH	HOLD	PL
6	CJV	COND JUMP VECTOR	X	PC	HOLD	D	HOLD	HOLD	VECT
7	JRP	COND JUMP R/PL	X	R	HOLD	D	HOLD	HOLD	PL
8	RFCT	REPEAT LOOP, CNTR ≠ 0	≠ 0	F	HOLD	F	HOLD	DEC	PL
			= 0	PC	POP	PC	POP	HOLD	PL
9	RPCT	REPEAT LOOP, CNTR ≠ 0	≠ 0	D	HOLD	D	HOLD	DEC	PL
			= 0	PC	HOLD	PC	HOLD	HOLD	PL
A	CRTN	COND RTN	X	PC	HOLD	F	POP	HOLD	PL
B	CJPP	COND JUMP PL & POP	X	PC	HOLD	D	POP	HOLD	PL
C	LDCT	LD CNTR & CONTINUE	X	PC	HOLD	PC	HOLD	LOAD	PL
D	LOOP	TEST END LOOP	X	F	HOLD	PC	POP	HOLD	PL
E	CONT	CONTINUE	X	PC	HOLD	PC	HOLD	HOLD	PL
F	TWB	THREE-WAY BRANCH	≠ 0	F	HOLD	PC	POP	DEC	PL
			= 0	D	POP	PC	POP	HOLD	PL

† If CCEN = LOW and CC = HIGH, hold; else load. X = Don't Care

Table 7 Pin Functions

Abbreviation	Name	Function
$D_i$	Direct Input Bit $i$	Direct input to register/counter and multiplexer. $D_0$ is LSB.
$I_i$	Instruction Bit $i$	Selects one-of-sixteen instructions for the AM 2910.
$\overline{CC}$	Condition Code	Used as test criterion. Pass test is a LOW on $\overline{CC}$ .
$\overline{CCEN}$	Condition Code Enable	Whenever the signal is HIGH, $\overline{CC}$ is ignored and the part operates as though $\overline{CC}$ were true (LOW).
$\overline{CI}$	Carry-In	Low order carry input to incrementer for microprogram counter.
$\overline{RLD}$	Register Load	When LOW forces loading of register/counter regardless of instruction or condition.
$\overline{OE}$	Output Enable	Three-state control of $Y_i$ outputs.
$\overline{CP}$	Clock Pulse	Triggers all internal state changes at LOW-to-HIGH edge.
$V_{cc}$	+5 Volts	
$GND$	Ground	
$Y_i$	Microprogram Address Bit $i$	Address to microprogram memory. $Y_0$ is LSB, $Y_{11}$ is MSB.
$\overline{FULL}$	FULL	Indicates that five items are on the stack.
$\overline{PL}$	Pipeline Address Enable	Can select #1 source (usually Pipeline Register) as direct input source.
$\overline{MAP}$	Map Address Enable	Can select #2 source (usually Mapping PROM or PLA) as direct input source.
$\overline{VECT}$	Vector Address Enable	Can select #3 source (for example, Interrupt Starting Address) as direct input source.

SUBROUTINE will not be executed; the contents of microprogram memory location 53 will be executed instead. Thus, the CONDITIONAL JUMP-TO-SUBROUTINE instruction at location 52 will cause the instruction either in location 90 or in location 53 to be executed next. If the TEST input is such that location 90 is selected, value 53 will be pushed onto the internal stack. This provides the return linkage for the machine when the subroutine beginning at location 90 is completed. In this example, the subroutine was completed at location 93 and a RETURN-FROM-SUBROUTINE was found at location 93.

*Instruction 2* is the JUMP MAP instruction. This is an unconditional instruction which causes the  $\overline{MAP}$  output to be enabled so that the next microinstruction location is determined by the address supplied via the mapping PROMs. Normally, the JUMP MAP instruction is used at the end of the instruction fetch sequence for the machine. In the example of Fig. 29, microinstructions at locations 50, 51, 52, and 53 might have been the fetch sequence, and at its completion at location 53, the jump map function would be contained in the pipeline register. This example shows the mapping PROM outputs to be 90; therefore, an unconditional jump to microprogram memory address 90 is performed.

*Instruction 3*, CONDITIONAL JUMP PIPELINE, derives its branch address from the pipeline register branch address value ( $BR_0$ - $BR_{11}$  in Fig. 28). This instruction provides a technique for

branching to various microprogram sequences depending upon the test condition inputs. Quite often, state machines are designed which simply execute tests on various inputs waiting for the condition to come true. When the true condition is reached, the machine then branches and executes a set of microinstructions to perform some function. This usually has the effect of resetting the input being tested until some point in the future. Figure 29 shows the conditional jump via the pipeline register address at location 52. When the contents of microprogram memory word 52 are in the pipeline register, the next address will be either location 53 or location 30 in this example. If the test is passed, the value currently in the pipeline register (3) will be selected. If the test fails, the next address selected will be contained in the microprogram counter, which in this example is 53.

*Instruction 4* is the PUSH/CONDITIONAL LOAD COUNTER instruction and is used primarily for setting up loops in microprogram firmware. In Figure 29, when instruction 52 is in the pipeline register, a PUSH will be made onto the stack and the counter will be loaded on the basis of the condition. When a PUSH occurs, the value pushed is always the next sequential instruction address. In this case, the address is 53. If the test fails, the counter is not loaded; if it is passed, the counter is loaded with the value contained in the pipeline register branch address field. Thus, a single microinstruction can be used to set up a loop to be executed a specific number of times. Instruction 8 will describe how to use the pushed value and the register/counter for looping.

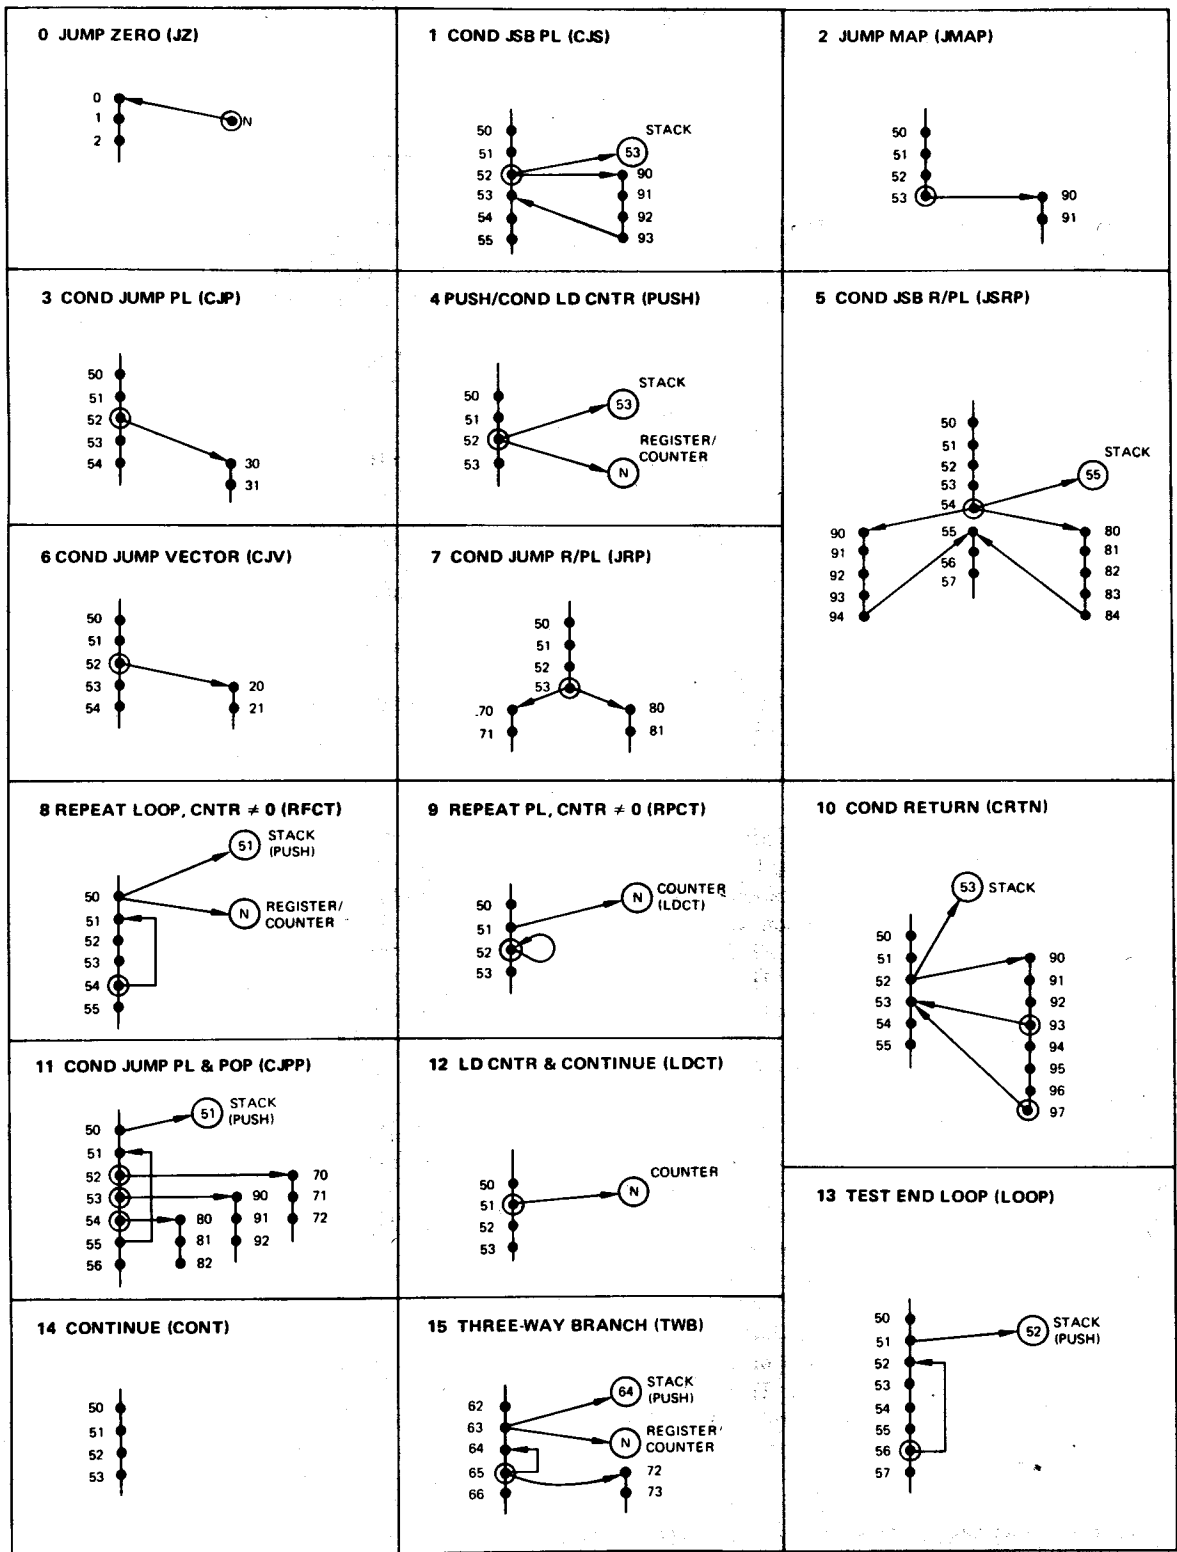


Fig. 29. Am2910 execution examples.

*Instruction 5* is a **CONDITIONAL JUMP-TO-SUBROUTINE** via the register/counter or the contents of the Pipeline register. As shown in Fig. 29, a **PUSH** is always performed and one of two subroutines executed. In this example, either the subroutine beginning at address 80 or the subroutine beginning at address 90 will be performed. A return-from subroutine (instruction 10) returns the microprogram flow to address 55. In order for this microinstruction control sequence to operate correctly, both the next-address fields of instruction 53 and the next-address fields of instruction 54 have to contain the proper value. Let us assume that the branch address fields of instruction 53 contain the value 90 so that it will be in the Am2910 register/counter when the contents of address 54 are in the pipeline register. This requires that the instruction at address 53 load the register/counter. Now, during the execution of instruction 5 (at address 54), if the test fails, the contents of the register (value = 90) will select the address of the next microinstruction. If the test input passes, the pipeline register contents (value = 80) will determine the address of the next microinstruction. Therefore, this instruction provides the ability to select one of two subroutines to be executed based on a test condition.

*Instruction 6* is a **CONDITIONAL JUMP VECTOR** instruction which provides the capability to take the branch address from a third source heretofore not discussed. In order for this instruction to be useful, the Am2910 output,  $\overline{\text{VECT}}$ , is used to control a three-state control input of a register, buffer, or PROM containing the next microprogram address. This instruction provides one technique for performing interrupt-type branching at the microprogram level. Since this instruction is conditional, a pass causes the next address to be taken from the vector source, while failure causes the next address to be taken from the microprogram counter. In the example of Fig. 29, if the **CONDITIONAL JUMP VECTOR** instruction is contained at location 52, execution will continue at vector address 20 if the **TEST** input is **HIGH** and the microinstruction at address 53 will be executed if the **TEST** input is **LOW**.

*Instruction 7* is a **CONDITIONAL JUMP** via the contents of the Am2910 register/counter or the contents of the pipeline register. This instruction is very similar to instruction 5, the **CONDITIONAL JUMP-TO-SUBROUTINE** via **R** or **PL**. The major difference between instruction 5 and instruction 7 is that no push onto the stack is performed with 7. Figure 29 depicts this instruction as a branch to one of two locations depending on the test condition. The example assumes the pipeline register contains the value 70 when the contents of address 52 are being executed. As the contents of address 53 are clocked into the pipeline register, the value 70 is loaded into the register/counter in the Am2910. The value 80 is available when the contents of address 53 are in

the pipeline register. Thus, control is transferred to either address 70 or address 80, depending on the test condition.

*Instruction 8* is the **REPEAT LOOP, COUNTER  $\neq$  ZERO** instruction. This microinstruction makes use of the decrementing capability of the register/counter. To be useful, some previous instruction, such as 4, must have loaded a count value into the register/counter. This instruction checks to see whether the register/counter contains a non-zero value. If so, the register/counter is decremented, and the address of the next microinstruction is taken from the top of the stack. If the register/counter contains zero, the loop exit condition is occurring; control falls through to the next sequential microinstruction by selecting  $\mu\text{PC}$ ; the stack is **POPPed** by decrementing the stack pointer, but the contents of the top of the stack are thrown away.

An example of the **REPEAT LOOP, COUNTER  $\neq$  ZERO** instruction is shown in Fig. 29. In this example, location 50 most likely would contain a **PUSH/CONDITIONAL LOAD COUNTER** instruction which would have caused address 51 to be **PUSHed** onto the stack and the counter to be loaded with the proper value for looping the desired number of times.

In this example, since the loop test is made at the end of the instructions to be repeated (microaddress 54), the proper value to be loaded by the instructions at address 50 is one less than the desired number of passes through the loop. This method allows a loop to be executed 1 to 4096 times. If it is desired to execute the loop from 0 to 4095 times, the firmware should be written to make the loop exit test immediately after loop entry.

Single-microinstruction loops provide a highly efficient capability for executing a specific microinstruction a fixed number of times. Examples include fixed rotates, byte swap, fixed-point multiply, and fixed-point divide.

*Instruction 9* is the **REPEAT PIPELINE REGISTER, COUNTER  $\neq$  ZERO** instruction. This instruction is similar to instruction 8 except that the branch address now comes from the pipeline register rather than the file. In some cases, this instruction may be thought of as a one-word file extension; that is, by using this instruction, a loop with the counter can still be performed when subroutines are nested five deep. This instruction's operation is very similar to that of instruction 8. The differences are that on this instruction, a failed test condition causes the source of the next microinstruction address to be the **D** inputs; and, when the test condition is passed, this instruction does not perform a **POP** because the stack is not being used.

In the example of Fig. 29, the **REPEAT PIPELINE, COUNTER  $\neq$  ZERO** instruction is instruction 52 and is shown as a single microinstruction loop. The address in the pipeline register would be 52. Instruction 51 in this example could be the **LOAD COUNTER AND CONTINUE** instruction (instruction 12). While

the example shows a single microinstruction loop, by simply changing the address in a pipeline register, multi-instruction loops can be performed in this manner for a fixed number of times as determined by the counter.

*Instruction 10* is the conditional RETURN-FROM-SUBROUTINE instruction. As the name implies, this instruction is used to branch from the subroutine back to the next microinstruction address following the subroutine call. Since this instruction is conditional, the return is performed only if the test is passed. If the test is failed, the next sequential microinstruction is performed. The example in Fig. 29 depicts the use of the conditional RETURN-FROM-SUBROUTINE instruction in both the conditional and the unconditional modes. This example first shows a JUMP-TO-SUBROUTINE at instruction location 52, where control is transferred to location 90. At location 93, a conditional RETURN-FROM-SUBROUTINE instruction is performed. If the test is passed, the stack is accessed and the program will transfer to the next instruction at address 53. If the test is failed, the next microinstruction at address 94 will be executed. The program will continue to address 97, where the subroutine is complete. To perform an unconditional RETURN-FROM-SUBROUTINE, the conditional RETURN-FROM-SUBROUTINE instruction is executed unconditionally; the microinstruction at address 97 is programmed to force  $\overline{\text{CCEN}} \text{ HIGH}$ , disabling the test, and the forced PASS causes an unconditional return.

*Instruction 11* is the CONDITIONAL JUMP PIPELINE register address and POP stack instruction. This instruction provides another technique for loop termination and stack maintenance. The example in Fig. 29 shows a loop being performed from address 55 back to address 51. The instructions at locations 52, 53, and 54 are all conditional JUMP and POP instructions. At address 52, if the TEST input is passed, a branch will be made to address 70 and the stack will be properly maintained via a POP. Should the test fail, the instruction at location 53 (the next sequential instruction) will be executed. Likewise, at address 53, either the instruction at 90 or 54 will be subsequently executed, depending on whether the test has been passed or failed. The instruction at 54 follows the same rules, going to either 80 or 55. An instruction sequence as described here, using the CONDITIONAL JUMP PIPELINE and POP instruction, is very useful when several inputs are being tested and the microprogram is looping waiting for any of the inputs being tested to occur before proceeding to another sequence of instructions. This provides the powerful jump-table programming technique at the firmware level.

*Instruction 12* is the LOAD COUNTER AND CONTINUE instruction, which simply enables the counter to be loaded with the value at its parallel inputs. These inputs are normally

connected to the pipeline branch address field which (in the architecture being described here) serves to supply either a branch address or a counter value, depending upon whether the microinstruction has been executed. There are altogether three ways of loading the counter: the explicit load by this instruction 12, the conditional load included as part of instruction 4, and the use of the  $\overline{\text{RLD}}$  input along with any instruction. The use of  $\overline{\text{RLD}}$  with any instruction overrides any counting or decrementation specified in the instruction, calling for a load instead. Its use provides additional microinstruction power, at the expense of one bit of microinstruction width. This instruction 12 is exactly equivalent to the combination of instruction 14 and  $\overline{\text{RLD}} \text{ LOW}$ . Its purpose is to provide a simple capability to load the register/counter in those implementations which do not provide microprogrammed control for  $\overline{\text{RLD}}$ .

*Instruction 13* is the TEST END-OF-LOOP instruction, which provides the capability of conditionally exiting a loop at the bottom; that is, this is a conditional instruction that will cause the microprogram to loop, via the file, if the test is failed or else to continue to the next sequential instruction. The example in Fig. 29 shows the TEST END-OF-LOOP microinstruction at address 56. If the test fails, the microprogram will branch to address 52. Address 52 is on the stack because a PUSH instruction has been executed at address 51. If the test is passed at instruction 56, the loop is terminated and the next sequential microinstruction at address 57 is executed, which also causes the stack to be POPped, thus accomplishing the required stack maintenance.

*Instruction 14* is the CONTINUE instruction, which simply causes the microprogram counter to increment so that the next sequential microinstruction is executed. This is the simplest microinstruction of all and should be the default instruction which the firmware requests whenever there is nothing better to do.

*Instruction 15*, THREE-WAY BRANCH, is the most complex. It provides for testing of both a data-dependent condition and the counter during one microinstruction and provides for selecting among one of three microinstruction addresses as the next microinstruction to be performed. Like instruction 8, a previous instruction will have loaded a count into the register/counter while pushing a microbranch address onto the stack. Instruction 15 performs a decrement-and-branch-until-zero function similar to instruction 8. The next address is taken from the top of the stack until the count reaches zero; then the next address comes from the pipeline register. The above action continues as long as the test condition fails. If at any execution of instruction 15 the test condition is passed, no branch is taken; the microprogram counter register furnishes the next address. When the loop is ended, either because the count has become zero or because the

conditional test has been passed, the stack is POPped by decrementing the stack pointer, since interest in the value contained at the top of the stack is then complete.

The application of instruction 15 can enhance performance of a variety of machine-level instructions, for instance: (1) a memory search instruction to be terminated either by finding a desired memory content or by reaching the search limit, (2) variable-field-length arithmetic terminated early upon finding that the content of the portion of the field still unprocessed is all zeros, (3) key search in a disc controller processing variable-length records, and (4) normalization of a floating-point number.

As one example, consider the case of a memory search instruction. As shown in Fig. 29, the instruction at microprogram address 63 can be instruction 4 (PUSH), which will push the value 64 onto the microprogram stack and load the number  $n$ , which is

one less than the number of memory locations to be searched before giving up. Location 64 contains a microinstruction which fetches the next operand from the memory area to be searched and compares it with the search key. Location 65 contains a microinstruction which tests the result of the comparison and also is a THREE-WAY BRANCH for microprogram control. If no match is found, the test fails and the microprogram goes back to location 64 for the next operand address. When the count becomes zero, the microprogram branches to location 72, which does whatever is necessary if no match is found. If a match occurs on any execution of the THREE-WAY BRANCH at location 65, control falls through to location 66, which handles this case. Whether the instruction ends by finding a match or not, the stack will have been POPped once, removing the value 64 from the top of the stack.