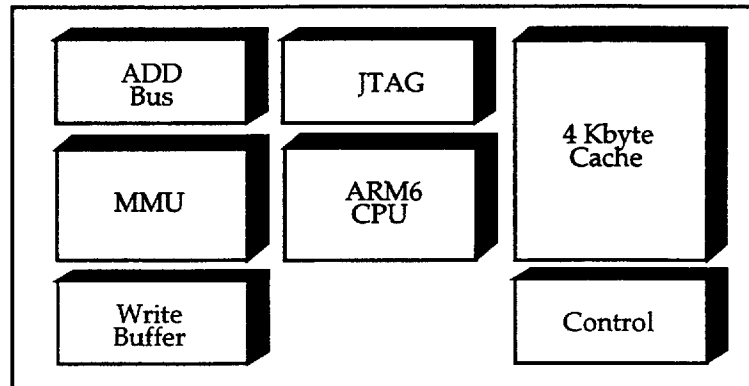


LH74610 RISC Processor

The LH74610 is a general purpose 32-bit microprocessor with 4 KByte cache, Write Buffer, and Memory Management Unit (MMU) combined in a single chip. It offers a high level of RISC performance yet it's fully-static design ensures minimal power consumption – making it ideal for portable, battery powered products.

The innovative MMU supports a conventional two-level page-table structure and a number of extensions which make it ideal for embedded control, UNIX, and Object Oriented systems. This results in a high instruction throughput and impressive real-time interrupt response from a small and cost-effective chip.



Features

- High-performance RISC – 15 MIPS sustained @ 20 MHz (20 MIPS peak)
- Memory Management Unit (MMU) support for virtual memory systems
- 4 KByte of instruction and data cache
- Write Buffer – enhancing performance
- Fully-static operation – low-power consumption – ideal for power-sensitive applications
- Fast sub microsecond interrupt response for real-time applications
- Excellent high-level language support
- Big and Little Endian operating modes
- IEEE 1149.1 Boundary Scan
- 144 Plastic Quad Flat Pack (PQFP) package

Applications

- Personal Computer Devices (i.e., PDAs)
- Portable Telecommunications
- Data Communications
- Consumer products
- Automotive

Introduction

The LH74610 architecture is based on 'Reduced Instruction Set Computers' (RISC) principles, and the instruction set and related decode mechanism are greatly simplified compared with microprogrammed 'Complex Instruction Set Computers' (CISC). The CPU within the LH74610 is the ARM6 and is software compatible with the ARM family.

The on-chip mixed data and instruction cache together with the Write Buffer substantially raise the average execution speed, and reduces the average amount of memory bandwidth required by the processor. This allows the external memory to support additional processors or Direct Memory Access (DMA) channels with minimal performance loss.

The MMU supports a conventional two-level page-table structure and a number of extensions which make it ideal for embedded control, UNIX and Object Oriented systems.

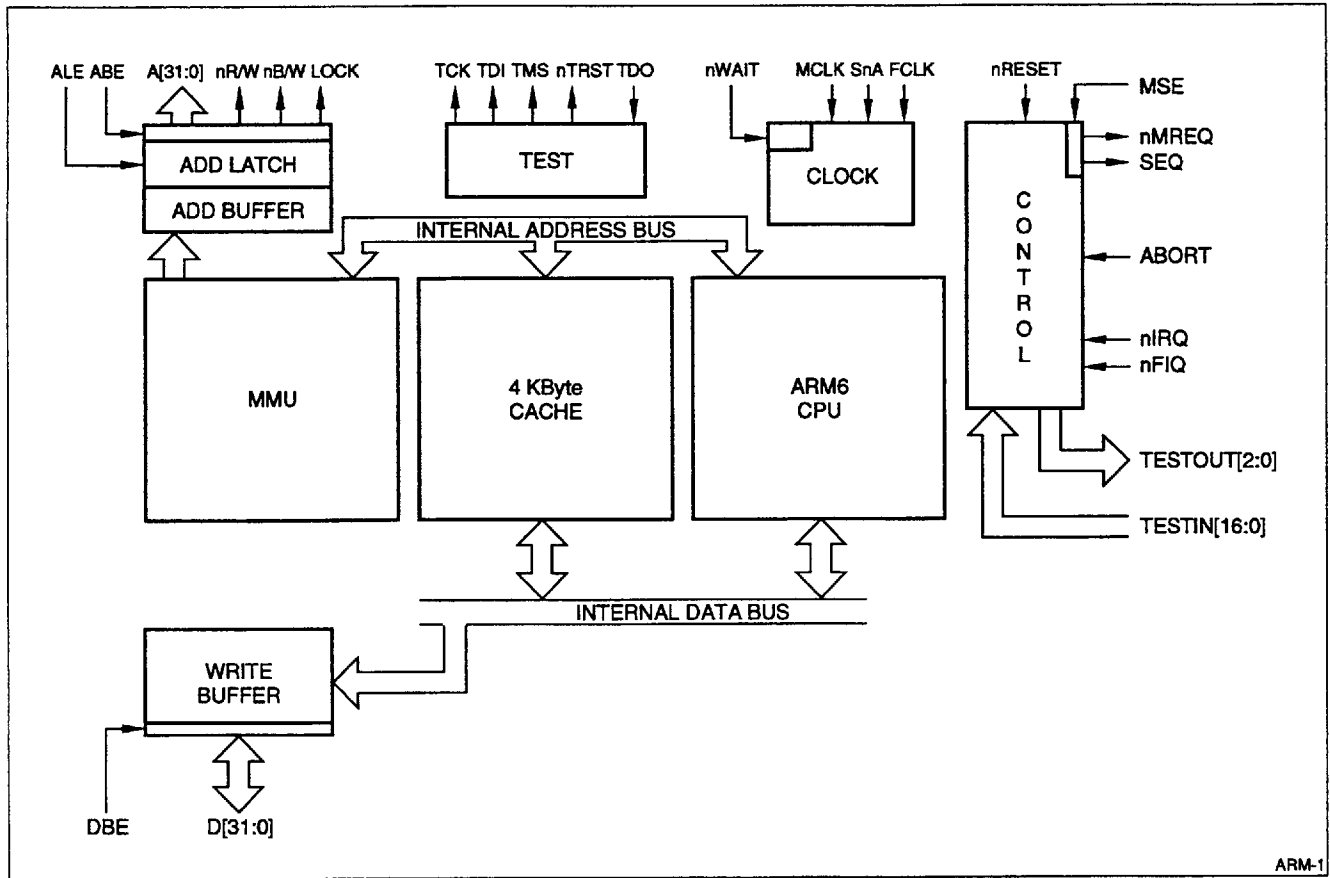
The instruction set comprises ten basic instruction types. Two of these make use of the on-chip arithmetic logic unit, barrel shifter and multiplier to perform high-speed operations on the data in a bank of 31 registers, each 32 bits wide. Three classes of instruction control the transfer of data between main memory and the register bank, one optimized for flexibility of addressing, another for rapid context switching and the third for swapping data. Two instructions control the flow and privilege level of execution, and another three types are dedicated to the control of external co-processors which allow the functionality of the instruction set to be extended off-chip in an open and uniform way.

The ARM instruction set has proved to be a good target for compilers of many different high-level languages. Where required for critical code segments, assembly code programming is also straightforward, unlike some RISC processors which depend on sophisticated compiler technology to manage complicated instruction interdependencies.

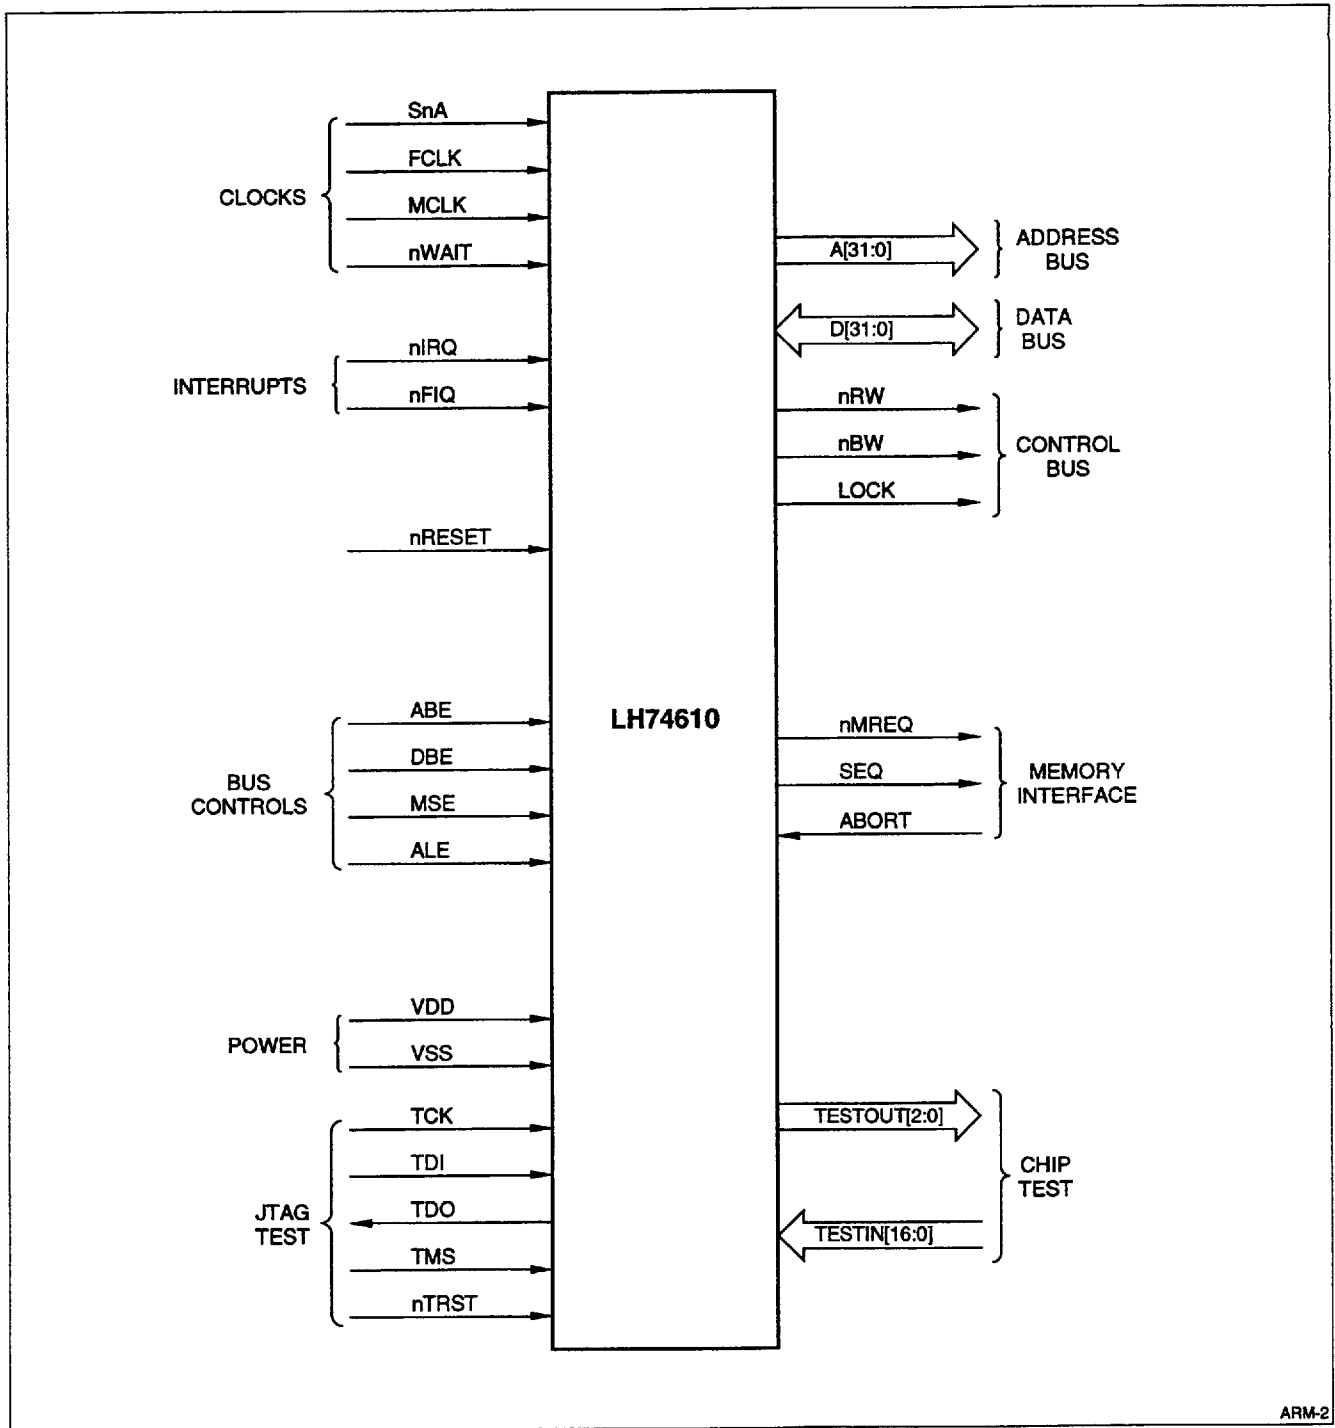
The memory interface has been designed to allow the performance potential to be realized without incurring high costs in the memory system. Speed-critical control signals are pipelined to allow system control functions to be implemented in standard low-power logic, and these control signals permit the exploitation of paged mode access offered by industry standard DRAMs.

The LH74610 is a fully static part and has been designed to minimize its power requirements. This makes it ideal for portable applications where both these features are essential.

Block Diagram



Functional Diagram



ARM-2

Description of Signals

Table 1. Description of Signals

NAME	TYPE	DESCRIPTION
A[31:0]	OCZ	Address Bus. This bus signals the address requested for memory accesses. Normally it changes during MCLK HIGH, but see ALE.
ABE	IT	Address bus enable. When this input is LOW, the address bus A[31:0], nR/W, nB/W and LOCK are put into a high impedance state (Note 1).
ABORT	IT	External abort. This input allows the memory system to tell the processor that a requested access has failed. This input is only monitored when the LH74610 is accessing external memory.
ALE	IT	Address latch enable. This input is used to control transparent latches on the address bus A[31:0], nB/W, nR/W and LOCK. Normally these signals change during MCLK HIGH, but they may be held by driving ALE LOW.
D[31:0]	ITOTZ	Data bus. These are bi-directional signal paths which are used for data transfers between the processor and external memory. For read operations (when nR/W is LOW), the input data must be valid before the falling edge of MCLK for write operations (when nR/W is HIGH), the output data will become valid while MCLK is LOW.
DBE	IT	Data bus enable. When this input is LOW, the data bus, D[31:0] is put into a high impedance state (Note 1). The drivers will always be high impedance except during write operations, and DBE must be driven HIGH in systems which do not require the data bus for DMA or similar activities.
FCLK	IT	Fast clock input. When the LH74610 CPU is accessing the cache or performing an internal cycle, it is clocked with the Fast clock, FCLK.
LOCK	OCZ	Locked operation. LOCK is driven HIGH, to signal a "locked" memory access sequence, and the memory manager should wait until LOCK goes LOW before allowing another device to access the memory. LOCK changes while MCLK is HIGH (but see ALE), and remains HIGH for the duration of the locked memory sequence.
MCLK	IT	Memory clock input. This clock times all LH74610 memory accesses. The LOW or HIGH period of MCLK may be stretched when accessing slow peripherals; alternatively, the nWAIT input may be used with a free-running MCLK to achieve similar effects.
MSE	IT	Memory request/sequential enable. When this input is LOW, the nMREQ and SEQ outputs are put into a high impedance state (Note 1).
nB/W	OCZ	Not byte/word. An output signal used by the processor to indicate to the external memory system when a data transfer of a byte length is required. nB/W is HIGH for word transfers and LOW for byte transfers, and is valid for both read and write operations. The signal changes while MCLK is HIGH, but see ALE.

Table 1. Description of Signals (Continued)

NAME	TYPE	DESCRIPTION
nFIQ	IT	Not fast interrupt request. If FIQs are enabled, the processor will respond to a LOW level on this input by taking the FIQ interrupt exception. This is an asynchronous, level-sensitive input, and must be held LOW until a suitable response is received from the processor.
nIRQ	IT	Not interrupt request. As nFIQ, but with lower priority. May be taken LOW asynchronously to interrupt the processor when the IRQ enable is active.
nMREQ	OCZ	Not memory request. This is a pipelined signal that changes while MCLK is LOW to indicate whether or not in the following cycle, the processor will be accessing external memory. When nMREQ is LOW, the processor will be accessing external memory
nRESET	IT	Not reset. This is a level sensitive input signal which is used to start the processor from a known address. A LOW level will cause the instruction being executed to terminate abnormally, and the on-chip caches, MMU, and write buffer to be disabled. When nRESET is driven HIGH, the processor will re-start from address 0. nRESET must remain LOW for at least 2 FCLK clock cycles, and 5 MCLK clock cycles while nWAIT is HIGH. During the LOW period the processor will perform idle cycles but with incrementing addresses.
nR/W	OCZ	Not read/write. When HIGH this signal indicates a processor write operation; when LOW, a read operation. The signal changes while MCLK is HIGH, but see ALE.
nTRST	IT	Test interface reset. Note this pin does NOT have an internal pullup resistor. This pin must be pulsed or driven LOW to achieve normal device operation, in addition to the normal device reset (nRESET).
nWAIT	IT	Not wait. When LOW this signal allows extra MCLK cycles to be inserted into memory accesses. It must change during the LOW phase of the MCLK cycle that is to be extended.
SEQ	OCZ	Sequential address. This signal is the inverse of nMREQ, and is provided for compatibility with existing ARM memory systems.
SnA	IT	This pin should be hard wired HIGH.
TESTIN[16:0]	IT	Test bus input. This bus is used for off-board testing of the device. When the device is fitted to a circuit all these pins must be tied LOW. The signals are further defined within the chip as: TESTIN[16], TESTIN[15:8], TESTIN[7:0]
TESTOUT[2:0]	OCZ	Test bus output. This bus is used for off-board testing of the device. When the device is fitted to a circuit and all the TESTIN[16:0] pins are driven LOW, these three outputs will be driven LOW. Note that these pins may not be tristated, except via the JTAG test port.
TCK	IT	Test interface reference Clock. This times all the transfers on the JTAG test interface.

Table 1. Description of Signals (Continued)

NAME	TYPE	DESCRIPTION
TDI	IT	Test interface data input. Note this pin does NOT have an internal pullup resistor.
TDO	OCZ	Test interface data output. Note this pin does NOT have an internal pullup resistor.
TMS	IT	Test interface mode select. Note this pin does NOT have an internal pullup resistor.
VDD		Positive supply. 15 pins are allocated to VDD in the 144 TQFP package.
VSS		Ground supply. 15 pins are allocated to VSS in the 144 TQFP package.
No Connect		Unallocated pins. There are 8 of these, which are reserved for future use. Note that these pins are not bonded, so no connection at the circuit board is necessary.

NOTES:

1. When output pads are placed in the high impedance state for long periods, care must be taken to ensure that they do not float to an undefined logic level, as this can dissipate power, especially in the pads.
2. Although the input pads have TTL thresholds, and will correctly interpret a TTL level input, it must be noted that unless all inputs are driven to the voltage rails, the input circuits will consume power.
3. Key to Signal Types
 - IT - Input, TTL threshold
 - OCZ - Output, CMOS levels, tri-stateable
 - IOTZ - Input/output tri-stateable, TTL thresholds

Configuration and Mode Selection

The LH74610 supports a variety of operating configurations. Some are controlled by register bits and are known as the register configurations. Others may be controlled by software and these are known as operating modes.

Register Configuration

The LH74610 processor provides 4 register configurations which may be changed while the processor is running. The actual register settings are given in detail in *Configuring the LH74610* but their effects are described below:

Two of the registers (data32 and prog32), allow one of three processor configurations to be selected as follows.

- (1) 26 bit program and data space - (prog32 LOW, data32 LOW). This configuration forces the LH74610 to operate like the earlier ARM processors with 26 bit address space. The programmer's model for these processors applies, but the new instructions to access the CPSR and SPSR registers operate as detailed later in this document. In this configuration it is impossible to select a 32 bit operating mode, and all exceptions (including address exceptions) enter the exception handler in the appropriate 26 bit mode.
- (2) 26 bit program space and 32 bit data space - (prog32 LOW, data32 HIGH). This is the same as the 26 bit program and data space configuration, but with address exceptions disabled to allow data transfer operations to access the full 32 bit address space.
- (3) 32 bit program and data space - (prog32 HIGH, data32 HIGH). This configuration extends the address space to 32 bits, introduces major changes in the programmer's model as described below and provides support for running existing 26 bit programs in the 32 bit environment.

The fourth processor configuration which is possible (26 bit data space and 32 bit program space) should not be selected.

The bigend signal controls whether the LH74610 treats words in memory as being stored in Big Endian or Little Endian format. Memory is viewed as a linear collection of bytes numbered upwards from zero. Bytes 0 to 3 hold the first stored word, bytes 4 to 7 the second and so on.

In the Little Endian scheme the lowest numbered byte in a word is considered to be the least significant byte of the word and the highest numbered byte is the most significant. Byte 0 of the memory system should be connected to data lines 7 through 0 (data[7:0] and dout[7:0]) in this scheme.

In the Big Endian scheme the most significant byte of a word is stored at the lowest numbered byte and the least significant byte is stored at the highest numbered byte. Byte 0 of the memory system should therefore be connected to data lines 31 through 24 (data[31:24] and dout[31:24])

The lateabt signal controls the processor's behavior when a data abort exception occurs. It only affects the behavior of LDR and STR instructions and is discussed more fully in the Programmer's Model and Instruction Set chapters.

Operating Mode Selection

When configured for 26 bit program space, the LH74610 is limited to operating in one of four modes known as the 26 bit modes. These modes correspond to the modes of the earlier ARM processors and are known as User26, FIQ26, IRQ26 and Supervisor26.

When using a 32 bit program space there are a total of 10 modes available. These are the four 26 bit modes described above plus 6 more known as the 32 bit modes. These are User32, FIQ32, IRQ32, Supervisor32, Abort32, and Undefined32. These are the normal operating modes in this configuration and the 26 bit modes are only provided for backwards compatibility to allow execution of programs originally written

NOTES

- (1) The LH74610 supports 32-bit program and data space configuration *only*.
- (2) The remainder of this document describes the LH74610 when configured for 32 bit program and data space and operating in one of the 32 bit modes. It is recommended that all new designs using the LH74610 should configure the processor in this way by setting prog32 and data32 HIGH and that all new code should be written to use only the 32 bit operating modes. It is also recommended that the lateabt register setting be set HIGH so that the Late Abort exception mechanism is used.
- (3) Because the original ARM instruction set has been modified to accommodate 32 bit operation there are certain additional restrictions which programmers must be aware of. These are indicated in the text by the words shall and shall not. Reference should also be made to the ARM Application Notes *Rules for ARM Code Writers* and *Notes for ARM Code Writers* available from your supplier.

Programmer's Model

The LH74610 has a 32 bit data bus and a 32 bit address bus. The data types the processor supports are Bytes (8 bits) and Words (32 bits), where words must be aligned to four byte boundaries. Instructions are exactly one word, and data operations (e.g. ADD) are only performed on word quantities. Load and store operations can transfer either bytes or words.

The LH74610 supports six modes of operation:

- (1) User mode: the normal program execution state
- (2) FIQ mode (fiq): designed to support a data transfer or channel process
- (3) IRQ mode (irq): used for general purpose interrupt handling
- (4) Supervisor mode (svc): a protected mode for the operating system
- (5) Abort mode (abt): entered after a data or instruction prefetch abort
- (6) Undefined mode (und): entered when an undefined instruction is executed

Mode changes may be made under software control or may be brought about by external interrupts or exception processing. Most application programs will execute in User mode. The other modes, known as privileged modes, will be entered to service interrupts or exceptions or to access protected resources.

Registers

The processor has a total of 37 registers made up of 31 general 32 bit registers and 6 status registers. At any one time 16 general registers (R0 to R15) and one or two status registers are visible to the programmer. The visible registers depend on the processor mode and the other registers (the banked registers) are switched in to support IRQ, FIQ, Supervisor, Abort and Undefined mode processing. The register bank organization is shown in Figure 1. The banked registers are shaded in the diagram.

In all modes 16 registers, R0 to R15, are directly accessible. All registers except R15 are general purpose and may be used to hold data or address values. Register R15 holds the Program Counter (PC). When R15 is read, bits [1:0] are zero and bits [31:2] contain the PC. A seventeenth register (the CPSR - Current Program Status Register) is also accessible. It contains condition code flags and the current mode bits and may be thought of as an extension to the PC.

R14 is used as the subroutine link register and receives a copy of R15 when a Branch and Link instruction is executed. It may be treated as a general purpose register at all other times. R14_svc, R14_irq, R14_fiq, R14_abt and R14_und are used similarly to hold the return values of R15 when interrupts and exceptions arise, or when Branch and Link instructions are executed within interrupt or exception routines.

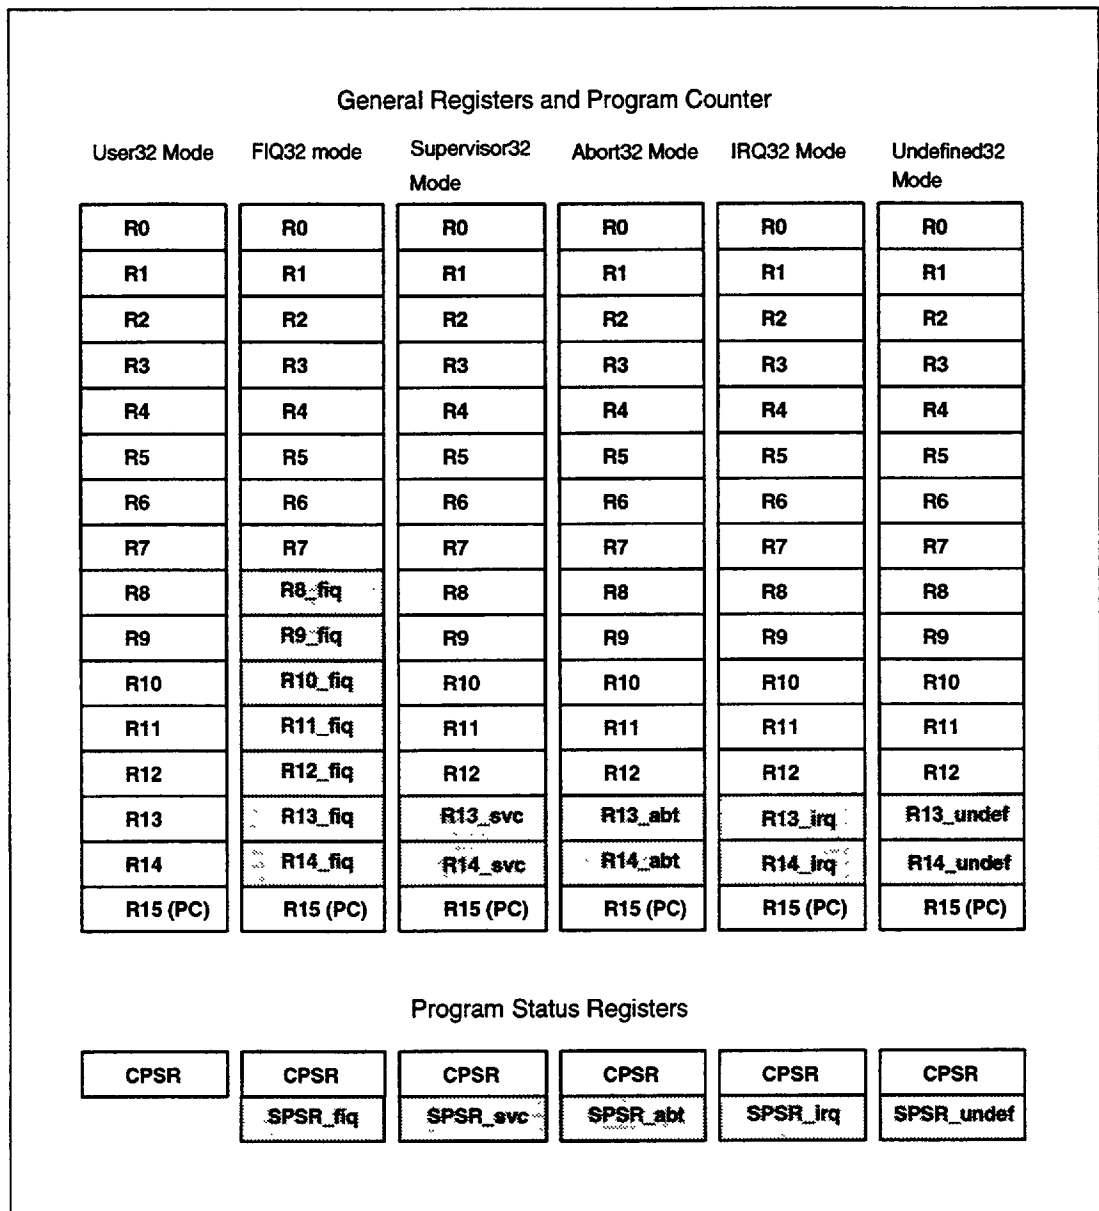


Figure 1. Register Organization

FIQ mode has seven banked registers mapped to R8-14 (R8_fiq-R14_fiq). Many FIQ programs will not need to save any registers. User mode, IRQ mode, Supervisor mode, Abort mode and Undefined mode each have two banked registers mapped to R13 and R14. The two banked registers allow these modes to each have a private stack pointer and link register. Supervisor, IRQ, Abort and Undefined mode programs which require more than these two banked registers are expected to save some or all of the caller's registers (R0 to R12) on their respective stacks. They are then free to use these registers which they will restore before returning to the caller. In addition there are also five SPSRs (Saved Program Status Registers) which are loaded with the CPSR when an exception occurs. There is one SPSR for each privileged mode.

The format of the Program Status Registers is shown in Figure 2. The N, Z, C and V bits are the condition code flags. The condition code flags in the CPSR may be changed as a result of arithmetic and logical operations in the processor and may be tested by all instructions to determine if the instruction is to be executed.

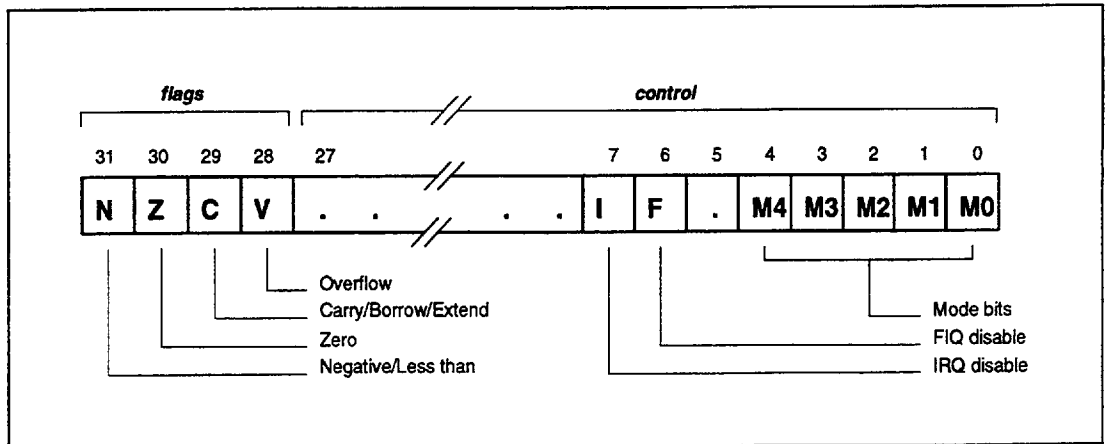


Figure 2. Format of the Program Status Registers (PSRs)

The I and F bits are the interrupt disable bits. The I bit disables IRQ interrupts when it is set and the F bit disables FIQ interrupts when it is set. The M0, M1, M2, M3 and M4 bits (M[4:0]) are the mode bits, and these determine the mode in which the processor operates. The interpretation of the mode bits is shown in Table 2. Not all combinations of the mode bits define a valid processor mode. Only those explicitly described shall be used.

The bottom 28 bits of a PSR (incorporating I, F and M[4:0]) are known collectively as the control bits. The control bits will change when an exception arises and in addition can be manipulated by software when the processor is in a privileged mode. Unused bits in the PSRs are reserved and their state shall be preserved when changing the flag or control bits. Programs shall not rely on specific values from the reserved bits when checking the PSR status, since they may read as one or zero in future processors.

Table 2. Mode Bits

M[4:0]	MODE	ACCESSIBLE REGISTER SET	
10000	usr	PC, R14..R0	CPSR
10001	fiq	PC, R14_fiq..R8_fiq, R7..R0	CPSR, SPSR_fiq
10010	irq	PC, R14_irq..R13_irq, R12..R0	CPSR, SPSR_irq
10011	svc	PC, R14_svc..R13_svc, R12..R0	CPSR, SPSR_svc
10111	abt	PC, R14_abt..R13_abt, R12..R0	CPSR, SPSR_abt
11011	und	PC, R14_und..R13_und, R12..R0	CPSR, SPSR_und

Exceptions

Exceptions arise whenever there is a need for the normal flow of program execution to be broken, so that (for example) the processor can be diverted to handle an interrupt from a peripheral. The processor state just prior to handling the exception must be preserved so that the original program can be resumed when the exception routine has completed. Many exceptions may arise at the same time.

The LH74610 handles exceptions by making use of the banked registers to save state. The old PC and CPSR contents are copied into the appropriate R14 and SPSR and the PC and mode bits in the CPSR bits are forced to a value which depends on the exception. Interrupt disable flags are set where required to prevent otherwise unmanageable nestings of exceptions. In the case of a re-entrant interrupt handler, R14 and the SPSR should be saved onto a stack in main memory before re-enabling the interrupt; when transferring the SPSR register to and from a stack, it is important to transfer the whole 32 bit value, and not just the flag or control fields. When multiple exceptions arise simultaneously, a fixed priority determines the order in which they are handled. The priorities are listed later in this chapter.

FIQ

The FIQ (Fast Interrupt reQuest) exception is externally generated by taking the Nfiq input LOW. This input can accept asynchronous transitions, and is delayed by one clock cycle for synchronization before it can affect the processor execution flow. It is designed to support a data transfer or channel process, and has sufficient private registers to remove the need for register saving in such applications (thus minimizing the overhead of context switching). The FIQ exception may be disabled by setting the F flag in the CPSR (but note that this is not possible from User mode). If the F flag is clear, the LH74610 checks for a LOW level on the output of the FIQ synchronizer at the end of each instruction.

When a FIQ is detected, the LH74610 performs the following:

- (1) Saves the address of the next instruction to be executed plus 4 in R14_fiq; saves CPSR in SPSR_fiq
- (2) Forces M[4:0]=%10001 (FIQ mode) and sets the F and I bits in the CPSR
- (3) Forces the PC to fetch the next instruction from address &1C

To return normally from FIQ, use SUBS PC, R14_fiq,#4 which will restore both the PC (from R14) and the CPSR (from SPSR_fiq) and resume execution of the interrupted code.

IRQ

The IRQ (Interrupt ReQuest) exception is a normal interrupt caused by a LOW level on the Nirq input. It has a lower priority than FIQ, and is masked out when a FIQ sequence is entered. Its effect may be masked out at any time by setting the I bit in the CPSR (but note that this is not possible from User mode). If the I flag is clear, the LH74610 checks for a LOW level on the output of the IRQ synchronizer at the end of each instruction. When an IRQ is detected, the LH74610 performs the following:

- (1) Saves the address of the next instruction to be executed plus 4 in R14_irq; saves CPSR in SPSR_irq
- (2) Forces M[4:0]=%10010 (IRQ mode) and sets the I bit in the CPSR
- (3) Forces the PC to fetch the next instruction from address &18

To return normally from IRQ, use SUBS PC,R14_irq,#4 which will restore both the PC and the CPSR and resume execution of the interrupted code.

Abort

The abort input comes from an external Memory Management system, and indicates that the current memory access cannot be completed. For instance, in a virtual memory system the data corresponding to the current address may have been moved out of memory onto a disc, and considerable processor activity may be required to recover the data before the access can be performed successfully. The LH74610 checks for abort during memory access (N and S) cycles. When successfully aborted, the LH74610 will respond in one of two ways:

- (1) If the abort occurred during an instruction prefetch (a *Prefetch Abort*), the prefetched instruction is marked as invalid but the abort exception does not occur immediately. If the instruction is not executed, for example as a result of a branch being taken while it is in the pipeline, no abort will occur. An abort will take place if the instruction reaches the head of the pipeline and is about to be executed.
- (2) If the abort occurred during a data access (a *Data Abort*), the action depends on the instruction type.
 - (a) Single data transfer instructions (LDR, STR) are aborted as though the instruction had not executed if the processor is configured for Early Abort. When configured for Late Abort, these instructions are able to write back modified base registers and the Abort handler must be aware of this.
 - (b) The swap instruction (SWP) is aborted as though it had not executed.
 - (c) Block data transfer instructions (LDM, STM) complete, and if write-back is set, the base is updated. If the instruction would normally have overwritten the base with data (i.e. LDM with the base in the transfer list), this overwriting is prevented. All register overwriting is prevented after the Abort is indicated, which means in particular that R15 (which is always last to be transferred) is preserved in an aborted LDM instruction.

When either a prefetch or data abort occurs, the LH74610 performs the following:

- (1) Saves the address of the aborted instruction plus 4 (for prefetch aborts) or 8 (for data aborts) in R14_abt; saves CPSR in SPSR_abt
- (2) Forces M[4:0]=%10111 (Abort mode) and sets the I bit in the CPSR.
- (3) Forces the PC to fetch the next instruction from either address &0C (prefetch abort) or address &10 (data abort)

To return after fixing the reason for the abort, use `SUBS PC,R14_abt,#4` (for a prefetch abort) or `SUBS PC,R14_abt,#8` (for a data abort). This will restore both the PC and the CPSR and retry the aborted instruction.

The abort mechanism allows a demand paged virtual memory system to be implemented when a suitable memory manager is available. The processor is allowed to generate arbitrary addresses, and when the data at an address is unavailable the memory manager signals an abort. The processor traps into system software which must work out the cause of the abort, make the requested data available, and retry the aborted instruction. The application program needs no knowledge of the amount of memory available to it, nor is its state in any way affected by the abort.

Software Interrupt

The software interrupt instruction (SWI) is used for getting into Supervisor mode, usually to request a particular supervisor function. When a SWI is executed, the LH74610 performs the following:

- (1) Saves the address of the SWI instruction plus 4 in `R14_svc`; saves CPSR in `SPSR_svc`
- (2) Forces `M[4:0]=%10011` (Supervisor mode) and sets the I bit in the CPSR
- (3) Forces the PC to fetch the next instruction from address `&08`

To return from a SWI, use `MOVS PC,R14_svc`. This will restore the PC and CPSR and return to the instruction following the SWI.

Undefined Instruction Trap

When the LH74610 comes across a coprocessor instruction which it can not handle (see Instruction Set chapter) then the LH74610 will take the undefined instruction trap.

The trap may be used for software emulation of a coprocessor in a system which does not have the coprocessor hardware, or for general purpose instruction set extension by software emulation.

When the LH74610 takes the undefined instruction trap it performs the following:

- (1) Saves the address of the Undefined or coprocessor instruction plus 4 in `R14_und`; saves CPSR in `SPSR_und`.
- (2) Forces `M[4:0]=%11011` (Undefined mode) and sets the I bit in the CPSR
- (3) Forces the PC to fetch the next instruction from address `&04`

To return from this trap after emulating the failed instruction, use `MOVS PC,R14_und`. This will restore the CPSR and return to the instruction following the undefined instruction.

Reset

When the Nreset signal goes LOW, the LH74610 abandons the currently executing instruction and then continues to fetch instructions from memory which it interprets as NOPs. When Nreset goes HIGH again, the LH74610 does the following:

- (1) Overwrites R14_svc and SPSR_svc by copying the current values of the PC and CPSR into them. The value of the saved PC and CPSR is not defined.
- (2) Forces M[4:0]=%10011 (Supervisor mode) and sets the I and F bits in the CPSR.
- (3) Forces the PC to fetch the next instruction from address &00

Vector Summary

Table 3. Vector Summary

ADDRESS	EXCEPTION	MODE ON ENTRY
&00000000	Reset	Supervisor
&00000004	Undefined instruction	Undefined
&00000008	Software interrupt	Supervisor
&0000000C	Abort (prefetch)	Abort
&00000010	Abort (data)	Abort
&00000014	-- reserved --	--
&00000018	IRQ	IRQ
&0000001C	FIQ	FIQ

These are byte addresses, and will normally contain a branch instruction pointing to the relevant routine.

The FIQ routine might reside at &1C onwards, and thereby avoid the need for (and execution time of) a branch instruction.

The reserved entry is for an Address Exception vector which is only operative when the processor is configured for a 26 bit program space.

Exception Priorities

When multiple exceptions arise at the same time, a fixed priority system determines the order in which they will be handled:

- (1) Reset (highest priority)
- (2) Data abort
- (3) FIQ
- (4) IRQ

- (5) Prefetch abort
- (6) Undefined Instruction, Software interrupt (lowest priority)

Note that not all exceptions can occur at once. Undefined instruction and software interrupt are mutually exclusive since they each correspond to particular (non-overlapping) decodings of the current instruction.

If a data abort occurs at the same time as a FIQ, and FIQs are enabled (i.e. the F flag in the CPSR is clear), the LH74610 will enter the data abort handler and then immediately proceed to the FIQ vector. A normal return from FIQ will cause the data abort handler to resume execution. Placing data abort at a higher priority than FIQ is necessary to ensure that the transfer error does not escape detection; the time for this exception entry should be added to worst case FIQ latency calculations.

Interrupt Latencies

The worst case latency for FIQ, assuming that it is enabled, consists of the longest time the request can take to pass through the synchronizer ($T_{syncmax}$), plus the time for the longest instruction to complete (T_{ldm} , the longest instruction is a LDM which loads all the registers including the PC), plus the time for the data abort entry (T_{exc}), plus the time for FIQ entry (T_{fiq}). At the end of this time the LH74610 will be executing the instruction at &1C.

$T_{syncmax}$ is 3 processor cycles, T_{ldm} is 20 cycles, T_{exc} is 3 cycles, and T_{fiq} is 2 cycles. The total time is therefore 28 processor cycles. This is just over 1.4 microseconds in a system which uses a continuous 20 MHz processor clock. The maximum IRQ latency calculation is similar, but must allow for the fact that FIQ has higher priority and could delay entry into the IRQ handling routine for an arbitrary length of time. The minimum latency for FIQ or IRQ consists of the shortest time the request can take through the synchronizer ($T_{syncmin}$) plus T_{fiq} . This is 4 processor cycles.

Instruction Set

The Condition Field

All LH74610 instructions are conditionally executed, which means that their execution may or may not take place depending on the values of the N, Z, C and V flags in the CPSR. The condition encoding is shown in Figure 3.

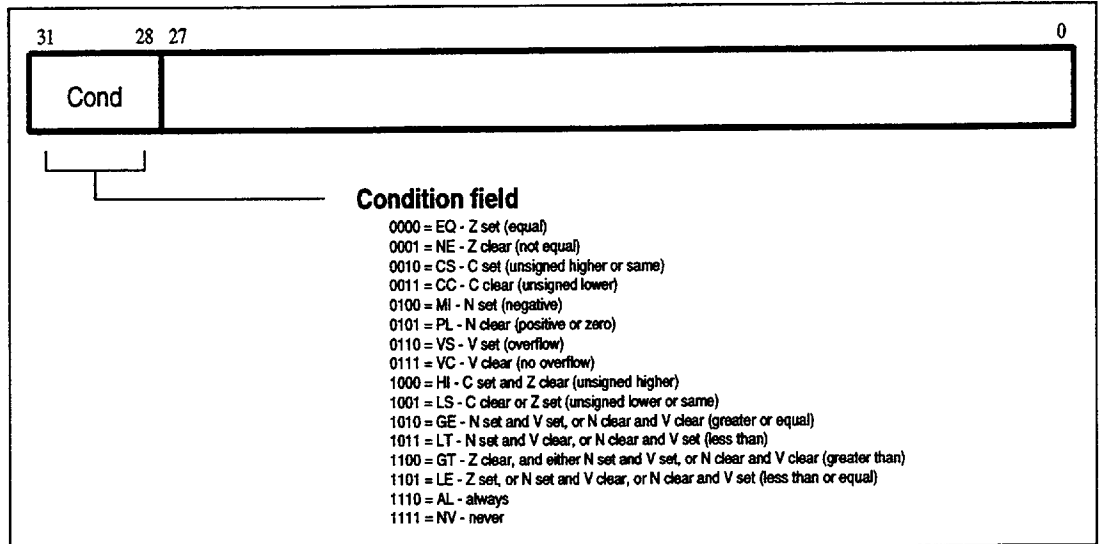


Figure 3. Condition Codes

If the always (AL) condition is specified, the instruction will be executed irrespective of the flags. The never (NV) class of condition codes should not be used as they will be redefined in future variants of the ARM architecture. If a NOP is required it is suggested that MOV R0,R0 be used. The assembler treats the absence of a condition code as though AL had been specified.

The other condition codes have meanings as detailed in Figure 5, for instance code 0000 (EQual) causes the instruction to be executed only if the Z flag is set. This would correspond to the case where a compare (CMP) instruction had found the two operands to be equal. If the two operands were different, the compare instruction would have cleared the Z flag and the instruction will not be executed.

Branch and Branch with Link (B, BL)

The instruction is only executed if the condition is true. The various conditions are defined at the beginning of this chapter. The instruction encoding is shown in Figure 4.

Branch instructions contain a signed 2's complement 24 bit offset. This is shifted left two bits, sign extended to 32 bits, and added to the PC. The instruction can therefore specify a branch of ± 32 Mbytes. The branch offset must take account of the prefetch operation, which causes the PC to be 2 words (8 bytes) ahead of the current instruction.

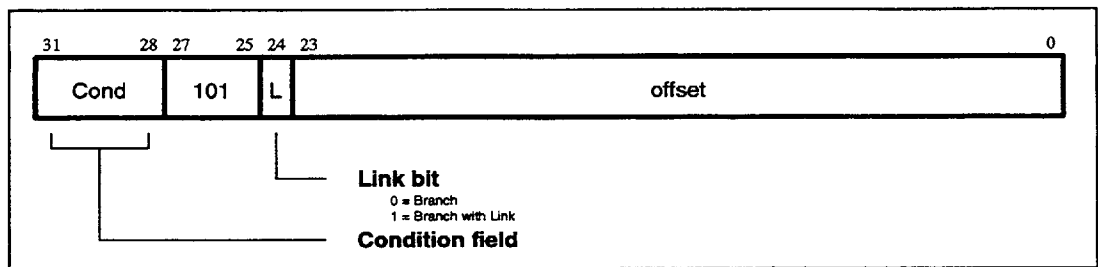


Figure 4. Branch Instructions

Branches beyond ± 32 Mbytes must use an offset or absolute destination which has been previously loaded into a register. In this case the PC should be manually saved in R14 if a Branch with Link type operation is required.

The Link Bit

Branch with Link (BL) writes the old PC into the link register (R14) of the current bank. The PC value written into R14 is adjusted to allow for the prefetch, and contains the address of the instruction following the branch and link instruction. Note that the CPSR is not saved with the PC.

To return from a routine called by Branch with Link use `MOV PC,R14` if the link register is still valid or `LDM Rn!,{..PC}` if the link register has been saved onto a stack pointed to by Rn.

Assembler Syntax

`B{L}{cond} <expression>`

`{L}` is used to request the Branch with Link form of the instruction. If absent, R14 will not be affected by the instruction.

`{cond}` is a two-char mnemonic as shown in Figure 3 (EQ, NE, VS etc.). If absent then AL (ALways) will be used.

`<expression>` is the destination. The assembler calculates the offset.

Items in `{}` are optional. Items in `<>` must be present.

Examples

```
here BAL here      ; assembles to &EAFFFFFFE (note effect of PC offset)
  B there          ; ALways condition used as default
```

```
  CMP R1,#0        ; compare R1 with zero and branch to fred if R1
  BEQ fred         ; was zero otherwise continue to next instruction
```

```
  BL sub + ROM     ; call subroutine at computed address
```

```
  ADDS R1,#1       ; add 1 to register 1, setting CPSR flags on the
  BLCC sub         ; result then call subroutine if the C flag is clear,
                  ; which will be the case unless R1 held &FFFFFFF
```

Data Processing

The instruction is only executed if the condition is true, defined at the beginning of this section. The instruction encoding is shown in Figure 5.

The instruction produces a result by performing a specified arithmetic or logical operation on one or two operands. The first operand is always a register (Rn). The second operand may be a shifted register (Rm) or a rotated 8 bit immediate value (Imm) according to the value of the I bit in the instruction. The condition codes in the CPSR may be preserved or updated as a result of this instruction, according to the value of the S bit in the instruction. Certain operations (TST, TEQ, CMP, CMN) do not write the result to Rd. They are used only to perform tests and to set the condition codes on the result and always have the S bit set. The instructions and their effects are listed in Table 4.

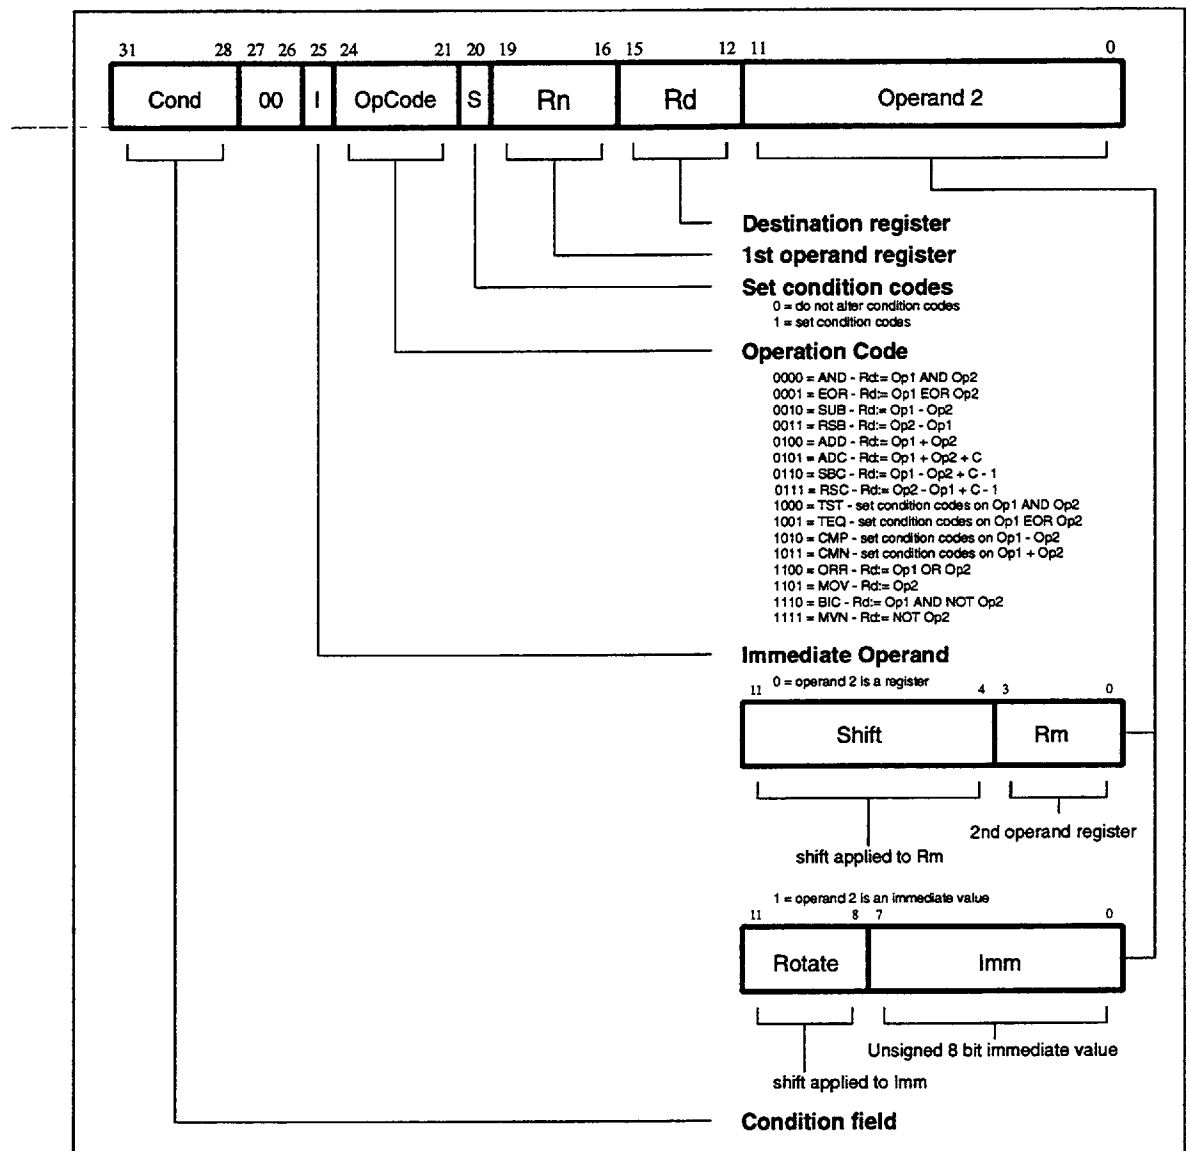


Figure 5. Data Processing Instructions

CPSR Flags

The data processing operations may be classified as logical or arithmetic. The logical operations (AND, EOR, TST, TEQ, ORR, MOV, BIC, MVN) perform the logical action on all corresponding bits of the operand or operands to produce the result. If the S bit is set (and Rd is not R15, see below) the V flag in the CPSR will be unaffected, the C flag will be set to the carry out from the barrel shifter (or preserved when the shift operation is LSL #0), the Z flag will be set if and only if the result is all zeros, and the flag will be set to the logical value of bit 31 of the result.

Table 4. ARM Data Processing Instruction

ASSEMBLER MNEMONIC	OPCODE	ACTION
AND	0000	operand1 AND operand2
EOR	0001	operand1 EOR operand2
SUB	0010	operand1 - operand2
RSB	0011	operand2 - operand1
ADD	0100	operand1 + operand2
ADC	0101	operand1 + operand2 + carry
SBC	0110	operand1 - operand2 + carry - 1
RSC	0111	operand2 - operand1 + carry - 1
TST	1000	as AND, but result is not written
TEQ	1001	as EOR, but result is not written
CMP	1010	as SUB, but result is not written
CMN	1011	as ADD, but result is not written
ORR	1100	operand1 OR operand2
MOV	1101	operand2 (operand1 is ignored)
BIC	1110	operand1 AND NOT operand2 (Bit clear)
MVN	1111	NOT operand2 (operand1 is ignored)

The arithmetic operations (SUB, RSB, ADD, ADC, SBC, RSC, CMP, CMN) treat each operand as a 32 bit integer (either unsigned or 2's complement signed, the two are equivalent). If the S bit is set (and Rd is not R15) the V flag in the CPSR will be set if an overflow occurs into bit 31 of the result; this may be ignored if the operands were considered unsigned, but warns of a possible error if the operands were 2's complement signed. The C flag will be set to the carry out of bit 31 of the ALU, the Z flag will be set if and only if the result was zero, and the N flag will be set to the value of bit 31 of the result (indicating a negative result if the operands are considered to be 2's complement signed).

Shifts

When the second operand is specified to be a shifted register, the operation of the barrel shifter is controlled by the Shift field in the instruction. This field indicates the type of shift to be performed (logical left or right, arithmetic right or rotate right). The amount by which the register should be shifted may be contained in an immediate field in the instruction, or in the bottom byte of another register (other than R15). The encoding for the different shift types is shown in Figure 6.

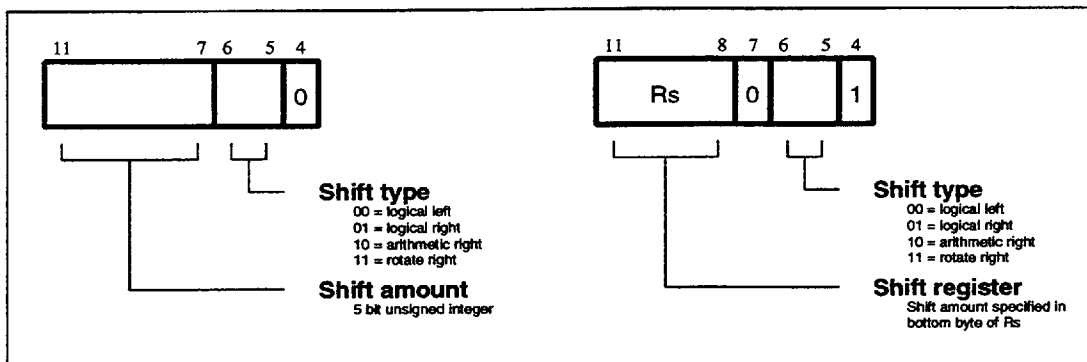


Figure 6. ARM Shift Operations

Instruction Specified Shift Amount

When the shift amount is specified in the instruction, it is contained in a 5 bit field which may take any value from 0 to 31. A logical shift left (LSL) takes the contents of Rm and moves each bit by the specified amount to a more significant position. The least significant bits of the result are filled with zeros, and the high bits of Rm which do not map into the result are discarded, except that the least significant discarded bit becomes the shifter carry output which may be latched into the C bit of the CPSR when the ALU operation is in the logical class (see above). For example, the effect of LSL #5 is shown in Figure 7.

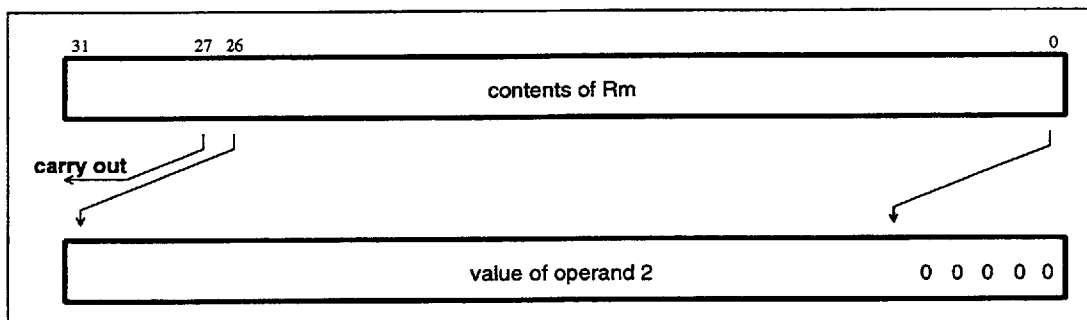


Figure 7. Logical Shift Left

Note that LSL #0 is a special case, where the shifter carry out is the old value of the CPSR C flag. The contents of Rm are used directly as the second operand. A logical shift right (LSR) is similar, but the contents of Rm are moved to less significant positions in the result. LSR #5 has the effect shown in Figure 8.

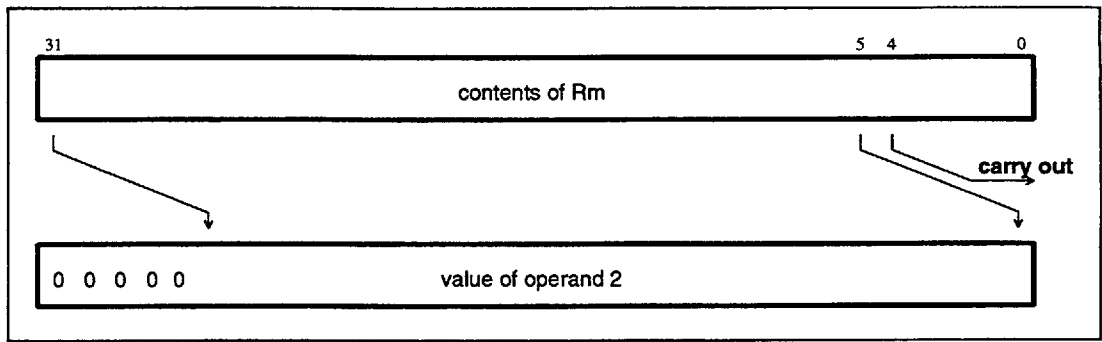


Figure 8. Logical Shift Right

The form of the shift field which might be expected to correspond to LSR #0 is used to encode LSR #32, which has a zero result with bit 31 of Rm as the carry output. Logical shift right zero is redundant as it is the same as logical shift left zero, so the assembler will convert LSR #0 (and ASR #0 and ROR #0) into LSL #0, and allow LSR #32 to be specified.

An arithmetic shift right (ASR) is similar to logical shift right, except that the high bits are filled with bit 31 of Rm instead of zeros. This preserves the sign in 2's complement notation. For example, ASR #5 is shown in Figure 9.

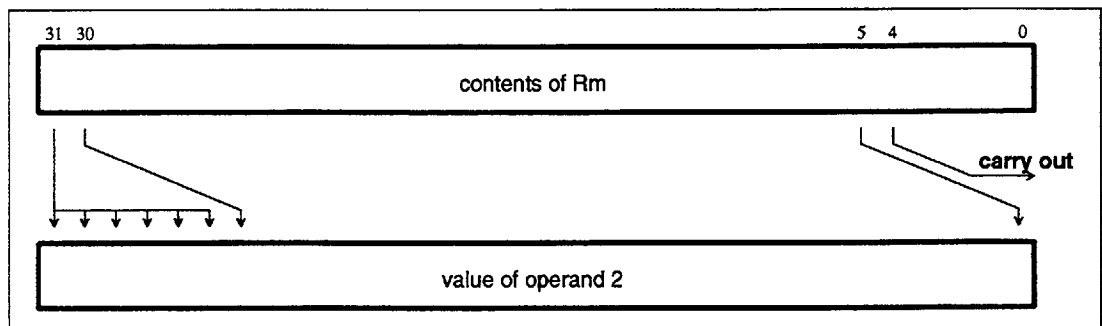


Figure 9. Arithmetic Shift Right

The form of the shift field which might be expected to give ASR #0 is used to encode ASR #32. Bit 31 of Rm is again used as the carry output, and each bit of operand 2 is also equal to bit 31 of Rm. The result is therefore all ones or all zeros, according to the value of bit 31 of Rm.

Rotate right (ROR) operations reuse the bits which 'overshoot' in a logical shift right operation by reintroducing them at the high end of the result, in place of the zeros used to fill the high end in logical right operations. For example, ROR #5 is shown in Figure 10.




Figure 10. Rotate Right

The form of the shift field which might be expected to give ROR #0 is used to encode a special function of the barrel shifter, rotate right extended (RRX). This is a rotate right by one bit position of the 33 bit quantity formed by appending the CPSR C flag to the most significant end of the contents of Rm as shown in Figure 11.

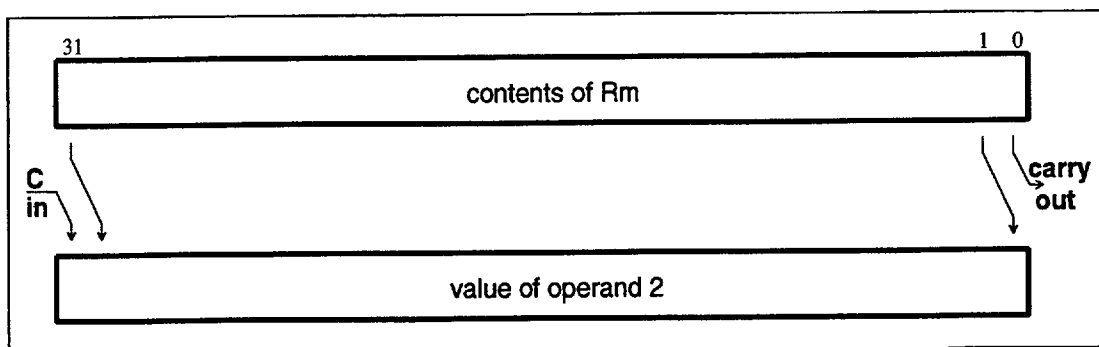


Figure 11. Rotate Right Extended

Register Specified Shift Amount

Only the least significant byte of the contents of Rs is used to determine the shift amount. Rs can be any general register other than R15.

If this byte is zero, the unchanged contents of Rm will be used as the second operand, and the old value of the CPSR C flag will be passed on as the shifter carry output.

If the byte has a value between 1 and 31, the shifted result will exactly match that of an instruction specified shift with the same value and shift operation.

If the value in the byte is 32 or more, the result will be a logical extension of the shifting processes described above:

- (1) LSL by 32 has result zero, carry out equal to bit 0 of Rm.
- (2) LSL by more than 32 has result zero, carry out zero.
- (3) LSR by 32 has result zero, carry out equal to bit 31 of Rm.

- (4) LSR by more than 32 has result zero, carry out zero.
- (5) ASR by 32 or more has result filled with and carry out equal to bit 31 of Rm.
- (6) ROR by 32 has result equal to Rm, carry out equal to bit 31 of Rm.
- (7) ROR by n where n is greater than 32 will give the same result and carry out as ROR by n-32; therefore repeatedly subtract 32 from n until the amount is in the range 1 to 32 and see above.

Note that the zero in bit 7 of an instruction with a register controlled shift is compulsory; a one in this bit will cause the instruction to be a multiply or undefined instruction.

Immediate Operand Rotates

The immediate operand rotate field is a 4 bit unsigned integer which specifies a shift operation on the 8 bit immediate value. The immediate value is zero extended to 32 bits, and then subject to a rotate right by twice the value in the rotate field. This enables many common constants to be generated, for example all powers of 2.

Writing to R15

When Rd is a register other than R15, the condition code flags in the CPSR may be updated from the ALU flags as described above.

When Rd is R15 and the S flag in the instruction is not set the result of the operation is placed in R15 and the CPSR is unaffected.

When Rd is R15 and the S flag is set the result of the operation is placed in R15 and the SPSR corresponding to the current mode is moved to the CPSR. This allows state changes which atomically restore both PC and CPSR. This form of instruction shall not be used in User mode.

Using R15 as an Operand

If R15 (the PC) is used as an operand in a data processing instruction the register is used directly.

The PC value will be the address of the instruction, plus 8 or 12 bytes due to instruction prefetching. If the shift amount is specified in the instruction, the PC will be 8 bytes ahead. If a register is used to specify the shift amount the PC will be 12 bytes ahead.

The TEQ, TST, CMP and CMN Opcodes

These instructions do not write the result of their operation but do set flags in the CPSR. An assembler shall always set the S flag for these instructions even if it is not specified in the mnemonic.

The TEQP form of the instruction used in earlier processors shall not be used in the 32 bit modes, the PSR transfer operations should be used instead. If used in these modes, its effect is to move SPSR_<mode> to CPSR if the processor is in a privileged mode and to do nothing if in User mode.

Assembler Syntax

- (1) MOV,MVN - single operand instructions

<opcode>{cond}{S} Rd,<Op2>

- (2) CMP,CMN,TEQ,TST - instructions which do not produce a result.

<opcode>{cond} Rn,<Op2>

- (3) AND,EOR,SUB,RSB,ADD,ADC,SBC,RSC,ORR,BIC

<opcode>{cond}{S} Rd,Rn,<Op2>

where <Op2> is Rm{,<shift>} or,<#expression>

{cond} - two-character condition mnemonic, see Figure 3

{S} - set condition codes if S present (implied for CMP, CMN, TEQ, TST).

Rd, Rn and Rm are expressions evaluating to a register number.

If <#expression> is used, the assembler will attempt to generate a shifted immediate 8-bit field to match the expression. If this is impossible, it will give an error.

<shift> is <shiftname> <register> or <shiftname> #expression, or RRX (rotate right one bit with extend).

<shiftname>s are: ASL, LSL, LSR, ASR, ROR.

(ASL is a synonym for LSL, the two assemble to the same code.)

Examples

```
ADDEQ R2,R4,R5      ; if the Z flag is set make R2:=R4+R5

TEQS R4,#3          ; test R4 for equality with 3
                   ; (the S is in fact redundant as the
                   ; assembler inserts it automatically)

SUB R4,R5,R7,LSR R2 ; logical right shift R7 by the number in
                   ; the bottom byte of R2, subtract the result
                   ; from R5, and put the answer into R4

MOV PC,R14          ; return from subroutine
MOVS PC,R14         ; return from exception & restore CPSR_<mode>
```

PSR Transfer (MRS, MSR)

The instruction is only executed if the condition is true. The various conditions are defined at the beginning of this chapter.

The MRS and MSR instructions are formed from a subset of the Data Processing operations and are implemented using the TEQ, TST, CMN and CMP instructions without the S flag set. The encoding is shown in Figure 12.

These instructions allow access to the CPSR and SPSR registers. The MRS instruction allows the contents of the CPSR or SPSR_<mode> to be moved to a general register. The MSR instruction allows the contents of a general register to be moved to the CPSR or SPSR_<mode> register.

The MSR instruction also allows an immediate value or register contents to be transferred to the condition code flags (N,Z,C and V) of CPSR or SPSR_<mode> without affecting the control bits. In this case, the top four bits of the specified register contents or 32 bit immediate value are written to the top four bits of the relevant PSR.

Operand Restrictions

In User mode, the control bits of the CPSR are protected from change, so only the condition code flags of the CPSR can be changed. In other (privileged) modes the entire CPSR can be changed.

The SPSR register which is accessed depends on the mode at the time of execution. For example, only SPSR_fiq is accessible when the processor is in FIQ mode.

R15 shall not be specified as the source or destination register.

A further restriction is that no attempt shall be made to access an SPSR in User mode, since no such register exists.

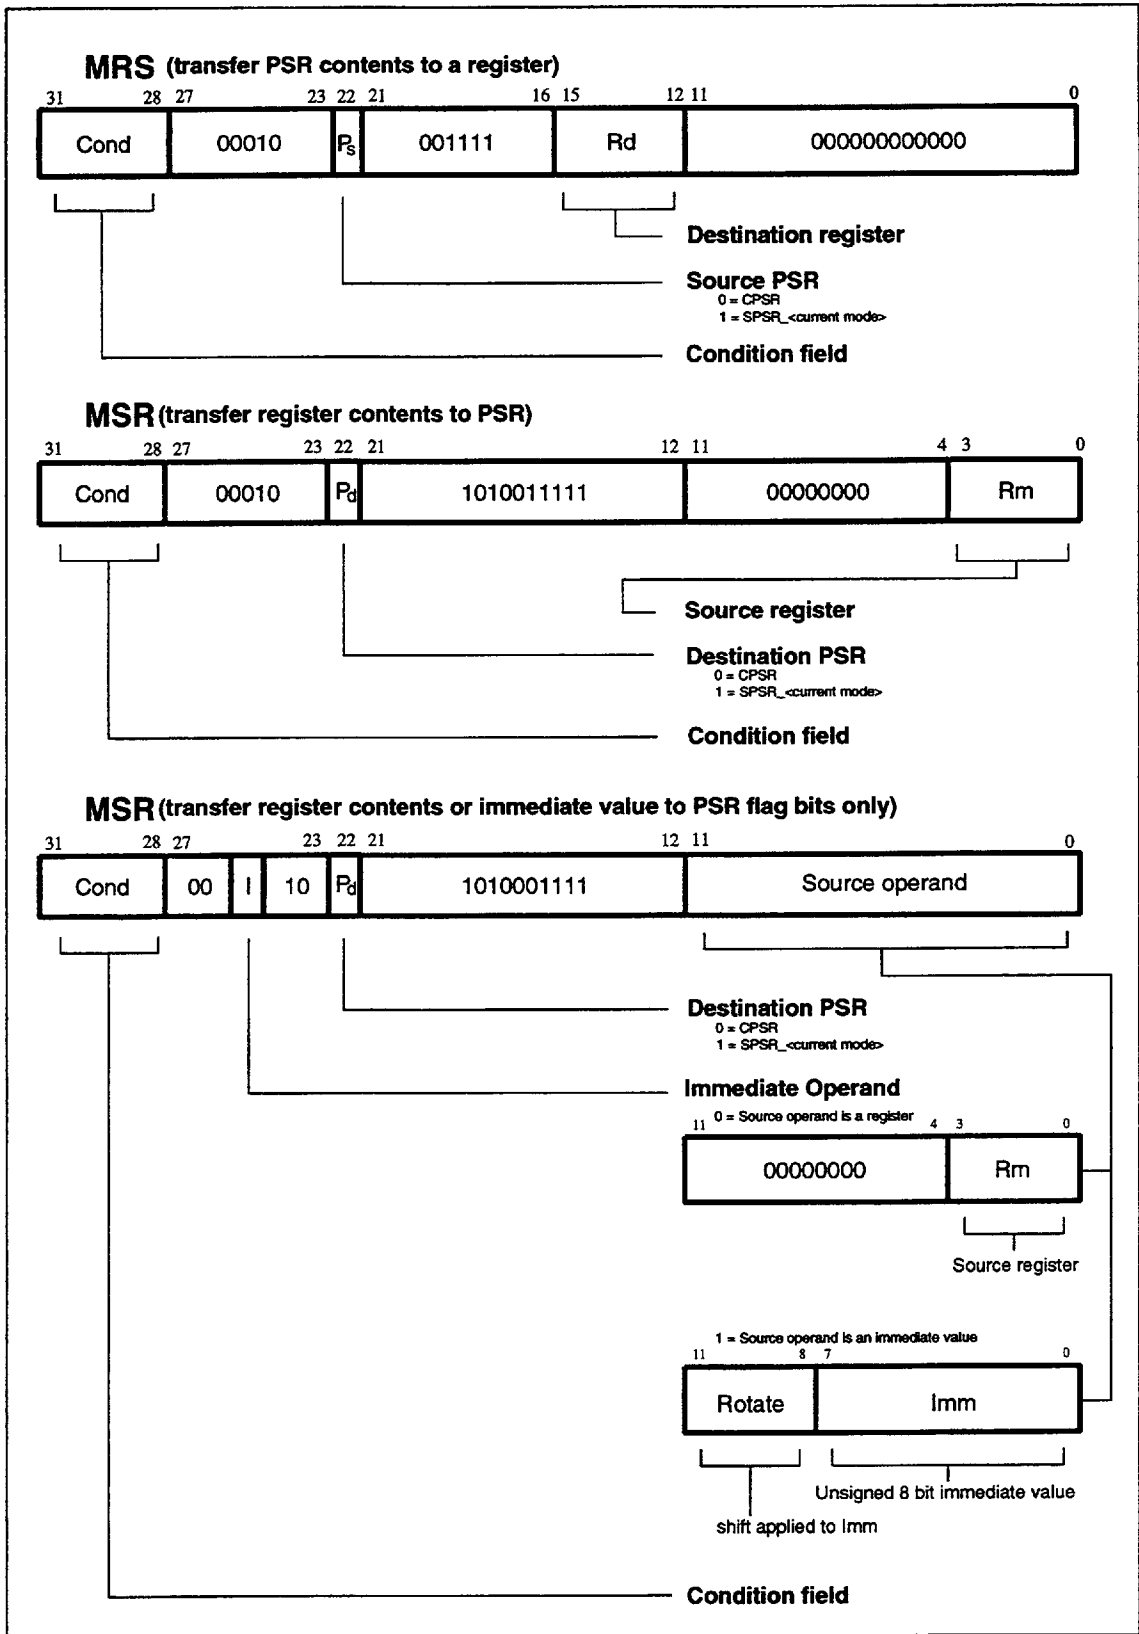


Figure 12. PSR Transfer Instructions

Reserved Bits

Only eleven bits of the PSR are defined in the LH74610 (N,Z,C,V,I,F & M[4:0]); the remaining bits (= PSR[27:8,5]) are reserved for use in future versions of the processor. To ensure the maximum compatibility between the LH74610 programs and future processors, the following rules should be observed:

- (1) The reserved bits shall be preserved when changing the value in a PSR.
- (2) Programs shall not rely on specific values from the reserved bits when checking the PSR status, since they may read as one or zero in future processors.

A read-modify-write strategy should therefore be used when altering the control bits of any PSR register; this involves transferring the appropriate PSR register to a general register using the MRS instruction, changing only the relevant bits and then transferring the modified value back to the PSR register using the MSR instruction.

e.g. The following sequence performs a mode change:

```
MRS Rtmp,CPSR           ; take a copy of the CPSR
BIC Rtmp,Rtmp,#&1F     ; clear the mode bits
ORR Rtmp,Rtmp,#new_mode ; select new mode
MSR CPSR,Rtmp          ; write back the modified CPSR
```

When the aim is simply to change the condition code flags in a PSR, an immediate value can be written directly to the flag bits without disturbing the control bits. e.g. The following instruction sets the N,Z,C & V flags:

```
MSR CPSR_flg,#&F0000000 ;set all the flags regardless of
                        ;their previous state (does not
                        ;affect any control bits)
```

No attempt shall be made to write an 8 bit immediate value into the whole PSR since such an operation cannot preserve the reserved bits.

Assembler Syntax

- (1) MRS - transfer PSR contents to a register

```
MRS{cond} Rd,<psr>
```

- (2) MSR - transfer register contents to PSR

```
MSR{cond} <psr>,Rm
```

- (3) MSR - transfer register contents to PSR flag bits only

```
MSR{cond} <psrf>,Rm
```

The most significant four bits of the register contents are written to the N,Z,C & V flags respectively.

(4) MSR - transfer immediate value to PSR flag bits only

MSR{cond} <psrf>,<#expression>

The expression should symbolize a 32 bit value of which the most significant four bits are written to the N,Z,C & V flags respectively.

{cond} - two-character condition mnemonic, see Figure 3.

Rd and Rm are expressions evaluating to a register number other than R15

<psrf> is CPSR, CPSR_all, SPSR or SPSR_all. (CPSR and CPSR_all are synonyms as are SPSR and SPSR_all)

<psrf> is CPSR_flg or SPSR_flg

Where <#expression> is used, the assembler will attempt to generate a shifted immediate 8-bit field to match the expression. If this is impossible, it will give an error.

Examples

In User mode the instructions behave as follows:

```
MSR  CPSR_all,Rm          ; CPSR[31:28] <- Rm[31:28]
MSR  CPSR_flg,Rm         ; CPSR[31:28] <- Rm[31:28]

MSR  CPSR_flg,#&A0000000 ; CPSR[31:28] <- &A
                               ; (i.e. set N,C; clear Z,V)

MRS  Rd,CPSR             ; Rd[31:0] <- CPSR[31:0]
```

In privileged modes the instructions behave as follows:

```
MSR  CPSR_all,Rm          ; CPSR[31:0] <- Rm[31:0]
MSR  CPSR_flg,Rm         ; CPSR[31:28] <- Rm[31:28]

MSR  CPSR_flg,#&50000000 ; CPSR[31:28] <- &5
                               ; (i.e. set Z,V; clear N,C)

MRS  Rd,CPSR             ; Rd[31:0] <- CPSR[31:0]

MSR  SPSR_all,Rm          ; SPSR_<mode>[31:0] <- Rm[31:0]
MSR  SPSR_flg,Rm         ; SPSR_<mode>[31:28] <- Rm[31:28]

MSR  SPSR_flg,#&C0000000 ; SPSR_<mode>[31:28] <- &C
                               ; (i.e. set N,Z; clear C,V)

MRS  Rd,SPSR            ; Rd[31:0] <- SPSR_<mode>[31:0]
```

Multiply and Multiply-Accumulate (MUL, MLA)

The instruction is only executed if the condition is true. The various conditions are defined at the beginning of this chapter. The instruction encoding is shown in Figure 13.

The multiply and multiply-accumulate instructions use a 2 bit Booth's algorithm to perform integer multiplication. They give the least significant 32 bits of the product of two 32 bit operands, and may be used to synthesize higher precision multiplications.

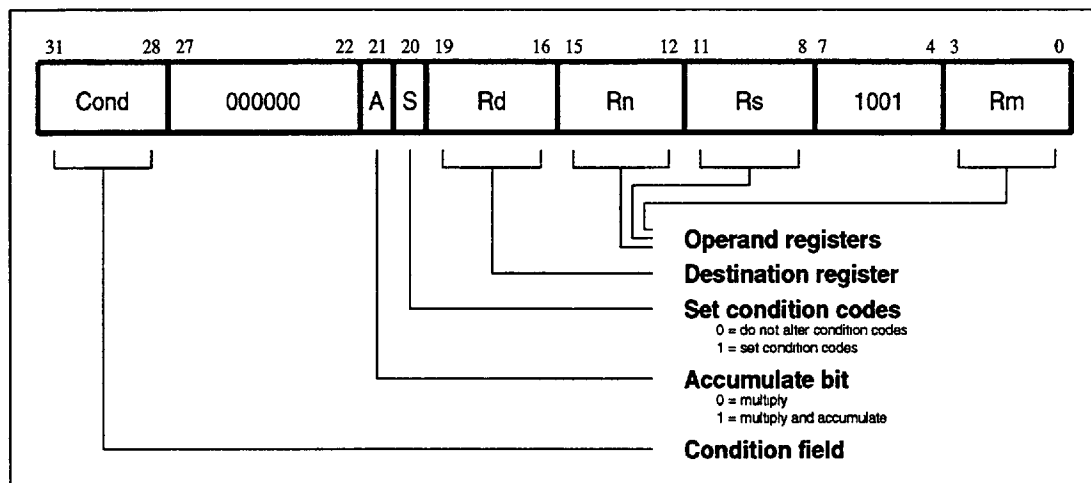


Figure 13. Multiply Instructions

The multiply form of the instruction gives $Rd := Rm * Rs$. Rn is ignored, and should be set to zero for compatibility with possible future upgrades to the instruction set.

The multiply-accumulate form gives $Rd := Rm * Rs + Rn$, which can save an explicit ADD instruction in some circumstances.

Both forms of the instruction work on operands which may be considered as signed (2's complement) or unsigned integers.

Operand Restrictions

Due to the way the Booth's algorithm has been implemented, certain combinations of operand registers should be avoided. (The assembler will issue a warning if these restrictions are overlooked.)

The destination register (Rd) should not be the same as the Rm operand register, as Rd is used to hold intermediate values and Rm is used repeatedly during the multiply. A MUL will give a zero result if $Rm = Rd$, and a MLA will give a meaningless result. R15 shall not be used as an operand or as the destination register.

All other register combinations will give correct results, and Rd, Rn and Rs may use the same register when required.

CPSR Flags

Setting the CPSR flags is optional, and is controlled by the S bit in the instruction. The N and Z flags are set correctly on the result (N is equal to bit 31 of the result, Z is set if and only if the result is zero), the V flag is unaffected by the instruction (as for logical data processing instructions), and the C flag is set to a meaningless value.

Assembler Syntax

MUL{cond}{S} Rd,Rm,Rs

MLA{cond}{S} Rd,Rm,Rs,Rn

{cond} - two-character condition mnemonic, see Figure 3.

{S} - set condition codes if S present

Rd, Rm, Rs and Rn are expressions evaluating to a register number other than R15.

Examples

```
MUL R1,R2,R3; R1:=R2*R3
MLAEQS R1,R2,R3,R4 ; conditionally R1:=R2*R3+R4,
; setting condition codes
```

The multiply instruction may be used to synthesize higher precision multiplications, for instance to multiply two 32 bit integers and generate a 64 bit result:

```
mul64
MOV a1,A,LSR #16 ; a1:= top half of A
MOV D,B,LSR #16 ; D := top half of B
BIC A,A,a1,LSL #16 ; A := bottom half of A
BIC B,B,D,LSL #16 ; B := bottom half of B
MUL C,A,B ; low section of result
MUL B,a1,B ; ) middle sections
MUL A,D,A ; ) of result
MUL D,a1,D ; high section of result
ADDSA,B,A ; add middle sections
; (couldn't use MLA as we need C correct)
ADDCSD,D,#&10000 ; carry from above add
ADDSC,C,A,LSL #16 ; C is now bottom 32 bits of product
ADC D,D,A,LSR #16 ; D is top 32 bits
```

A, B are registers containing the 32 bit integers; C, D are registers for the 64 bit result; a1 is a temporary register. A and B are overwritten during the multiply.

Single Data Transfer (LDR, STR)

The instruction is only executed if the condition is true. The various conditions are defined at the beginning of this chapter. The instruction encoding is shown in Figure 14.

The single data transfer instructions are used to load or store single bytes or words of data. The memory address used in the transfer is calculated by adding an offset to or subtracting an offset from a base register.

The result of this calculation may be written back into the base register if 'auto-indexing' is required.

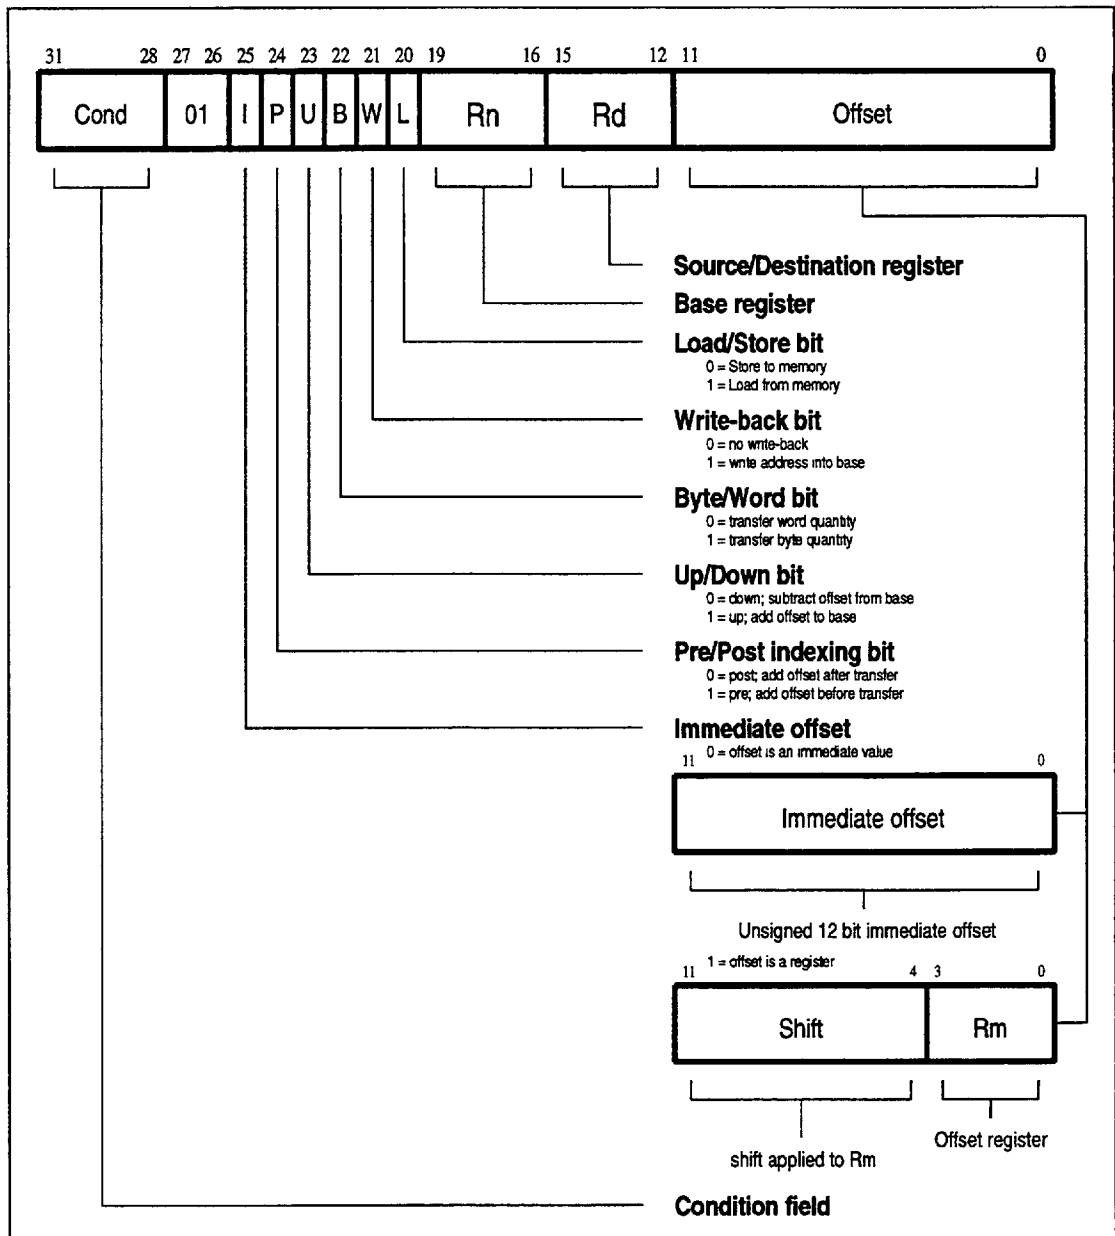


Figure 14. Single Data Transfer Instructions

Offsets and Auto-Indexing

The offset from the base may be either a 12 bit unsigned binary immediate value in the instruction, or a second register (possibly shifted in some way). The offset may be added to ($U=1$) or subtracted from ($U=0$) the base register R_n . The offset modification may be performed either before (pre-indexed, $P=1$) or after (post-indexed, $P=0$) the base is used as the transfer address.

The W bit gives optional auto increment and decrement addressing modes. The modified base value may be written back into the base ($W=1$), or the old base value may be kept ($W=0$). In the case of post-indexed addressing, the write back bit is redundant and is always set to zero, since the old base value can be retained by setting the offset to zero. Therefore post-indexed data transfers always write back the modified base. The only use of the W bit in a post-indexed data transfer is in privi-

leged mode code, where setting the W bit forces the Ntrans signal to go LOW for the transfer, allowing the operating system to generate a user address in a system where the memory management hardware makes suitable use of this hardware.

Shifted Register Offset

The 8 shift control bits are described in the data processing instructions section. However, the register specified shift amounts are not available in this instruction class.

Bytes and Words

This instruction class may be used to transfer a byte (B=1) or a word (B=0) between a LH74610 register and memory.

The action of LDR(B) and STR(B) instructions is influenced by the bigend configuration signal to the processor. The two possible configurations are described below.

Little Endian Configuration

A byte load (LDRB) expects the data on data bus inputs 7 through 0 if the supplied address is on a word boundary, on data bus inputs 15 through 8 if it is a word address plus one byte, and so on. The selected byte is placed in the bottom 8 bits of the destination register, and the remaining bits of the register are filled with zeros.

A byte store (STRB) repeats the bottom 8 bits of the source register four times across data bus outputs 31 through 0. The external memory system should activate the appropriate byte subsystem to store the data.

A word load (LDR) will normally use a word aligned address. However, an address offset from a word boundary will cause the data to be rotated into the register so that the addressed byte occupies bits 0 to 7. This means that half-words accessed at offsets 0 and 2 from the word boundary will be correctly loaded into bits 0 through 15 of the register. Two shift operations are then required to clear or to sign extend the upper 16 bits.

A word store (STR) should generate a word aligned address. The word presented to the data bus is not affected if the address is not word aligned. That is, bit 31 of the register being stored always appears on data bus output 31.

Big Endian Configuration

A byte load (LDRB) expects the data on data bus inputs 31 through 24 if the supplied address is on a word boundary, on data bus inputs 23 through 16 if it is a word address plus one byte, and so on. The selected byte is placed in the bottom 8 bits of the destination register and the remaining bits of the register are filled with zeros.

A byte store (STRB) repeats the bottom 8 bits of the source register four times across data bus outputs 31 through 0. The external memory system should activate the appropriate byte subsystem to store the data.

A word load (LDR) should generate a word aligned address. An address offset of 0 or 2 from a word boundary will cause the data to be rotated into the register so that the addressed byte occupies bits 31 through 24. This means that half-words accessed at these offsets will be correctly loaded into bits 16 through 31 of the register. A shift operation is then required to move (and optionally sign extend) the data into the bottom 16 bits. An address offset of 1 or 3 from a word boundary will cause the data to be rotated into the register so that the addressed byte occupies bits 15 through 8.

A word store (STR) should generate a word aligned address. The word presented to the data bus is not affected if the address is not word aligned. That is, bit 31 of the register being stored always appears on data bus output 31.

Use of R15

Write-back shall not be specified if R15 is specified as the base register (Rn). When using R15 as the base register one must remember that it contains an address 8 bytes on from the address of the current instruction.

R15 shall not be specified as the register offset (Rm).

When R15 is the source register (Rd) of a register store (STR) instruction, the stored value will be address of the instruction plus 12.

Restriction on the Use of Base Register

When configured for late aborts, the following code is very difficult to unwind as the base register, Rn, gets updated before the abort handler is entered. In certain circumstances it may be impossible to calculate the initial value.

```
<LDR | STR> Rd, [Rn],{+/-}Rn{,<shift>}
```

A post-indexed LDR | STR where Rm=Rn shall not be used.

Data Aborts

A transfer to or from a legal address may cause problems for a memory management system. For instance, in a system which uses virtual memory the required data may be absent from main memory. The memory manager can signal a problem by taking the processor abort input HIGH, whereupon the data transfer instruction will be prevented from changing the processor state and the Data Abort trap will be taken. It is up to the system software to resolve the cause of the problem, then the instruction can be restarted and the original program continued.

The LH74610 supports two types of Data Abort processing depending on the lateabt configuration input. When configured for Early Aborts, any base register write-back which would have occurred is prevented from happening in the event of an abort. When configured for Late Aborts, this write-back is allowed to take place and the Abort handler must correct this before allowing the instruction to be re-executed.

Assembler Syntax

```
<LDR | STR>{cond}{B}{T} Rd,<Address>
```

LDR - load from memory into a register

STR - store from a register into memory

{cond} - two-character condition mnemonic, see Figure 3.

{B} - if B is present then byte transfer, otherwise word transfer

{T} - if T is present the W bit will be set in a post-indexed instruction, causing the Ntrans signal to go LOW for the transfer cycle. T is not allowed when a pre-indexed addressing mode is specified or implied.

Rd is an expression evaluating to a valid register number.

<Address> can be:

- (i) An expression which generates an address:

<expression>

The assembler will attempt to generate an instruction using the PC as a base and a corrected immediate offset to address the location given by evaluating the expression. This will be a PC relative, pre-indexed address. If the address is out of range, an error will be generated.

- (ii) A pre-indexed addressing specification:

[Rn] offset of zero

[Rn,<#expression>](!) offset of <expression> bytes

[Rn,{+/-}Rm{,<shift>}](!) offset of +/- contents of index register, shifted by <shift>

- (iii) A post-indexed addressing specification:

[Rn],<#expression> offset of <expression> bytes

[Rn]{+/-}Rm{,<shift>} offset of +/- contents of index register, shifted as by <shift>.

Rn and Rm are expressions evaluating to a valid register number. NOTE if Rn is R15 then the assembler will subtract 8 from the offset value to allow for the LH74610 pipelining. In this case base write-back shall not be specified.

<shift> is a general shift operation (see section on data processing instructions) but note that the shift amount may not be specified by a register.

(!) write back the base register (set the W bit) if ! is present.

Examples

```
STR R1, [BASE, INDEX]!           ; store R1 at BASE+INDEX (both of which are
                                ; registers) and write back address to BASE

STR R1, [BASE], INDEX           ; store R1 at BASE and write back
                                ; BASE+INDEX to BASE

LDR R1, [BASE, #16]             ; load R1 from contents of BASE+16
                                ; Don't write back

LDR R1, [BASE, INDEX, LSL#2]    ; load R1 from contents of BASE+INDEX*4

LDREQB R1, [BASE, #5]          ; conditionally load byte at BASE+5 into
                                ; R1 bits 0 to 7, filling bits 8 to 31
                                ; with zeros

STR R1, PLACE                   ; generate PC relative offset to address
                                ; PLACE
                                .
                                .
PLACE
```

Block Data Transfer (LDM, STM)

The instruction is only executed if the condition is true. The various conditions are defined at the beginning of this chapter. The instruction encoding is shown in Figure 15.

Block data transfer instructions are used to load (LDM) or store (STM) any subset of the currently visible registers. They support all possible stacking modes, maintaining full or empty stacks which can grow up or down memory, and are very efficient instructions for saving or restoring context, or for moving large blocks of data around main memory.

The Register List

The instruction can cause the transfer of any registers in the current bank (and non-user mode programs can also transfer to and from the user bank, see below). The register list is a 16 bit field in the instruction, with each bit corresponding to a register. A 1 in bit 0 of the register field will cause R0 to be transferred, a 0 will cause it not to be transferred; similarly bit 1 controls the transfer of R1, and so on.

Any subset of the registers, or all the registers, may be specified. The only restriction is that the register list should not be empty.

Whenever R15 is stored to memory the stored value is the address of the STM instruction plus 12.

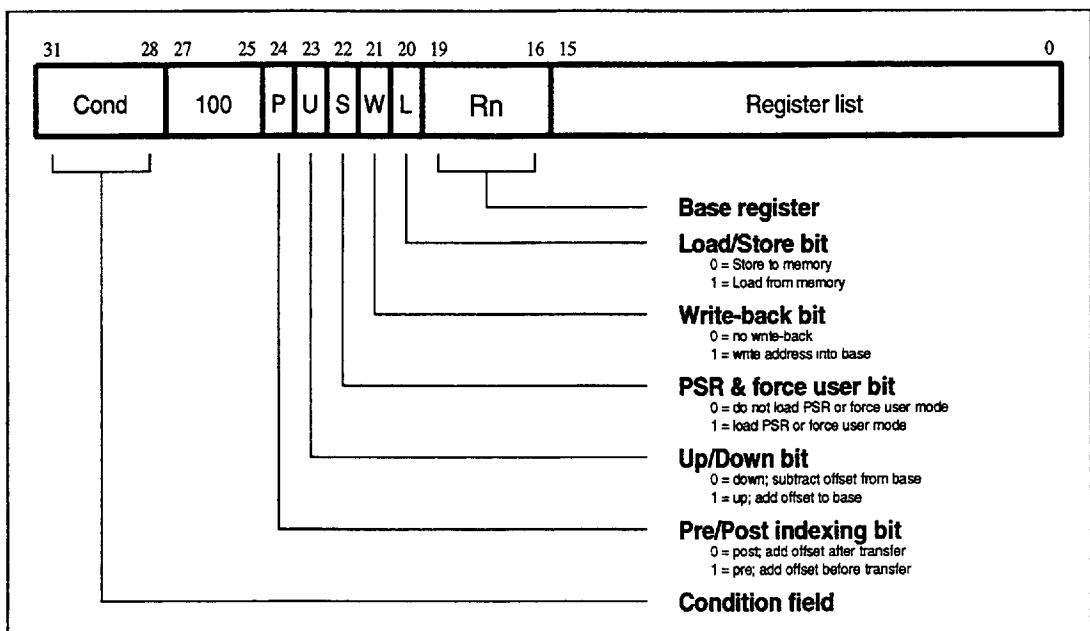


Figure 15. Block Data Transfer Instructions

Addressing Modes

The transfer addresses are determined by the contents of the base register (Rn), the pre/post bit (P) and the up/down bit (U). The registers are transferred in the order lowest to highest, so R15 (if in the list) will always be transferred last. The lowest register also gets transferred to/from the lowest memory address. By way of illustration, consider the transfer of R1, R5 and R7 in the case where Rn=1000H and write back of the modified base is required (W=1). Figures 16-19 show the sequence of register transfers, the addresses used, and the value of Rn after the instruction has completed.

In all cases, had write back of the modified base not been required (W=0), Rn would have retained its initial value of 1000H unless it was also in the transfer list of a load multiple register instruction, when it would have been overwritten with the loaded value.

Address Alignment

The address should normally be a word aligned quantity and non-word aligned addresses do not affect the instruction. However, the bottom 2 bits of the address will appear on a[1:0] and might be interpreted by the memory system.

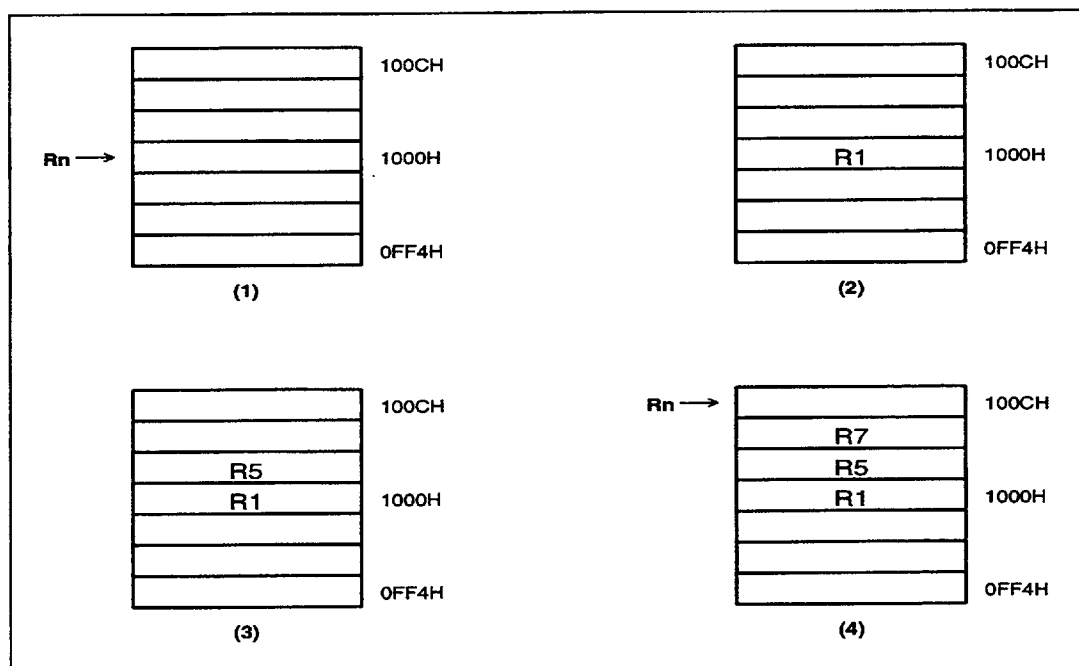


Figure 16. Post-Increment Addressing

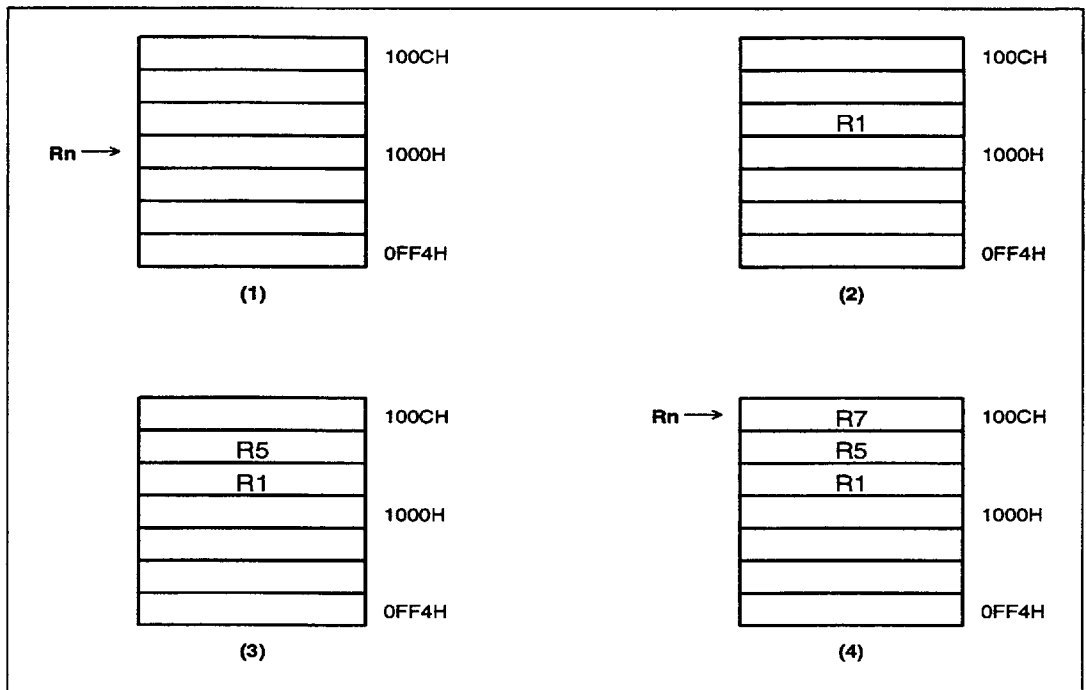


Figure 17. Pre-Increment Addressing

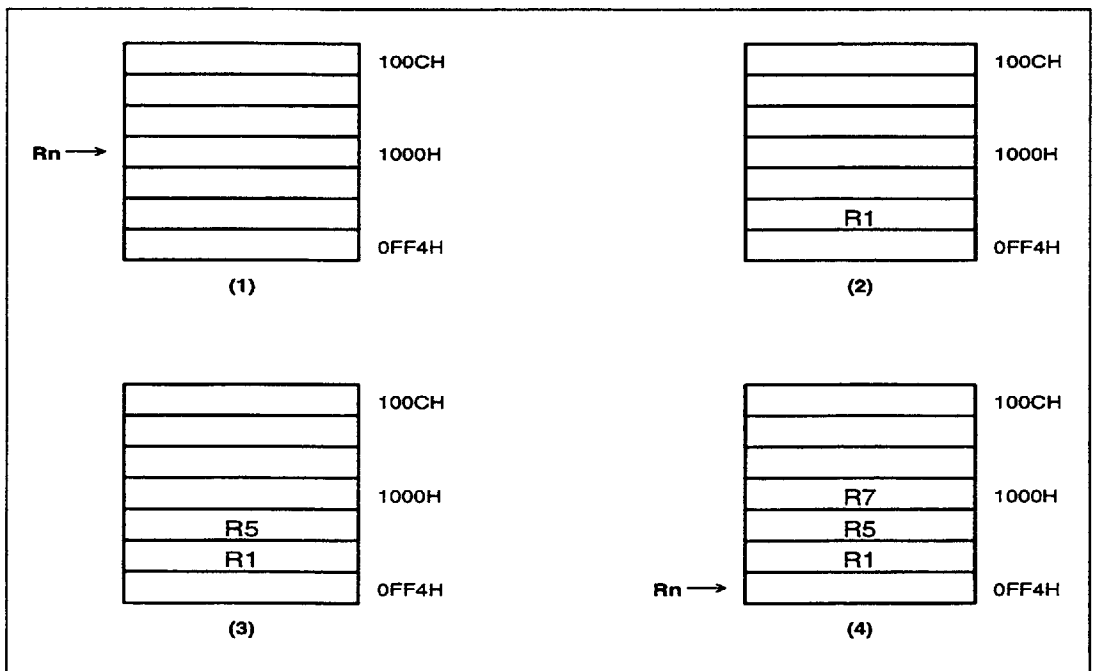


Figure 18. Post-Decrement Addressing

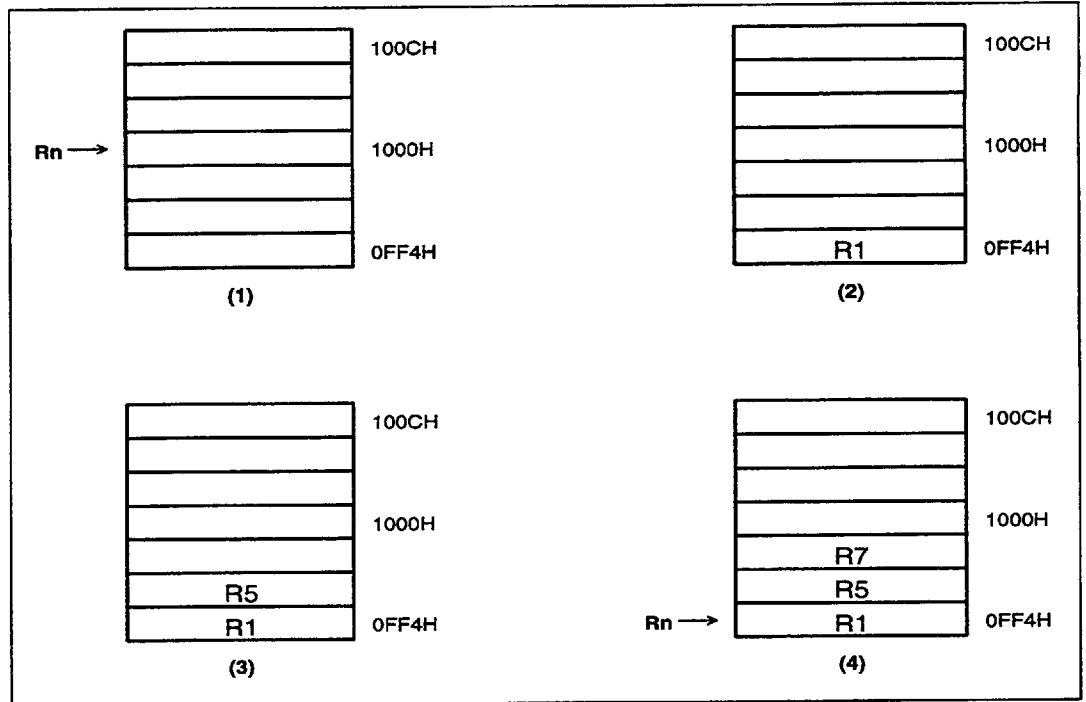


Figure 19. Pre-Decrement Addressing

Use of the S Bit

When the S bit is set in a LDM/STM instruction its meaning depends on whether or not R15 is in the transfer list and on the type of instruction. The S bit should only be set if the instruction is to execute in a privileged mode.

LDM With R15 in Transfer List and S Bit Set (Mode Changes)

If the instruction is a LDM then SPSR_<mode> is transferred to CPSR at the same time as R15 is loaded.

STM With R15 in Transfer List and S Bit Set (User Bank Transfer)

The registers transferred are taken from the User bank rather than the bank corresponding to the current mode. This is useful for saving the user state on process switches. Base write-back shall not be used when this mechanism is employed.

R15 Not in List and S Bit Set (User Bank Transfer)

For both LDM and STM instructions, the User bank registers are transferred rather than the register bank corresponding to the current mode. This is useful for saving the user state on process switches. Base write-back shall not be used when this mechanism is employed.

When the instruction is LDM, care must be taken not to read from a banked register during the following cycle (inserting a NOP after the LDM will ensure safety).

Use of R15 as the Base

R15 shall not be used as the base register in any LDM or STM instruction.

Inclusion of the Base in the Register List

When write-back is specified, the base is written back at the end of the second cycle of the instruction. During a STM, the first register is written out at the start of the second cycle. A STM which includes storing the base, with the base as the first register to be stored, will therefore store the unchanged value, whereas with the base second or later in the transfer order, will store the modified value. A LDM will always overwrite the updated base if the base is in the list.

Data Aborts

Some legal addresses may be unacceptable to a memory management system, and the memory manager can indicate a problem with an address by taking the abort signal HIGH. This can happen on any transfer during a multiple register load or store, and must be recoverable if the LH74610 is to be used in a virtual memory system.

The state of the lateabt configuration input does not affect the behavior of LDM and STM instructions in the event of an Abort exception.

Aborts During STM Instructions

If the abort occurs during a store multiple instruction, the LH74610 takes little action until the instruction completes, whereupon it enters the data abort trap. The memory manager is responsible for preventing erroneous writes to the memory. The only change to the internal state of the processor will be the modification of the base register if write-back was specified, and this must be reversed by software (and the cause of the abort resolved) before the instruction may be retried.

Aborts During LDM Instructions

When the LH74610 detects a data abort during a load multiple instruction, it modifies the operation of the instruction to ensure that recovery is possible.

- (i) Overwriting of registers stops when the abort happens. The aborting load will not take place but earlier ones may have overwritten registers. The PC is always the last register to be written and so will always be preserved.
- (ii) The base register is restored, to its modified value if write-back was requested. This ensures recoverability in the case where the base register is also in the transfer list, and may have been overwritten before the abort occurred.

The data abort trap is taken when the load multiple has completed, and the system software must undo any base modification (and resolve the cause of the abort) before restarting the instruction.

Assembler Syntax

<LDM | STM>{cond}<FD | ED | FA | EA | IA | IB | DA | DB> Rn{!},<Rlist>{^}

{cond} - two character condition mnemonic, see Figure 3.

Rn is an expression evaluating to a valid register number

<Rlist> is a list of registers and register ranges enclosed in {} (eg {R0,R2-R7,R10}).

{!} if present requests write-back (W=1), otherwise W=0

{^} if present set S bit to load the CPSR along with the PC, or force transfer of user bank when in privileged mode

Addressing Mode Names

There are different assembler mnemonics for each of the addressing modes, depending on whether the instruction is being used to support stacks or for other purposes. The equivalences between the names and the values of the bits in the instruction are shown in the following Table 5.

Table 5. Addressing Mode Names

NAME	STACK	OTHER	L BIT	P BIT	U BIT
pre-increment load	LDMED	LDMIB	1	1	1
post-increment load	LDMFD	LDMIA	1	0	1
pre-decrement load	LDMEA	LDMDB	1	1	0
post-decrement load	LDMFA	LDMDA	1	0	0
pre-increment store	STMFA	STMIB	0	1	1
post-increment store	STMEA	STMIA	0	0	1
pre-decrement store	STMFD	STMDB	0	1	0
post-decrement store	STMED	STMDA	0	0	0

FD, ED, FA, EA define pre/post indexing and the up/down bit by reference to the form of stack required. The F and E refer to a "full" or "empty" stack, i.e. whether a pre-index has to be done (full) before storing to the stack. The A and D refer to whether the stack is ascending or descending. If ascending, a STM will go up and LDM down, if descending, vice-versa.

IA, IB, DA, DB allow control when LDM/STM are not being used for stacks and simply mean Increment After, Increment Before, Decrement After, Decrement Before.

Examples

```
LDMFD SP!, {R0,R1,R2} ; unstack 3 registers

STMIA BASE, {R0-R15} ; save all registers

LDMFD SP!, {R15} ; R15 <- (SP), CPSR unchanged
LDMFD SP!, {R15} ; R15 <- (SP), CPSR <- SPSR_mode (allowed only
; in privileged modes)
STMFD R13, {R0-R14} ; Save user mode regs on stack (allowed only
; in privileged modes)
```

These instructions may be used to save state on subroutine entry, and restore it efficiently on return to the calling routine:

```
STMED SP!, {R0-R3,R14} ; save R0 to R3 to use as workspace
; and R14 for returning

BL somewhere ; this nested call will overwrite R14

LDMED SP!, {R0-R3,R15} ; restore workspace and return
```

Single Data Swap (SWP)

The instruction is only executed if the condition is true. The various conditions are defined at the beginning of this chapter. The instruction encoding is shown in Figure 20.

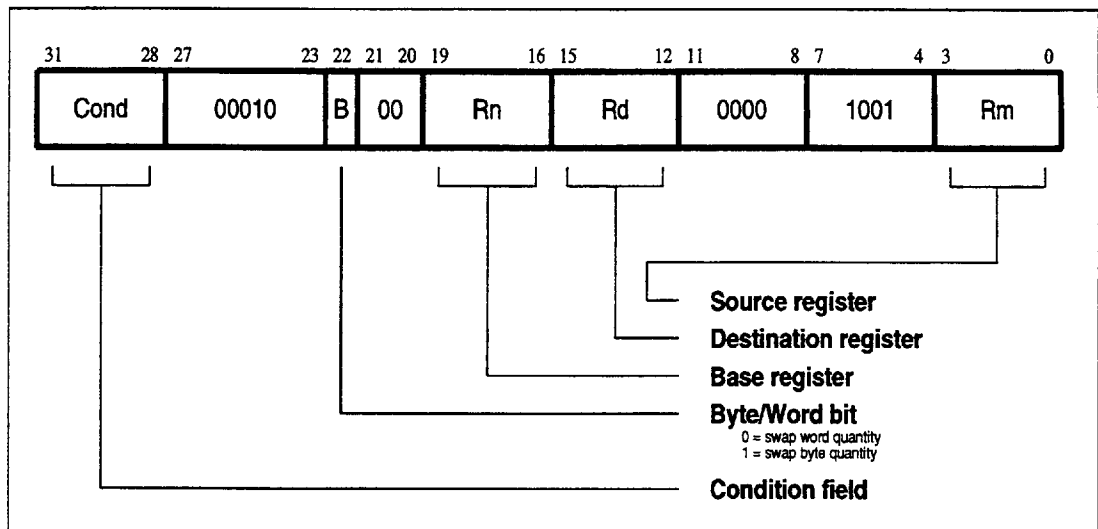


Figure 20. Swap Instruction

The data swap instruction is used to swap a byte or word quantity between a register and external memory. This instruction is implemented as a memory read followed by a memory write which are "locked" together (the processor cannot be interrupted until both operations have completed, and the memory manager is warned to treat them as inseparable). This class of instruction is particularly useful for implementing software semaphores.

The swap address is determined by the contents of the base register (Rn). The processor first reads the contents of the swap address. Then it writes the contents of the source register (Rm) to the swap address, and stores the old memory contents in the destination register (Rd). The same register may be specified as both the source and destination.

The **lock** output goes HIGH for the duration of the read and write operations to signal to the external memory manager that they are locked together, and should be allowed to complete without interruption. This is important in multi-processor systems where the swap instruction is the only indivisible instruction which may be used to implement semaphores; control of the memory must not be removed from a processor while it is performing a locked operation.

Bytes and Words

This instruction class may be used to swap a byte (B=1) or a word (B=0) between a LH74610 register and memory. The SWP instruction is implemented as a LDR followed by a STR and the action of these is as described in the section on single data transfers. In particular, the description of Big and Little Endian configuration applies to the SWP instruction.

Use of R15

R15 shall not be used as an operand (Rd, Rn or Rs) in a SWP instruction.

Return From the Supervisor

The PC is saved in R14_svc upon entering the software interrupt trap, with the PC adjusted to point to the word after the SWI instruction. `MOVS PC,R14_svc` will return to the calling program and restore the CPSR.

Note that the link mechanism is not re-entrant, so if the supervisor code wishes to use software interrupts within itself it must first save a copy of the return address and SPSR.

Comment Field

The bottom 24 bits of the instruction are ignored by the processor, and may be used to communicate information to the supervisor code. For instance, the supervisor may look at this field and use it to index into an array of entry points for routines which perform the various supervisor functions.

Assembler Syntax

`SWI{cond} <expression>`

{cond} - two character condition mnemonic, see Figure 3

<expression> is evaluated and placed in the comment field (which is ignored by the LH74610).

Examples

```
SWI  ReadC           ; get next character from read stream
SWI  WriteI+"k"     ; output a "k" to the write stream
SWINE 0             ; conditionally call supervisor
                       ; with 0 in comment field
```

The above examples assume that suitable supervisor code exists, for instance:

```
&08 B Supervisor   ; SWI entry point

EntryTable         ; addresses of supervisor routines
  & ZeroRtn
  & ReadCRtn
  & WriteIRtn
  ...

Zero      * 0
ReadC    * 256
WriteI    * 512
Supervisor

; SWI has routine required in bits 8-23 and data (if any) in bits 0-7.
; Assumes R13_svc points to a suitable stack

STM R13, {R0-R2,R14} ; save work registers and return address
LDR R0, [R14, #-4]   ; get SWI instruction
BIC R0, R0, #&FF000000 ; clear top 8 bits
MOV R1, R0, LSR#8    ; get routine offset
ADR R2, EntryTable  ; get start address of entry table
LDR R15, [R2, R1, LSL#2] ; branch to appropriate routine

WriteIRtn         ; enter with character in R0 bits 0-7
.....
LDM R13, {R0-R2,R15}^ ; restore workspace and return
```

Coprocessor Instructions on the LH74610

NOTE: The LH74610 does not have a coprocessor interface. The LH74610 only directly supports Coprocessor Register Transfer instructions (CPRT) to coprocessor #15 which are used for programming internal registers.

All other coprocessor instructions will cause the undefined instruction trap to be taken.

Coprocessor Data Operations (CDP)

The instruction is only executed if the condition is true. The various conditions are defined at the beginning of this chapter. The instruction encoding is shown in Figure 22.




Figure 22. Coprocessor Data Operation Instruction

This class of instruction is used to tell a coprocessor to perform some internal operation. No result is communicated back to the LH74610, and it will not wait for the operation to complete. The coprocessor could contain a queue of such instructions awaiting execution, and their execution can overlap other LH74610 activity allowing the coprocessor and the LH74610 to perform independent tasks in parallel.

Use of the CDP instruction on the LH74610 will cause an undefined instruction trap to be taken.

Coprocessor Data Transfers (LDC, STC)

The instruction is only executed if the condition is true. The various conditions are defined at the beginning of this chapter. The instruction encoding is shown in Figure 23.




Figure 23. Coprocessor Data Transfer Instructions

This class of instruction is used to load (LDC) or store (STC) a subset of the coprocessor's registers directly to memory. The LH74610 is responsible for supplying the memory address, and the coprocessor supplies or accepts the data and controls the number of words transferred.

Use of the CDP instruction on the LH74610 will cause an undefined instruction trap to be taken.

Coprocessor Register Transfers (MRC, MCR)

The instruction is only executed if the condition is true. The various conditions are defined at the beginning of this chapter. The instruction encoding is shown in Figure 24.




Figure 24. Coprocessor Register Transfer Instructions

The instruction is implemented specifically on the LH74610 for coprocessor #15. For other coprocessors the instruction results in an undefined instruction trap.

This class of instruction is used to communicate information directly between the processor core and the control registers.

An example of a MRC instruction would be a read of the LH74610 i-d register. An example of a MCR instruction is the writing of the LH74610 control register.

The Coprocessor Fields

The CP# field [11:8] should be set to #15 (1111) as shown in fig 26.

Cond denotes Condition field

CRn denotes LH74610 register

Rd denotes ARM core register

Assembler Syntax

<MCR | MRC>{cond} CP#15,0,Rd,CRn,C0

MRC - move to LH74610 register (L=1)

MCR - move from LH74610 register (L=0)

{cond} - two character condition mnemonic, see Figure 3

CP#15 - the unique number of the required coprocessor

Rd is an expression evaluating to a valid LH74610 register number

CRn is an expression evaluating to a LH74610 register number

Examples

MRC CP#15,0,R0,MMUID,C0 ; get the LH74610 i-d

MCREQ CP#15,0,R14,MMUCR,C0 ; conditionally configure the control register

Undefined Instruction

The instruction is only executed if the condition is true. The various conditions are defined at the beginning of this chapter. The instruction format is shown in Figure 25.




Figure 25. Undefined Instruction

If the condition is true, the undefined instruction trap will be taken.

Note that the undefined instruction mechanism involves offering this instruction to any coprocessors which may be present, and all coprocessors must refuse to accept it by driving cpa and cpb HIGH.

Assembler Syntax

At present the assembler has no mnemonics for generating this instruction. If it is adopted in the future for some specified use, suitable mnemonics will be added to the assembler. Until such time, this instruction shall not be used.

Instruction Set Summary

A summary of the LH74610 instruction set is shown in Figure 26.

NOTE: Some instruction codes are not defined but do not cause the Undefined instruction trap to be taken, for instance a Multiply instruction with bit 5 or bit 6 changed to a 1. These instructions shall not be used, as their action may change in future ARM implementations.

(3) multiplication by 4, 5 or 6 (run time)

```

MOV   Rc,Ra,LSL#2    ;multiply by 4
CMP   Rb,#5          ;test value
ADDCS Rc,Rc,Ra       ;complete multiply by 5
ADDHI Rc,Rc,Ra       ;complete multiply by 6

```

(4) combining discrete and range tests

```

TEQ   Rc,#127        ;discrete test
CMPNE Rc,#" -1       ;range test
MOVLS Rc,#"."        ;IF Rc<=" " OR Rc=ASCII(127)
                        ;THEN Rc:="."

```

(5) division and remainder

```

; enter with numbers in Ra and Rb
;
MOV   Rcnt,#1        ;bit to control the division
Div1 CMP   Rb,#&80000000 ;move Rb until greater than Ra
      CMPCC Rb,Ra
      MOVCC Rb,Rb,ASL#1
      MOVCC Rcnt,Rcnt,ASL#1
      BCC   Div1
      MOV   Rc,#0
Div2 CMP   Ra,Rb      ;test for possible subtraction
      SUBCS Ra,Ra,Rb  ;subtract if ok
      ADDCS Rc,Rc,Rcnt ;put relevant bit into result
      MOVS  Rcnt,Rcnt,LSR#1 ;shift control bit
      MOVNE Rb,Rb,LSR#1 ;halve unless finished
      BNE   Div2
;
; divide result in Rc
; remainder in Ra

```

Pseudo Random Binary Sequence Generator

It is often necessary to generate (pseudo-) random numbers and the most efficient algorithms are based on shift generators with exclusive or feedback rather like a cyclic redundancy check generator. Unfortunately the sequence of a 32 bit generator needs more than one feedback tap to be maximal length (i.e. $2^{32}-1$ cycles before repetition), so this example uses a 33 bit register with taps at bits 33 and 20. The basic algorithm is newbit:=bit33 eor bit20, shift left the 33 bit number and put in newbit at the bottom; this operation is performed for all the newbits needed (i.e. 32 bits). The entire operation can be done in 55 cycles:

```

; enter with seed in Ra (32 bits), Rb (1 bit in Rb lsb), uses Rc
;
TST   Rb,Rb,LSR#1    ;top bit into carry
MOVS  Rc,Ra,RRX      ;33 bit rotate right
ADC   Rb,Rb,Rb       ;carry into lsb of Rb
EOR   Rc,Rc,Ra,LSL#12 ;(involved!)
EOR   Ra,Rc,Rc,LSR#20 ;(similarly involved!)
;
; new seed in Ra, Rb as before

```

Multiplication by Constant Using the Barrel Shifter

- (1) Multiplication by 2^n (1,2,4,8,16,32..)

```
MOV Ra, Rb, LSL #n
```

- (2) Multiplication by 2^{n+1} (3,5,9,17..)

```
ADD Ra,Ra,Ra,LSL #n
```

- (3) Multiplication by 2^{n-1} (3,7,15..) .display

```
RSB Ra,Ra,Ra,LSL #n
```

- (4) Multiplication by 6 .display

```
ADD Ra,Ra,Ra,LSL #1 ;multiply by 3
MOV Ra,Ra,LSL#1 ;and then by 2
```

- (5) Multiply by 10 and add in extra number

```
ADD Ra,Ra,Ra,LSL#2 ;multiply by 5
ADD Ra,Rc,Ra,LSL#1 ;multiply by 2 and add in next digit
```

- (6) General recursive method for $Rb := Ra * C$, C a constant:

- (a) If C even, say $C = 2^n * D$, D odd:

```
D=1: MOV Rb,Ra,LSL #n
D<>1: {Rb := Ra*D}
      MOV Rb,Rb,LSL #n
```

- (b) If $C \text{ MOD } 4 = 1$, say $C = 2^n * D + 1$, D odd, $n > 1$:

```
D=1: ADD Rb,Ra,Ra,LSL #n
D<>1: {Rb := Ra*D}
      ADD Rb,Ra,Rb,LSL #n
```

- (c) If $C \text{ MOD } 4 = 3$, say $C = 2^n * D - 1$, D odd, $n > 1$:

```
D=1: RSB Rb,Ra,Ra,LSL #n
D<>1: {Rb := Ra*D}
      RSB Rb,Ra,Rb,LSL #n
```

This is not quite optimal, but close. An example of its non-optimality is multiply by 45 which is done by:

```
RSB  Rb,Ra,Ra,LSL#2 ;multiply by 3
RSB  Rb,Ra,Rb,LSL#2 ;multiply by 4*3-1 = 11
ADD  Rb,Ra,Rb,LSL#2 ;multiply by 4*11+1 = 45
```

rather than by:

```
ADD  Rb,Ra,Ra,LSL#3 ;multiply by 9
ADD  Rb,Rb,Rb,LSL#2 ;multiply by 5*9 = 45
```

Loading a Word From an Unknown Alignment

```
; enter with address in Ra (32 bits)
; uses Rb, Rc; result in Rd.
; Note d must be less than c e.g. 0,1
;
BIC  Rb,Ra,#3           ;get word aligned address
LDMIA Rb,{Rd,Rc}       ;get 64 bits containing answer
AND  Rb,Ra,#3           ;correction factor in bytes
MOVS Rb,Rb,LSL#3        ;...now in bits and test if aligned
MOVNE Rd,Rd,LSR Rb     ;produce bottom of result word
; (if not aligned)
RSBNE Rb,Rb,#32        ;get other shift amount
ORRNE Rd,Rd,Rc,LSL Rb  ;combine two halves to get result
```

Loading a Halfword (Little Endian)

```
LDR  Ra, [Rb,#2]       ; Get halfword to bits 15:0
MOV  Ra,Ra,LSL #16     ; move to top
MOV  Ra,Ra,LSR #16     ; and back to bottom
; use ASR to get sign extended version
```

Loading a Halfword (Big Endian)

```
LDR  Ra, [Rb,#2]       ; Get halfword to bits 31:16
MOV  Ra,Ra,LSR #16     ; and back to bottom
; use ASR to get sign extended version
```

Configuring the LH74610

The operation and configuration of the LH74610 is controlled both directly via coprocessor instructions and indirectly via the Memory Management Page tables. The coprocessor instructions manipulate a number of on-chip registers which control the configuration of the Cache, Write Buffer, MMU and a number of other configuration options.

To ensure backwards compatibility of future CPUs, all reserved or unused bits in registers and coprocessor instructions should be programmed to '0'. Invalid registers must not be read/written. The following bits shall be programmed to '0'.

- Register 1 bits[30:9]
- Register 2 bits[13:0]
- Register 5 bits[31:0]
- Register 6 bits[11:0]
- Register 7 bits[31:0]

Internal Coprocessor Instructions

The on-chip registers may be read using MRC instructions and written using MCR instructions. These operations are only allowed in non-user modes and the undefined instruction trap will be taken if accesses are attempted in user mode.

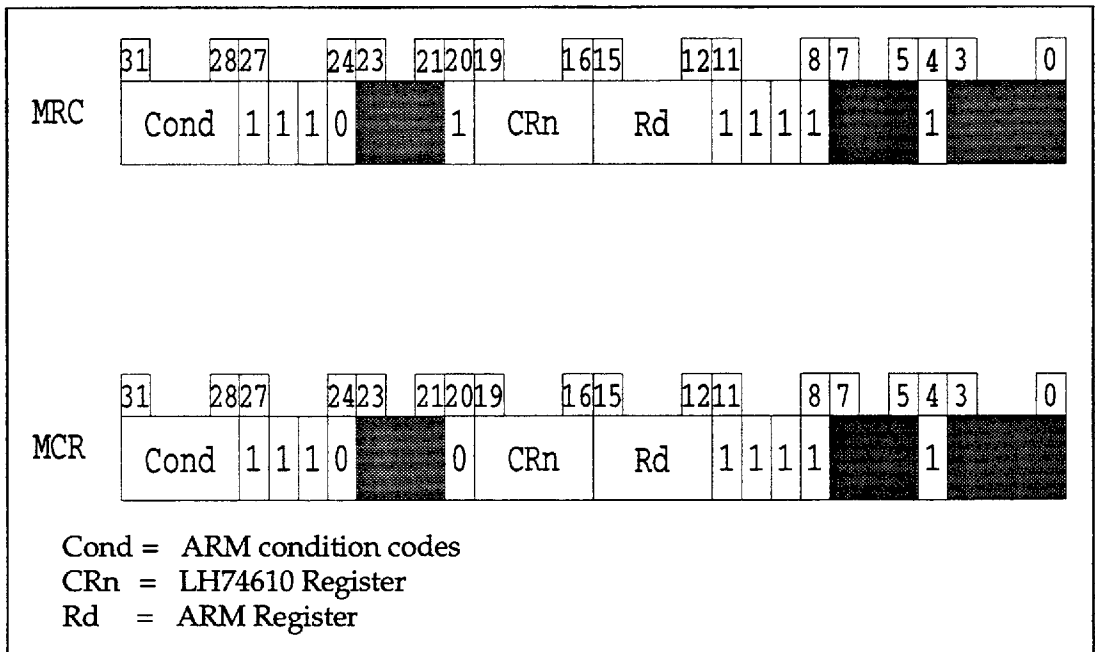


Figure 27. Internal Coprocessor Instructions

LH74610 Registers

The LH74610 contains registers which control the cache and MMU operation. These registers are accessed using CPRT instructions to Coprocessor #15 with the processor in a privileged mode. Only some of registers 0-7 are valid: an access to an invalid register will cause neither the access nor an undefined instruction trap, and therefore should never be carried out; an access to any of the registers 8-15 will cause the undefined instruction trap to be taken.

Register Reads

	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	41								56								06				14											
1	RESERVED																															
2	RESERVED																															
3	RESERVED																															
4	RESERVED																															
5																					0	0	0	0	Domain	Status						
6	Fault Address																															
7	RESERVED																															
8-15	RESERVED																															

Register Writes


	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	RESERVED																															
1	0																					S	B	L	D	P	W	C	A	M		
2	Translation Table Base																															
3	Domain Access Control																															
	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0																
4	RESERVED																															
5	FLUSH TLB (data - don't care)																															
6	Purge TLB (data = purge address)																															
7	FLUSH IDC (data - don't care)																															
8-15	RESERVED																															

S Bit 8 System

This bit controls the LH74610 permission system selection. Refer to Section Descriptor

Register 2 Translation Table Base

Register 2 is a write-only register which holds the base of the currently active Level One page table.



Register 3 Domain Access Control

Register 3 is a write-only register which holds the current access control for domains 0 to 15. See the section on the MMU for more details.

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0																

Domain




Register 4 Reserved

Register 4 is Reserved. Accessing this register has no effect, but should never be attempted.

Register 5

Read: Fault Status

Reading register 5 returns the status of the last data fault. It is not updated for a prefetch fault. See the section on the MMU for more details. Note that only the bottom 12 bits are returned. The upper 20 bits will be the last value on the internal data bus, and therefore will have no meaning. Bits 11:8 are always returned as zero.




Write: Translation Lookaside Buffer Flush

Writing Register 5 flushes the TLB. (The data written is discarded).

Register 6


Read: Fault Address

Reading register 6 returns the virtual address of the last data fault.



Write: TLB Purge

Writing Register 6 purges the TLB; the data is treated as an address and the TLB is searched for a corresponding page table descriptor. If a match is found, the corresponding entry is marked as invalid. This allows the page table descriptors in main memory to be updated and invalid entries in the on-chip TLB to be purged without requiring the entire TLB to be flushed.



Register 7 IDC Flush

Register 7 is a write-only register. The data written to this register is discarded and the IDC is flushed.

Registers 8-15 Reserved

Accessing any of these registers will cause the undefined instruction trap to be taken.

Instruction and Data Cache (IDC)

The LH74610 contains a 4kByte mixed instruction and data cache; the IDC has 256 lines of 16 bytes (4 words), organized as 4 blocks of 64 lines (making it 64-way set associative), and uses the virtual addresses generated by the processor core. The IDC is always reloaded a line at a time (four words). It may be enabled or disabled via the LH74610 Control Register and is disabled on RESET. The operation of the cache is further controlled by two bits: **Cacheable** and **Updateable**, which are stored in the Memory Management Page Tables. For this reason, in order to use the IDC, the MMU must be enabled. The two functions may however be enabled simultaneously, with a single write to the Control Register.

The Cacheable Bit

The **Cacheable** bit determines whether data being read may be placed in the IDC and used for subsequent read operations. Typically main memory will be marked as Cacheable to improve system performance, and I/O space as Non-cacheable to stop the data being stored in the LH74610's cache. [For example if the processor is polling a hardware flag in I/O space, it is important that the processor is forced to read data from the external peripheral, and not a copy of initial data held in the cache]. The Cacheable bit can be configured for both pages and sections.

The Updateable Bit

The **Updateable** bit determines whether the data in the cache should be updated during a write operation to maintain consistency with the external memory. [In certain cases automatic updating of cached data is not required: for instance, when using the MEMC1a memory manager, a read operation in the address space between 3400000H-3FFFFFFH would access the ROMs, but a write operation in the same address space would change a MEMC register, and should not affect the cached ROM data]. The Updateable bit can only be configured by the Level One descriptor: that is an entire section or all the pages for a single Level One descriptor share the same configuration.

IDC Operation

When the processor performs a read or write operation, the translation entry for that address is inspected and the state of the Cacheable and Updateable bits determines the subsequent action.

Cacheable Reads $C = 1$

The cache is searched for the relevant data; if found in the cache, the data is fed to the processor using a fast clock cycle (from FCLK). If the data is not found in the cache, an external memory access is initiated to read the appropriate line of data (4 words) from external memory and it is stored in a pseudo-randomly chosen entry in the cache (a linefetch operation).

Uncacheable Reads $C = 0$

The cache is not searched for the relevant data; instead an external memory access is initiated. No linefetch operation is performed, and the cache is not updated.

Updateable Writes $U = 1$

An external memory access is initiated, and the cache is searched; if the cache holds a copy of the data from the address being written to, then the cache data is simultaneously updated.

Non-Updateable Writes $U = 0$

An external memory access is initiated, but the cache is not searched and the contents of the cache are not affected.

IDC Validity

The IDC operates with virtual addresses, so care must be taken to ensure that its contents remain consistent with the virtual to physical mappings performed by the Memory Management Unit. If the Memory Mappings are changed, the IDC validity must be ensured.

Software IDC Flush

The entire IDC may be marked as invalid by writing to the LH74610 IDC Flush Register (Register 7). The cache will be flushed immediately the register is written, but note that the following two instruction fetches may come from the cache before the register is written.

Doubly Mapped Space

Since the cache works with virtual addresses, it is assumed that every virtual address maps to a different physical address. If the same physical location is accessed by more than one virtual address, the cache cannot maintain consistency, since each virtual address will have a separate entry in the cache, and only one entry will be updated on a processor write operation. To avoid any cache inconsistencies, both doubly-mapped virtual addresses should be marked as uncacheable.

Read-Lock-Write

The IDC treats the Read-Locked-Write instruction as a special case. The read phase always forces a read of external memory, regardless of whether the data is contained in the cache. The write phase is treated as a normal write operation (and if marked as Updateable, and the data is already in the cache, the cache will be updated). Externally the two phases are flagged as indivisible by asserting the LOCK signal.

IDC Enable/Disable and Reset

The IDC is automatically disabled and flushed on RESET. Once enabled, cacheable read accesses will cause lines to be placed in the cache. If subsequently disabled, no new lines will be placed in the cache, and the cache is not searched, but, Updateable write operations will continue to operate, thus maintaining consistency with the external memory. If the cache is subsequently re-enabled, it must be flushed if data already in the cache no longer matches that in external memory.

To enable the IDC

- Ensure the MMU is enabled (set bit 0 in Control Register)
- Enable the IDC (set bit 2 in Control Register)

NOTE: The MMU and IDC may be enabled simultaneously.

To disable the IDC

- Disable the IDC (clear bit 2 in Control Register)

NOTE: Updateable writes continue but no linefetches are performed. To fully inhibit the cache's operation it should be disabled and then flushed to ensure it contains no valid entries.

Write Buffer (WB)

The LH74610 Write Buffer is provided to improve system performance. It can buffer up to 8 words of data, and 2 independent addresses. It may be enabled or disabled via the W bit (bit 3) in the LH74610 Control Register and the buffer is disabled and flushed on RESET. The operation of the Write Buffer is further controlled by one bit, B, or Bufferable, which is stored in the Memory Management Page Tables. For this reason, in order to use the Write Buffer, the MMU must be enabled. The two functions may however be enabled simultaneously, with a single write to the Control Register. For a write to use the Write Buffer, both the W bit in the Control Register, and the B bit in the corresponding page table must be set.

Bufferable Bit

This bit controls whether a write operation may or may not use the Write Buffer. [Typically main memory will be bufferable and I/O space unbufferable]. The Bufferable bit can be configured for both pages and sections.

Write Buffer Operation

When the CPU performs a write operation, the translation entry for that address is inspected and the state of the B bit determines the subsequent action. If the Write Buffer is disabled via the LH74610 Control Register, bufferable writes are treated in the same way as unbuffered writes.

Bufferable Write

If the Write Buffer is enabled and the processor performs a write to a bufferable area, the data is placed in the Write Buffer at FCLK speeds and the CPU continues execution. The Write Buffer then performs the external write in parallel. If however the Write Buffer is full (either because there are already 8 words of data in the buffer, or because there is no slot for the new address) then the processor is stalled, until there is sufficient space in the buffer.

Unbufferable Writes

If the Write Buffer is disabled or the CPU performs a write to an unbufferable area, the processor is stalled until the write completes externally, which may require synchronization and several external clock cycles.

Read-Lock-Write

The write phase of a read-lock-write sequence is treated as an Unbuffered write, even if it is marked as buffered.

Note that a single write requires one address slot and one data slot in the write buffer; a sequential write of n words requires one address slot and n data slots. The total of 8 data slots in the buffer may be used as required. So for instance there could be one non-sequential write and one sequential write of 7 words in the buffer, and the processor could continue as normal: a third write or an eighth word in the second write would stall the processor until the first write had completed.

To enable the Write Buffer

- Ensure the MMU is enabled (set bit 0 in Control Register)
- Enable the Write Buffer (set bit 3 in Control Register)

NOTE: The MMU and Write Buffer may be enabled simultaneously

To disable the Write Buffer

- Disable the Write Buffer (clear bit 3 in Control Register)

NOTE: Any writes already in the write buffer will complete normally.

External Coprocessors

The LH74610 has no external coprocessor bus, so it is not possible to add external coprocessors to this device. If this is required, then the ARM600 should be used.

The LH74610 does still have an internal coprocessor designated #15 for internal control of the device. If a coprocessor other than #15 is accessed, then the CPU will take the undefined instruction trap.

Memory Management Unit (MMU)

The MMU performs two primary functions: it translates virtual addresses into physical addresses, and it controls memory access permissions. The MMU hardware required to perform these functions consists of a Translation Look-aside Buffer (TLB), access control logic, and translation table walking logic.

The MMU supports memory accesses based on Sections or Pages. Sections are comprised of 1MB blocks of memory. Two different page sizes are supported: Small Pages consist of 4kB blocks of memory and Large Pages consist of 64kB blocks of memory. (Large Pages are supported to allow mapping of a large region of memory while using only a single entry in the TLB). Additional access control mechanisms are extended within Small Pages to 1kB Sub-Pages and within Large Pages to 16kB Sub-Pages.

The MMU also supports the concept of domains - areas of memory that can be defined to possess individual access rights. The Domain Access Control Register is used to specify access rights for up to 16 separate domains.

The TLB caches 32 translated entries. During most memory accesses, the TLB provides the translation information to the access control logic.

If the TLB contains a translated entry for the virtual address, the access control logic determines whether access is permitted. If access is permitted, the MMU outputs the appropriate physical address corresponding to the virtual address. If access is not permitted, the MMU signals the CPU to abort.

If the TLB misses (it does not contain a translated entry for the virtual address), the translation table walk hardware is invoked to retrieve the translation information from a translation table in physical memory. Once retrieved, the translation information is placed into the TLB, possibly overwriting an existing value. The entry to be overwritten is chosen cyclically.

When the MMU is turned off (as happens on RESET), the virtual address is output directly onto the physical address bus.

MMU Program Accessible Registers

The LH74610 Processor provides several 32-bit registers which determine the operation of the MMU. The format for these registers is shown in Figure 28. A brief description of the registers is provided below. Each register will be discussed in more detail within the section that describes its use.

Data is written to and read from the MMU's registers using the ARM CPU's MRC and MCR coprocessor instructions.

	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
1 write	0																								S	B	L	D	P	W	C	A	M
2 write	Translation Table Base																																
3 write	Domain Access Control																																
	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0																	
5 read	0 0 0 0 Domain Status																																
5 write	FLUSH TLB (data - don't care)																																
6 read	Fault Address																																
6 write	Purge TLB (data = purge address)																																

Figure 28. MMU Register Summary

- The Translation Table Base Register holds the physical address of the base of the translation table maintained in main memory. Note that this base must reside on 16kB boundaries.
- The Domain Access Control Register consists of sixteen 2-bit fields, each of which defines the access permissions for one of the sixteen Domains (D15-D0).
- The Fault Status Register indicates the domain and type of access being attempted when an abort occurred. Bits 7:4 specify which of the sixteen domains (D15-D0) was being accessed when a fault occurred. Bits 3:1 indicate the type of access being attempted. The encoding of these bits is different for internal and external faults (as indicated by bit 0 in the register) and is shown in Table 6. A write to this register flushes the TLB.
- The Fault Address Register holds the virtual address of the access which was attempted when a fault occurred. A write to this register causes the data written to be treated as an address and, if it is found in the TLB, the entry is marked as invalid. (This operation is known as a TLB purge). The Fault Status Register and Fault Address Register are only updated for data faults, not for prefetch faults.

Address Translation

The MMU translates virtual addresses generated by the CPU into physical addresses to access external memory, and also derives and checks the access permission. Translation information, which consists of both the address translation data and the access permission data, resides in a translation table located in physical memory. The MMU provides the logic needed to traverse this translation table, obtain the translated address, and check the access permission.

There are three routes by which the address translation (and hence permission check) takes place. The route taken depends on whether the address in question has been marked as a section-mapped access or a page-mapped access; and there are two sizes of page-mapped access (large pages and small pages). However, the translation process always starts out in the same way, as described below, with a Level One fetch. A section-mapped access only requires a Level One fetch, but a page-mapped access also requires a Level Two fetch.

The Translation Process

Translation Table Base

The translation process is initiated when the on-chip TLB does not contain an entry for the requested virtual address. The Translation Table Base (TTB) Register points to the base of a table in physical memory which contains Section and/or Page descriptors. The 14 low-order bits of the TTB Register are set to zero as illustrated in Figure 29; the table must reside on a 16 kB boundary.




Figure 29. Translation Table Base Register

Level One Fetch

Bits 31:14 of the Translation Table Base register are concatenated with bits 31:20 of the virtual address to produce a 30-bit address as illustrated in Figure 30. This address selects a four-byte translation table entry which is a First Level Descriptor for either a Section or a Page (bit1 of the descriptor returned specifies whether it is for a Section or Page).

Level One Descriptor

The Level One Descriptor returned is either a Page Table Descriptor or a Section Descriptor, and its format varies accordingly. Figure 31 illustrates the format of Level One Descriptors.

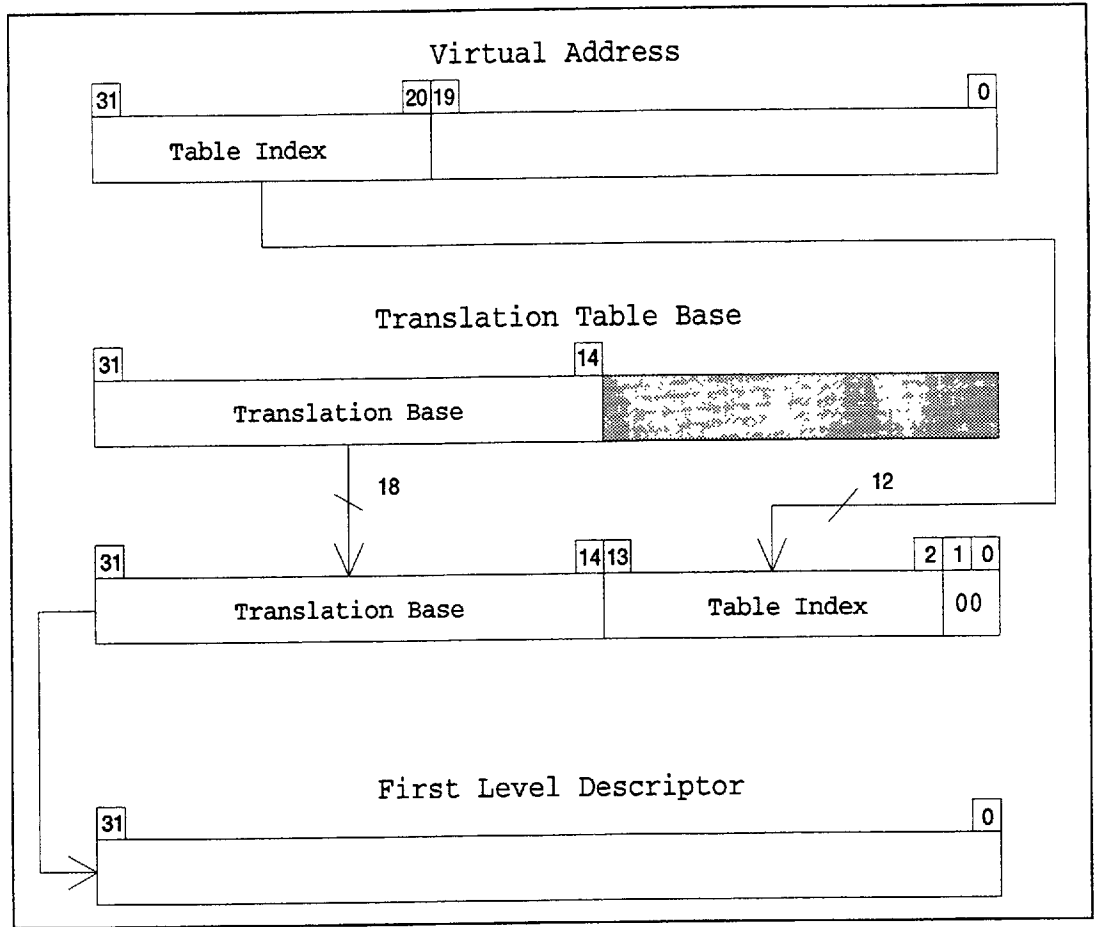


Figure 30. Accessing the Translation Table First Level Descriptors

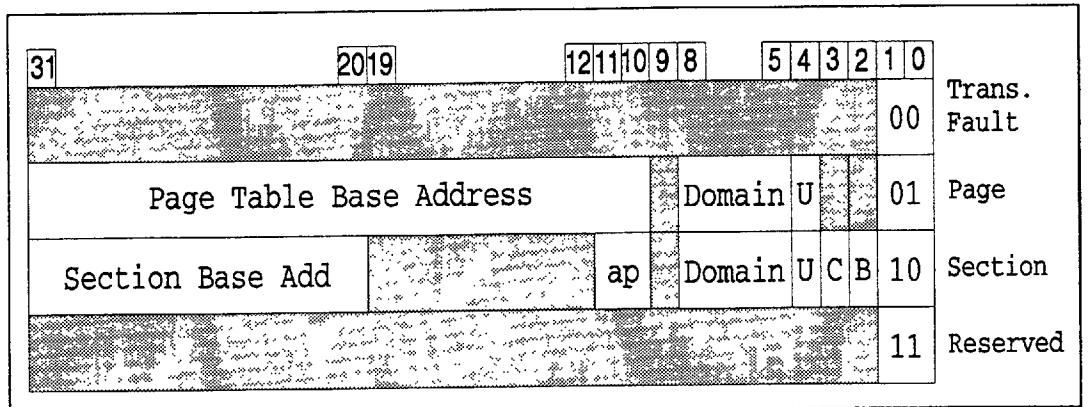


Figure 31. Level One Descriptors

The two least significant bits indicate the descriptor type and validity, and are interpreted as shown in Table 6.

Table 6. Interpreting Level One Descriptor Bits [1:0]

VALUE	MEANING	NOTES
0 0	Invalid	Generates a Section Translation Fault.
0 1	Page	Indicates that this is a Page Descriptor
1 0	Section	Indicates that this is a Section Descriptor
1 1	Reserved	Reserved for future use. (currently as for invalid)

Page Table Descriptor

Bits 3:2 are always written as 0.

Bit 4 Updateable: indicates that the data in the cache should be updated during a write operation to maintain consistency with external memory (if the cache is enabled).

Bits 8:5 specify one of the sixteen possible domains (held in the Domain Access Control Register) that contain the primary access controls.

Bits 31:10 form the base for referencing the Page Table Entry. (The page table index for the entry is derived from the virtual address as illustrated in Figure 5).

If a Page Table Descriptor is returned from the Level One fetch, a Level Two fetch is initiated as described below.

Section Descriptor

Bits 4:2 (U, C, & B) control the cache- and write-buffer-related functions as follows:

U - Updateable: indicates that the data in the cache should be updated during a write operation to maintain consistency with external memory (if the cache is enabled).

C - Cacheable: indicates that data at this address will be placed in the cache (if the cache is enabled).

B - Bufferable: indicates that data at this address will be written through the write buffer (if the write buffer is enabled).

Bits 8:5 specify one of the sixteen possible domains (held in the Domain Access Control Register) that contain the primary access controls.

Bits 11:10 (a) specify the access permissions for this section and are interpreted as shown in Table 7. Their interpretation is dependent upon the setting of the S bit (Control Register bit 8). Note that the Domain Access Control specifies the primary access control; the ap bits only have an effect in client mode. Refer to section on access permissions.

Table 7. Interpreting Access Permission (ap) Bits

ap	S	PERMISSIONS		NOTES
		SUPERVISOR	USER	
00	0	No Access	No Access	Any access generates a permission fault
00	1	Read Only	No Access	Supervisor read only permitted
01	x	Read/Write	No Access	Access allowed only in Supervisor mode
10	x	Read/Write	Read Only	Writes in User mode cause permission fault
11	x	Read/Write	Read/Write	All access types permitted in both modes.

Bits 19:12 are always written as 0.

Bits 31:20 form the corresponding bits of the physical address for the 1MByte section.

Translating Section References

Figure 32 illustrates the complete Section translation sequence. Note that the access permissions contained in the Level One descriptor must be checked before the physical address is generated. The sequence for checking access permissions is described below.

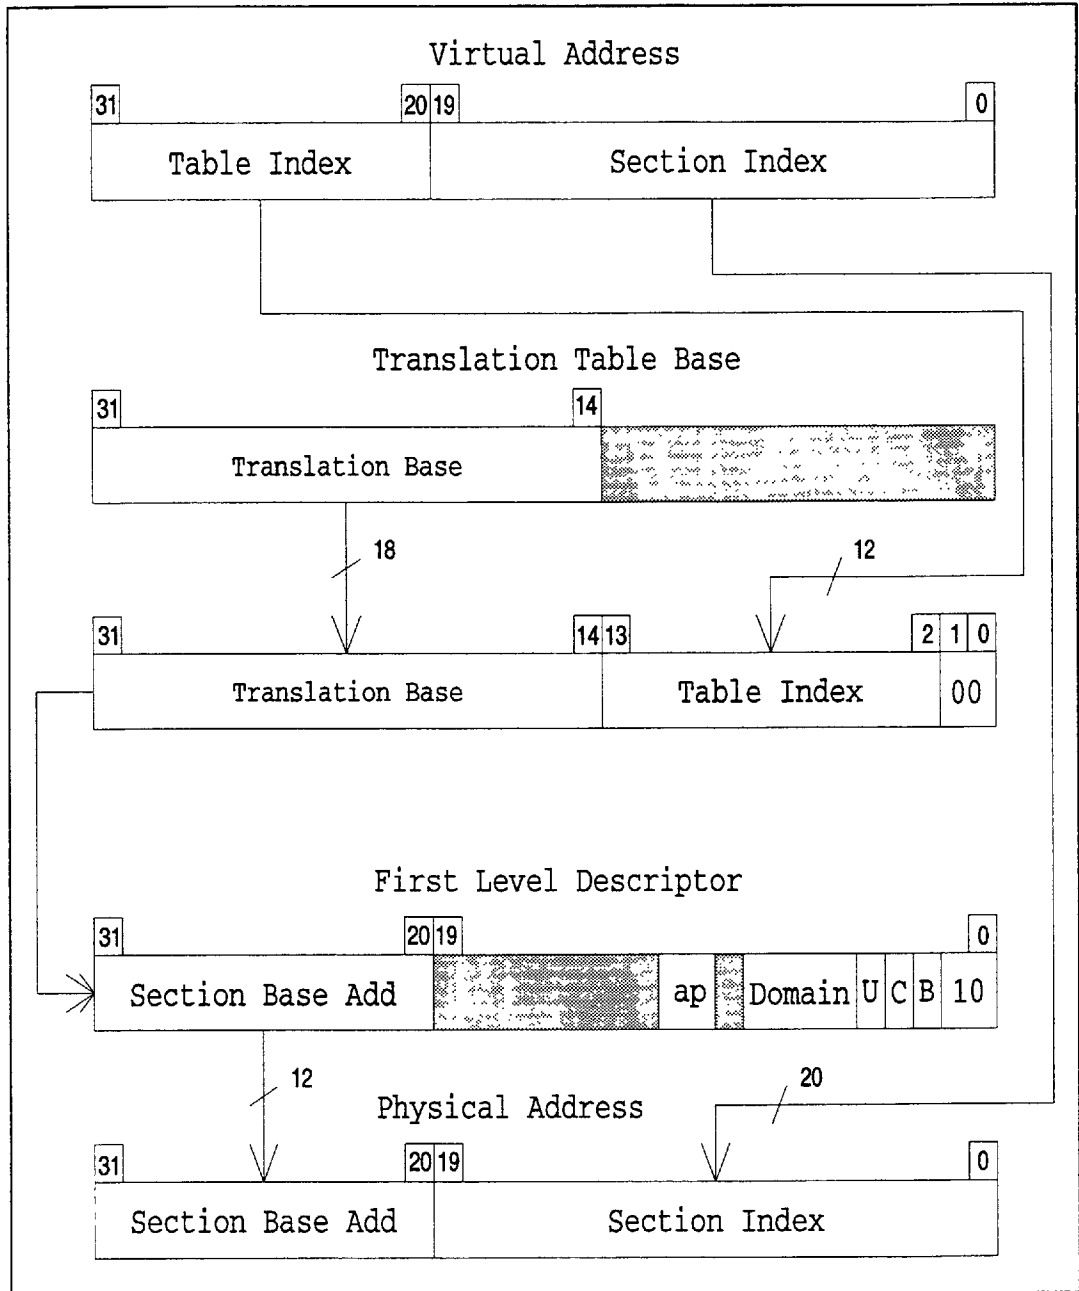


Figure 32. Section Translation

Level Two Descriptor

If the Level One fetch returns a Page Table Descriptor, this provides the base address of the page table to be used. The page table is then accessed as described in Figure 5, and a Page Table Entry, or Level Two Descriptor, is returned. This in turn may define either a Small Page or a Large Page access. Figure 33 illustrates the format of Level Two Descriptors.

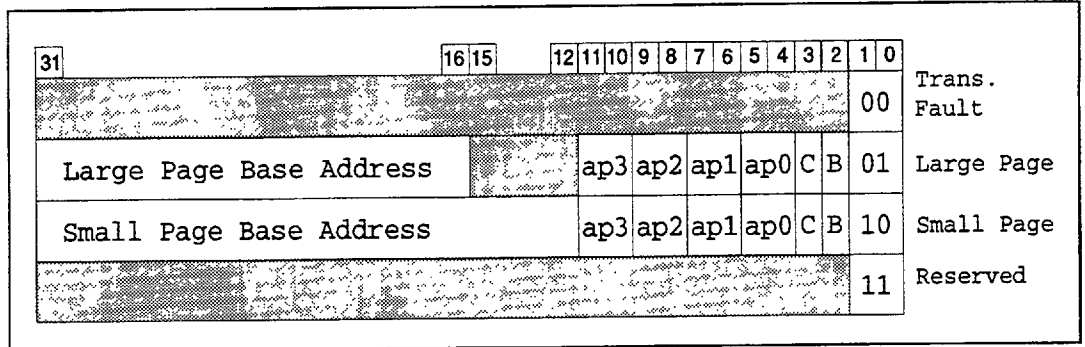


Figure 33. Page Table Entry (Level Two Descriptor)

The two least significant bits indicate the page size and validity, and are interpreted as shown in Table 8.

Table 8. Interpreting Page

VALUE	MEANING	NOTES
0 0	Invalid	Generates a Page Translation Fault.
0 1	Large Page	Indicates that this is a 64 kB Page
1 0	Small Page	Indicates that this is a 4 kB Page
1 1	Reserved	Reserved for future use. (currently as for invalid)

Bit 2 B - Bufferable: indicates that data at this address will be written through the Write Buffer (if the Write Buffer is enabled).

Bit 3 C - Cacheable: indicates that data at this address will be placed in the IDC (if the cache is enabled).

Bits 11:4 specify the access permissions (ap3 - ap0) for the four sub-pages and interpretation of these bits is described earlier in Table 4.

For large pages, bits 15:12 are programmed as 0

Bits 31:12 (small pages) or **bits 31:16** (large pages) are used to form the corresponding bits of the physical address - the physical page number. (The page index is derived from the virtual address as illustrated in Figures 34 and 35).

Translating Small Page References

Figure 5 illustrates the complete translation sequence for a 4kB Small Page. Page translation involves one additional step beyond that of a section translation: the Level One descriptor is the Page Table descriptor, and this is used to point to the Level Two descriptor, or Page Table Entry. (Note that the access permissions are now contained in the Level Two descriptor and must be checked before the physical address is generated. The sequence for checking access permissions is described later).

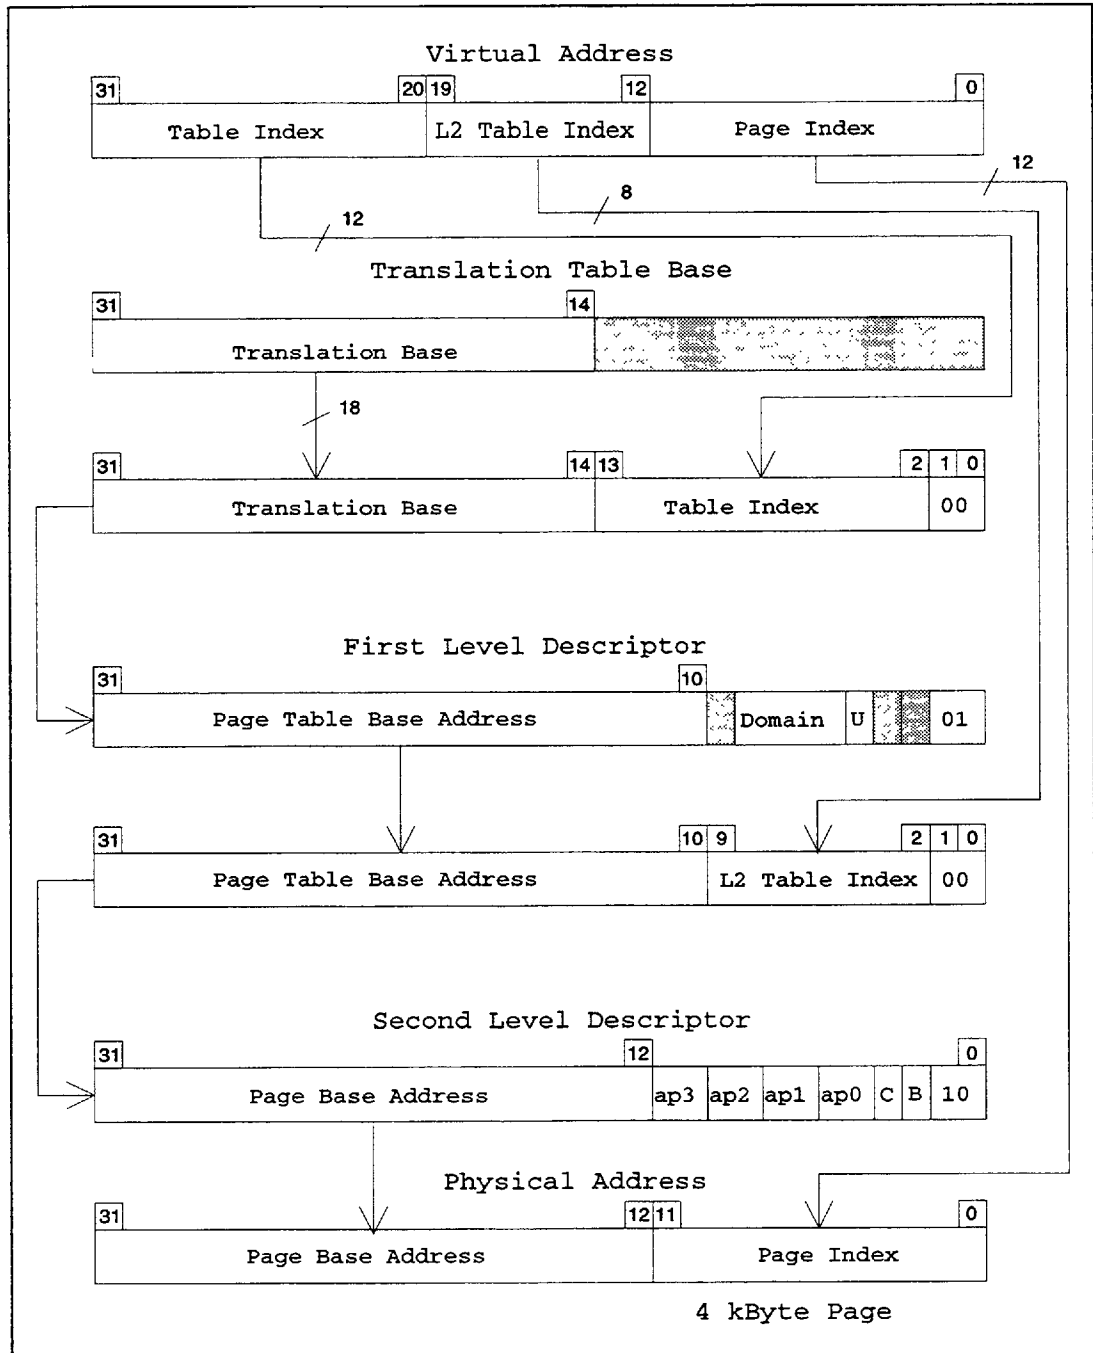


Figure 34. Small Page Translation

Translating Large Page References

Figure 35 illustrates the complete translation sequence for a 64 kB Large Page. Note that since the upper four bits of the Page Index and low-order four bits of the Page Table index overlap, each Page Table Entry for a Large Page must be duplicated 16 times (in consecutive memory locations) in the Page Table.

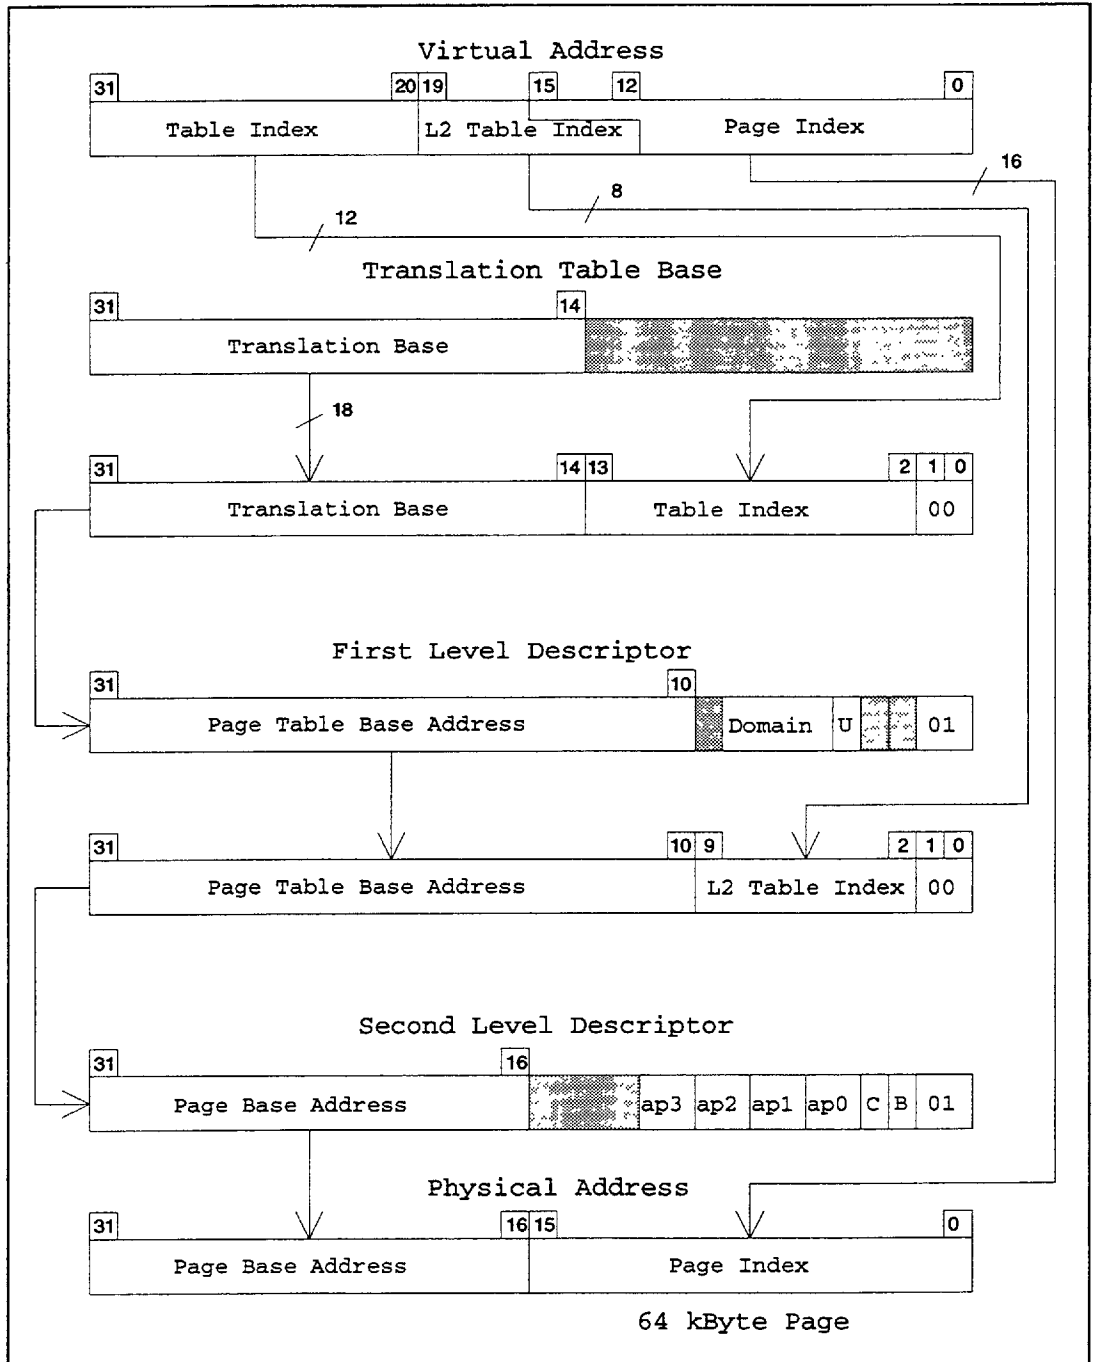


Figure 35. Large Page Translation

MMU Faults and CPU Aborts

The MMU generates four types of faults:

- Alignment Fault.
- Translation Fault.
- Domain Fault.
- Permission Fault.

In addition, an external abort may be raised on external data access.

The access control mechanisms of the MMU detect the conditions that produce these faults. If a fault is detected as the result of a memory access, the MMU will abort the access and signal the fault condition to the CPU. The MMU is also capable of retaining status and address information about the abort. The CPU recognizes two types of abort: data aborts and prefetch aborts, and these are treated differently by the MMU.

If the MMU detects an access violation, it will do so before the external memory access takes place, and it will therefore inhibit the access. External aborts will not necessarily inhibit the external access, as described in the section on external aborts.

Fault Address and Fault Status Registers (FAR and FSR)

Aborts resulting from data accesses (data aborts) are acted upon by the CPU immediately, and the MMU places an encoded 4 bit value FS[3:0], along with the 4bit encoded Domain number, in the Fault Status Register (FSR). In addition, the virtual processor address which caused the data abort is latched into the Fault Address Register (FAR). If an access violation simultaneously generates more than one source of abort, they are encoded in the priority given in Table 6.

CPU instructions on the other hand are prefetched, so a prefetch abort simply flags the instruction as it enters the instruction pipeline. Only when (and if) the instruction is executed does it cause an abort; an abort is not acted upon if the instruction is not used (i.e. it is branched around). Because instruction prefetch aborts may or may not be acted upon, the MMU status information is not preserved for the resulting CPU abort; for a prefetch abort, the MMU does not update the FSR or FAR.

The sections that follow describe the various access permissions and controls supported by the MMU and detail how these are interpreted to generate faults.

Table 9. Priority Encoding of Fault Status

		SOURCE	FS[32:10]	DOMAIN[3:0]	FAR		
Highest		Write Buffer	00x0	x	NOTE2		
		Bus Error (linefetch)	Section	0100	valid	valid	NOTE 3
			Page	0110	valid	valid	
		Bus Error (other)	Section	1000	valid	valid	
			Page	1010	valid	valid	
		Alignment		00x1	x	valid	
		Bus Error (translation)	level1	1100	x	valid	
			level2	1110	valid	valid	
		Translation	Section	0101	valid	valid	
			Page	0111	valid	valid	
	Domain	Section	1001	valid	valid		
		Page	1011	valid	valid		
Lowest		Permission	Section	1101	valid	valid	
			Page	1111	valid	valid	

NOTES:

x is undefined: may read as 0 or 1

1. Any abort masked by the priority encoding may be regenerated by fixing the primary abort and restarting the instruction.
2. The Write Buffer Bus Error is asynchronous and the Fault Address Register reflects the first data operation that could be aborted. This instruction must be restarted, (and it could of course generate an abort of its own).
3. This assumes that the error was flagged on word 0 of the linefetch
4. In fact this register will contain bits[8:5] of the Level 1 entry which are undefined, but would encode the domain in a valid entry.

Access Control

MMU accesses are primarily controlled via domains. There are 16 domains, and each has a 2-bit field to define it. Two basic kinds of users are supported: Clients and Managers. Clients use a domain; Managers control the behavior of the domain. The domains are defined in the Domain Access Control Register. Figure 36 illustrates how the 32 bits of the register are allocated to define the sixteen 2-bit domains.

Table 10 defines how the bits within each domain are interpreted to specify the access permissions.

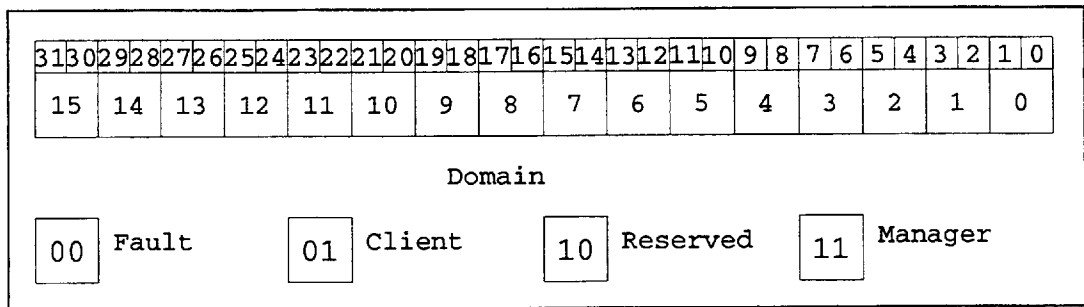


Figure 36. Domain Access Control Register Format

Table 10. Interpreting Access Bits in Domain Access Control Register

VALUE	MEANING	NOTES
00	No Access	Any access will generate a Domain Fault.
01	Client	Accesses are checked against the access permission bits in the Section or Page descriptor.
10	Reserved	Reserved. Currently behaves like the no access mode.
11	Manager	Accesses are NOT checked against the access Permission bits so a Permission fault cannot be generated.

Fault Checking Sequence

The sequence by which the MMU checks for access faults is slightly different for Sections and Pages. Figure 37 illustrates the sequence for both types of accesses. The sections and figures that follow describe the conditions that generate each of the faults.

Alignment Fault

If Alignment fault is enabled (bit 1 in Control Register set), the MMU will generate an alignment fault on any data word access the address of which is not word-aligned irrespective of whether the MMU is enabled or not; in other words, if either of virtual address bits [1:0] are not 0. Alignment fault will not be generated on any instruction fetch, nor on any byte access. Note that if the access generates an alignment fault, the access sequence will abort without reference to further permission checks.

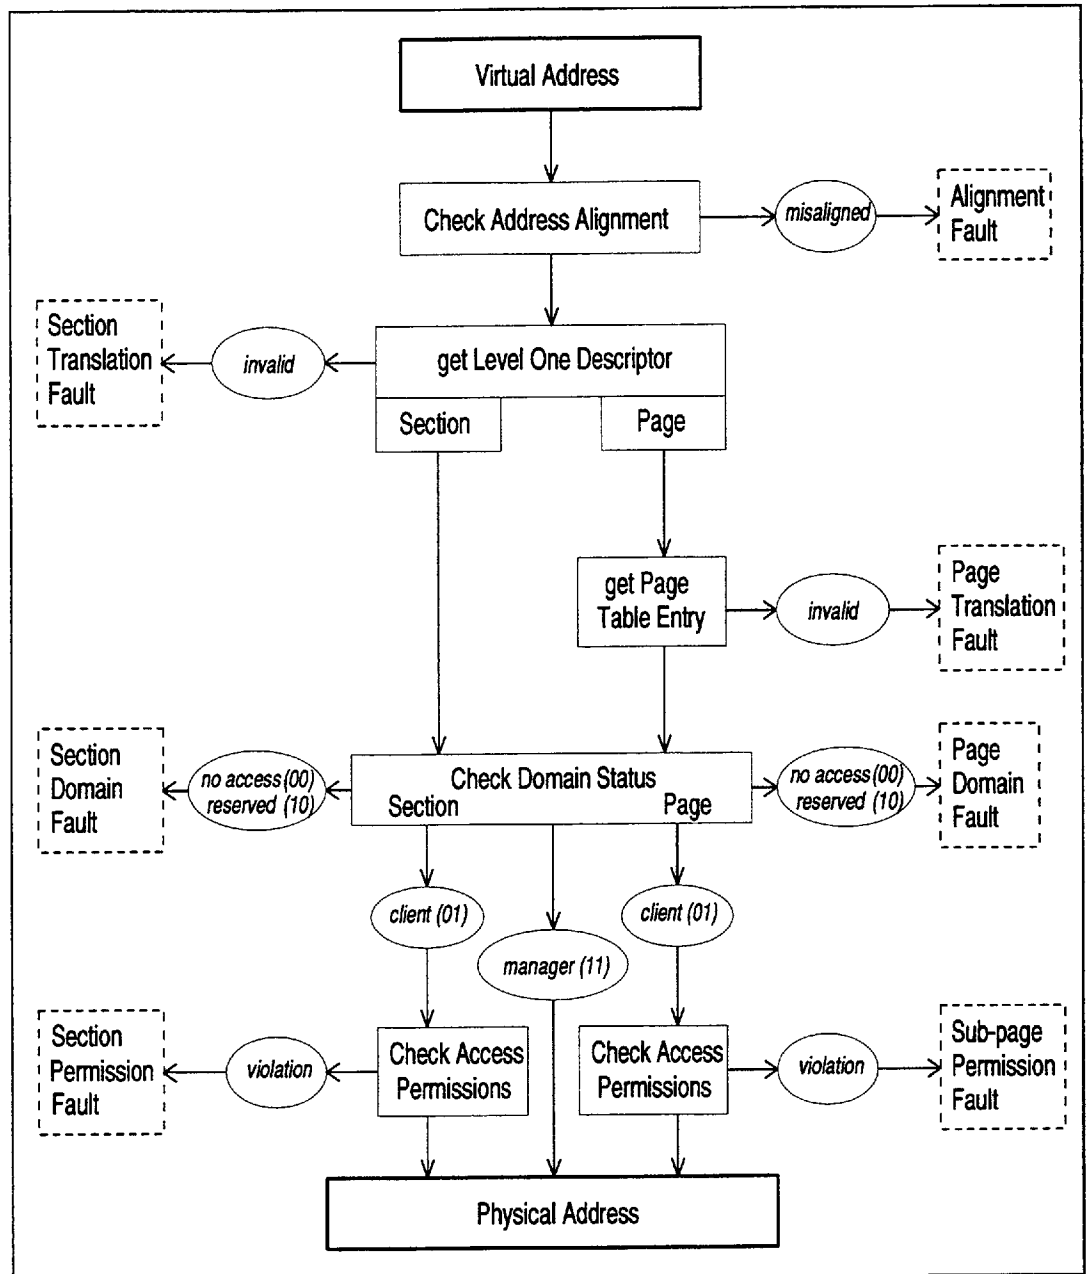


Figure 37. Sequence for Checking Faults

Translation Fault

There are two types of translation fault: section & page.

A Section Translation Fault is generated if the Level One descriptor is marked as invalid. This happens if bits[1:0] of the descriptor are both 0 or both 1.

A Page Translation Fault is generated if the Page Table Entry is marked as invalid. This happens if bits[1:0] of the entry are both 0 or both 1.

Domain Fault

There are two types of domain fault: section & page. In both cases the Level One descriptor holds the 4-bit Domain field which selects one of the sixteen 2-bit domains in the Domain Access Control Register. The two bits of the specified domain are then checked for access permissions as detailed in Table 7. In the case of a section, the domain is checked once the Level One descriptor is returned, and in the case of a page, the domain is checked once the Page Table Entry is returned.

If the specified access is either No Access (00) or Reserved (10) then either a Section Domain Fault or Page Domain Fault occurs.

Permission Fault

There are two types of permission fault: section & sub-page. Permission fault is checked at the same time as Domain fault. If the 2-bit domain field returns client (01), then the permission access check is invoked as follows:

section:

If the Level One descriptor defines a section-mapped access, then the ap bits of the descriptor define whether or not the access is allowed according to Table 7 which is repeated below. Their interpretation is dependent upon the setting of the S bit (Control Register bit 8). If the access is not allowed, then a Section Permission fault is generated.

sub-page:

If the Level One descriptor defines a page-mapped access, then the Level Two descriptor specifies four access permission fields (ap3..ap0) each corresponding to one quarter of the page. Hence for small pages, ap3 is selected by the top 1kB of the page, and ap0 is selected by the bottom 1kB of the page; for large pages, ap3 is selected by the top 6kB of the page, and ap0 is selected by the bottom 16kB of the page. The selected ap bits are then interpreted in exactly the same way as for a section (see table below), the only difference being that the fault generated is a sub-page permission fault.

Table 11. Interpreting Access Permission (ap) Bits

ap	S	PERMISSIONS		NOTES
		SUPERVISOR	USER	
00	0	No Access	No Access	Any access generates a permission fault
00	1	Read Only	No Access	Supervisor read only permitted
01	x	Read/Write	No Access	Supervisor read or write only permitted
10	x	Read/Write	Read Only	Writes in User mode cause permission fault
11	x	Read/Write	Read/Write	All access types permitted in both modes

External Aborts

In addition to the MMU-generated aborts, the LH74610 has an external abort pin which may be used to flag an error on an external memory access. However, some accesses aborted in this way are not restartable, so this pin must be used with great care. The following section describes the restrictions.

- Uncacheable reads
- Unbuffered writes
- Level One descriptor fetch
- Level Two descriptor fetch
- read-lock-write sequence

These accesses may be aborted and restarted safely. If any of the above are aborted the external access will cease on the following cycle. In the case of a read-lock-write sequence in which the read aborts, the write will not happen.

Cacheable reads (linefetches)

A linefetch may be aborted safely provided that the abort is flagged on word 0. In this case, the IDC will not be updated or corrupted, and the access will be restartable. It is not advisable to flag an abort on any word other than word 0 of a linefetch, as the IDC will contain a corrupt line, and the instruction may not be restartable. On the external bus, a linefetch which is externally aborted will continue to the end as though it had not aborted.

Buffered writes

Buffered writes cannot be safely externally aborted. Because the processor will have moved on before the external abort is received, this class of abort is not restartable. If the system does flag this type of abort, then the Fault Status Register will record the fact, but this is a non-recoverable error, and the machine must be reset. Therefore, the system should be configured such that it does not do buffered writes to areas of memory which are capable of flagging an external abort. If a buffered write burst is externally aborted, then the external write will continue to the end.

Interaction of the MMU, IDC and Write Buffer

The MMU, IDC and WB may be enabled/disabled independently. However there are only five valid combinations. There are no hardware interlocks on these restrictions, so invalid combinations will cause undefined results.

Table 12. Interaction of MMU, IDC and Write Buffer

MMU	IDC	WB
off	off	off
on	off	off
on	on	off
on	off	on
on	on	on

The following procedures must be observed.

To enable the MMU

- Program the Translation Table Base and Domain Access Control Registers
- Program Level 1 and Level 2 page tables as required
- Enable the MMU by setting bit 0 in the Control Register.

NOTE: Care must be taken if the translated address differs from the untranslated address as the two instructions following the enabling of the MMU will have been fetched using "flat translation" and enabling the MMU may be considered as a branch with delayed execution. A similar situation occurs when the MMU is disabled. Consider the following code sequence:

```
MOV R1, #&1
MCR 15,0,R1,0,0; Enable MMU
Fetch Flat
Fetch Flat
Fetch Translated
```

To disable the MMU

- Disable the WB by clearing bit 3 in the Control Register.
- Disable the IDC by clearing bit 2 in the Control Register.
- Disable the MMU by clearing bit 0 in the Control Register.

NOTE: If the MMU is enabled, then disabled and subsequently re-enabled the contents of the TLB will have been preserved. If these are now invalid, the TLB should be flushed before re-enabling the MMU. Disabling of all three functions may be done simultaneously.

Effect of Reset

At the end of the reset sequence, the MMU is disabled, so forces "flat" translation (i.e. the physical address is the virtual address, and there is no permission checking); alignment faults are also disabled; the cache is disabled and flushed; the write buffer is disabled and flushed; the ARM6 CPU core is put into 26 bit data and address mode, with early abort timing and little-endian mode.

LH74610 Bus Interface

Overview

The LH74610 bus interface is designed to operate in synchronous mode. In this mode, there is a tightly defined relationship between FCLK and MCLK. MCLK may only make transitions on the falling edge of FCLK. An amount of jitter between the two clocks is permitted, and the device will function correctly, but MCLK must not be later than FCLK. MCLK may lead FCLK by up to the difference in MCLK and FCLK periods without affecting the critical paths.

Cycle Speed

The speed of the memory bus interface may be controlled in two ways:

1. The LOW and HIGH Phases of MCLK may be stretched
2. The nWAIT pin can be used to insert entire MCLK cycles into any access. When LOW this signal inserts an extra MCLK cycle into the current access. It must change while MCLK is LOW.

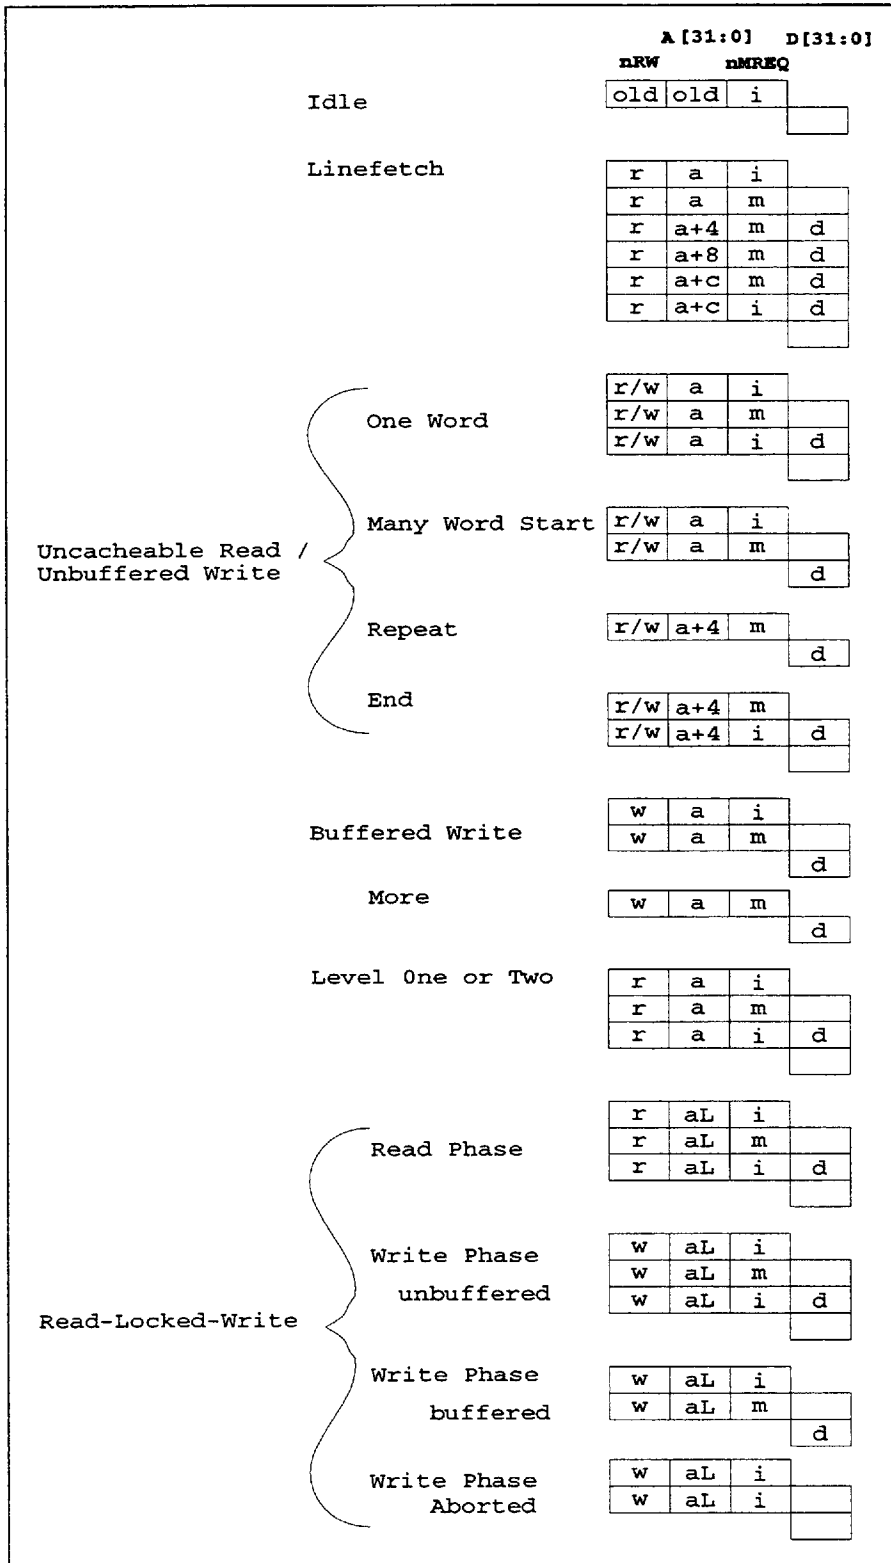
Cycle Types

The LH74610 can perform many different bus cycles, and they may be combined in any order.

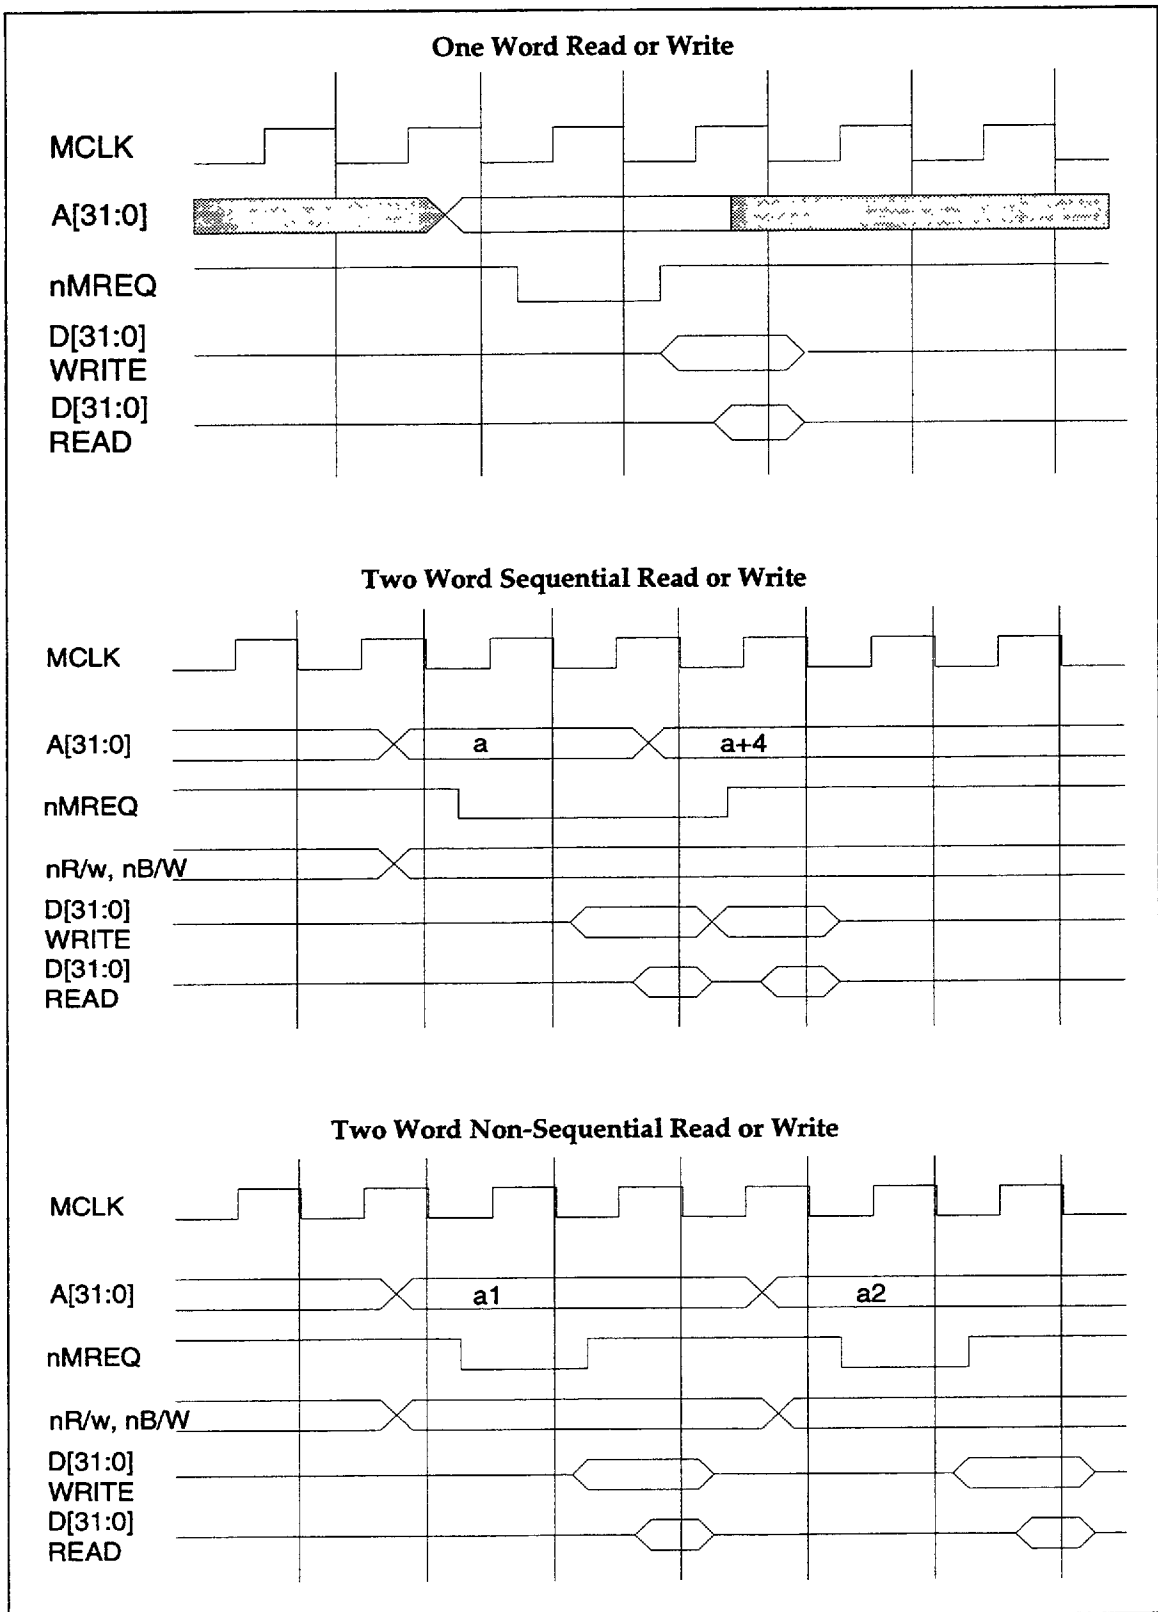
- Idle
- Line Fetch
- Unbuffered Write
- Uncacheable Read
- Buffered Write
- Translation Level 1 Fetch
- Translation Level 2 Fetch
- Read-Locked-Write

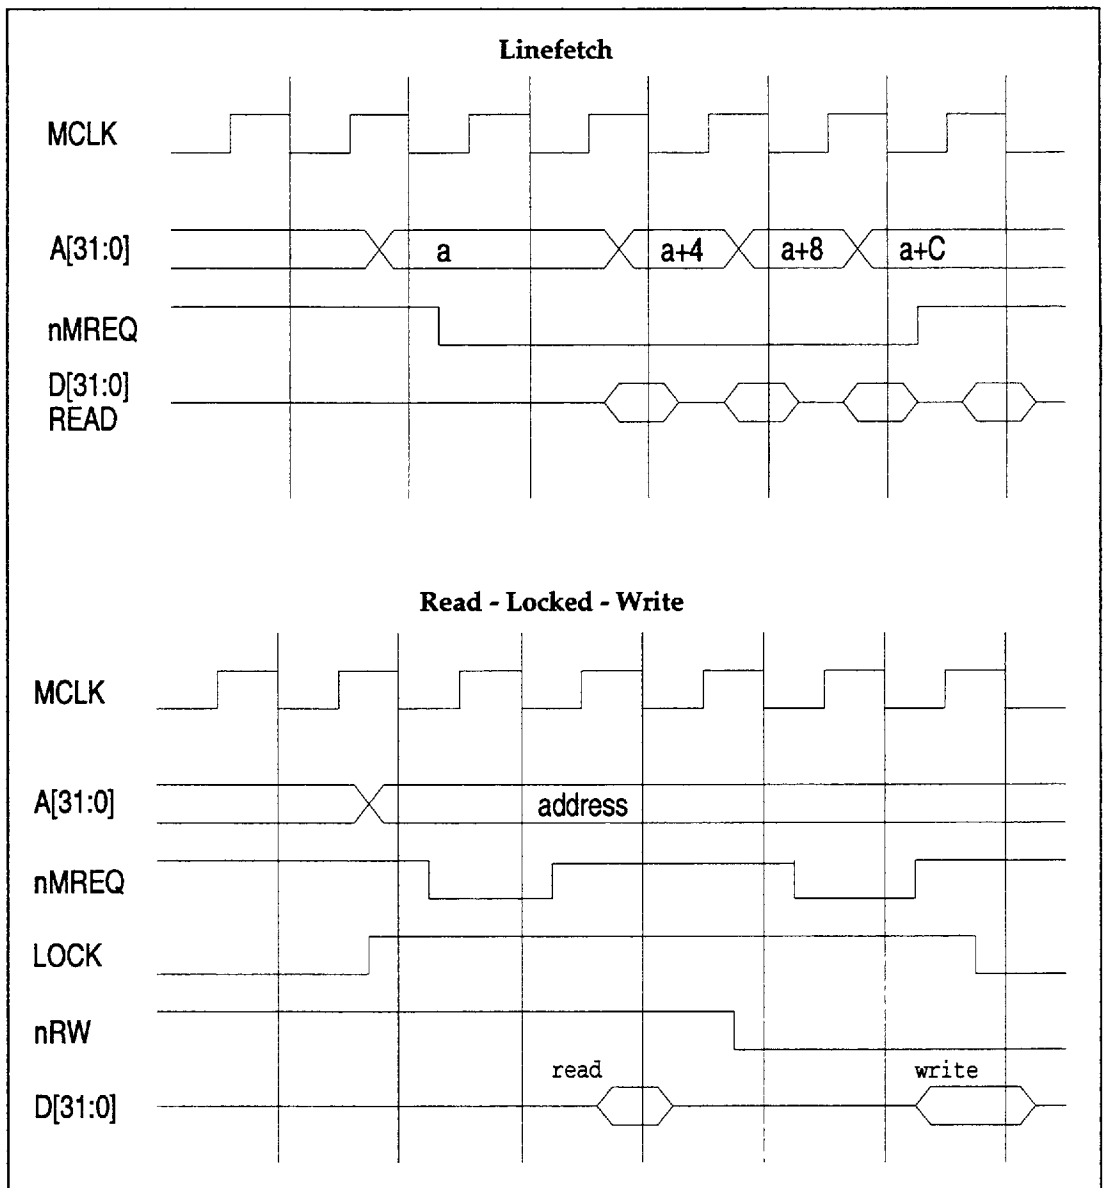
Linefetch, unbuffered writes, uncacheable reads and buffered writes all have very similar bus cycles. They differ only in the number of idle cycles at the end, which in turn depends on which cycle follows. The following table defines this for each bus cycle. The following diagrams give examples of typical sequences.

Cycle Type Details




Example Cycles



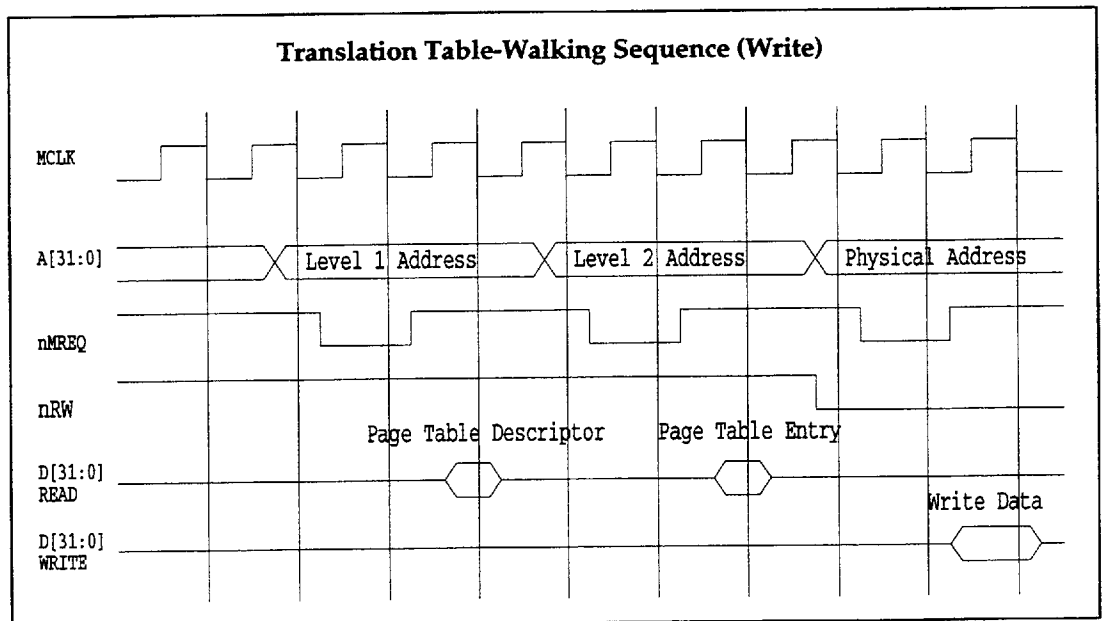


Note that the address bus and nR/W do not change at the end of the sequence, unless the following cycle is to a different address, or is a read. LOCK, however, always returns to a LOW level signifying the end of the locked sequence.

Example Cycles (cont'd)



In this example the $nWAIT$ pin has been used to introduce an extra MCLK cycle to ease the system timing. This allows all control signals to be sampled on the rising edges of MCLK. More sophisticated systems may sample control signals on both rising and falling edges of MCLK to improve performance.

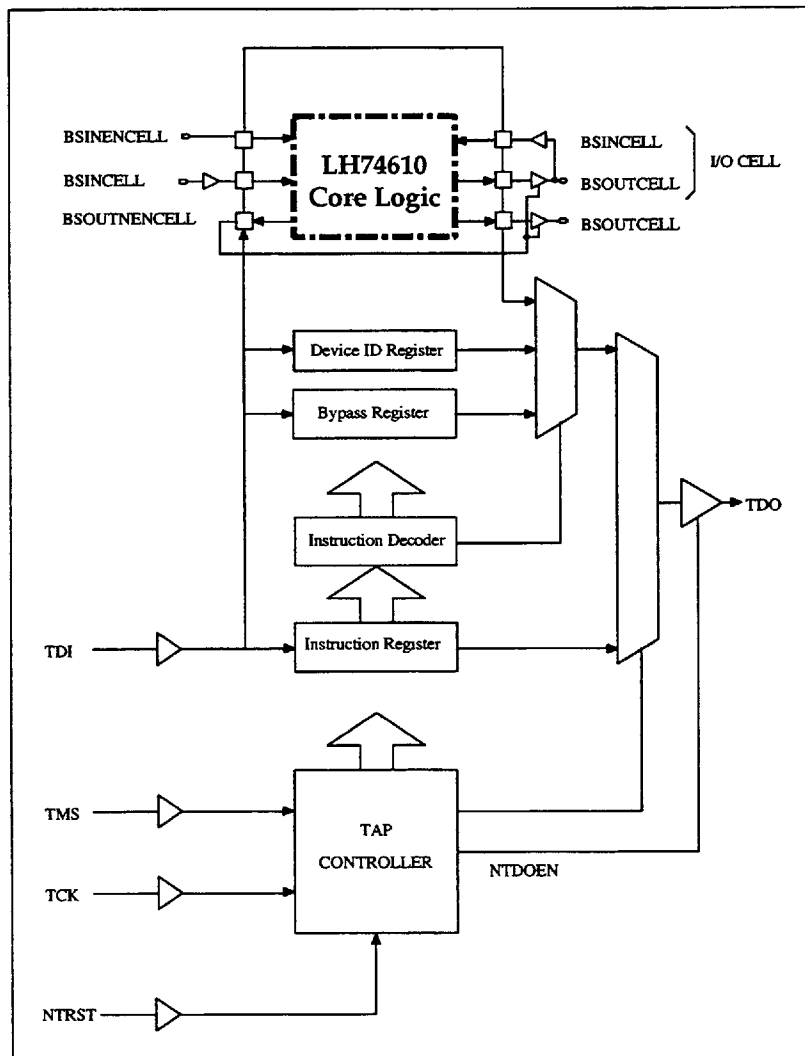


Boundary Scan Test Interface

The boundary-scan interface conforms to the IEEE Std. 1149.1- 1990, Standard Test Access Port and Boundary-Scan Architecture (please refer to this document for an explanation of the terms used in this section and for a description of the TAP controller states.)

Overview

The boundary-scan interface provides a means of testing the core of the device when it is fitted to a circuit board, and a means of driving and sampling all the external pins of the device irrespective of the core state. This latter function permits testing of both the device's electrical connections to the circuit board, and (in conjunction with other devices on the circuit board having a similar interface) testing the integrity of the circuit board connections between devices. The interface intercepts all external connections within the device, and each such "cell" is then connected together to form a serial register (the boundary scan register). The whole interface is controlled via 5 dedicated pins: TDI, TMS, TCK, nTRST and TDO.



Reset

The boundary-scan interface includes a state-machine controller (the TAP controller). In order to force the TAP controller into the correct state after power-up of the device, a reset pulse must be applied to the pin. If the boundary scan interface is to be used, then nTRST must be driven LOW, and then HIGH again. If the boundary scan interface is not to be used, then the nTRST pin may be tied permanently LOW. Note that a clock on TCK is not necessary to reset the device.

The action of reset (either a pulse or a DC level) is as follows:

System mode is selected (i.e. the boundary scan chain does NOT intercept any of the signals passing between the pads and the core).

IDcode mode is selected. If TCK is pulsed, the contents of the ID register will be clocked out of TDO.

Pullup Resistors

The 1149.1 standard effectively requires that TDI, TMS, and nTRST should have internal pullup resistors. In order to allow the LH74610 to consume zero static current, these resistors are NOT fitted to this device. Accordingly, the 4 inputs to the test interface (the above 3 signals plus TCK) must all be driven to good logic levels to achieve normal circuit operation.

Instruction Register

The instruction register is 4 bits in length. There is no parity bit.

The fixed value loaded into the instruction register during the CAPTURE-IR controller state is: 0001

Public Instructions

The following public instructions are supported:

INSTRUCTION	BINARY CODE
BYPASS	1111
SAMPLE/PRELOAD	0011
EXTEST	0000
INTEST	1100
IDCODE	1110
HIGHZ	0111
CLAMP	0101
CLAMPZ	1001

In the descriptions that follow, TDI and TMS are sampled on the rising edge of TCK and all output transitions on TDO occur as a result of the falling edge of TCK.

BYPASS (1111)

The BYPASS instruction connects a 1 bit shift register (the BYPASS register) between TDI and TDO.

When the BYPASS instruction is loaded into the instruction register, all the boundary-scan cells are placed in their normal (system) mode of operation. This instruction has no effect on the system pins. In the CAPTURE-DR state, a logic 0 is captured by the bypass register. In the SHIFT-DR state, test data is shifted into the bypass register via TDI and out via TDO after a delay of one TCK cycle. Note that the first bit shifted out will be a zero. The bypass register is not affected in the UPDATE-DR state.

SAMPLE/PRELOAD (0011)

The BS (boundary-scan) register is placed in test mode by the SAMPLE/PRELOAD instruction.

The SAMPLE/PRELOAD instruction connects the BS register between TDI and TDO.

When the instruction register is loaded with the SAMPLE/PRELOAD instruction, all the boundary-scan cells are placed in their normal system mode of operation.

In the CAPTURE-DR state, a snapshot of the signals at the boundary-scan cells is taken on the rising edge of TCK. Normal system operation is unaffected. In the SHIFT-DR state, the sampled test data is shifted out of the BS register via the TDO pin, whilst new data is shifted in via the TDI pin to preload the BS register parallel input latch. In the UPDATE-DR state, the preloaded data is transferred into the BS register parallel output latch. Note that this data is not applied to the system logic or system pins while the SAMPLE/PRELOAD instruction is active. This instruction should be used to preload the boundary-scan register with known data prior to selecting the INTEST or EXTEST instructions (see the table below for appropriate guard values to be used for each boundary-scan cell).

EXTEST (0000)

The BS (boundary-scan) register is placed in test mode by the EXTEST instruction.

The EXTEST instruction connects the BS register between TDI and TDO.

When the instruction register is loaded with the EXTEST instruction, all the boundary-scan cells are placed in their test mode of operation.

In the CAPTURE-DR state, inputs from the system pins and outputs from the boundary-scan output cells to the system pins are captured by the boundary-scan cells. In the SHIFT-DR state, the previously captured test data is shifted out of the BS register via the TDO pin, whilst new test data is shifted in via the TDI pin to the BS register parallel input latch. In the UPDATE-DR state, the new test data is transferred into the BS register parallel output latch. Note that this data is applied immediately to the system logic and system pins. The first EXTEST vector should be clocked into the boundary-scan register, using the SAMPLE/PRELOAD instruction, prior to selecting INTEST to ensure that known data is applied to the system logic.

INTEST (1100)

The BS (boundary-scan) register is placed in test mode by the INTEST instruction.

The INTEST instruction connects the BS register between TDI and TDO.

When the instruction register is loaded with the INTEST instruction, all the boundary-scan cells are placed in their test mode of operation.

In the CAPTURE-DR state, the complement of the data supplied to the core logic from input boundary-scan cells is captured, while the true value of the data that is output from the core logic to output boundary-scan cells is captured. Note that CAPTURE-DR captures the complemented value of the input cells for testability reasons.

In the SHIFT-DR state, the previously captured test data is shifted out of the BS register via the TDO pin, whilst new test data is shifted in via the TDI pin to the BS register parallel input latch. In the UPDATE-DR state, the new test data is transferred into the BS register parallel output latch. Note that this data is applied immediately to the system logic and system pins. The first INTEST vector should be clocked into the boundary-scan register, using the SAMPLE/PRELOAD instruction, prior to selecting INTEST to ensure that known data is applied to the system logic.

Single-step operation is possible using the INTEST instruction.

IDCODE (1110)

The IDCODE instruction connects the device identification register (or ID register) between TDI and TDO. The ID register is a 32-bit register that allows the manufacturer, part number and version of a component to be determined through the TAP.

When the instruction register is loaded with the IDCODE instruction, all the boundary-scan cells are placed in their normal (system) mode of operation.

In the CAPTURE-DR state, the device identification code (specified at the end of this section) is captured by the ID register. In the SHIFT-DR state, the previously captured device identification code is shifted out of the ID register via the TDO pin, whilst data is shifted in via the TDI pin into the ID register. In the UPDATE-DR state, the ID register is unaffected.

HIGHZ (0111)

The HIGHZ instruction connects a 1 bit shift register (the BYPASS register) between TDI and TDO.

When the HIGHZ instruction is loaded into the instruction register, all outputs are placed in an inactive drive state.

In the CAPTURE-DR state, a logic 0 is captured by the bypass register. In the SHIFT-DR state, test data is shifted into the bypass register via TDI and out via TDO after a delay of one TCK cycle. Note that the first bit shifted out will be a zero. The bypass register is not affected in the UPDATE-DR state.

CLAMP (0101)

The CLAMP instruction connects a 1 bit shift register (the BYPASS register) between TDI and TDO.

When the CLAMP instruction is loaded into the instruction register, the state of all output signals is defined by the values previously loaded into the boundary-scan register. A guarding pattern (specified for this device at the end of this section) should be pre-loaded into the boundary-scan register using the SAMPLE/PRELOAD instruction prior to selecting the CLAMP instruction.

In the CAPTURE-DR state, a logic 0 is captured by the bypass register. In the SHIFT-DR state, test data is shifted into the bypass register via TDI and out via TDO after a delay of one TCK cycle. Note that the first bit shifted out will be a zero. The bypass register is not affected in the UPDATE-DR state.

CLAMPZ (1001)

The CLAMPZ instruction connects a 1 bit shift register (the BYPASS register) between TDI and TDO.

When the CLAMPZ instruction is loaded into the instruction register, all outputs are placed in an inactive drive state, but the data supplied to the disabled output drivers is derived from the boundary-scan cells.

The purpose of this instruction is to ensure, during production testing, that each output driver can be disabled when its data input is either a 0 or a 1.

A guarding pattern (specified for this device at the end of this section) should be pre-loaded into the boundary-scan register using the SAMPLE/PRELOAD instruction prior to selecting the CLAMPZ instruction.

In the CAPTURE-DR state, a logic 0 is captured by the bypass register. In the SHIFT-DR state, test data is shifted into the bypass register via TDI and out via TDO after a delay of one TCK cycle. Note that the first bit shifted out will be a zero. The bypass register is not affected in the UPDATE-DR state.

Test Data Registers

Bypass Register

Purpose: This is a single bit register which can be selected as the path between TDI and TDO to allow the device to be bypassed during boundary-scan testing.

Length: 1 bit

Operating Mode: When the BYPASS instruction is the current instruction in the instruction register, serial data is transferred from TDI to TDO in the SHIFT-DR state with a delay of one TCK cycle.

There is no parallel output from the bypass register.

A logic 0 is loaded from the parallel input of the bypass register in the CAPTURE-DR state.

LH74610 Device Identification (ID) Code Register

Purpose: This register is used to read the 32-bit device identification code. No programmable supplementary identification code is provided.

Length: 32 bits

The format of the ID register is as follows:

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
version				part number																manufacturer identity											
00046061H																															

Operating Mode: When the IDCODE instruction is current, the ID register is selected as the serial path between TDI and TDO.

There is no parallel output from the ID register.

The 32-bit device identification code is loaded into the ID register from its parallel inputs during the CAPTURE-DR state.

LH74610 Boundary Scan (BS) Register

Purpose: The BS register consists of a serially connected set of cells around the periphery of the device, at the interface between the core logic and the system input/output pads. This register can be used to isolate the core logic from the pins and then apply tests to the core logic, or conversely to isolate the pins from the core logic and then drive or monitor the system pins.

Operating modes: The BS register is selected as the register to be connected between TDI and TDO only during the SAMPLE/PRELOAD, EXTEST and INTEST instructions. Values in the BS register are used, but are not changed, during the CLAMP and CLAMPZ instructions.

In the normal (system) mode of operation, straight-through connections between the core logic and pins are maintained and normal system operation is unaffected.

In TEST mode (i.e. when either EXTEST or INTEST is the currently selected instruction), values can be applied to the core logic or output pins independently of the actual values on the input pins and core logic outputs respectively. On the LH74610 all of the boundary scan cells include an update register and thus all of the pins can be controlled in the above manner. Additional boundary-scan cells are interposed in the scan chain in order to control the enabling of tristateable buses.

The correspondence between boundary-scan cells and system pins, system direction controls and system output enables is as shown on the following page. The cells are listed in the order in which they are connected in the boundary-scan register, starting with the cell closest to nTDI. All boundary-scan register cells at input pins can apply tests to the on-chip core logic. The EXTEST guard values specified in the table below should be clocked into the boundary-scan register (using the SAMPLE/PRELOAD instruction) before the EXTEST instruction is selected, to ensure that known data is applied to the core logic during the test. The INTEST guard values shown in the table below should be clocked into the boundary-scan register (using the SAMPLE/PRELOAD instruction) before the INTEST instruction is selected to ensure that all outputs are disabled. These guard values should also be used when new EXTEST or INTEST vectors are clocked into the boundary-scan register.

Output Enable Boundary-Scan Cells

The boundary-scan register cells Nendout, Nabe, Ntbe, and Nmse control the output drivers of tristate outputs as shown in the table below. In the case of OUTEN0 enable cells (Nendout, Ntbe), loading a 1 into the cell will place the associated drivers into the tristate state, while in the case of type INEN1 enable cells (Nabe, Nmse), loading a 0 into the cell will tristate the associated drivers.

To put all LH74610 tristate outputs into their high impedance state, a logic 1 should be clocked into the output enable boundary-scan cells Nendout and Ntbe, and a logic 0 should be clocked into Nabe and Nmse. Alternatively, the HIGHZ instruction can be used.

If the on-chip core logic causes the drivers controlled by Nendout, for example, to be tristate, (i.e. by driving the signal Nendout HIGH), then a 1 will be observed on this cell if the SAMPLE/PRELOAD or INTEST instructions are active.

Single-Step Operation

The LH74610 is a static design and there is no minimum clock speed. It can therefore be single-stepped while the INTEST instruction is selected. This can be achieved by serialising a parallel stimulus and clocking the resulting serial vectors into the boundary-scan register. When the boundary-scan register is updated, new test stimuli are applied to the core logic inputs; the effect of these stimuli can then be observed on the core logic outputs by capturing them in the boundary-scan register.

NO.	CELL NAME	PIN	TYPE	OUTPUT ENABLE BS CELL	INTEST	EXTEST/ CLAMP	NO.	CELL NAME	PIN	TYPE	OUTPUT ENABLE BS CELL	INTEST	EXTEST/ CLAMP
from tdi													
guard value													
1	din23	D[23]	IN	-			69	testin[4]	TESTIN[4]	IN	-		0
2	dout2	D[23]	OUT	Nendout			70	testin[3]	TESTIN[3]	IN	-		0
3	din22	D[22]	IN	-			71	testin[2]	TESTIN[2]	IN	-		0
4	dout22	D[22]	OUT	Nendout			72	testin[1]	TESTIN[1]	IN	-		0
5	din21	D[21]	IN	-			73	testin[0]	TESTIN[0]	IN	-		0
6	dout21	D[21]	OUT	Nendout			74	Ntbe		-	OUTEN0	-	1
7	din20	D[20]	IN	-			75	ale	ALE	IN	-		
8	dout20	D[20]	OUT	Nendout			76	a31	A[31]	OUT	Nabe		
9	din19	D[19]	IN	-			77	a30	A[30]	OUT	Nabe		
10	dout19	D[19]	OUT	Nendout			78	a29	A[29]	OUT	Nabe		
11	din18	D[18]	IN	-			79	a28	A[28]	OUT	Nabe		
12	dout18	D[18]	OUT	Nendout			80	a27	A[27]	OUT	Nabe		
13	din17	D[17]	IN	-			81	a26	A[26]	OUT	Nabe		
14	dout17	D[17]	OUT	Nendout			82	a25	A[25]	OUT	Nabe		
15	din16	D[16]	IN	-			83	a24	A[24]	OUT	Nabe		
16	dout16	D[16]	OUT	Nendout			84	a23	A[23]	OUT	Nabe		
17	din15	D[15]	IN	-			85	a22	A[22]	OUT	Nabe		
18	dout15	D[15]	OUT	Nendout			86	a21	A[21]	OUT	Nabe		
19	din14	D[14]	IN	-			87	a20	A[20]	OUT	Nabe		
20	dout14	D[14]	OUT	Nendout			88	a19	A[19]	OUT	Nabe		
21	din13	D[13]	IN	-			89	a18	A[18]	OUT	Nabe		
22	dout13	D[13]	OUT	Nendout			90	a17	A[17]	OUT	Nabe		
23	din12	D[12]	IN	-			91	a16	A[16]	OUT	Nabe		
24	dout12	D[12]	OUT	Nendout			92	a15	A[15]	OUT	Nabe		
25	din11	D[11]	IN	-			93	a14	A[14]	OUT	Nabe		
26	dout11	D[11]	OUT	Nendout			94	a13	A[13]	OUT	Nabe		
27	din10	D[10]	IN	-			95	a12	A[12]	OUT	Nabe		
28	dout10	D[10]	OUT	Nendout			96	a11	A[11]	OUT	Nabe		
29	din9	D[9]	IN	-			97	a10	A[10]	OUT	Nabe		
30	dout9	D[9]	OUT	Nendout			98	a09	A[09]	OUT	Nabe		
31	Nendout	-	OUTEN0	-	1		99	a08	A[08]	OUT	Nabe		
32	din8	D[8]	IN	-			100	a07	A[07]	OUT	Nabe		
33	dout8	D[8]	OUT	Nendout			101	a06	A[06]	OUT	Nabe		
							102	a05	A[05]	OUT	Nabe		

NO.	CELL NAME	PIN	TYPE	OUTPUT ENABLE BS CELL	INTEST	EXTEST/ CLAMP	NO.	CELL NAME	PIN	TYPE	OUTPUT ENABLE BS CELL	INTEST	EXTEST/ CLAMP
34	din7	D[7]	IN	-			103	a04	A[04]	OUT	Nabe		
35	dout7	D[7]	OUT	Nendout			104	a03	A[03]	OUT	Nabe		
36	din6	D[6]	IN	-			105	a02	A[02]	OUT	Nabe		
37	dout6	D[6]	OUT	Nendout			106	a01	A[01]	OUT	Nabe		
38	din5	D[5]	IN	-			107	a00	A[00]	OUT	Nabe		
39	dout5	D[5]	OUT	Nendout			108	Nabe	ABE	INEN1	-	0	
40	din4	D[4]	IN	-			109	rlw	LOCK	OUT	Nabe		
41	dout4	D[4]	OUT	Nendout			110	Nbw	nBW	OUT	Nabe		
42	din3	D[3]	IN	-			111	Nrw	nRW	OUT	Nabe		
43	dout3	D[3]	OUT	Nendout			112	testin[15]	TESTIN[15]	IN	-	0	
44	din2	D[2]	IN	-			113	testin[14]	TESTIN[14]	IN	-	0	
45	dout2	D[2]	OUT	Nendout			114	testin[13]	TESTIN[13]	IN	-	0	
46	din1	D[1]	IN	-			115	testin[12]	TESTIN[12]	IN	-	0	
47	dout1	D[1]	OUT	Nendout			116	testin[11]	TESTIN[11]	IN	-	0	
48	din0	D[0]	IN	-			117	testin[10]	TESTIN[10]	IN	-	0	
49	dout0	D[0]	OUT	Nendout			118	testin[09]	TESTIN[09]	IN	-	0	
50	dbe	DBE	IN	-			119	testin[08]	TESTIN[08]	IN	-	0	
51	seq	SEQ	OUT	Nmse			120	din31	D[31]	IN	-		
52	Nimreq	nMREQ	OUT	Nmse			121	dout31	D[31]	OUT	Nendout		
53	Nmse	MSE	INEN1	-	0		122	din30	D[30]	IN	-		
54	sNa	SnA	IN	-			123	dout30	D[30]	OUT	Nendout		
55	Nwait	nWAIT	IN	-			124	din29	D[29]	IN	-		
56	mclk	MCLK	IN	-	0		125	dout29	D[29]	OUT	Nendout		
57	fcik	FCLK	IN	-	0		126	din28	D[28]	IN	-		
58	abort	ABORT	IN	-			127	dout28	D[28]	OUT	Nendout		
59	Nreset	nRESET	IN	-			128	din27	D[27]	IN	-		
60	testin[16]	TESTIN[16]	IN	-	0		129	dout27	D[27]	OUT	Nendout		
61	testout[2]	TESTOUT[2]	OUT	Nitbe			130	din26	D[26]	IN	-		
62	testout[1]	TESTOUT[1]	OUT	Nitbe			131	dout26	D[26]	OUT	Nendout		
63	testout[0]	TESTOUT[0]	OUT	Nitbe			132	din25	D[25]	IN	-		
64	Nirq	nIRQ	IN	-			133	dout25	D[25]	OUT	Nendout		
65	Nfiq	nFIQ	IN	-			134	din24	D[24]	IN	-		
67	testin[6]	TESTIN[6]	IN	-	0		135	dout24	D[24]	OUT	Nendout		
68	testin[5]	TESTIN[5]	IN	-	0								to TDO

Type Key
 IN Input pad OUT Output pad INEN1 Input enable active high OUTENO Output enable active low

DC Parameters

Absolute Maximum Ratings

Table 13. LH74610 DC Maximum Ratings

SYMBOL	PARAMETER	MIN	MAX	UNITS	NOTE
VDD	Supply voltage	VSS-0.3	VSS+6.0	V	1
Vip	Voltage applied to any pin	VSS-0.3	VDD+0.3	V	1
Ts	Storage temperature	-40	125	deg C	1

NOTE:

These are stress ratings only. Exceeding the absolute maximum ratings may permanently damage the device. Operating the device at absolute maximum ratings for extended periods may affect device reliability.

DC Operating Conditions

Table 14. LH74610 DC Operating Conditions

SYMBOL	PARAMETER	MIN	TYP	MAX	UNITS	NOTES
VDD	Supply voltage	4.50	5.0	5.50	V	
Vihc	IC input HIGH voltage	3.5		VDD	V	1,2
Vilc	IC input LOW voltage	0.0		1.5	V	1,2
Viht	IT/ITP input HIGH voltage	2.4		VDD	V	1,3,4
Vilt	IT/ITP input LOW voltage	0.0		0.8	V	1,3,4
Ta	Ambient operating temperature	-10		70	deg.C	

NOTES:

Voltages measured with respect to VSS.

IC - CMOS-level inputs

IT - TTL-level inputs (includes IT and ITOTZ pin types)

DC Characteristics

NOTE: Nominal values shown are derived from transient analysis simulations.

Table 15. LH74610 DC Characteristics

SYMBOL	PARAMETER	MIN	TYP	MAX	UNIT
IDD	Operating current			110	mA
	Standby current			20	μ A
Iin	IT input leakage current			± 1	μ A
Iinp	ITP input leakage current			± 5	μ A
Vol	Output HIGH voltage: Ioh = -8 mA				V
	Output HIGH voltage: Ioh = -4 mA	4.1			V
Voh	Output LOW voltage: Iol = 8 mA			0.4	V
Vihk	CK input HIGH voltage	VDD - 0.4			V
Vilk	CK input LOW voltage			0.4	V
Viht	IT input HIGH voltage	2.2			V
Vilt	IT input LOW voltage			0.8	V
Vohc	OC output HIGH voltage	4.1		VDD	V
Volc	OC output LOW voltage	0.0		0.4	V
Cin	Input capacitance			15	pF
Co	Output capacitance			20	pF
Cio	I/O capacitance			20	pF

AC Parameters

Test Conditions

The AC timing diagrams presented in this section assume that the outputs of the LH74610 have been loaded with the capacitive loads shown in the 'Test Load' column of the table below; these loads have been chosen as typical of the system in which the LH74610 might be employed. The output pads of the LH74610 are CMOS drivers which exhibit a propagation delay that increases linearly with the increase in load capacitance. An 'Output derating' figure is given for each output pad, showing the approximate rate of increase of output time with increasing load capacitance.

Table 16. LH74610 AC Test Conditions

OUTPUT SIGNAL	TEST LOAD (pF)	OUTPUT DERATING (ns/pF)
A[25:0]	50	0.072
D[31:0]	50	0.072
nR/W	50	0.072
nB/W	50	0.072
LOCK	50	0.072
nMREQ	50	0.072
SEQ	50	0.072

Relationship Between FCLK & MCLK




Table 17. LH74610 fclk and mclk Relationship

SYMBOL	PARAMETER	MIN	TYP	MAX	UNIT	NOTE
T_{fckl}	FCLK LOW time	23			ns	1
T_{fckh}	FCLK HIGH time	23			ns	1
T_{fmh}	FCLK - MCLK hold time	25			ns	
T_{mfs}	MCLK - FCLK setup	3			ns	

NOTE:

1. FCLK timings measured at 50% of Vdd.

Main Bus Signals

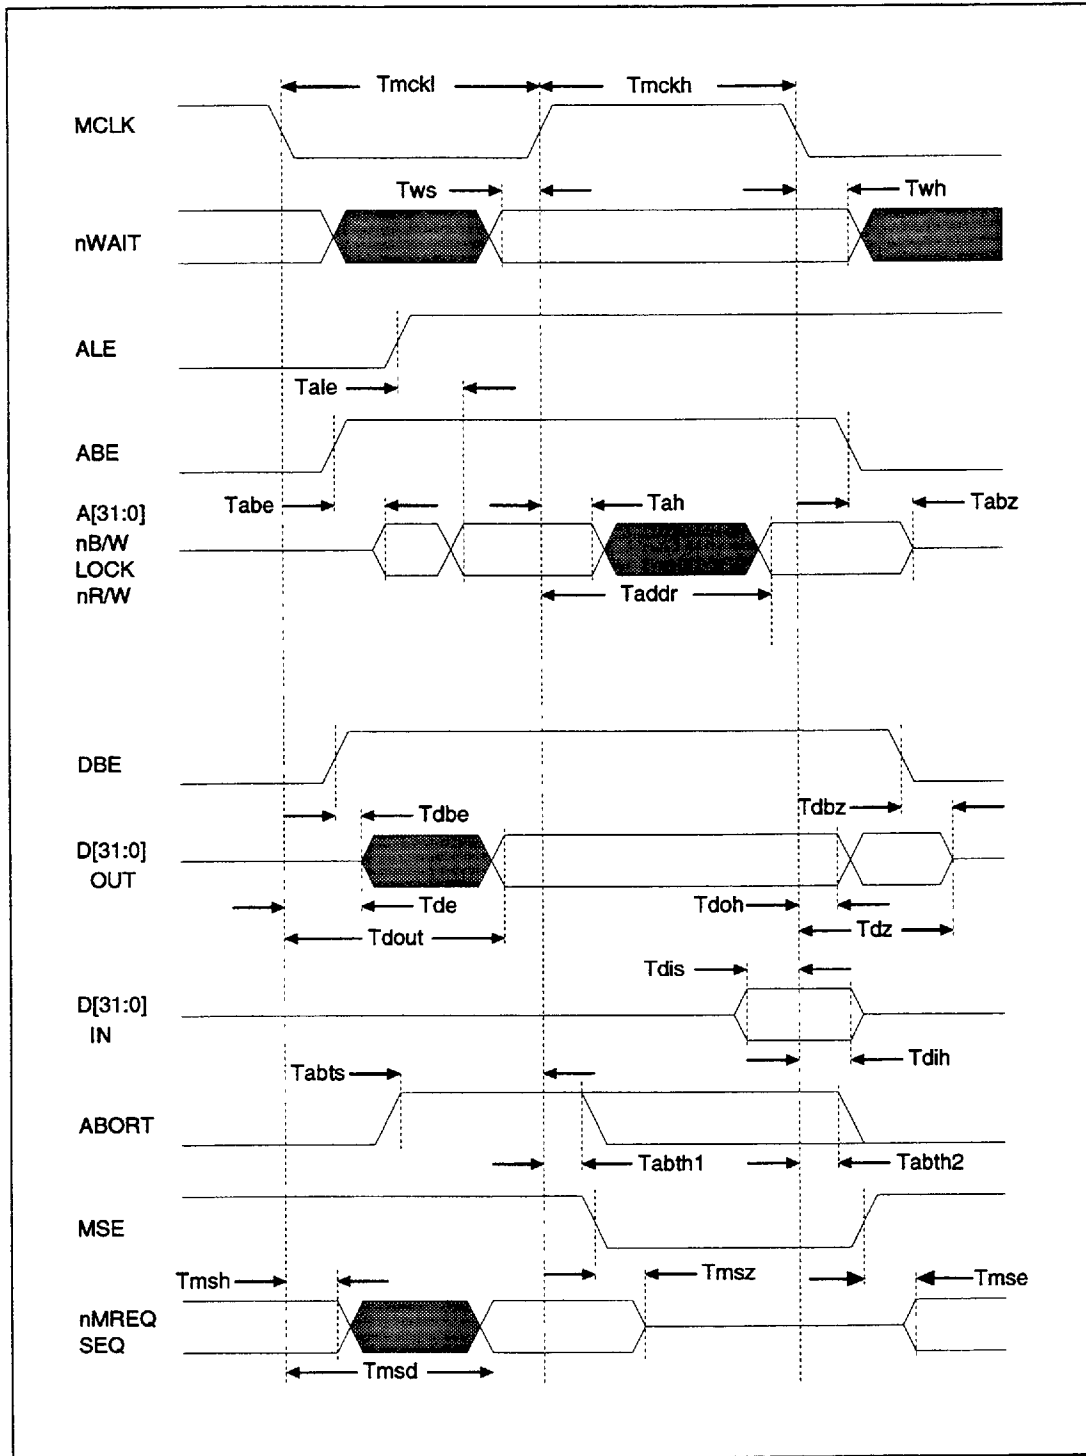


Table 18. LH74610 Bus Timing

SYMBOL	PARAMETER	MIN	TYP	MAX	UNIT	NOTE
Tmckl	MCLK LOW time	30			ns	1
Tmckh	MCLK HIGH time	30			ns	
Tws	nWAIT setup to MCLK	5			ns	
Twh	nWAIT hold from MCLK	5			ns	
Tale	address latch open			18	ns	2
Tabe	address bus enable			15	ns	2
Tabz	address bus disable			25	ns	
Taddr	MCLK to address delay			28	ns	2
Tah	address hold time	5			ns	2
Tdbe	DBE to data enable			22	ns	2
Tde	MCLK to data enable	10			ns	2
Tdbz	DBE to data disable			25	ns	
Tdz	MCLK to data disable			25	ns	
Tdout	data out delay			40	ns	2
Tdoh	data out hold	5			ns	2
Tdis	data in setup	2			ns	
Tdih	data in hold	10			ns	
Tabts	ABORT setup time	10			ns	
Tabth1	ABORT hold time	5			ns	3
Tabth2	ABORT hold time	5			ns	3
Tmse	nMREQ & SEQ enable			10	ns	
Tmsz	nMREQ & SEQ disable			20	ns	
Tmsd	nMREQ & SEQ delay			38	ns	
Tmsh	nMREQ & SEQ hold	5			ns	

NOTES:

1. MCLK timings measured between clock edges at 50% of Vdd.
2. The timings of these buses are measured to TTL levels.
3. Tabth1 is a requirement for LH74610. To ensure compatibility with future processors, designs should meet Tabth2. Tabth2 is not tested on the LH74610.


Table 19. LH74610 Boundary Scan Interface Timing

SYMBOL	PARAMETER	MIN	TYP	MAX	UNITS	NOTES
Tbscl	tck low period	50			ns	1
Tbsch	tck high period	50			ns	1
Tbsis	tdi,tms setup to [TCr]	10			ns	
Tbsih	tdi,tms hold from [TCr]	10			ns	
Tbsod	TCf to tdo valid			40	ns	2
Tbsoh	tdo hold time	5			ns	2
Tbsoe	tdo enable time	5			ns	2,3
Tbsod	tdo disable time			40	ns	2,4
Tbsss	I/O signal setup to [TCr]	5			ns	5
Tbssh	I/O signal hold from [TCr]	20			ns	5
Tbsdd	TCf to data output valid			40	ns	
Tbsdh	data output hold time	5			ns	6
Tbsde	data output enable time	5			ns	6,7
Tbsdd	data output disable time			40	ns	6,8
Tbsr	Reset period	30			ns	
Tbsrs	tms setup to [TRr]	10			ns	9
Tbsrh	tms hold from [TRr]	10			ns	9


NOTES:

- tck may be stopped indefinitely in either the low or high phase.
- Assumes a 25pF load on tdo. Output timing derates at 0.072ns/pF of extra load applied.
- tdo enable time applies when the TAP controller enters the Shift-DR or Shift-IR states.
- tdo disable time applies when the TAP controller leaves the Shift-DR or Shift-IR states.
- For correct data latching, the I/O signals (from the core and the pads) must be setup and held with respect to the rising edge of tck in the CAPTURE-DR state of the SAMPLE/PRELOAD, INTEST and EXTEST instructions.
- Assumes that the data outputs are loaded with the AC test loads (see AC parameter specification).
- Data output enable time applies when the boundary scan logic is used to enable the output drivers.
- Data output disable time applies when the boundary scan logic is used to disable the output drivers.
- The tms input must be held high as Ntrst is taken high at the end of the boundary-scan reset sequence.

Physical Details



LH74610 PQFP Package = 0.5 mm Lead Pitch



Pinout

PIN	SIGNAL	TYPE	PIN	SIGNAL	TYPE	PIN	SIGNAL	TYPE	PIN	SIGNAL	TYPE
1	MSE	i	37	D[24]	i/o	73	LOCK	o	109	A[26]	o
2	nMREQ	o	38	D[25]	i/o	74	ABE	i	110	A[27]	o
3	SEQ	o	39	D[26]	i/o	75	A[0]	o	111	A[28]	o
4	DBE	i	40	Vss1	-	76	A[1]	o	112	Vdd2	-
5	Vss2	-	41	Vss2	-	77	A[2]	o	133	Vss2	-
6	Vdd2	-	42	Vdd2	-	78	Vss2	-	114	A[29]	o
7	D[0]	i/o	43	D[27]	i/o	79	Vdd2	-	115	A[30]	o
8	D[1]	i/o	44	D[28]	i/o	80	A[3]	o	116	A[31]	o
9	D[2]	i/o	45	D[29]	i/o	81	A[4]	o	117	ALE	i
10	D[3]	i/o	46	D[30]	i/o	82	A[5]	o	118	n/c	
11	D[4]	i/o	47	D[31]	i/o	83	A[6]	o	119	n/c	
12	D[5]	i/o	48	TDO	o	84	A[7]	o	120	n/c	
13	D[6]	i/o	49	TDI	i	85	A[8]	o	121	Vss1	-
14	D[7]	i/o	50	nTRST	i	86	A[9]	o	122	Vdd1	-
15	D[8]	i/o	51	Vdd1	-	87	A[10]	o	123	TESTIN[7]	i
16	Vss2	-	52	TMS	i	88	A[11]	o	124	TESTIN[6]	i
17	Vdd2	-	53	TCK	i	89	A[12]	o	125	TESTIN[5]	i
18	Vss1	-	54	n/c	-	90	Vdd2	-	126	TESTIN[4]	i
19	Vdd1	-	55	n/c	-	91	Vss1	-	127	TESTIN[3]	i
20	D[9]	i/o	56	n/c	-	92	Vdd1	-	128	TESTIN[2]	i
21	D[10]	i/o	57	n/c	-	93	Vss2	-	129	TESTIN[1]	i
22	D[11]	i/o	58	n/c	-	94	A[13]	o	130	TESTIN[0]	i
23	D[12]	i/o	59	TESTIN[8]	i	95	A[14]	o	131	nFIQ	
24	D[13]	i/o	60	TESTIN[9]	i	96	A[15]	o	132	nIRQ	
25	D[14]	i/o	61	Vdd1	-	97	A[16]	o	133	TESTOUT[0]	o
26	D[15]	i/o	62	Vss1	-	98	A[17]	o	134	TESTOUT[1]	o
27	D[16]	i/o	63	TESTIN[10]	i	99	A[18]	o	135	TESTOUT[2]	o
28	D[17]	i/o	64	TESTIN[11]	i	100	A[19]		136	TESTIN[16] i	
29	D[18]	i/o	65	TESTIN[12]	i	101	A[20]	o	137	nRESET	i
30	D[19]	i/o	66	TESTIN[13]	i	102	Vdd2	-	138	ABORT	i
31	Vdd2	-	67	TESTIN[14]	i	103	Vss2	-	139	FCLK	i
32	Vss2	-	68	TESTIN[15]	i	104	A[21]	o	140	MCLK	i
33	D[20]	i/o	69	Vss2	-	105	A[22]	o	141	Vdd2	-
34	D[21]	i/o	70	Vdd2	-	106	A[23]	o	142	Vss2	-
35	D[22]	i/o	71	nR/W	o	107	A[24]	o	143	nWAIT	i
36	D[23]	i/o	72	nB/W	o	108	A[25]	o	144	SnA	i