



Section 2. CPU

HIGHLIGHTS

This section of the manual contains the following topics:

2.1	Introduction	2-2
2.2	Programmer's Model.....	2-4
2.3	Software Stack Pointer (SSP)	2-7
2.4	CPU Register Descriptions	2-10
2.5	Arithmetic Logic Unit (ALU).....	2-13
2.6	Multiplication and Divide Support.....	2-14
2.7	Compiler Friendly Architecture.....	2-17
2.8	Multi-Bit Shift Support	2-17
2.9	Instruction Flow Types	2-18
2.10	Program Flow Loop Control	2-21
2.11	Address Register Dependencies	2-23
2.12	Register Maps.....	2-26
2.13	Related Application Notes.....	2-27
2.14	Revision History	2-28

2.1 INTRODUCTION

The PIC24H CPU has a 16-bit (data) modified Harvard architecture with an enhanced instruction set. The CPU has a 24-bit instruction word with a variable length opcode field. The Program Counter (PC) is 24 bits wide and addresses up to 4M x 24 bits of user program memory space.

A single-cycle instruction prefetch mechanism is used to help maintain throughput and provides predictable execution. All instructions execute in a single cycle, except the instructions that change the program flow, the double-word move (`MOV.D`) instruction, table instructions and also instructions accessing Program Space Visibility (PSV) take more than one cycle. Overhead-free program loop constructs are supported using the `REPEAT` instructions, which are interruptible at any point.

2.1.1 Registers

The PIC24H devices have sixteen 16-bit working registers in the programmer's model. Each working register can act as a data, address or address offset register. The sixteenth working register (W15) operates as a Software Stack Pointer (SSP) for interrupts and calls. The fifteenth working register (W14) can be used as a Stack Frame Pointer when used with `LNK` and `UNLK` instructions.

The data space can be addressed as 32K words or 64 Kbytes. The upper 32 Kbytes of the data space memory map can optionally be mapped into program space at any 16K word program boundary defined by the 8-bit Program Space Visibility Page (PSVPAG) register. The data to Program Space Mapping feature lets any instruction access program space as if it were data space. Refer to **Section 4. "Program Memory"** (DS70238), for more details on Program Space Visibility.

2.1.2 Addressing Modes

The CPU supports the following Addressing modes:

- Inherent (no operand)
- Relative
- Literal
- Memory Direct
- Register Direct
- Register Indirect

Each instruction is associated with a predefined addressing mode group depending on its functional requirements. As many as six addressing modes are supported for each instruction. For most instructions, the PIC24H CPU can execute the following functions in a single instruction cycle:

- Data (or program data) memory read
- Working register (data) read
- Data memory write
- Program (instruction) memory read

Therefore, three operand instructions can be supported allowing $A + B = C$ operations to be executed in a single cycle.

2.1.3 Arithmetic and Logic Unit

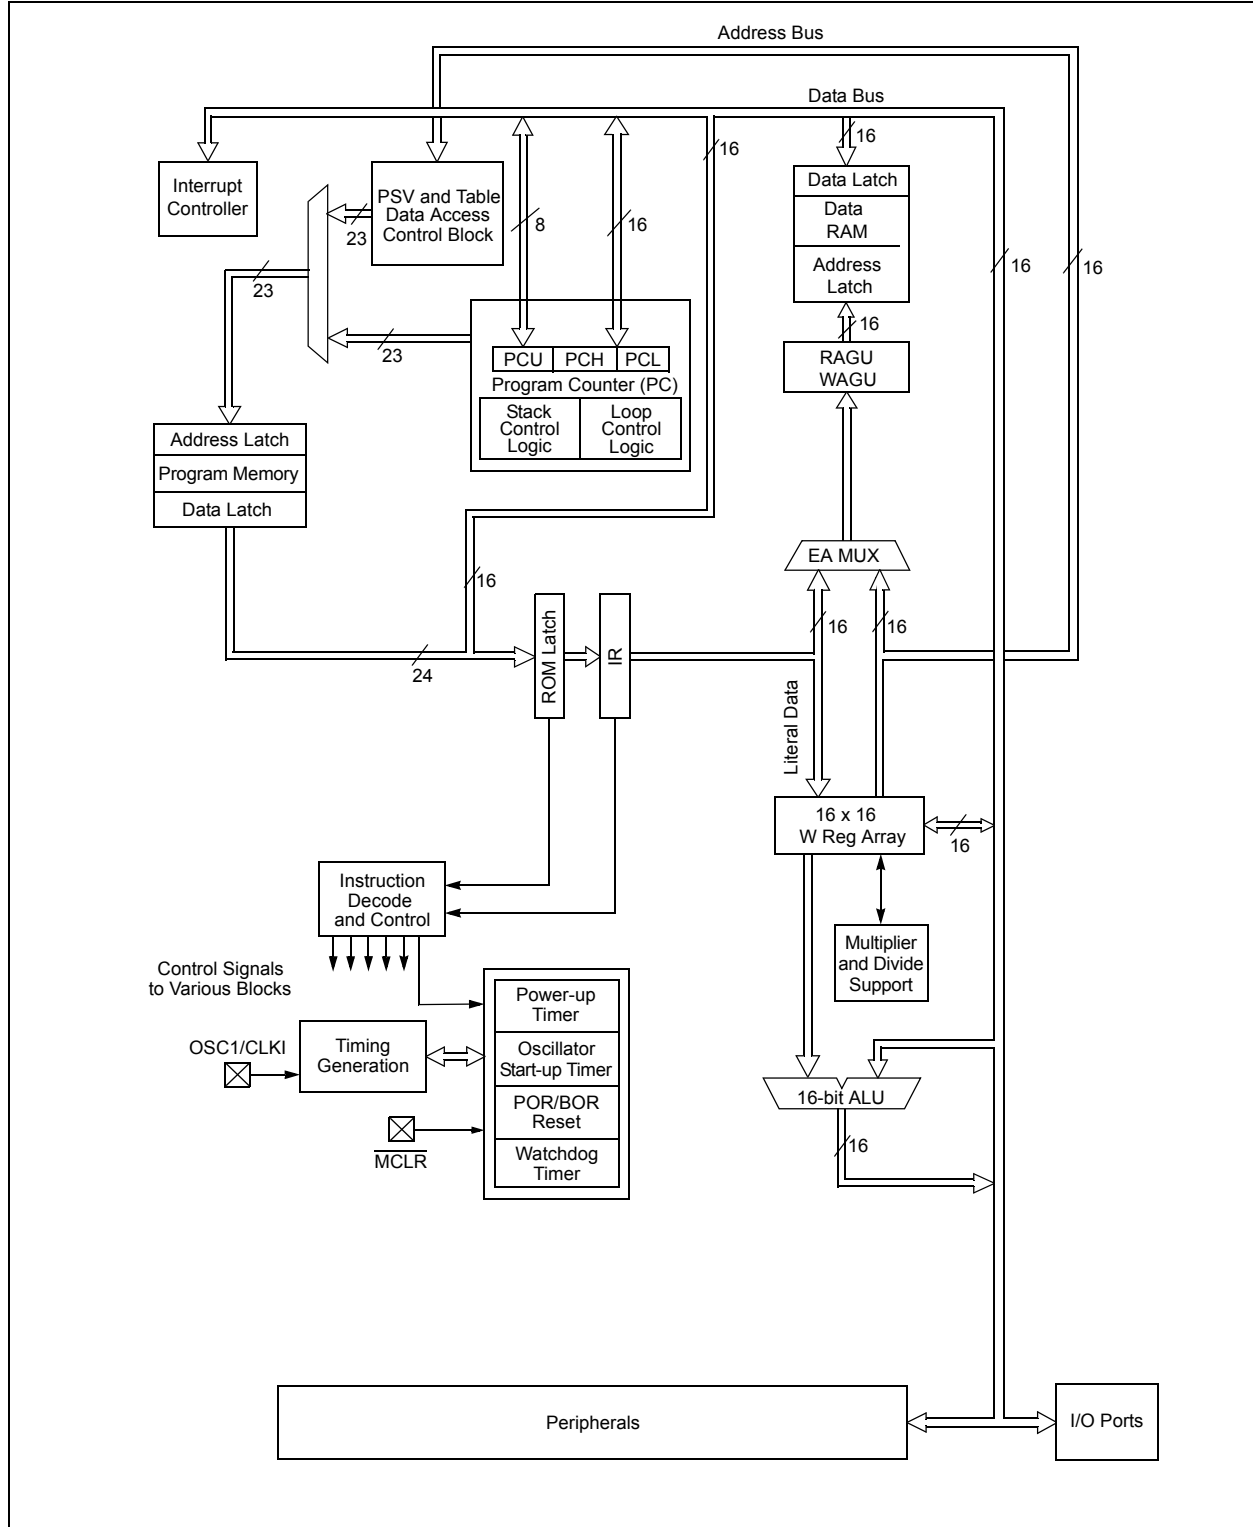
A high-speed, 17-bit by 17-bit multiplier is included to significantly enhance the core arithmetic capability and throughput. The multiplier supports Signed, Unsigned and Mixed mode, 16-bit by 16-bit or 8-bit by 8-bit integer multiplication. All multiply instructions execute in a single cycle.

The 16-bit Arithmetic Logic Unit (ALU) is enhanced with integer divide assist hardware that supports an iterative non-restoring divide algorithm. It operates in conjunction with the `REPEAT` instruction looping mechanism, and a selection of iterative divide instructions, to support 32-bit (or 16-bit) divided by 16-bit integer signed and unsigned division. All divide operations require 19 cycles to complete, but are interruptible at any cycle boundary.

2.1.4 Exception Processing

The PIC24H CPU has a vectored exception scheme with up to eight sources of non-maskable traps and up to 118 interrupt sources. Each interrupt source can be assigned to one of the seven priority levels. Figure 2-1 shows the PIC24H CPU block diagram.

Figure 2-1: PIC24H CPU Block Diagram



PIC24H Family Reference Manual

2.2 PROGRAMMER'S MODEL

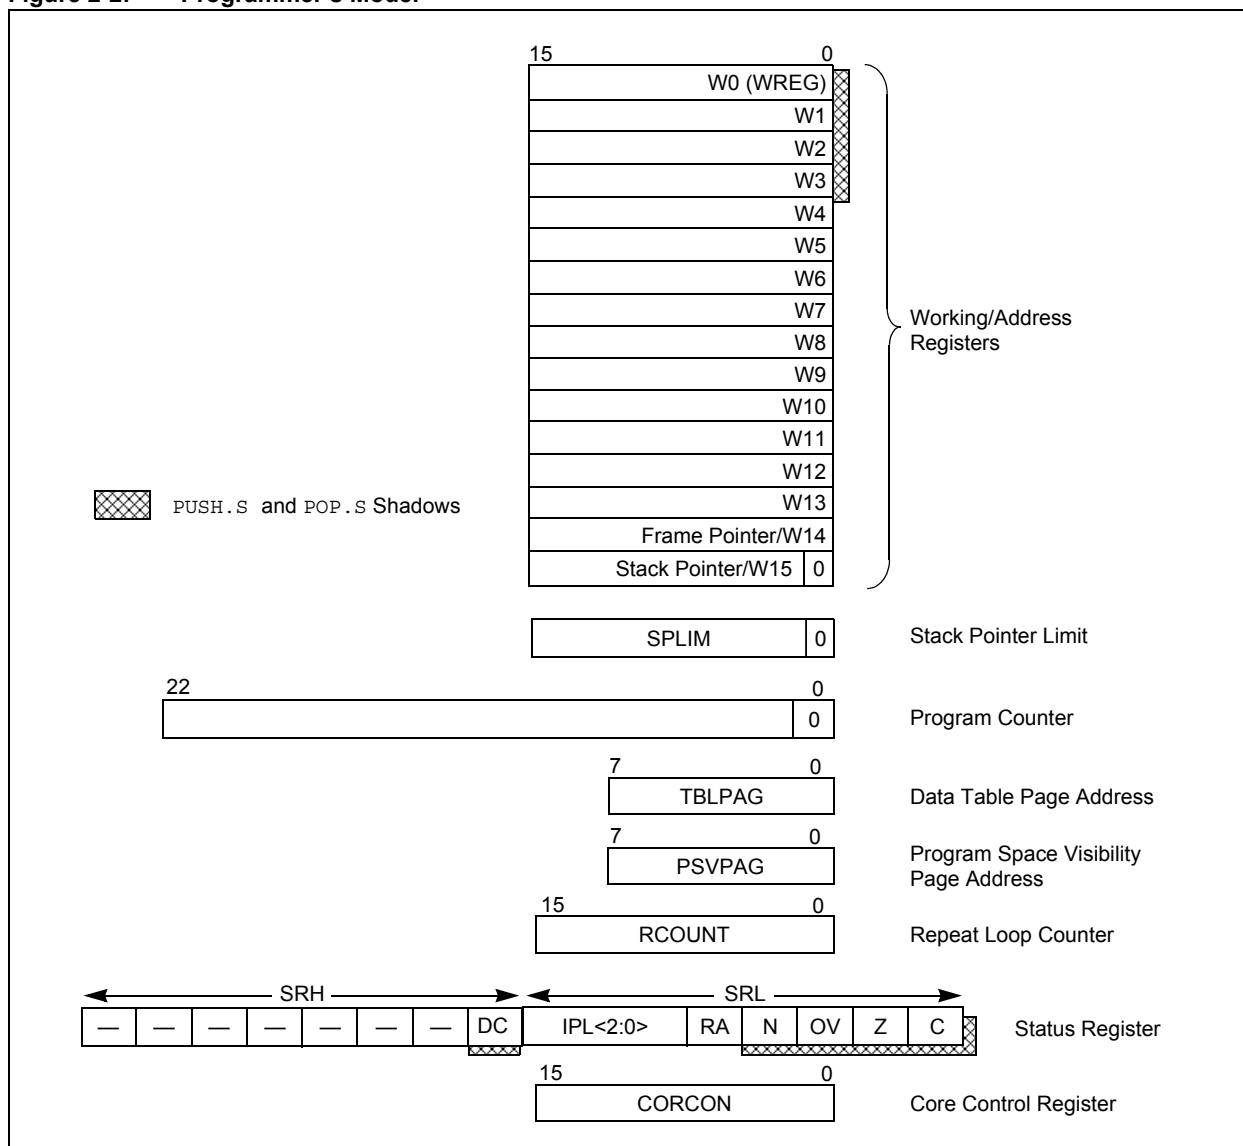
Figure 2-2 shows the programmer's model for the PIC24H CPU. All registers in the programmer's model are memory mapped and can be manipulated directly by instructions. Table 2-1 provides a description of each register.

Table 2-1: Programmer's Model Register Descriptions

Register(s) Name	Description
W0 through W15	Working register array
PC	23-bit PC
SR	ALU STATUS register
SPLIM	Stack Pointer Limit Value register
TBLPAG	Table Memory Page Address register
PSVPAG	Program Space Visibility Page Address register
RCOUNT	Repeat Loop Counter register
CORCON	CPU Control register

Table 2-7 shows how all registers associated with the programmer's model are memory mapped.

Figure 2-2: Programmer's Model



2.2.1 Working Register Array

The 16 working (W) registers can function as data, address or address offset registers. The function of a W register is determined by the addressing mode of the instruction that accesses it.

The PIC24H instruction set can be divided into two instruction types:

- Register instructions
- File register instructions

2.2.1.1 REGISTER INSTRUCTIONS

Register instructions can use each W register as a data value or an address offset value, as shown in Example 2-1.

Example 2-1: Register Instructions

```
MOV    W0, W1           ; move contents of W0 to W1
MOV    W0, [W1]         ; move W0 to address contained in W1
ADD    W0, [W4], W5     ; add contents of W0 to contents pointed
                        ; to by W4. Place result in W5.
```

2.2.1.2 FILE REGISTER INSTRUCTIONS

W0 is a special working register as it is the only working register that can be used in file register instructions. File register instructions operate on a specific memory address contained in the instruction opcode and W0. W1-W15 cannot be specified as a target register in file register instructions.

The file register instructions provide backward compatibility with existing PIC® MCU devices that have only one W register. The label “WREG” is used in the assembler syntax to denote W0 in a file register instruction, as shown in Example 2-2.

Example 2-2: File Register Instructions

```
MOV    WREG, 0x0100     ; move contents of W0 to address 0x0100
ADD    0x0100, WREG     ; add W0 to address 0x0100, store in W0
```

Note: Refer to the “*dsPIC30F/33F Programmer’s Reference Manual*” (DS70157), for complete descriptions of addressing modes and instruction syntax.

2.2.1.3 W REGISTER MEMORY MAPPING

Since the W registers are memory mapped, it is possible to access a W register in a file register instruction, as shown in Example 2-3.

Example 2-3: Accessing W Register in File Register Instruction

```
MOV    0x0004, W10      ; equivalent to MOV W2, W10
```

where:

0x0004 is the memory address in W2

In addition, it is also possible to execute an instruction that attempts to use a W register as both an address pointer and operand destination, as shown in Example 2-4.

Example 2-4: W Register Used as Address Pointer and Operand Destination

```
MOV    W1, [W2++]
```

where:

W1 = 0x1234

W2 = 0x0004 ; [W2] addresses W2

In Example 2-4, the contents of W2 are 0x0004. Since W2 is used as an address pointer, it points to location 0x0004 in memory. W2 is also mapped to this address in memory. Even though this is an unlikely event, it is impossible to detect until run time. The PIC24H CPU ensures that the data write dominates, resulting in W2 = 0x1234 in this example.

2.2.1.4 W REGISTERS AND BYTE MODE INSTRUCTIONS

The byte instructions, which target the W register array, affect only the Least Significant Byte (LSB) of the target register. Since the working registers are memory mapped, the LSBs and the Most Significant Bytes (MSBs) can be manipulated through byte wide data memory space access.

2.2.2 Shadow Registers

Table 2-7 shows how some of the registers have a shadow register associated with them. The shadow register is used as a temporary holding register and can transfer its contents to or from its host register upon an occurring event. The shadow registers are not directly accessible. The `PUSH.S` and `POP.S` shadow rule is applied to register transfer into and out of shadows.

The `PUSH.S` and `POP.S` instructions are useful for fast context save/restore during a function call or Interrupt Service Routine (ISR). The `PUSH.S` instruction transfers the following register values into their respective shadow registers:

- W0...W3
- SR (N, OV, Z, C, DC bits only)

The `POP.S` instruction restores the values from the shadow registers into these register locations. Example 2-5 shows a code example using the `PUSH.S` and `POP.S` instructions.

Example 2-5: `PUSH.S` and `POP.S` Instructions

```
MyFunction:
    PUSH.S           ; Save W registers, MCU status
    MOV    #0x03, W0 ; load a literal value into W0
    ADD    RAM100    ; add W0 to contents of RAM100
    BTSC   SR, #Z    ; is the result 0?
    BSET   Flags, #IsZero ; Yes, set a flag
    POP.S           ; Restore W regs, MCU status
    RETURN
```

The `PUSH.S` instruction overwrites the contents previously saved in the shadow registers. The shadow registers are only one level in depth. Therefore, care must be taken if the shadow registers are to be used for multiple software tasks.

The user-assigned application must ensure that any task using the shadow registers is not interrupted by a higher priority task that also uses the shadow registers. If the higher priority task is allowed to interrupt the lower priority task, the contents of the shadow registers saved in the lower priority task are overwritten by the higher priority task.

2.2.3 Uninitialized W Register Reset

The W register array (with the exception of W15) is cleared during all resets and is considered uninitialized until written to. An attempt to use an uninitialized register as an address pointer resets the device.

A word write must be performed to initialize a W register. A byte write will not affect the initialization detection logic.

2.3 SOFTWARE STACK POINTER (SSP)

W15 serves as a dedicated SSP and is automatically modified by exception processing, subroutine calls and returns. However, W15 can be referenced by any instruction in the same manner as all other W registers. This simplifies reading, writing and manipulating the Stack Pointer (SP). For example, creating stack frames.

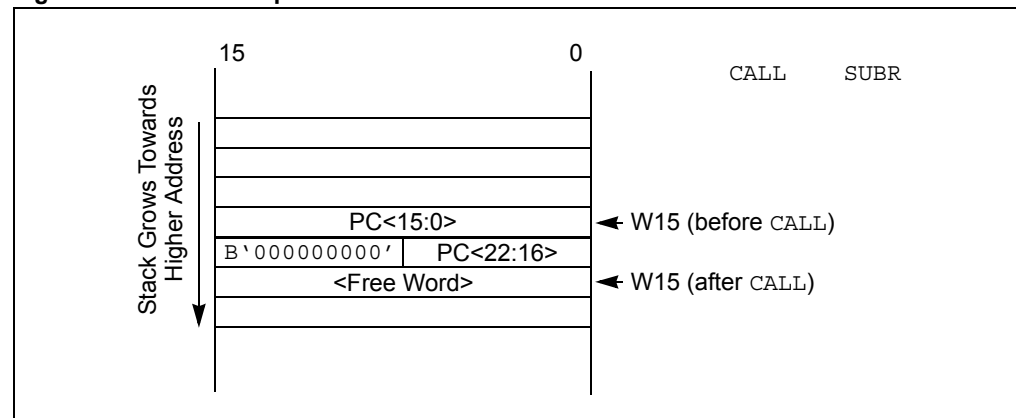
Note: To protect against misaligned stack access, W15<0> is set to '0' by the hardware.

W15 is initialized to 0x0800 during all resets. This address ensures that the SP points to valid RAM in all PIC24H devices and permits stack availability for non-maskable trap exceptions, which may occur before the SP is initialized by the user-assigned application software. The user-assigned application may reprogram the SP during initialization to any location within data space.

The SP always points to the first available free word and fills the software stack, working from lower towards higher addresses. Figure 2-3 shows how it pre-decrements for a stack pop (read) and post-increments for a stack push (writes).

When the PC is pushed onto the stack, the PC<15:0> bits are pushed onto the first available stack word, and then the PC<22:16> bits are pushed onto the second available stack location. Figure 2-3 shows that for a PC push during any CALL instruction, the MSB of the PC is zero-extended before the push. During exception processing, the MSB of the PC is concatenated with the lower eight bits of the CPU STATUS register, SR. This allows the contents of SRL to be preserved automatically during interrupt processing.

Figure 2-3: Stack Operation for a CALL Instruction



2.3.1 Software Stack Examples

The software stack is manipulated using the PUSH and POP instructions. The PUSH and POP instructions are the equivalent of a MOV instruction with W15 used as the destination pointer. The contents of W0 can be pushed onto the stack, as shown in Example 2-6.

Example 2-6: Software Stack Example

```
PUSH    W0
This syntax is equivalent to:
        MOV    W0, [W15++]

The contents of the Top-of-Stack (TOS) can be returned to W0 by:
        POP    W0
This syntax is equivalent to:
        MOV    [--W15], W0
```

PIC24H Family Reference Manual

Figure 2-4 through Figure 2-7 show examples of how the software stack is used. Figure 2-4 shows the software stack at device initialization. W15 is initialized to 0x0800. Furthermore, this example assumes that the values 0x5A5A and 0x3636 are written to W0 and W1, respectively. The stack is pushed for the first time in Figure 2-5 and the value contained in W0 is copied to the stack. W15 is automatically updated to point to the next available stack location (0x0802). In Figure 2-6, the contents of W1 are pushed onto the stack. In Figure 2-7, the stack is popped and the Top-of-Stack (TOS) value (previously pushed from W1) is written to W3.

Figure 2-4: Stack Pointer at Device Reset

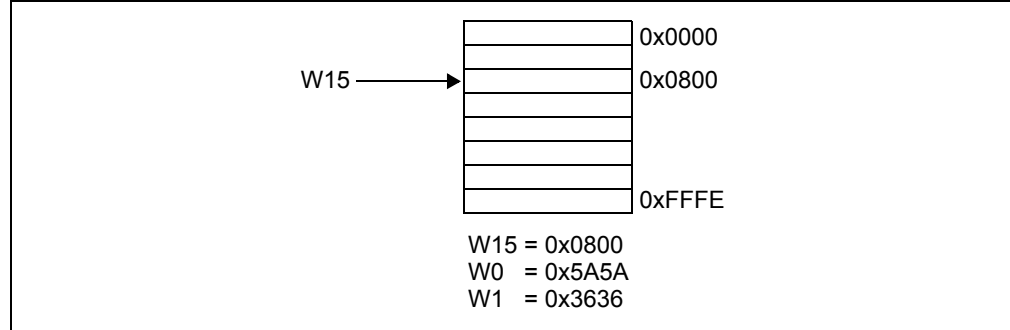


Figure 2-5: Stack Pointer After the First PUSH Instruction

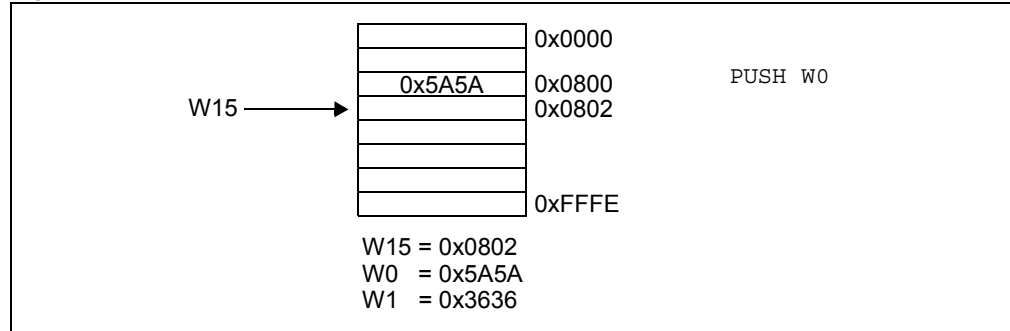


Figure 2-6: Stack Pointer After the Second PUSH Instruction

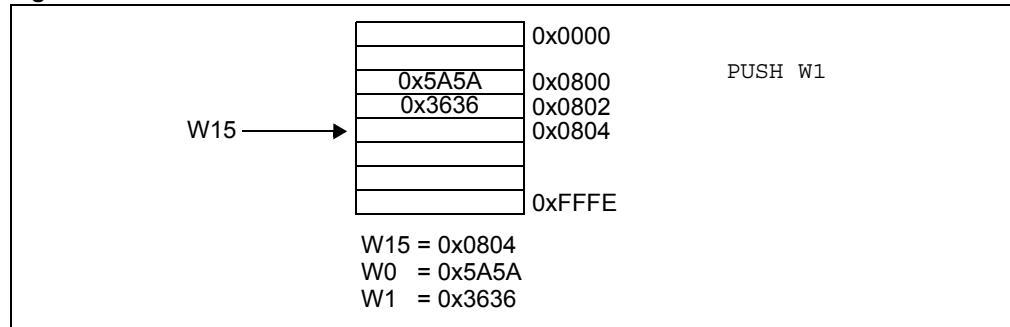
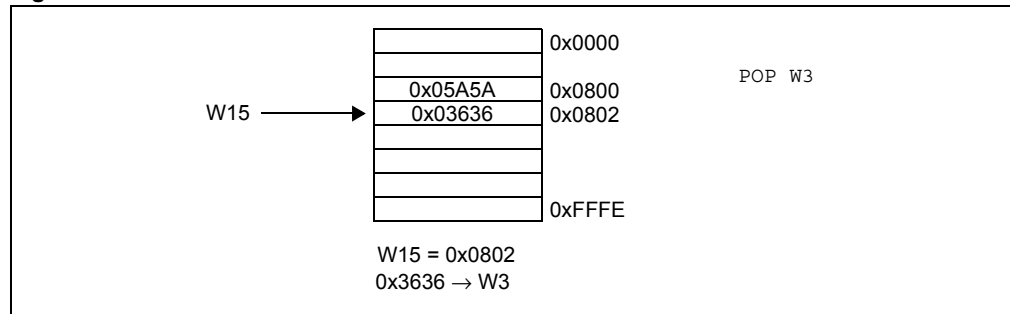


Figure 2-7: Stack Pointer After a POP Instruction



2.3.2 W14 Software Stack Frame Pointer

A frame is a user-defined section of memory in the stack used by a single subroutine. W14 is a special working register because it can be used as a Stack Frame Pointer with the `LNK` (link) and `ULNK` (unlink) instructions. W14 can be used in a normal working register by instructions when it is not used as a Stack Frame Pointer.

For software examples that use W14 as a Stack Frame Pointer, refer to the “*dsPIC30F/33F Programmer’s Reference Manual*” (DS70157).

2.3.3 Stack Pointer Overflow

There is a Stack Pointer Limit (SPLIM) register associated with the Stack Pointer that is reset to 0x0000. SPLIM is a 16-bit register, but SPLIM<0> is set to ‘0’ because all stack operations must be word-aligned.

The stack overflow check is not enabled until a word write to the SPLIM register occurs, after this time it can only be disabled by a device Reset. All effective addresses generated using W15 as a source or destination are compared against the value in SPLIM. If the contents of the Stack Pointer (W15) are greater than the contents of the SPLIM register by 2, and a Push operation is performed, a stack error trap will not occur. The stack error trap will occur on a subsequent Push operation. Thus, for example, if it is desirable to cause a stack error trap when the stack grows beyond address 0x2000 in RAM, initialize the SPLIM register with the value, 0x1FFE.

Note: A stack error trap can be caused by any instruction that uses the contents of the W15 register to generate an Effective Address (EA). Thus, if the contents of W15 are greater than the contents of the SPLIM register by 2, and a `CALL` instruction is executed, or an interrupt occurs, a stack error trap is generated.

If stack overflow checking is enabled, a stack error trap also occurs if the W15 effective address calculation wraps over the end of data space (0xFFFF).

Note: A write to the SPLIM register should not be followed by an indirect read operation using W15.

Refer to **Section 6. “Interrupts”** (DS70224), for more details on stack error traps.

2.3.4 Stack Pointer Underflow

The stack is initialized to 0x0800 during a Reset. A stack error trap is initiated, if the Stack Pointer address is ever less than 0x0800.

Note: Locations in data space between 0x0000 and 0x07FF are, in general, reserved for core and peripheral Special Function Registers (SFRs).

2.4 CPU REGISTER DESCRIPTIONS

2.4.1 SR: CPU Status Register

The PIC24H CPU has a 16-bit Status Register (SR), the LSB of which is referred to as the lower Status register (SRL). The upper byte of SR is referred to as higher Status register (SRH). Register 2-1 shows a detailed description of SR.

The SRL register contains all of the MCU ALU operation Status flags, the CPU Interrupt Priority Level Status bits (IPL<2:0>), and the REPEAT Loop Active Status bit, RA (SR<4>). During exception processing, SRL is concatenated with the MSB of the PC, to form a complete word value, which is then stacked.

SRH contains only the Digit Carry bit, DC (SR<8>). The SR bits are readable/writable with the following exceptions:

- The RA bit (SR<4>): RA is a read-only bit
- IPL<2:0>: When the SRL register is disabled (NSTDIS = 1), the IPL<2:0> bits become read-only

Note: A description of the SR bits affected by each instruction is provided in the <i>dsPIC30F/33F Programmer's Reference Manual</i> (DS70157).
--

2.4.2 CORCON: Core Control Register

The CORCON register contains the IPL3 status bit, which is concatenated with IPL<2:0> (SR<7:5>), to form the CPU Interrupt Priority Level, as well as the PSV Enable bit.

Register 2-1: SR: CPU Status Register

U-0	U-0	U-0	U-0	U-0	U-0	U-0	R/W-0
—	—	—	—	—	—	—	DC
bit 15							bit 8
R/W-0 ^(1,2)	R/W-0 ^(1,2)	R/W-0 ^(1,2)	R-0	R/W-0	R/W-0	R/W-0	R/W-0
IPL<2:0>			RA	N	OV	Z	C
bit 7							bit 0

Legend:

R = Readable bit	W = Writable bit	U = Unimplemented bit, read as '0'
-n = Value at POR	'1' = Bit is set	'0' = Bit is cleared
		x = Bit is unknown

- bit 15-9 **Unimplemented:** Read as '0'
- bit 8 **DC:** MCU ALU Half Carry/Borrow bit
 - 1 = A carry out from the 4th low-order bit (for byte-sized data) or 8th low-order bit (for word-sized data) of the result occurred
 - 0 = No carry out from the 4th low-order bit (for byte-sized data) or 8th low-order bit (for word-sized data) of the result occurred
- bit 7-5 **IPL<2:0>:** CPU Interrupt Priority Level Status bits^(1,2)
 - 111 = CPU interrupt priority level is 7 (15), user interrupts disabled
 - 110 = CPU interrupt priority level is 6 (14)
 - 101 = CPU interrupt priority level is 5 (13)
 - 100 = CPU interrupt priority level is 4 (12)
 - 011 = CPU interrupt priority level is 3 (11)
 - 010 = CPU interrupt priority level is 2 (10)
 - 001 = CPU interrupt priority level is 1 (9)
 - 000 = CPU interrupt priority level is 0 (8)
- bit 4 **RA:** REPEAT Loop Active bit
 - 1 = REPEAT loop in progress
 - 0 = REPEAT loop not in progress
- bit 3 **N:** MCU ALU Negative bit
 - 1 = Result was negative
 - 0 = Result was non-negative (zero or positive)
- bit 2 **OV:** MCU ALU Overflow bit
 - This bit is used for signed arithmetic (2's complement). It indicates an overflow of the magnitude which causes the sign bit to change state.
 - 1 = Overflow occurred for signed arithmetic (in this arithmetic operation)
 - 0 = No overflow occurred
- bit 1 **Z:** MCU ALU Zero bit
 - 1 = Last operation resulted in zero
 - 0 = Last operation did not result in zero
- bit 0 **C:** MCU ALU Carry/Borrow bit
 - 1 = A carry out from the Most Significant bit (MSb) of the result occurred
 - 0 = No carry out from the Most Significant bit (MSb) of the result occurred

Note 1: The IPL<2:0> bits are concatenated with the IPL<3> bit (CORCON<3>) to form the CPU interrupt priority level. The value in parentheses indicates the IPL if IPL<3> = 1. User interrupts are disabled when IPL<3> = 1.

2: The IPL<2:0> Status bits are read-only when NSTDIS = 1 (INTCON1<15>).

PIC24H Family Reference Manual

Register 2-2: CORCON: Core Control Register

U-0	U-0	U-0	U-0	U-0	U-0	U-0	U-0
—	—	—	—	—	—	—	—
bit 15						bit 8	

U-0	U-0	U-0	U-0	R/C-0	R/W-0	U-0	U-0
—	—	—	—	IPL3 ⁽¹⁾	PSV	—	—
bit 7						bit 0	

Legend:	C = Clearable bit		
R = Readable bit	W = Writable bit	U = Unimplemented bit, read as '0'	
-n = Value at POR	'1' = Bit is set	'0' = Bit is cleared	x = Bit is unknown

- bit 15-4 **Unimplemented:** Read as '0'
- bit 3 **IPL3:** CPU Interrupt Priority Level Status bit⁽¹⁾
 - 1 = CPU interrupt priority level is greater than 7
 - 0 = CPU interrupt priority level is 7 or less
- bit 2 **PSV:** Program Space Visibility in Data Space Enable bit
 - 1 = Program space visible in data space
 - 0 = Program space not visible in data space
- bit 1-0 **Unimplemented:** Read as '0'

Note 1: User interrupts are disabled when IPL3 = 1.

2.4.3 Other PIC24H CPU Control Registers

This section provides brief descriptions of additional registers that are associated with the PIC24H CPU. These registers are described in greater detail in other sections of the “PIC24H Family Reference Manual”.

2.4.3.1 TBLPAG: TABLE PAGE ADDRESS POINTER

The TBLPAG register holds the upper eight bits of a program memory address during table read and write operations. Table instructions transfer the data between program memory space and data memory space. Refer to **Section 4. “Program Memory”** (DS70238), for more details.

2.4.3.2 PSVPAG: PROGRAM MEMORY VISIBILITY PAGE ADDRESS POINTER

Program Space Visibility allows the user to map a 32 Kbyte section of the program memory space into the upper 32 Kbytes of data address space. This feature allows transparent access of constant data through instructions that operate on data memory. The PSVPAG register selects the 32 Kbyte region of program memory space mapped to the data address space. Refer to **Section 4. “Program Memory”** (DS70238), for more details on the PSVPAG register.

2.4.3.3 DISICNT: DISABLE INTERRUPTS COUNTER REGISTER

The DISICNT register is used by the `DISI` instruction to disable interrupts of priority 1-6 for the specified number of cycles. Refer to **Section 6. “Interrupts”** (DS70224), for more details.

2.5 ARITHMETIC LOGIC UNIT (ALU)

The PIC24H ALU is 16 bits wide and is capable of addition, subtraction, single bit shifts and logic operations. Unless specified, arithmetic operations are 2’s complement in nature. Depending on the operation, the ALU can affect the values of these status bits in the SR register:

- Carry/Borrow (C)
- Zero (Z), Negative (N)
- Overflow (OV)
- Half Carry/Borrow (DC)

The C and DC Status bits operate as $\overline{\text{Borrow}}$ and $\overline{\text{Digit Borrow}}$ bits, respectively, for subtraction operations.

The ALU can perform 8-bit or 16-bit operations depending on the mode of the instruction used. Data for the ALU operation can come from the W register array, or data memory, depending on the addressing mode of the instruction. Likewise, output data from the ALU can be written to the W register array or a data memory location.

Refer to the “*dsPIC30F/33F Programmer’s Reference Manual*” (DS70157), for more details on the SR bits affected by each instruction, addressing modes and 8-bit/16-bit instruction modes.

- Note 1:** Byte operations use the 16-bit ALU and can produce results in excess of eight bits. However, to maintain backward compatibility with PIC MCU devices, the ALU result from all of the byte operations is written back as a byte (i.e., MSB not modified), and the CPU STATUS register, SR, is updated based only on the LSB state of the result.
- 2:** All register instructions performed in Byte mode only affect the LSB of the W registers. The MSB of any W register can be modified by using file register instructions that access the memory mapped contents of the W registers.

2.6 MULTIPLICATION AND DIVIDE SUPPORT

2.6.1 Overview

The PIC24H CPU contains a 17-bit x 17-bit multiplier and is capable of unsigned, signed or mixed sign operation with the following multiplication modes:

- 16-bit x 16-bit Signed
- 16-bit x 16-bit Unsigned
- 16-bit Signed x 5-bit (literal) Unsigned
- 16-bit Unsigned x 16-bit Unsigned
- 16-bit Unsigned x 5-bit (literal) Unsigned
- 16-bit Unsigned x 16-bit Signed
- 8-bit Unsigned x 8-bit Unsigned

The divide block is capable of supporting 32-bit/16-bit and 16-bit/16-bit signed and unsigned integer divide operation with the following data sizes:

- 32-bit Signed/16-bit Signed divide
- 32-bit Unsigned/16-bit Unsigned divide
- 16-bit Signed/16-bit Signed divide
- 16-bit Unsigned/16-bit Unsigned divide

2.6.2 Multiplier

Figure 2-8 shows a block diagram of the multiplier, which supports the multiply instructions that include integer 16-bit signed, unsigned and mixed-sign multiplies. All multiply instructions only support Register Direct Addressing mode for the result. A 32-bit result (from all multiplies other than `MULWF`) is written to any two aligned consecutive W register pairs, except `W15:W14`, which are not allowed.

The file register `MUL` instruction may be directed to use byte-sized or word-sized operands. The destination is always the `W3:W2` register pair in the W array. Byte multiplicands direct a 16-bit result to `W2` (`W3` is not changed), and word multiplicands direct a 32-bit result to `W3:W2`.

<p>Note: The destination register pair for multiply instructions must be 'aligned' (i.e., odd:even), where 'odd' contains the most significant result word and 'even' contains the least significant result word. For example, <code>W3:W2</code> is acceptable, whereas <code>W4:W3</code> is not acceptable.</p>

The multiplicands for all multiply instructions are derived from the W array (first word) and data space (second word). File register `MUL` instructions derive these multiplicands from `W2` (first word or byte) and data space (second word or byte) using a zero-extended, 13-bit absolute address.

Figure 2-8 shows additional data paths are provided to allow these instructions to write the result back into the W array and data bus (via the W array).

2.6.2.1 MCU MULTIPLY INSTRUCTIONS

The same multiplier supports the MCU multiply instructions, which include integer 16-bit signed, unsigned, and mixed sign multiplies as shown in Table 2-2. All multiplications performed by the MUL instruction produce integer results. The MUL instruction can be directed to use byte-sized or word-sized operands. Byte input operands produce a 16-bit result and word input operands produce a 32-bit result to the specified register(s) in the W array.

Table 2-2: MCU Instructions that Utilize the Multiplier

MCU Instruction	Description
MUL/MUL.UU	Multiply two unsigned integers
MUL.SS	Multiply two signed integers
MUL.SU/MUL.US	Multiply a signed integer with an unsigned integer

- Note 1:** MCU instructions using the multiplier operate only in Integer mode.
Note 2: Result of an MCU multiply is 32 bits long and is stored in a pair of W registers.

Table 2-3: Multiplication Options

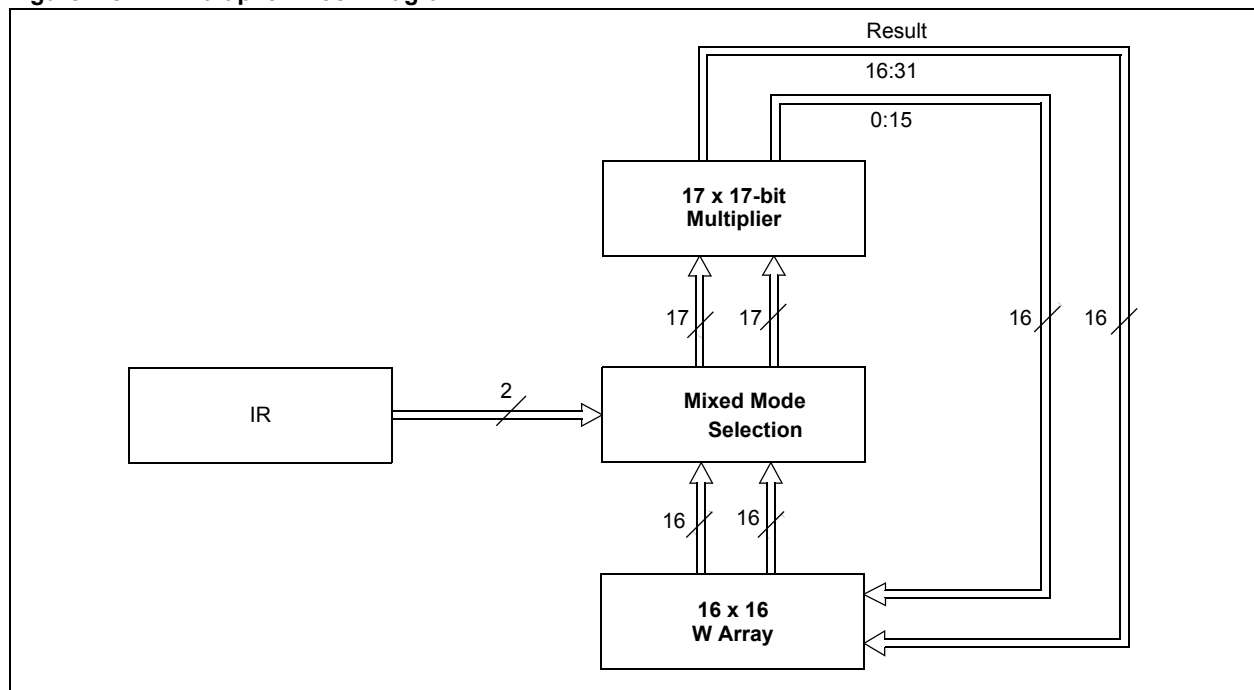
Instruction Class	Signed x Signed	Unsigned x Unsigned	Unsigned x Signed	Signed x Unsigned
MUL (MCU Multiplication)	Yes (Integer only)	Yes (Integer only)	Yes (Integer only)	Yes (Integer only)

2.6.2.2 SINGLE AND MIXED MODE INTEGERS

Simple data preprocessing logic, either zero or sign, extends all operands to 17 bits such that the unsigned, signed or mixed sign multiplications can be executed as signed values. All unsigned operands are always zero-extended into the seventeenth bit of the multiplier input value. All signed operands are always sign-extended into the seventeenth bit of the multiplier input value.

- For unsigned 16-bit multiplies, the multiplier produces a 32-bit, unsigned result
- For signed 16-bit multiplies, the multiplier produces 30 bits of data and two bits of sign
- For 16-bit Mixed mode (signed/unsigned) multiplies, the multiplier produces 31 bits of data and one bit of sign

Figure 2-8: Multiplier Block Diagram



2.6.3 Divider

The PIC24H CPU features both 32-bit/16-bit and 16-bit/16-bit signed and unsigned, integer divide operations implemented as single instruction iterative divides.

The quotient for all divide instructions ends up in W0 and the remainder in W1. 16-bit signed and unsigned DIV instructions can specify any W register for both the 16-bit divisor (Wn) and any W register (aligned) pair [W(m + 1):Wm] for the 32-bit dividend. The divide algorithm takes one cycle per bit of divisor. Therefore, both 32-bit/16-bit and 16-bit/16-bit instructions take the same number of cycles to execute.

The divide instructions must be executed within a REPEAT loop. Any other form of execution (for example, a series of discrete divide instructions) will not function correctly because the instruction flow function is conditional on RCOUNT. The divide flow does not automatically set up the REPEAT, which must be explicitly executed with the correct operand value as shown in Table 2-4 (REPEAT will execute the target instruction {operand value + 1} time).

Table 2-4: Divide Execution Time

Instruction	Description	Iterations	REPEAT Operand Value	Total Execution Time (including REPEAT)
DIV.SD	Signed divide: W(m + 1):Wm/Wn → W0; Rem → W1	18	17	19
DIV.SW	Signed divide: Wm/Wn → W0; Rem → W1	18	17	19
DIV.UD	Unsigned divide: W(m + 1):Wm/Wn → W0; Rem → W1	18	17	19
DIV.UW	Unsigned divide: Wm/Wn → W0; Rem → W1	18	17	19

All intermediate data is saved in W1:W0 after each iteration. The N, C and Z Status flags convey control information between iterations. Consequently, although the divide instructions are listed as 19 cycle operations, the divide iterative sequence is interruptible, just like any other REPEAT loop.

Dividing by zero initiates an arithmetic error trap. The divisor is evaluated during the first cycle of the divide instruction, so the first cycle executes prior to the start of exception processing for the trap. Refer to **Section 6. “Interrupts”** (DS70224), for more details.

2.7 COMPILER FRIENDLY ARCHITECTURE

The PIC24H CPU architecture is designed to produce an efficient (code size and speed) C compiler.

- For most instructions, the core is capable of executing a data (or program data) memory read, a working register (data) read, a data memory write and a program (instruction) memory read per instruction cycle. As a result, three parameter instructions are supported, allowing $A + B = C$ operations to be executed in a single cycle.
- Instruction addressing modes are very flexible and are matched closely to compiler needs.
- There are sixteen, 16 x 16-bit working register arrays, each of which can act as data, address or offset registers. One working register (W15) operates as a software stack for interrupts and calls.
- Linear indirect access of all data space is supported, plus the memory direct address range is extended to 8 Kbytes, with the addition of 16-bit direct address load and store instructions.
- Linear indirect access of 32K word (64 Kbyte) pages within program space (user and test space) is supported using any working register via new table read and write instructions.
- Part of the data space can be mapped into program space, allowing constant data to be accessed as if it were in data space using PSV mode.

2.8 MULTI-BIT SHIFT SUPPORT

The PIC24H CPU supports single-cycle, multi-bit arithmetic and logic shifts using a shifter block. It also supports single bit shifts through the ALU. The multi-bit shifter is capable of performing up to a 15-bit arithmetic right shift, or up to a 15-bit left shift, in a single cycle. Table 2-5 provides a full summary of instructions that use the shift operation.

Table 2-5: Instructions Using Single and Multi-Bit Shift Operations

Instruction	Description
ASR	Arithmetic shift right source register by one or more bits.
SL	Shift left source register by one or more bits.
LSR	Logical shift right source register by one or more bits.

All multi-bit shift instructions only support Register Direct Addressing mode for both the operand source and result destination.

2.9 INSTRUCTION FLOW TYPES

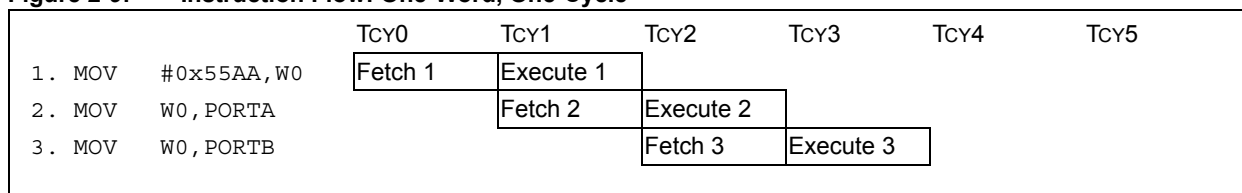
Most instructions in the PIC24H architecture occupy a single word of program memory and execute in a single cycle. An instruction prefetch mechanism facilitates single cycle (1 Tcy) execution. However, some instructions take two or three instruction cycles to execute. Consequently, there are seven different types of instruction flow in the PIC24H architecture as listed below and are described in this section:

- 1 Instruction Word, 1 Instruction Cycle
- 1 Instruction Word, 2 Instruction Cycles
- 1 Instruction Word, 2 or 3 Instruction Cycles (Program Flow Changes)
- 1 Instruction Word, 3 Instruction Cycles (RETFIE, RETURN, RETLW)
- Table Read/Write Instructions
- 2 Instruction Words, 2 Instruction Cycles
- Address Register Dependencies

2.9.1 1 Instruction Word, 1 Instruction Cycle

These instructions take one instruction cycle to execute as shown in Figure 2-9. Most instructions are one-word, one-cycle instructions.

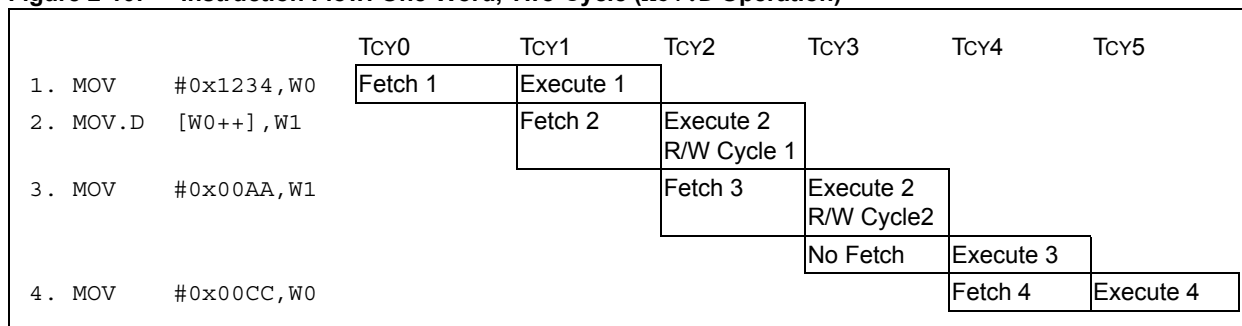
Figure 2-9: Instruction Flow: One-Word, One-Cycle



2.9.2 1 Instruction Word, 2 Instruction Cycles

In these instructions, there is no prefetch flush. The only instructions of this type are the MOV.D instructions (load and store double-word). Two cycles are required to complete these instructions, as shown in Figure 2-10.

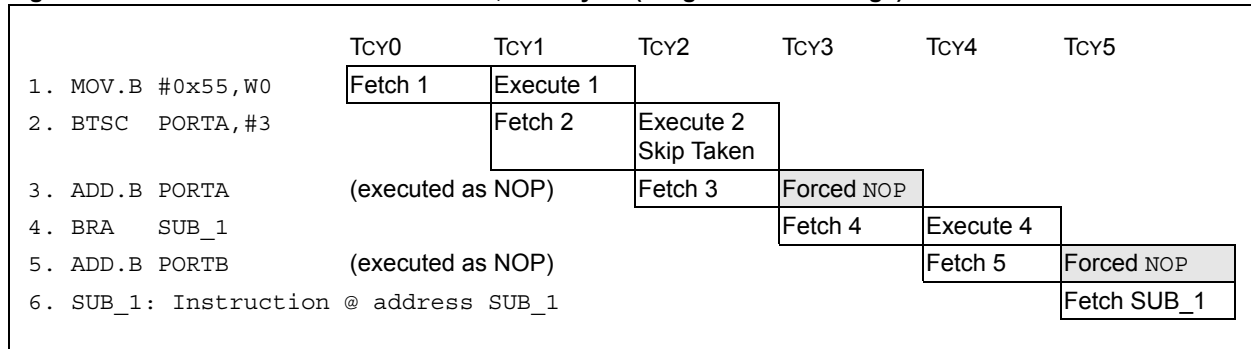
Figure 2-10: Instruction Flow: One-Word, Two-Cycle (MOV.D Operation)



2.9.3 1 Instruction Word, 2 or 3 Instruction Cycles (Program Flow Changes)

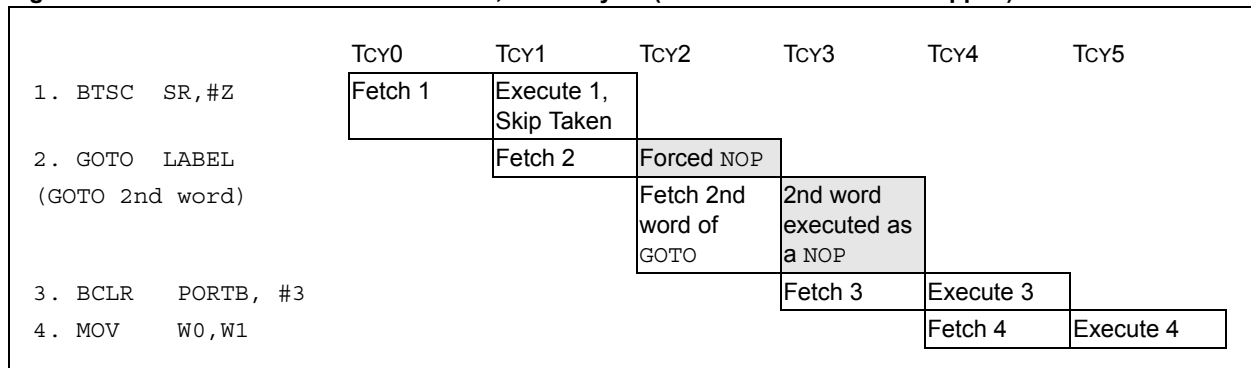
These instructions include relative call and branch instructions, and skip instructions. When an instruction changes the PC (other than to increment it), the program memory prefetch data must be discarded. This makes the instruction take two effective cycles to execute, as shown in Figure 2-11.

Figure 2-11: Instruction Flow: One-Word, Two-Cycle (Program Flow Change)



Three cycles are required when a two-word instruction is skipped. In this case, the program memory prefetch data is discarded and the second word of the two-word instruction is fetched. Figure 2-12 shows the second word of the instruction is executed as a NOP.

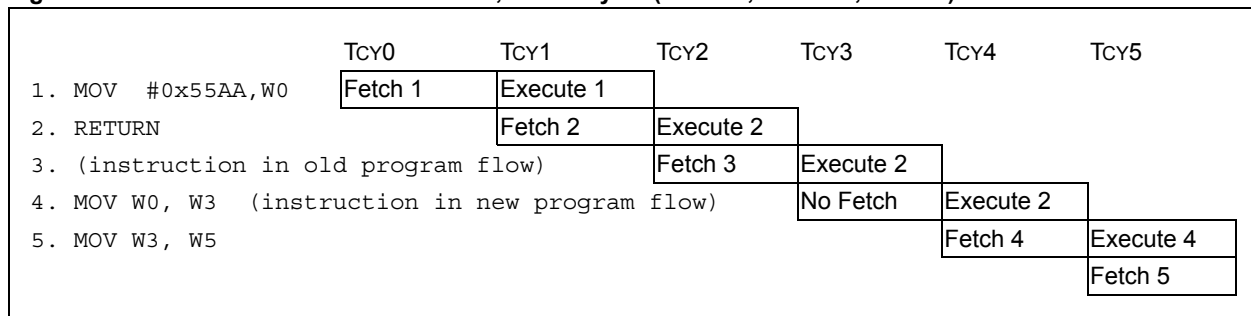
Figure 2-12: Instruction Flow: One-Word, Three-Cycle (Two-Word Instruction Skipped)



2.9.4 1 Instruction Word, 3 Instruction Cycles (RETFIE, RETURN, RETLW)

Figure 2-13 shows the RETFIE, RETURN and RETLW instructions that are used to return from a subroutine call or an Interrupt Service Routine (ISR) and take three instruction cycles to execute.

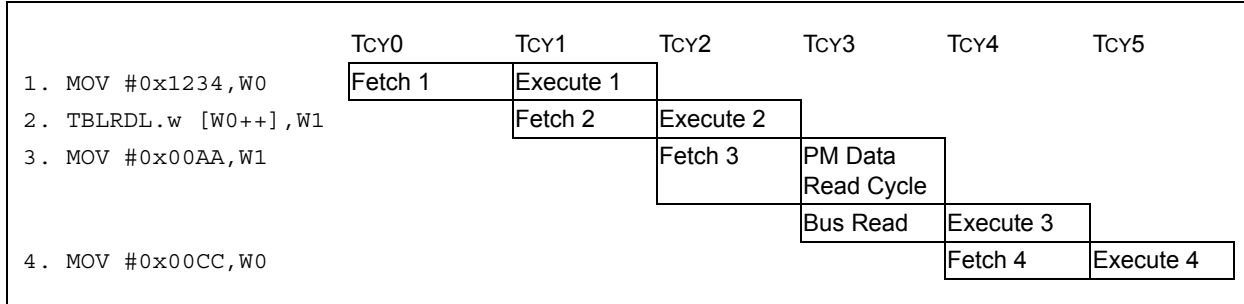
Figure 2-13: Instruction Flow: One-Word, Three-Cycle (RETURN, RETFIE, RETLW)



2.9.5 Table Read/Write Instructions

These instructions suspend fetching to insert a read or write cycle to the program memory. Figure 2-14 shows the instruction that is fetched while executing the table operation that is saved for one cycle, and executed in the cycle immediately after the table operation.

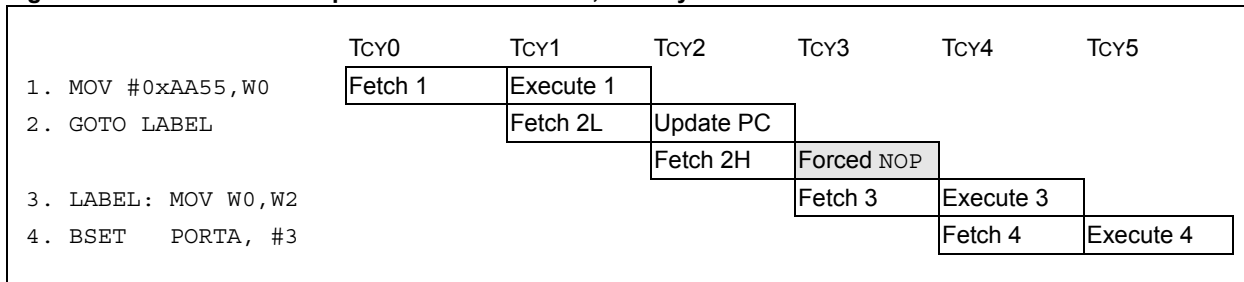
Figure 2-14: Instruction Pipeline Flow: Table Operations



2.9.6 2 Instruction Words, 2 Instruction Cycles

In these instructions, the fetch after the instruction contains data. This results in a two-cycle instruction, as shown in Figure 2-15. The second word of a two-word instruction is encoded so that it executes as a NOP if it is fetched by the CPU, when the CPU did not first fetch the first word of the instruction. This is important when a two-word instruction is skipped by a skip instruction (refer to Figure 2-12).

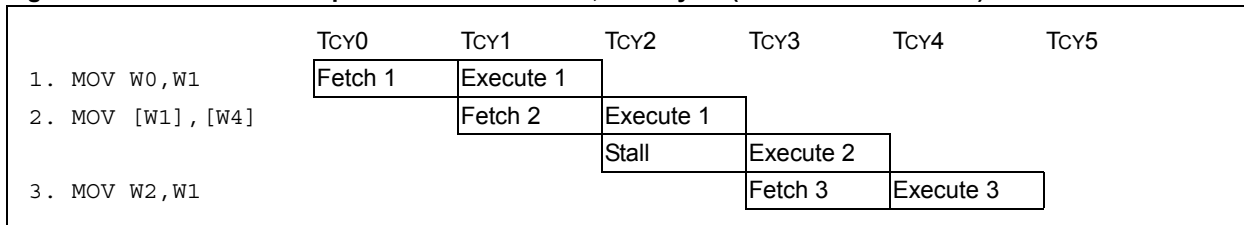
Figure 2-15: Instruction Pipeline Flow: Two-Word, Two-Cycle



2.9.7 Address Register Dependencies

These are instructions that are subjected to a stall due to data address dependency between the X-data space read and write operations. An additional cycle is inserted to resolve the resource conflict, as discussed in 2.11 “Address Register Dependencies”.

Figure 2-16: Instruction Pipeline Flow: One-Word, One-Cycle (With Instruction Stall)



Note: If the RETURN instruction is placed at the end of the program memory, the illegal address error trap will be generated by the device during the run-time. This is due to the prefetch operation that will try to preload the next two instructions from the memory location, which in this case do not exist. The solution is to leave two extra instruction words available after the RETURN instruction, so that the compiler can place NOP and RESET instructions at the end of the program memory.

2.10 PROGRAM FLOW LOOP CONTROL

The PIC24H CPU supports the REPEAT instruction construct to provide unconditional automatic program loop control. The REPEAT instruction implements a single instruction program loop. The instruction uses control bits within the CPU STATUS register, SR, to temporarily modify CPU operation.

2.10.1 REPEAT Loop

The REPEAT instruction causes the instruction that follows it to be repeated a number of times. A literal value contained in the instruction, or a value in one of the W registers, can be used to specify the REPEAT count value. The W register option enables the loop count to be a software variable.

An instruction in a REPEAT loop executes at least once. The number of iterations for a REPEAT loop is the 14-bit literal value + 1, or Wn + 1. The syntax for the two forms of the REPEAT instruction is shown in Example 2-7.

Example 2-7: REPEAT Instruction Syntax

```

REPEAT #lit14          ; RCOUNT <-- lit14
(Valid target Instruction)

or

REPEAT Wn              ; RCOUNT <-- Wn
(Valid target Instruction)
    
```

2.10.1.1 REPEAT OPERATION

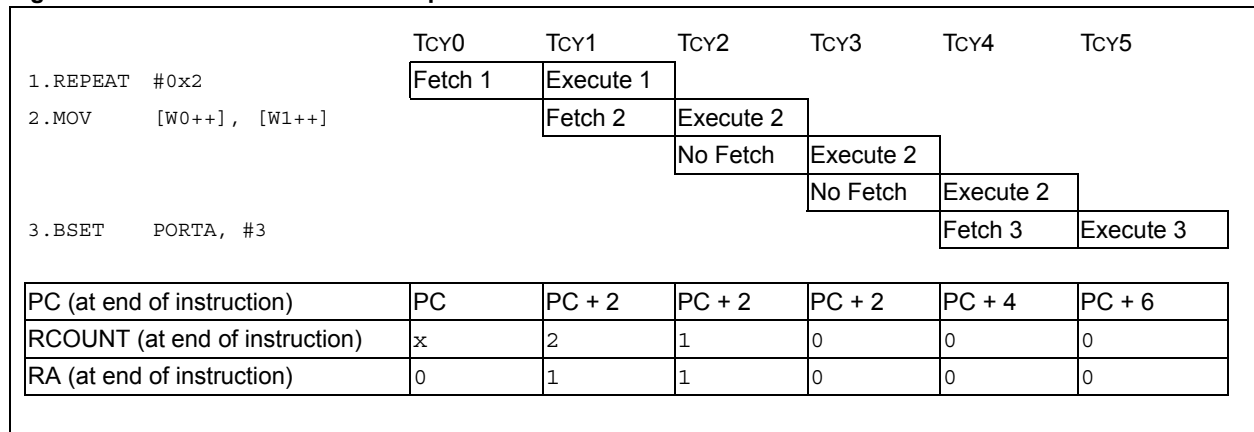
The loop count for REPEAT operations is held in the 14-bit RCOUNT register, which is memory mapped. RCOUNT is initialized by the REPEAT instruction. The REPEAT instruction sets the Repeat Active Status bit, RA (SR<4>), to '1' if the RCOUNT is a non-zero value.

RA is a read-only bit and cannot be modified through software. For REPEAT loop count values greater than '0', the PC is not incremented. Furthermore, PC increments are inhibited until RCOUNT = 0. For an instruction flow example of a REPEAT loop, refer to Figure 2-17.

For a loop count value equal to '0', REPEAT has the effect of a NOP and the RA (SR<4>) bit is not set. The REPEAT loop is essentially disabled before it begins, allowing the target instruction to execute only once while prefetching the subsequent instruction (i.e., normal execution flow).

Note: The instruction immediately following the REPEAT instruction (i.e., the target instruction) is always executed at least one time. It is always executed one time more than the value specified in the 14-bit literal or the W register operand.

Figure 2-17: REPEAT Instruction Pipeline Flow



2.10.1.2 INTERRUPTING A REPEAT LOOP

A `REPEAT` instruction loop may be interrupted any time. The RA state is preserved on the stack during exception processing to allow the user to execute further `REPEAT` loops from within any number of nested interrupts. After the SRL register is stacked, the RA Status bit is cleared to restore normal execution flow within the ISR.

Note 1: If a `REPEAT` loop is interrupted and an ISR is being processed, the user must stack the RCOUNT (Repeat Loop Counter) register prior to executing another `REPEAT` instruction within an ISR.

2: If `REPEAT` is used within an ISR, the user must unstack RCOUNT prior to executing `RETFIE`.

Returning into a `REPEAT` loop from an ISR using `RETFIE` requires no special handling. Interrupts prefetch the repeated instruction during the third cycle of the `RETFIE`. The stacked RA bit is restored when the SRL register is popped, and if set, the interrupted `REPEAT` loop is resumed.

Note: If the repeated instruction (target instruction in the `REPEAT` loop) is accessing data from Program Space (PS) using PSV, it requires two instruction cycles the first time it is executed after a return from an exception. Similar to the first iteration of a loop, timing limitations will not allow the first instruction to access data residing in PS in a single instruction cycle.

2.10.1.2.1 Early Termination of a REPEAT Loop

An interrupted `REPEAT` loop can be terminated earlier than normal in the ISR by clearing the RCOUNT register in software.

2.10.1.3 RESTRICTIONS ON THE REPEAT INSTRUCTION

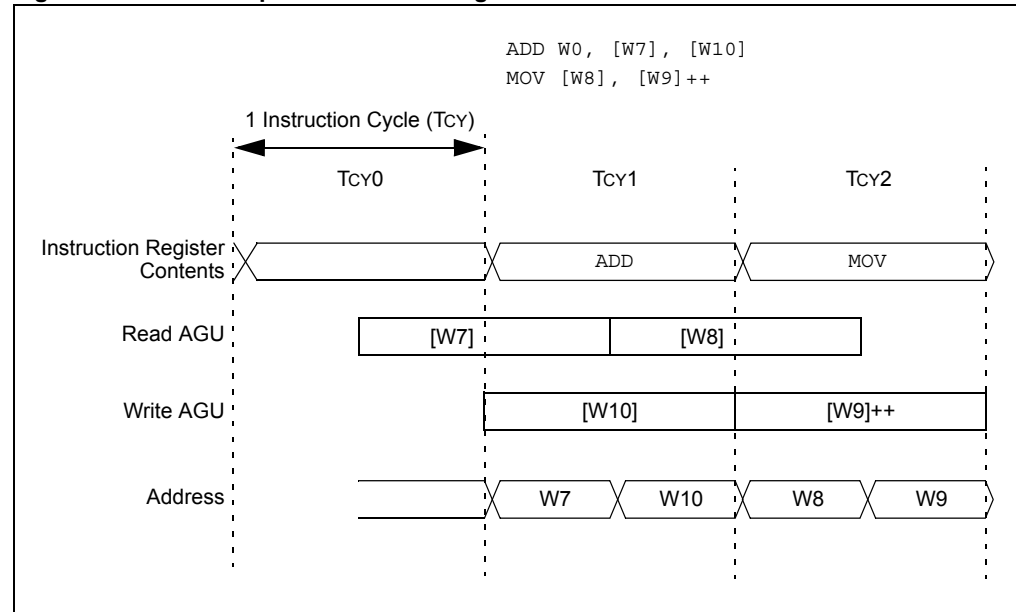
Any instruction can immediately follow a `REPEAT` except for the following:

- Program flow control instructions (any branch, compare and skip, subroutine calls, returns, etc.)
- Another `REPEAT` instruction
- `DISI`, `ULNK`, `LNK`, `PWRSV`, `RESET` instructions
- `MOV.D` instruction

2.11 ADDRESS REGISTER DEPENDENCIES

The PIC24H architecture supports a data space read (source) and a data space write (destination) for most instructions. The Effective Address (EA) calculation by the Address Generator Unit (AGU), and subsequent data space read or write, each take a period of one instruction cycle to complete. Figure 2-18 shows how the timing causes the data space read and write operations for each instruction to partially overlap. Because of this overlap, a Read-After-Write (RAW) data dependency can occur across instruction boundaries. RAW data dependencies are detected and handled at run time by the PIC24H CPU.

Figure 2-18: Data Space Access Timing



2

CPU

2.11.1 Read-After-Write (RAW) Dependency Rules

If a working register, W_n , is used as a write operation destination in the current instruction, and the same working register, W_n , that is read in the prefetched instruction are the same, the following rules apply:

- If the destination write (current instruction) does not modify the contents of W_n , no stalls occur, or
- If the source read (prefetched instruction) does not calculate an EA using W_n , no stalls occur

During each instruction cycle, the PIC24H hardware automatically checks to see if a RAW data dependency is about to occur. If the conditions specified above are not satisfied, the CPU automatically adds one instruction cycle delay before executing the prefetched instruction. The instruction stall provides enough time for the destination W register write to take place before the next (prefetched) instruction has to use the written data.

Table 2-6: Read-After-Write (RAW) Dependency Summary

Destination Addressing Mode Using Wn	Source Addressing Mode Using Wn	Status	Examples (Wn = W2)
Direct	Direct	Allowed	ADD.w W0, W1, W2 MOV.w W2, W3
Direct	Indirect	Stall	ADD.w W0, W1, W2 MOV.w [W2], W3
Direct	Indirect with modification	Stall	ADD.w W0, W1, W2 MOV.w [W2++], W3
Indirect	Direct	Allowed	ADD.w W0, W1, [W2] MOV.w W2, W3
Indirect	Indirect	Allowed	ADD.w W0, W1, [W2] MOV.w [W2], W3
Indirect	Indirect with modification	Allowed	ADD.w W0, W1, [W2] MOV.w [W2++], W3
Indirect with modification	Direct	Allowed	ADD.w W0, W1, [W2++] MOV.w W2, W3
Indirect	Indirect	Stall	ADD.w W0, W1, [W2] MOV.w [W2], W3 ; W2=0x0004 (mapped W2)
Indirect	Indirect with modification	Stall	ADD.w W0, W1, [W2] MOV.w [W2++], W3 ; W2=0x0004 (mapped W2)
Indirect with modification	Indirect	Stall	ADD.w W0, W1, [W2++] MOV.w [W2], W3
Indirect with modification	Indirect with modification	Stall	ADD.w W0, W1, [W2++] MOV.w [W2++], W3

2.11.2 Instruction Stall Cycles

An instruction stall is essentially a one instruction cycle wait period appended in front of the read phase of an instruction that allows the prior write to complete before the next read operation. For the purpose of interrupt latency, it should be noted that the stall cycle is associated with the instruction following the instruction where it was detected (i.e., stall cycles always precede instruction execution cycles).

If a RAW data dependency is detected, the PIC24H CPU begins an instruction stall. During an instruction stall, the following events occur:

- The write operation underway (for the previous instruction) is allowed to complete as normal.
- Data space is not addressed until after the instruction stall.
- PC increment is inhibited until after the instruction stall.
- Further instruction fetches are inhibited until after the instruction stall.

2.11.2.1 INSTRUCTION STALL CYCLES AND INTERRUPTS

When an interrupt event coincides with two adjacent instructions that cause an instruction stall, one of two possible outcomes could occur:

1. The interrupt could be coincident with the first instruction. In this situation, the first instruction is allowed to complete and the second instruction is executed after the ISR completes. In this case, the stall cycle is eliminated from the second instruction because the exception process provides time for the first instruction to complete the write phase.
2. The interrupt may be coincident with the second instruction. In this situation, the second instruction and the appended stall cycle is allowed to execute prior to the ISR. In this case, the stall cycle associated with the second instruction executes normally. However, the stall cycle is effectively absorbed into the exception process timing. The exception process proceeds as if an ordinary one-cycle instruction or two-cycle instruction is interrupted.

2.11.2.2 INSTRUCTION STALL CYCLES AND FLOW CHANGE INSTRUCTIONS

The `CALL` and `RCALL` instructions write to the stack using `W15` and may, therefore, force an instruction stall prior to the next instruction if the source read of the next instruction uses `W15`.

The `RETFIE` and `RETURN` instructions can never force an instruction stall prior to the next instruction because they only perform read operations. However, the user should note that the `RETLW` instruction can force a stall because it writes to a `W` register during the last cycle.

The `GOTO` and branch instructions can never force an instruction stall because they do not perform write operations.

2.11.2.3 INSTRUCTION STALLS AND REPEAT LOOPS

Other than the addition of instruction stall cycles, RAW data dependencies will not affect the operation of `REPEAT` loops.

The prefetched instruction within a `REPEAT` loop does not change until the loop is complete or an exception occurs. Although register dependency checks occur across instruction boundaries, the PIC24H CPU effectively compares the source and destination of the same instruction during a `REPEAT` loop.

2.11.2.4 INSTRUCTION STALLS AND PROGRAM SPACE VISIBILITY (PSV)

When the PSV is enabled and the Effective Address (EA) falls within the visible PSV window, the read cycle is redirected to the address in program space. Accessing data from program space takes up to three instruction cycles.

Instructions operating in PSV address space are subject to instruction stalls, just like any other instruction. Although the instruction stall and PSV cycles both occur at the beginning of an instruction, it is not possible to combine them. If a stall occurs coincidentally with a PSV cycle, the stall cycle is forced first, and then the PSV cycle, and finally, the instruction cycle. Consider the code segment shown in Example 2-8.

Example 2-8: Code Example

```
ADD    W0, [W1], [W2++]      ; PSV = 1, W1=0x8000, PSVPAG=0xAA
MOV    [W2], [W3]
```

This sequence of instructions would take five instruction cycles to execute. Two instruction cycles are added to perform the PSV access via `W1`. Furthermore, an instruction stall cycle is inserted to resolve the RAW data dependency caused by `W2`.

2.12 REGISTER MAPS

A summary of the registers associated with the PIC24H CPU core is provided in Table 2-7.

Table 2-7: PIC24H CPU SFR Memory Map

Name	Bit 15	Bit 14	Bit 13	Bit 12	Bit 11	Bit 10	Bit 9	Bit 8	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0	All Resets
W0	Working Register 0																0000
W1	Working Register 1																0000
W2	Working Register 2																0000
W3	Working Register 3																0000
W4	Working Register 4																0000
W5	Working Register 5																0000
W6	Working Register 6																0000
W7	Working Register 7																0000
W8	Working Register 8																0000
W9	Working Register 9																0000
W10	Working Register 10																0000
W11	Working Register 11																0000
W12	Working Register 12																0000
W13	Working Register 13																0000
W14	Working Register 14																0000
W15	Working Register 15																0800
SPLIM	Stack Pointer Limit																xxxx
PCL	Program Counter, Low Word																0000
PCH	—	—	—	—	—	—	—	—	Program Counter, High Byte								0000
TBLPAG	—	—	—	—	—	—	—	—	Table Page Address Pointer								0000
PSVPAG	—	—	—	—	—	—	—	—	Program Memory Visibility Page Address Pointer								0000
RCOUNT	Repeat Loop Counter																xxxx
SR	—	—	—	—	—	—	—	DC	IPL2	IPL1	IPL0	RA	N	OV	Z	C	0000
CORCON	—	—	—	—	—	—	—	—	—	—	—	—	IPL3	PSV	—	—	0000
DISICNT	—	—	Disable Interrupts Counter													xxxx	

Legend: x = unknown value on Reset, — = unimplemented, read as '0'. Reset values are shown in hexadecimal.

Note: Refer to the device data sheet for specific core register map details.

2.13 RELATED APPLICATION NOTES

This section lists application notes that are related to this section of the manual. These application notes may not be written specifically for the PIC24H device family, but the concepts are pertinent and could be used with modification and possible limitations. The current application notes related to the CPU are:

Title	Application Note #
No related application notes at this time.	N/A

Note: For additional Application Notes and code examples for the PIC24H device family, visit the Microchip web site (www.microchip.com).

2.14 REVISION HISTORY

Revision A (May 2007)

This is the initial released version of this document.

Revision B (September 2009)

This revision includes the following updates:

- Note:
 - The note reference for bit 5, bit 6 and bit 7 has been updated in Register 2-1.
 - Added a note on program memory error in **2.9.7 “Address Register Dependencies”**.
- Sections:
 - Updated the **2.1 “Introduction”** section with the following data: All instructions execute in a single cycle, except the instructions that change the program flow, double-word move (MOV.D) instruction, table instructions and also the instructions accessing Program Space Visibility (PSV) take more than one cycle.
 - Updated the exception process in **2.11.2.1 “Instruction Stall Cycles and Interrupts”** as follows: The exception process proceeds as if an ordinary one-cycle instruction or two-cycle instruction is interrupted.
 - Updated the read or write cycle as “read cycle” in **2.11.2.4 “Instruction Stalls and Program Space Visibility (PSV)”**.
- Tables:
 - Added a table on MCU Instructions that utilize the Multiplier (see Table 2-2) in **2.6.2.1 “MCU Multiply Instructions”**.
 - Added a table on Multiplication Options (see Table 2-3) in **2.6.2.1 “MCU Multiply Instructions”**.
- Updated the incorrect references to external documents specified in this document.
- Additional minor corrections such as language and formatting updates have been incorporated throughout the document.