
ARM Datasheet

Part No 1 85250 0360 0
Issue No 1.0
17 March 1987

© Copyright Acorn Computers Limited 1987

Neither the whole nor any part of the information contained in, or the product described in, this manual may be adapted or reproduced in any material form except with the prior written permission of the copyright holder.

The product described in this manual is subject to continuous developments and improvements. All particulars of the product and its use contained in this manual are given by Acorn Computers in good faith. However, all warranties implied or expressed, including but not limited to implied warranties or merchantability or fitness for purpose, are excluded.

This manual is intended only to assist the reader in the use of the product. Acorn Computers shall not be liable for any loss or damage arising from the use of any information in this manual, or any error or omission in such information, or any incorrect use of the product.

ISBN 1 85250 026 3

Published by:

Acorn Computers Limited, Fulbourn Road, Cherry Hinton, Cambridge CB1 4JN, UK

Contents

1. Introduction	1
2. Block Diagram	2
3. Functional Diagram	3
4. Description of Signals	4
5. Programmers' Model	7
5.1 Introduction	7
5.2 Registers	7
5.3 Exceptions	7
5.3.1 FIQ	9
5.3.2 IRQ	9
5.3.3 Address exception trap	10
5.3.4 Abort	10
5.3.5 Software interrupt	11
5.3.6 Undefined instruction trap	11
5.3.7 Reset	12
5.3.8 Vector Summary	12
5.3.9 Exception Priorities	12
5.3.10 Interrupt Latencies	13
6. Instruction Set	14
6.1 The condition field	14
6.2 Branch and branch with link (B, BL)	15
6.2.1 The link bit	15
6.2.2 Assembler syntax	15
6.2.3 Examples	16
6.3 Data processing	17
6.3.1 Operations	18
6.3.2 PSR flags	18
6.3.3 Shifts	18
Instruction specified shift amount	19
Register specified shift amount	20
6.3.4 Immediate operand rotates	21
6.3.5 Writing to R15	21
6.3.6 Using R15 as an operand	21
6.3.7 Assembler syntax	21
6.3.8 Examples	22
6.4 Multiply and multiply-accumulate (MUL, MLA)	23
6.4.1 Operand restrictions	23
6.4.2 PSR flags	23
6.4.3 Writing to R15	24
6.4.4 Using R15 as an operand	24
6.4.5 Assembler syntax	24
6.4.6 Examples	24
6.5 Single data transfer (LDR, STR)	25
6.5.1 Offsets and auto-indexing	25
6.5.2 Shifted register offset	26
6.5.3 Bytes and words	26
6.5.4 Use of R15	26
6.5.5 Address exceptions	26
6.5.6 Data Aborts	27

6.5.7 Assembler syntax	27
6.5.8 Examples	28
6.6 Block data transfer (LDM, STM)	29
6.6.1 The register list	29
6.6.2 Addressing modes	29
Post-increment addressing	30
Pre-increment addressing	30
Post-decrement addressing	31
Pre-decrement addressing	31
6.6.3 Transfer of R15	32
6.6.4 Forcing transfer of the user bank	32
6.6.5 Use of R15 as the base	32
6.6.6 Inclusion of the base in the register list	32
6.6.7 Address exceptions	32
6.6.8 Data Aborts	32
Aborts during STM instructions	33
Aborts during LDM instructions	33
6.6.9 Assembler syntax	33
Addressing mode names	33
6.6.10 Examples	34
6.7 Software interrupt	35
6.7.1 Return from the supervisor	35
6.7.2 Comment field	35
6.7.3 Assembler syntax	35
6.7.4 Examples	36
6.8 Co-Processor data operations	37
6.8.1 The Co-Processor fields	37
6.8.2 Assembler syntax	37
6.8.3 Examples	38
6.9 Co-Processor data transfers	39
6.9.1 The Co-Processor fields	39
6.9.2 Addressing modes	39
6.9.3 Use of R15	40
6.9.4 Forcing address translation	40
6.9.5 Address exceptions	40
6.9.6 Data aborts	40
6.9.7 Assembler syntax	40
6.9.8 Examples	41
6.10 Co-Processor register transfers	42
6.10.1 The Co-Processor fields	42
6.10.2 Transfers to R15	42
6.10.3 Transfers from R15	42
6.10.4 Assembler syntax	43
6.10.5 Examples	43
6.11 Undefined instructions	44
6.11.1 Assembler syntax	44
6.12 Instruction set summary	45
6.13 Instruction set examples	46
6.13.1 Using the conditional instructions	46
6.13.2 Pseudo random binary sequence generator	46
6.13.3 Multiplication by constant using the barrel shifter	47
6.13.4 Loading a word from an unknown alignment	48
6.13.5 Sign/zero extension of a half word	48

6.13.6 Return setting condition codes	48
7. Memory Interface	49
7.1 Cycle types	49
7.2 Byte addressing	51
7.3 Address timing	51
7.4 Memory management	52
7.5 Use of MEMC	52
8. Co-Processor Interface	53
8.1 Interface signals	53
8.1.1 Co-Processor present/absent	53
8.1.2 Busy-waiting	53
8.1.3 Pipeline following	53
8.2 Data transfer cycles	53
8.3 Register transfer cycle	54
8.4 Privileged instructions	54
8.5 Idempotency	54
8.6 Undefined instructions	55
8.7 Use of MEMC	55
9. Instruction Cycle Operations	56
9.1 Branch and branch with link	56
9.2 Data operations	56
9.3 Multiply and multiply accumulate	57
9.4 Load register	58
9.5 Store register	59
9.6 Store multiple registers	59
9.7 Load multiple registers	60
9.8 Software interrupt and exception entry	61
9.9 Co-Processor data operation	61
9.10 Co-Processor data transfer (from memory to Co-Processor)	62
9.11 Co-Processor data transfer (from Co-Processor to memory)	63
9.12 Co-Processor register transfer (Load from Co-Processor)	63
9.13 Co-Processor register transfer (Store to Co-Processor)	64
9.14 Undefined instructions and Co-Processor absent	64
9.15 Unexecuted instructions	64
9.16 Instruction speeds	65
10. DC Parameters	66
10.1 Absolute Maximum Ratings	66
10.2 DC Operating Conditions	67
10.3 DC Characteristics	68
11. AC Parameters	69
12. Packaging	73
13. Compatibility with Prototype ARMs	75
13.1 Plug-in compatibility	75
13.2 Bug fixes	75
13.3 Design differences	75

1. Introduction

The ARM (Acorn RISC Machine) is a general purpose 32-bit single-chip microprocessor. The architecture is based on Reduced Instruction Set Computer (RISC) principles, and the instruction set and related decode mechanism are greatly simplified compared with microprogrammed Complex Instruction Set Computers. This simplification results in a high instruction throughput and a good real-time interrupt response from a small and cost-effective chip.

The instruction set comprises nine basic instruction types. Two of these make use of the on-chip arithmetic logic unit (ALU), barrel shifter and multiplier to perform high-speed operations on the data in a bank of 27 registers, each 32 bits wide. Two instruction types control the transfer of data between main memory and the register bank, one optimised for flexibility of addressing and the other for rapid context switching. Two instructions control the flow and privilege level of execution, and the remaining 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.

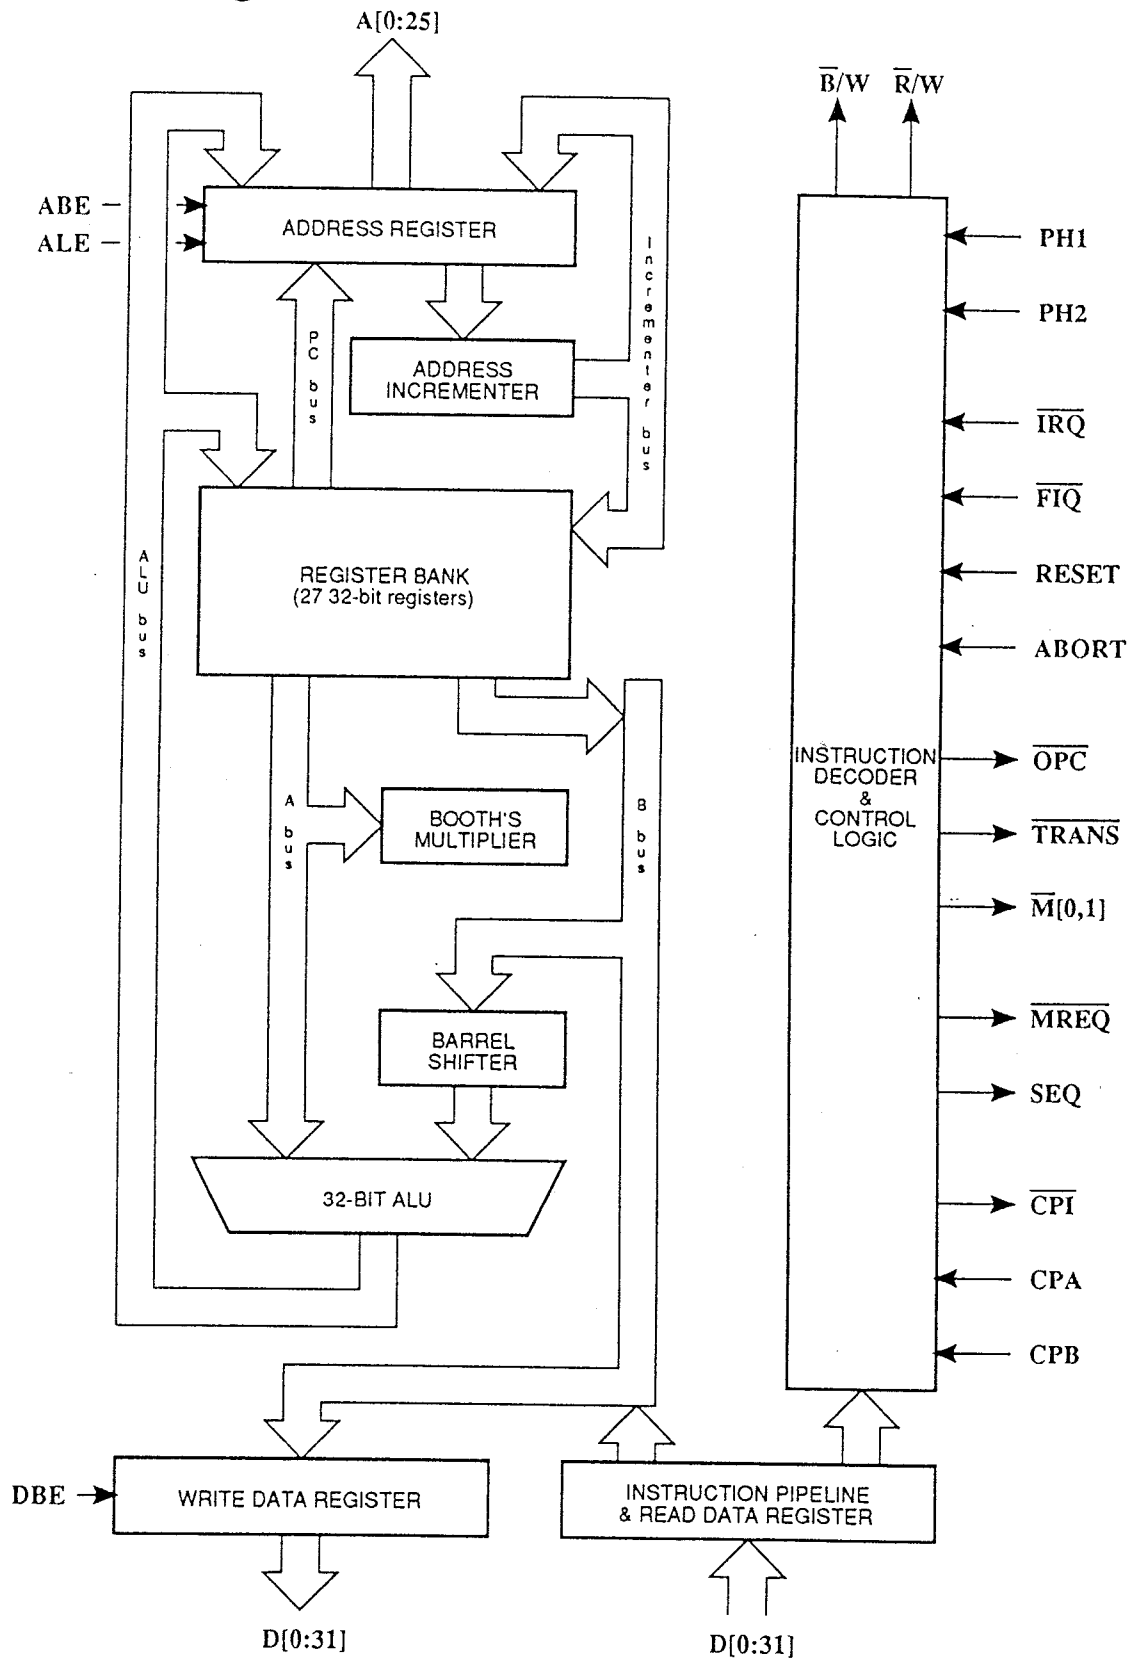
Pipelining is employed so that all parts of the processing and memory systems can operate continuously. Typically, while one instruction is being executed, its successor is being decoded, and a third instruction is being fetched from memory.

The memory interface has been designed to allow the performance potential to be realised 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 facilitate the exploitation of the fast local access modes offered by industry standard dynamic random access memories (DRAMs).

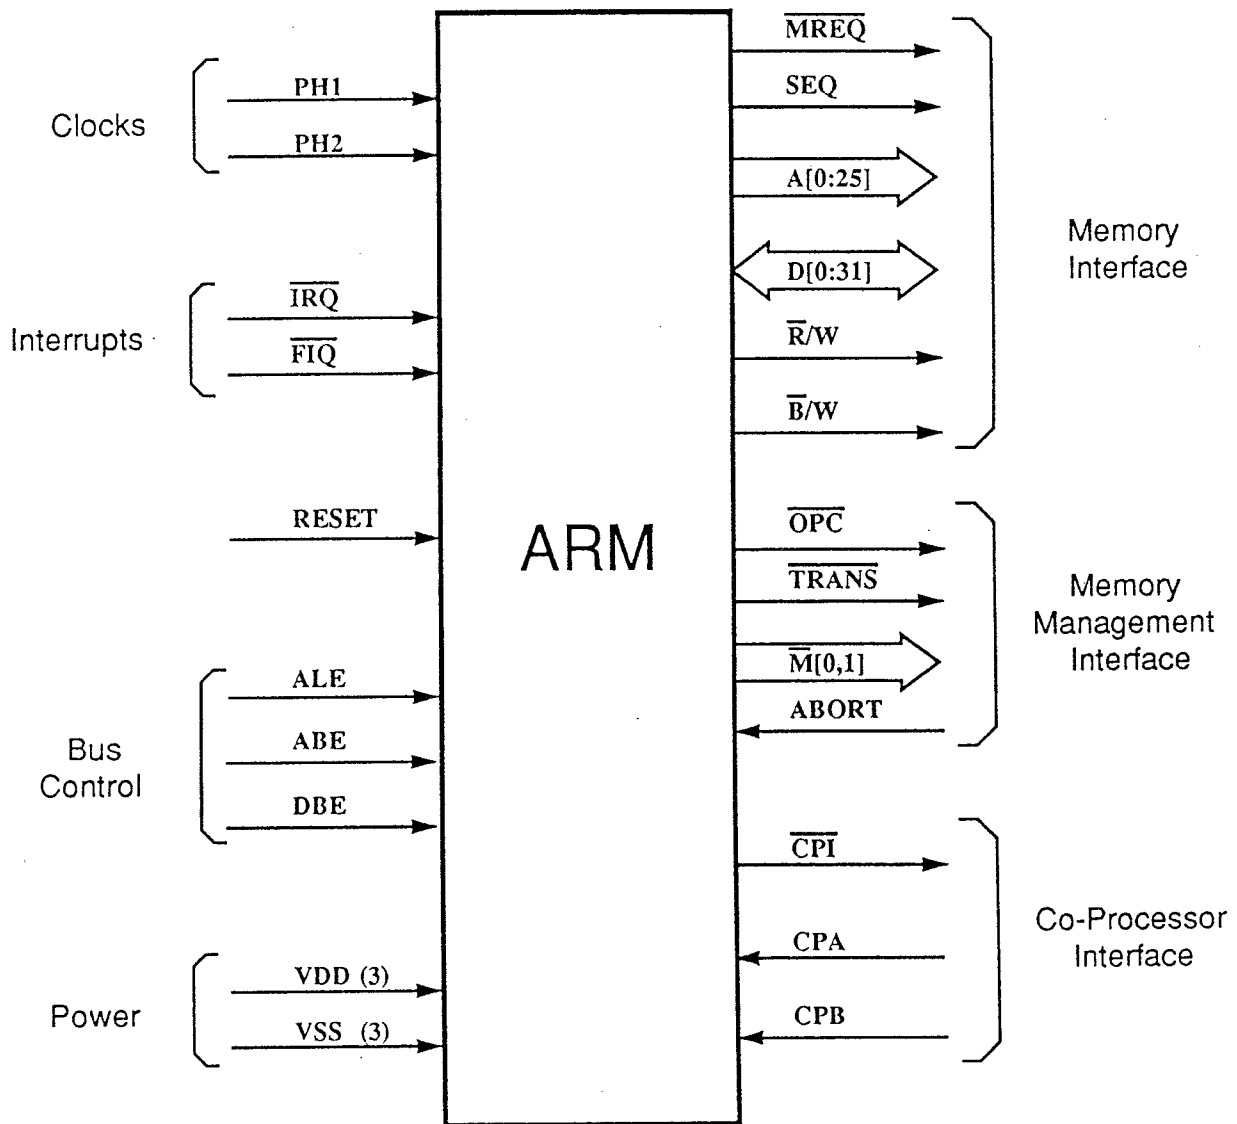
FEATURES

- * 32-bit data bus
- * 26-bit address bus giving a 64-MByte uniform address space
- * Support for virtual memory systems
- * Simple but powerful instruction set
- * Co-Processor interface for instruction set extension
- * Good high-level language compiler support
- * Peak execution rate of 10 million instructions per second (MIPS)
- * Fast interrupt response for real-time applications
- * Low power consumption (0.1 W typical) with a single +5 V supply
- * 84-pin JEDEC B leadless chip carrier or plastic leaded chip carrier

2. Block Diagram



3. Functional Diagram



4. Description of Signals

Name	Pin	Type	Description
PH2	1	ICk	Phase two clock.
PH1	2	ICk	Phase one clock.
\overline{R}/W	3	OC	Not read / write. When HIGH this signal indicates a processor write cycle; when LOW, a read cycle. It becomes valid during phase 2 of the cycle before that to which it refers, and remains valid to the end of phase 1 of the referenced cycle.
\overline{OPC}	4	OC	Not op-code fetch. When LOW this signal indicates that the processor is fetching an instruction from memory; when HIGH data (if anything) is being transferred. The signal becomes valid during phase 2 of the previous cycle, remaining valid through phase 1 of the referenced cycle.
\overline{MREQ}	5	OC	Not memory request. This signal, when LOW, indicates that the processor requires memory access during the following cycle. The signal becomes valid during phase 1, remaining valid through phase 2 of the cycle preceding that to which it refers.
ABORT	6	IT	Memory abort. This is an input which allows the memory system to tell the processor that a requested access is not allowed. The signal must be valid before the end of phase 1 of the cycle during which the memory transfer is attempted.
\overline{IRQ}	7	IT	Not interrupt request. This is an asynchronous interrupt request to the processor which causes it to be interrupted if taken LOW when the appropriate enable in the processor is active. The signal is level sensitive and must be held LOW until a suitable response is received from the processor.
\overline{FIQ}	8	IT	Not fast interrupt request. As \overline{IRQ} , but with higher priority. May be taken LOW asynchronously to interrupt the processor when the appropriate enable is active.
RESET	9	IT	Reset. This is a level sensitive input signal which is used to start the processor from a known address. A HIGH level will cause the instruction being executed to terminate abnormally. When RESET becomes LOW for at least one clock cycle, the processor will re-start from address 0. RESET must remain HIGH for at least two clock cycles, and during the HIGH period the processor will perform dummy instruction fetches with the address incrementing from the point where reset was activated. The address value will overflow to zero if RESET is held beyond the maximum address limit.
\overline{TRANS}	10	OC	Not memory translate. When this signal is LOW it indicates that the processor is in user mode, or that the supervisor is using a single transfer instruction with the force translate bit active. It may be used to tell memory management hardware when translation of the addresses should be turned on, or as an indicator of non-user mode activity.
VDD	11,32,55	PWR	Supply.

VSS	33,54,75	PWR	Supply.
$\overline{M}[1,0]$	13,14	OC	Not processor mode. These are output signals which are the inverses of the internal status bits indicating the processor operation mode.
SEQ	15	OC	Sequential address. This is an output signal. It will become HIGH when either: <ul style="list-style-type: none"> the address for the next cycle is being generated in the address incrementer, so will be equal to the present address (in bytes) plus 4, or during a cycle which did not use memory (\overline{MREQ} inactive), when the next cycle will use memory and the address will be the same as the current address. <p>The signal becomes valid during phase 1 and remains so through phase 2 of the cycle before the cycle whose address it anticipates. It may be used, in combination with the low-order address lines, to indicate that the next cycle can use a fast memory mode (for example DRAM page mode) and/or to by-pass the address translation system.</p>
ALE	16	IT	Address latch enable. This input to the processor is used to control transparent latches on the address outputs. Normally the addresses change during phase 2 to the value required during the next cycle, but for direct interfacing to ROMs they are required to be stable to the end of phase 2. Taking ALE LOW until the end of phase 2 will ensure that this happens. If the system does not require address lines to be held in this way, ALE may be held permanently HIGH. The ALE latch is dynamic, and ALE should not be held LOW indefinitely.
A[25:0]	17-31,34-44	OCZ	Addresses. This is the processor address bus. If ALE (address latch enable) is HIGH, the addresses become valid during phase 2 of the cycle before the one to which they refer and remain so during phase 1 of the referenced cycle. Their stable period may be controlled by ALE as described above.
ABE	45	IC	Address bus enable. This is an input signal which, when LOW, puts the address bus drivers into a high impedance state. ABE may be tied HIGH when there is no system requirement to turn off the address drivers.
D[0:31]	46-53,56-74, 77-81	IOTZ	Data Bus. These are bi-directional signal paths which are used for data transfers between the processor and external memory, as follows: <ul style="list-style-type: none"> during read cycles (when $\overline{R}/W = 0$), the input data must be valid before the end of phase 2 of the transfer cycle during write cycles (when $\overline{R}/W = 1$), the output data will become valid during phase 1 and remain so throughout phase 2 of the transfer cycle.
DBE	83	IT	Data bus enable. This is an input signal which, when LOW, forces data bus drivers into a high impedance state. (The drivers will always be high impedance except during write cycles, and DBE may be tied HIGH in systems which do not require the data bus for DMA or similar activities.)
\overline{B}/W	84	OC	Not byte / word. This is an output signal used by the processor to indicate to the external memory system when a data transfer of a byte

length is required. The signal is HIGH for word transfers and LOW for byte transfers and is valid for both read and write cycles. The signal will become valid during phase 2 of the cycle before the one during which the transfer will take place. It will remain stable throughout phase 1 of the transfer cycle.

$\overline{\text{CPI}}$	82	OC	Co-Processor instruction. When ARM executes a Co-Processor instruction, it will take this output LOW and wait for a response from the Co-Processor. The action taken will depend on this response, which the Co-Processor signals on the CPA and CPB inputs.
CPB	12	IT	Co-Processor busy. A Co-Processor which is capable of performing the operation which ARM is requesting (by asserting $\overline{\text{CPI}}$), but cannot commit to starting it immediately, should indicate this by letting CPB float HIGH. When the Co-Processor is ready to start it should take CPB LOW. ARM samples CPB at the end of phase 1 of the cycle when $\overline{\text{CPI}}$ is LOW.
CPA	76	IT	Co-Processor absent. A Co-Processor which is capable of performing the operation which ARM is requesting (by asserting $\overline{\text{CPI}}$) should take CPA LOW immediately. If CPA is HIGH at the end of phase 1 of the cycle when $\overline{\text{CPI}}$ is LOW, ARM will abort the Co-Processor handshake and take the undefined instruction trap. If CPA is LOW and remains LOW, ARM will busy-wait until CPB is LOW and then complete the Co-Processor instruction.

Key to Signal Types

ICK	Unbuffered clock inputs
IT	Input with TTL compatible levels
OC	Output with CMOS compatible levels
OCZ	3-state output with CMOS compatible levels
IOTZ	Bi-directional 3-state input/output with TTL compatible levels
PWR	Power pins

5. Programmers' Model

5.1 Introduction

ARM has a 32 bit data bus and a 26 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.

ARM supports four modes of operation, including protected supervisor and interrupt handling modes.

5.2 Registers

The processor has 27 32-bit registers, 16 of which are visible to the programmer at any time. The visible subset depends on the processor mode; special registers are switched in to support interrupt and supervisor processing. The register bank organisation is shown in figure 1.

User mode is the normal program execution state; registers R0-15 are directly accessible.

All registers are general purpose and may be used to hold data or address values, except that register R15 contains the Program Counter (PC) and the Processor Status Register (PSR). Special bits in some instructions allow the PC and PSR to be treated together or separately as required. Figure 2 shows the allocation of bits within R15.

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 and R14_fiq are used similarly to hold the return values of R15 when interrupts and exceptions arise, or when Branch and Link instructions are executed within supervisor or interrupt routines.

The FIQ processing state (described in the Exceptions section) has seven private registers mapped to R8-14 (R8_fiq-R14_fiq). Many FIQ programs will not need to save any registers.

The IRQ processing state has two private registers mapped to R13 and R14 (R13_irq and R14_irq).

Supervisor mode (entered on SWI instructions and other traps) has two private registers mapped to R13 and R14 (R13_svc and R14_svc).

The two private registers allow the IRQ and supervisor modes each to have a private stack pointer and link register. Supervisor and IRQ mode programs are expected to save the User state on their respective stacks and then use the User registers, remembering to restore the User state before returning.

In User mode only the N, Z, C and V bits of the PSR may be changed. The I, F and Mode flags will change only when an exception arises. In supervisor and interrupt modes all flags may be manipulated directly.

5.3 Exceptions

Exceptions arise whenever there is a need for the normal flow of program execution to be broken, so that (for instance) 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.

ARM handles exceptions by making use of the banked registers to save state. The old PC and PSR are copied into the appropriate R14, and the PC and processor mode bits are forced to a value which depends on the exception. Interrupt disable flags are set where required to prevent otherwise unmanageable nestings

user mode		svc mode	irq mode	fiq mode
R0				
R1				
R2				
R3				
R4				
R5				
R6				
R7				
R8			R8_fiq	
R9			R9_fiq	
R10			R10_fiq	
R11			R11_fiq	
R12			R12_fiq	
R13	R13_svc	R13_irq	R13_fiq	
R14	R14_svc	R14_irq	R14_fiq	
R15 (PC/PSR)				

Figure 1: Register Organisation

of exceptions. In the case of a re-entrant interrupt handler, R14 should be saved onto a stack in main memory before re-enabling the interrupt. When multiple exceptions arise simultaneously a fixed priority determines the order in which they are handled.

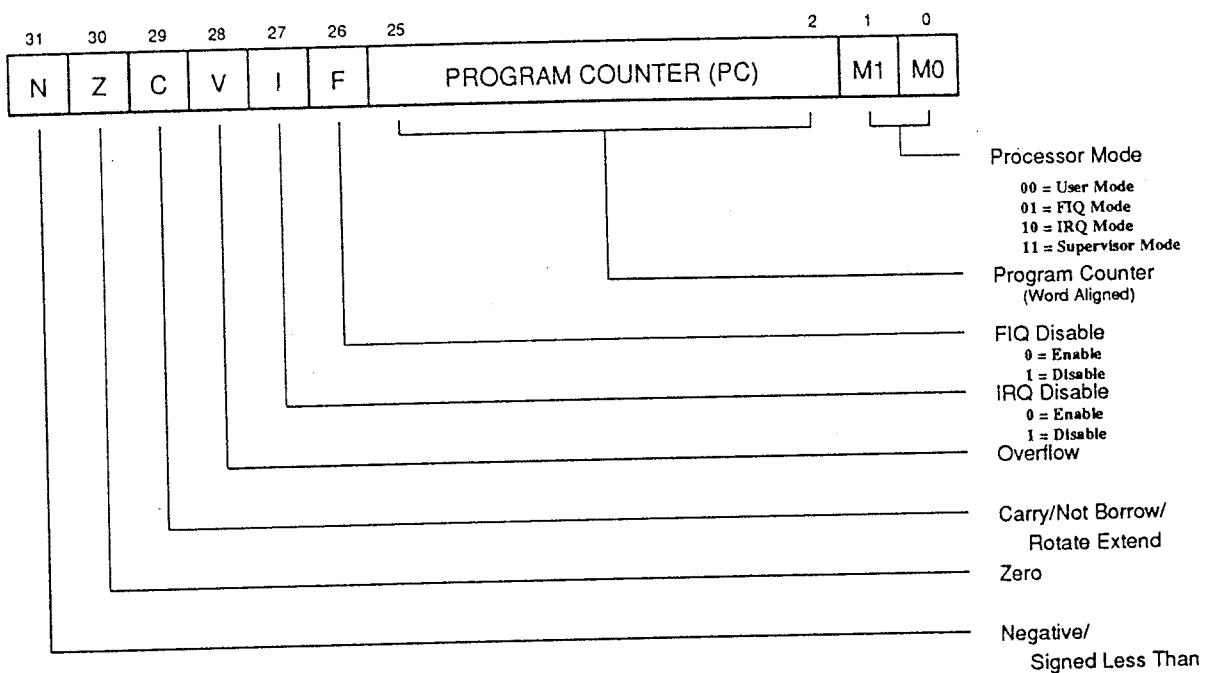


Figure 2: The Program Counter (PC) and Processor Status Register (PSR)

5.3.1 FIQ

The FIQ (Fast Interrupt reQuest) exception is externally generated by taking the $\overline{\text{FIQ}}$ pin LOW. This input can accept asynchronous transitions, and is delayed by one clock cycle for synchronisation 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, so that the overhead of context switching is minimised. The FIQ exception may be disabled by setting the F flag in the PSR (but note that this is not possible from user mode). If the F flag is clear ARM checks for a LOW level on the output of the FIQ synchroniser at the end of each instruction. When ARM is FIQed it will:

- (1) save R15 in R14_fiq;
- (2) force M0, M1 to FIQ mode and set the F and I bits in the PC word;
- (3) force the PC to fetch the next instruction from address 1CH.

To return normally from FIQ use SUBS PC,R14_fiq,#4. This will resume execution of the interrupted code sequence, and restore the original mode and interrupt enable state.

5.3.2 IRQ

The IRQ (Interrupt ReQuest) exception is a normal interrupt caused by a LOW level on the $\overline{\text{IRQ}}$ pin. 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 PC (but note that this is not possible from user mode). If the I flag is clear ARM checks for a LOW level on the output of the IRQ synchroniser at the end of each instruction. When successfully IRQed ARM will:

- (1) save R15 in R14_irq;
- (2) force M0, M1 to IRQ mode and set the I bit in the PC word;
- (3) force the PC to fetch the next instruction from address 18H.

To return normally from IRQ use SUBS PC,R14_irq,#4. This will restore the original processor state and thereby re-enable IRQ.

5.3.3 Address exception trap

An address exception arises whenever a data transfer is attempted with a calculated address above 3FFFFFFH. The ARM address bus is 26 bits wide, and an address calculation will have a 32-bit result. If this result has a logic "1" in any of the top 6 bits it is assumed that the address overflow is an error, and the address exception trap is taken.

Note that a branch cannot cause an address exception, and a block data transfer instruction which starts in the legal area but increments into the illegal area will not trap. The check is performed only on the address of the first word to be transferred.

When an address exception is seen ARM will:

- (1) if the data transfer was a store, force it to load. (This protects the memory from spurious writing.)
- (2) complete the instruction, but prevent internal state changes where possible. The state changes are the same as if the instruction had aborted on the data transfer.
- (3) save R15 in R14_svc;
- (4) force M0, M1 to supervisor mode and set the I bit in the PC word;
- (5) force the PC to fetch the next instruction from address 14H.

Normally an address exception is caused by erroneous code, and it is inappropriate to resume execution. If a return is required from this trap, use SUBS PC,R14_svc,#4. This will return to the instruction after the one causing the trap.

5.3.4 Abort

The Abort signal 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. ARM checks for an Abort at the end of the first phase of each bus cycle. When successfully Aborted ARM will respond in one of three ways:

- (i) if the abort occurred during an instruction prefetch (a *Prefetch Abort*), the prefetched instruction is marked as invalid; when it comes to execution, it is reinterpreted as below. (If the instruction is not executed, for example as a result of a branch being taken while it is in the pipeline, the abort will have no effect.)
- (ii) if the abort occurred during a data access (a *Data Abort*), the action depends on the instruction type. Data transfer instructions (LDR, STR) are aborted as though the instruction had not executed. The LDM and STM instructions complete, and if writeback 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.
- (iii) if the abort occurred during an internal cycle it is ignored.

Then, in cases (i) and (ii):

- (1) save R15 in R14_svc;
- (2) force M0, M1 to supervisor mode and set the I bit in the PC word;
- (3) force the PC to fetch the next instruction from address 0CH for Prefetch Abort, 10H for Data Abort.

To continue after a Prefetch Abort use SUBS PC,R14_svc,#4. This will attempt to re-execute the aborting instruction (which will only be effective if action has been taken to remove the cause of the original abort). A Data Abort requires any auto-indexing to be reversed before returning to re-execute the offending instruction, the return being done by SUBS PC,R14_svc,#8.

The abort mechanism allows a *demand paged virtual memory system* to be implemented when a suitable memory management unit (such as MEMC) 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.

5.3.5 Software interrupt

The software interrupt is used for getting into supervisor mode, usually to request a particular supervisor function. ARM will:

- (1) save R15 in R14_svc;
- (2) force M0, M1 to supervisor mode and set the I bit in the PC word;
- (3) force the PC to fetch the next instruction from address 08H.

To return from a SWI, use MOVS PC,R14_svc. This returns to the instruction following the SWI.

5.3.6 Undefined instruction trap

When ARM executes a Co-Processor instruction or an Undefined instruction, it offers it to any Co-Processors which may be present. If a Co-Processor can perform this instruction but is busy at that moment, ARM will wait until the Co-Processor is ready. If no Co-Processor can handle the instruction ARM will take the undefined instruction trap.

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

When the undefined instruction trap is taken ARM will:

- (1) save R15 in R14_svc;
- (2) force M0, M1 to supervisor mode and set the I bit in the PC word;
- (3) force the PC to fetch the next instruction from address 04H.

To return from this trap (after performing a suitable emulation of the required function), use MOVS PC,R14_svc. This will return to the instruction following the undefined instruction.

5.3.7 Reset

When Reset goes HIGH ARM will:

- (1) stop the currently executing instruction and start executing no-ops. When Reset goes LOW again it will;
- (2) save R15 in R14_svc;
- (3) force M0, M1 to supervisor mode and set the F and I bits in the PC word;
- (4) force the PC to fetch the next instruction from address 0H.

5.3.8 Vector Summary

Address	
00000000	Reset
00000004	Undefined instruction
00000008	Software interrupt
0000000C	Abort (prefetch)
00000010	Abort (data)
00000014	Address exception
00000018	IRQ
0000001C	FIQ

These are byte addresses, and will normally contain a branch instruction pointing to the relevant routine. The FIQ routine might reside at 000001CH onwards, and thereby avoid the need for (and execution time of) a branch instruction.

5.3.9 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) Address exception, 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. Address exception and data abort are mutually exclusive, since if an address is illegal the ARM will ignore the ABORT input. Undefined instruction and software interrupt are also mutually exclusive since they each correspond to particular (non-overlapping) decodings of the current instruction.

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

5.3.10 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 synchroniser ($T_{syncmax}$), plus the time for the longest instruction to complete (T_{ldm} , the longest instruction is load multiple registers), plus the time for address exception or data abort entry (T_{exc}), plus the time for FIQ entry (T_{fiq}). At the end of this time ARM will be executing the instruction at 1CH.

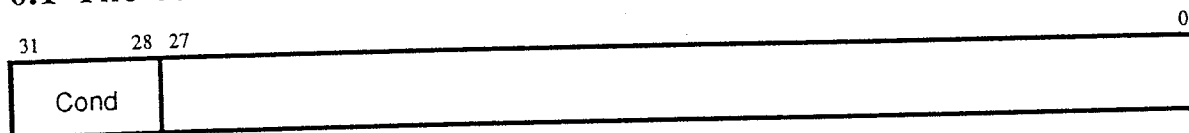
$T_{syncmax}$ is 2.5 processor cycles, T_{ldm} is 18 cycles, T_{exc} is 3 cycles, and T_{fiq} is 2 cycles. The total time is therefore 25.5 processor cycles, which is just over 2.5 microseconds in a system which uses a continuous 10 MHz processor clock. In a DRAM based system running at 4 and 8 MHz, for example using MEMC, this time becomes 4.5 microseconds, and if bus bandwidth is being used to support video or other DMA activity, the time will increase accordingly.

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 synchroniser ($T_{syncmin}$) plus T_{fiq} . This is 3.5 processor cycles.

6. Instruction Set

6.1 The condition field



Condition field

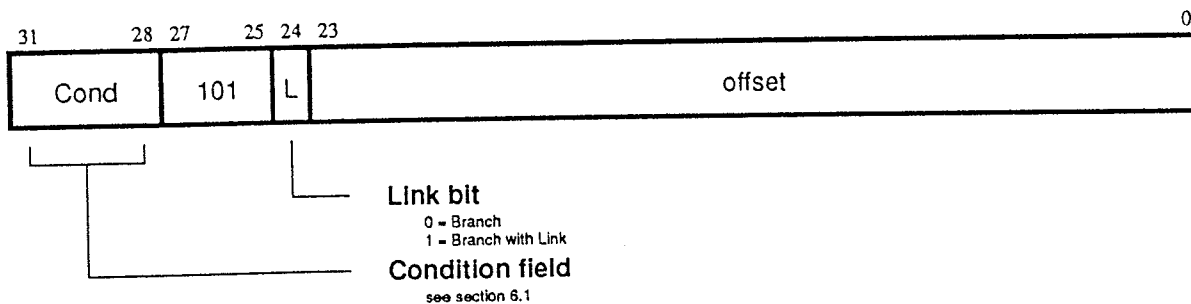
0000 = EQ - Z set (equal)
0001 = NE - Z clear (not equal)
0010 = CS - C set (unsigned higher or same)
0011 = CC - C clear (unsigned lower)
0100 = MI - N set (negative)
0101 = PL - N clear (positive or zero)
0110 = VS - V set (overflow)
0111 = VC - V clear (no overflow)
1000 = HI - C set and Z clear (unsigned higher)
1001 = LS - C clear or Z set (unsigned lower or same)
1010 = GE - N set and V set, or N clear and V clear (greater or equal)
1011 = LT - N set and V clear, or N clear and V set (less than)
1100 = GT - Z clear, and either N set and V set, or N clear and V clear (greater than)
1101 = LE - Z set, or N set and V clear, or N clear and V set (less than or equal)
1110 = AL - always
1111 = NV - never

All ARM 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 PSR at the end of the preceding instruction.

If the ALways condition is specified, the instruction will be executed irrespective of the flags, and likewise the NeVer condition will cause it not to be executed (it will be a no-op, ie take one cycle and have no effect on the processor state).

The other condition codes have meanings as detailed above, 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.

6.2 Branch and branch with link (B, BL)



The instruction is only executed if the condition specified in the condition field is true (see section 6.1).

All branches take a 24 bit offset. This is shifted left two bits and added to the PC, with any overflow being ignored. The branch can therefore reach any word aligned address within the address space. The branch offset must take account of the prefetch operation, which causes the PC to be 2 words ahead of the current instruction.

6.2.1 The link bit

Branch with Link writes the old PC and PSR into R14 of the current bank. The PC value written into the link register (R14) is adjusted to allow for the prefetch, and contains the address of the instruction following the branch and link instruction.

To return and restore the PSR use `MOVS PC,R14` if the link register is still valid or `LDM Rn!,(PC)^` if the link register has been saved onto a stack. To return without restoring the PSR 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.

6.2.2 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 section 6.1 (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.

Chapter 6

6.2.3 Examples

```
here BAL here ; assembles to EAffffffE
               ; (note effect of PC offset)

B there       ; ALways condition used as default

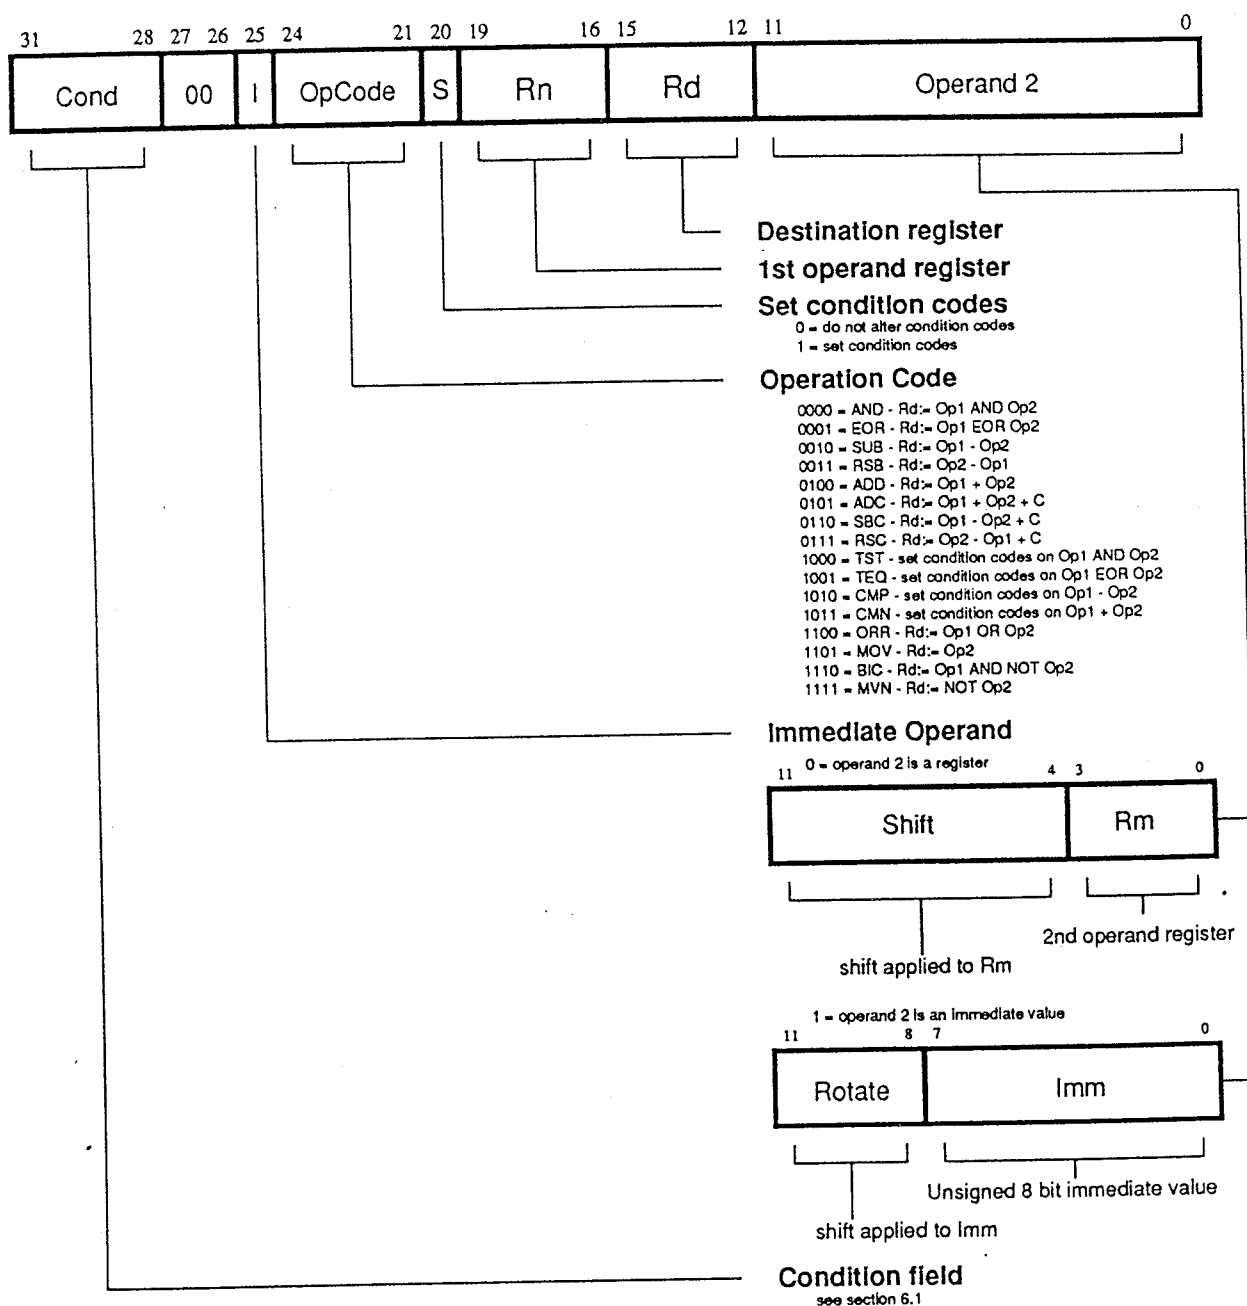
CMP R1,#0     ; compare register 1 with zero
BEQ fred      ; branch to fred if register 1 was zero
               ; otherwise continue to next instruction

BL sub + ROM  ; unconditionally call subroutine at computed address

ADDS R1,#1    ; add 1 to register 1, setting PSR flags on the result
BLCC sub      ; call subroutine if the C flag is clear, which will be
               ; the case unless R1 contained FFFFFFFFH
               ; otherwise continue to next instruction

BLNV sub      ; NeVer call subroutine (this is a NO-OP)
```

6.3 Data processing



The instruction is only executed if the condition is true. The various conditions are defined in section 6.1.

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 PSR 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 therefore should always have the S bit set. (The assembler treats TST, TEQ, CMP and CMN as TSTS, TEQS, CMPS and CMNS by default.)

6.3.1 Operations

The operations supported are:

Assembler Mnemonic	OpCode	Action
AND	0000	Bit-wise logical AND of operands
EOR	0001	Bit-wise logical EOR of operands
SUB	0010	Subtract operand 2 from operand 1
RSB	0011	Subtract operand 1 from operand 2
ADD	0100	Add operands
ADC	0101	Add operands plus carry (PSR C flag)
SBC	0110	Subtract operand 2 from operand 1 plus carry
RSC	0111	Subtract operand 1 from operand 2 plus carry
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	Bit-wise logical OR of operands
MOV	1101	Move operand 2 (operand 1 is ignored)
BIC	1110	Bit clear (bit-wise logical AND of operand 1 and NOT operand 2)
MVN	1111	Move NOT operand 2 (operand 1 is ignored)

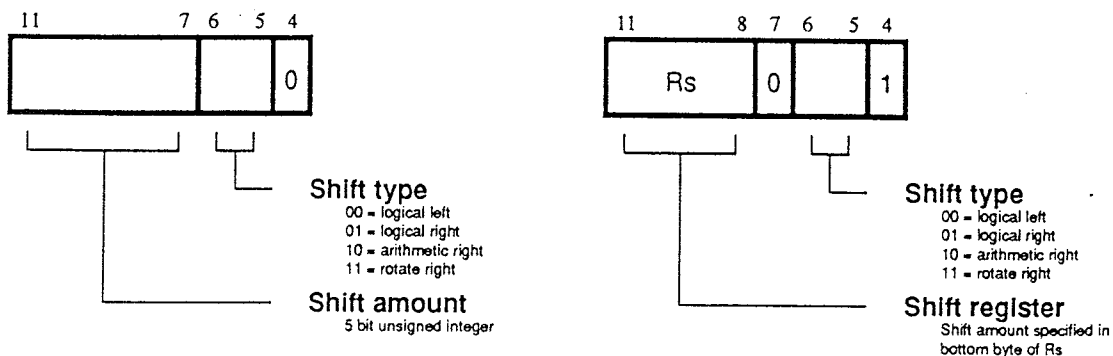
6.3.2 PSR flags

The 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 PSR 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 zeroes, and the N flag will be set to the logical value of bit 31 of the result.

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 PSR 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).

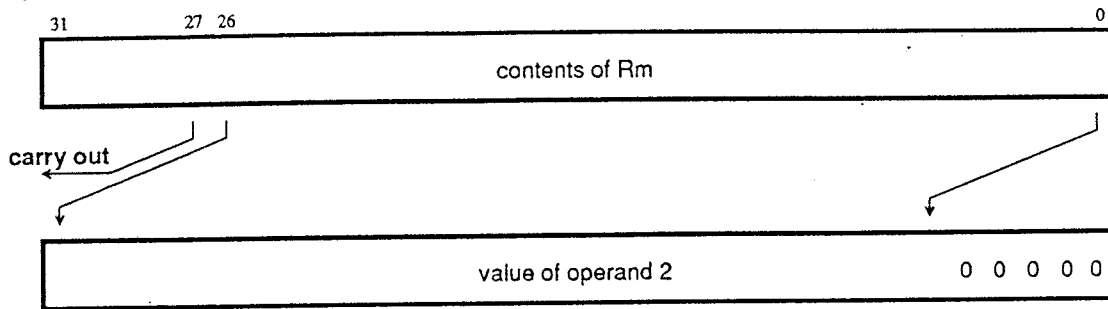
6.3.3 Shifts

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:



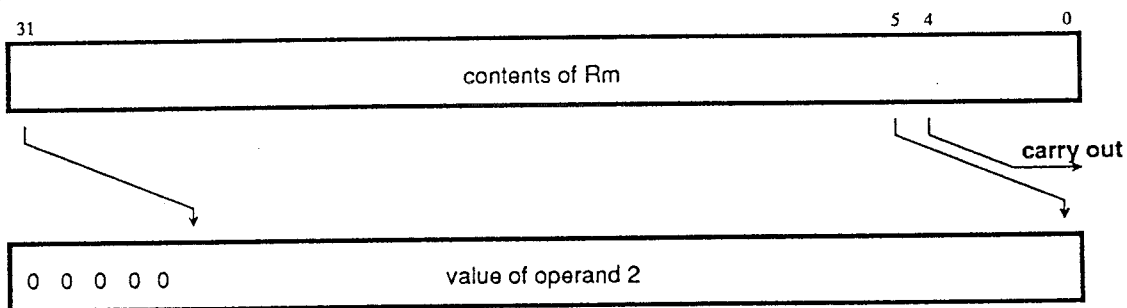
Instruction specified shift amount

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 zeroes, 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 PSR when the ALU operation is in the logical class (see above). For example, the effect of LSL #5 is:



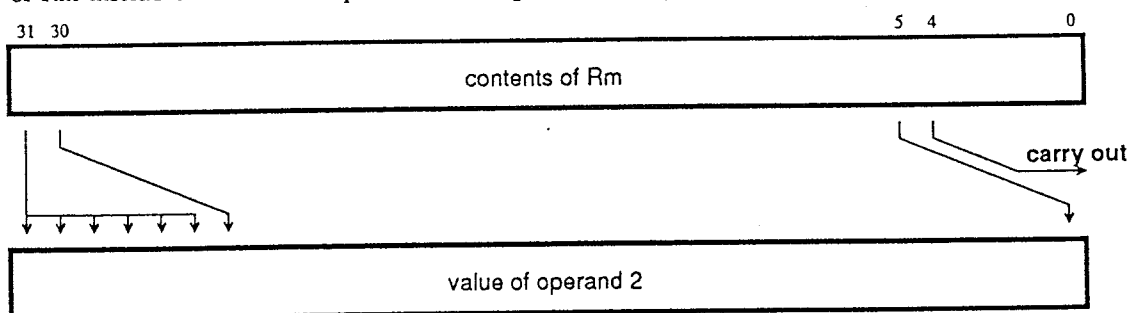
Note that LSL #0 is a special case, where the shifter carry out is the old value of the PSR C flag. The contents of Rm are used directly as the second operand.

A logical shift right (LSR) is similar, but the contents of Rm are moved to less significant positions in the result. LSR #5 has this effect:



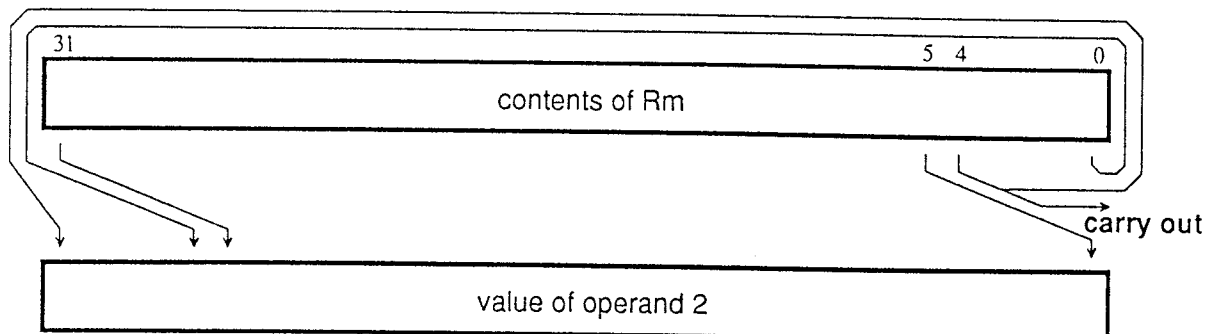
The form of the shift field which might be expected to correspond to LSR #0 is used to encode LSR #32, which has a zero result with bit 31 of Rm as the carry output. Logical shift right zero is redundant as it is the same as logical shift left zero, so the assembler will convert LSR #0 (and ASR #0 and 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 zeroes. This preserves the sign in 2's complement notation. For example, ASR #5:

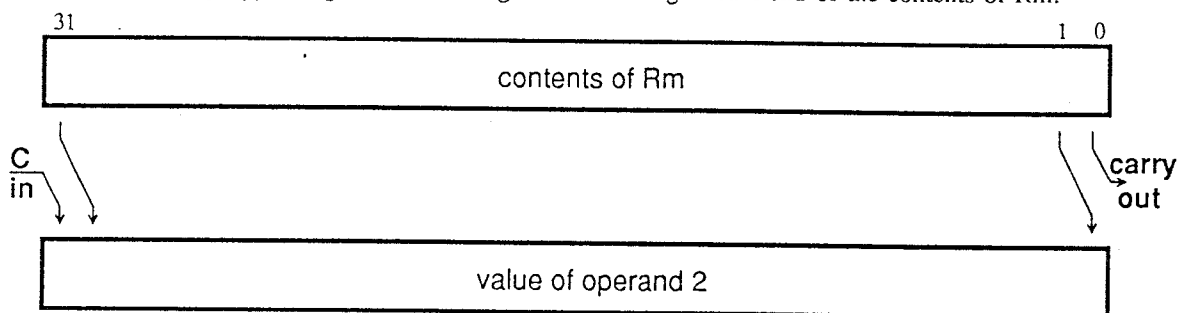


The form of the shift field which might be expected to give ASR #0 is used to encode ASR #32. Bit 31 of Rm is again used as the carry output, and each bit of operand 2 is also equal to bit 31 of Rm. The result is therefore all ones or all zeroes, according to the value of bit 31 of Rm.

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 zeroes used to fill the high end in logical right operations. For example, ROR #5:



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 PSR C flag to the most significant end of the contents of Rm:



Register specified shift amount

Only the least significant byte of the contents of Rs is used to determine the shift amount.

If this byte is zero, the unchanged contents of Rm will be used as the second operand, and the old value of the PSR 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:

- * LSL by 32 has result zero, carry out equal to bit 0 of Rm.
- * LSL by more than 32 has result zero, carry out zero.
- * LSR by 32 has result zero, carry out equal to bit 31 of Rm.
- * LSR by more than 32 has result zero, carry out zero.
- * ASR by 32 or more has result filled with and carry out equal to bit 31 of Rm.
- * ROR by 32 has result equal to Rm, carry out equal to bit 31 of Rm.
- * 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 an undefined instruction.

6.3.4 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. Another example is that the 8 bit constant may be aligned with the PSR flags (bits 0, 1, and 26 to 31). All the flags can thereby be initialised in one TEQP instruction (see section 6.2.5).

6.3.5 Writing to R15

When Rd is a register other than R15, the condition code flags in the PSR may be updated from the ALU flags as described above. When Rd is R15 and the S flag in the instruction is set, the PSR is overwritten by the corresponding bits in the ALU result, so bit 31 of the result goes to the N flag, bit 30 to the Z flag, bit 29 to the C flag and bit 28 to the V flag. In user mode the other flags (I, F, M1, M0) are protected from direct change, but in non-user modes these will also be affected, accepting copies of bits 27, 26, 1 and 0 of the result respectively.

When one of these instructions is used to change the processor mode (which is only possible in a non-user mode), the following instruction should not access a banked register (R8-R14) during its first cycle. A no-op should be inserted if the next instruction must access a banked register. Accesses to the unbanked registers (R0-R7 and R15) are safe.

If the S flag is clear when Rd is R15, only the 24 PC bits of R15 will be written. Conversely, if the instruction is of a type which does not normally produce a result (CMP, CMN, TST, TEQ) but Rd is R15 and the S bit is set, the result will be used in this case to update those PSR flags which are not protected by virtue of the processor mode.

6.3.6 Using R15 as an operand

If R15 is used as an operand in a data processing instruction it can present different values depending on which operand position it occupies. It will always contain the value of the PC. It may or may not contain the values of the PSR flags as they were at the completion of the previous instruction.

When R15 appears in the Rm position it will give the value of the PC together with the PSR flags to the barrel shifter.

When R15 appears in either of the Rn or Rs positions it will give the value of the PC alone, with the PSR bits replaced by zeroes.

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 8 bytes ahead when used as Rs, and 12 bytes ahead when used as Rn or Rm.

6.3.7 Assembler syntax

- * MOV,MVN - single operand instructions
`<opcode>{cond}{S} Rd,<Op2>`
- * CMP,CMN,TEQ,TST - instructions which do not produce a result.
`<opcode>{cond}{P} Rn,<Op2>`
- * 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 section 6.1.

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

{P} - make Rd = R15 in instructions where Rd is not specified, otherwise Rd will default to R0. (Used for changing the PSR directly from the ALU result.)

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.)

6.3.8 Examples

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

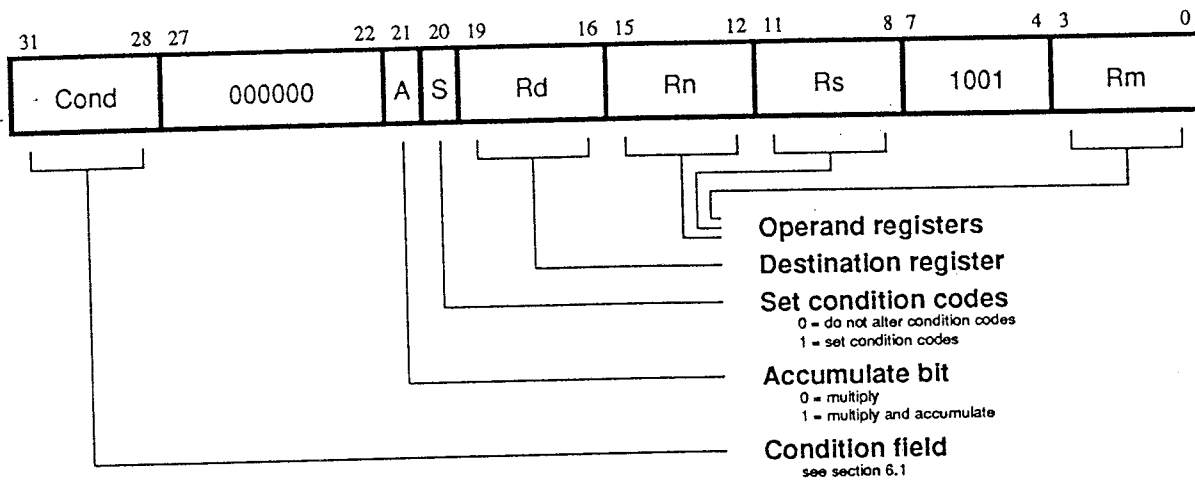
TEQS R4,#3           ; test R4 for equality with 3
                     ; (the S is in fact redundant as the
                     ; assembler inserts it automatically)

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

                     ; assume non-user mode here
TEQP R15,#0          ; Change to user mode and clear N,Z,C,V,I,F
                     ; NB R15 is here in the Rn position, so it
                     ; comes without the PSR flags
MOVNV R0,R0          ; no-op to avoid mode change hazard
MOV PC,R14           ; return from subroutine (R14 is a banked register)

MOVS PC,R14          ; return from subroutine and restore the PSR
```

6.4 Multiply and multiply-accumulate (MUL, MLA)



The instruction is only executed if the condition is true. The various conditions are defined in section 6.1.

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.

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.

6.4.1 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.

The destination register (Rd) should also not be $R15$. $R15$ is protected from modification by these instructions, so the instruction will have no effect, except that it will put meaningless values in the PSR flags if the S bit is set.

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

6.4.2 PSR flags

Setting the PSR 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.

6.4.3 Writing to R15

As mentioned above, R15 must not be used as the destination register (Rd). If it is so used, the instruction will have no effect except possibly to scramble the PSR flags.

6.4.4 Using R15 as an operand

R15 may be used as one or more of the operands, though the result will rarely be useful. When used as Rs the PC bits will be used without the PSR flags, and the PC value will be 8 bytes on from the address of the multiply instruction. When used as Rn, the PC bits will be used along with the PSR flags, and the PC will again be 8 bits on from the address of the instruction. When used as Rm, the PC bits will be used together with the PSR flags, but the PC will be the address of the instruction plus 12 bytes in this case.

6.4.5 Assembler syntax

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

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

{cond} - two-character condition mnemonic, see section 6.1.

{S} - set condition codes if S present.

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

(Rd must not be R15 and must not be the same as Rm.)

6.4.6 Examples

```
MUL R1,R2,R3      ; R1:=R2*R3

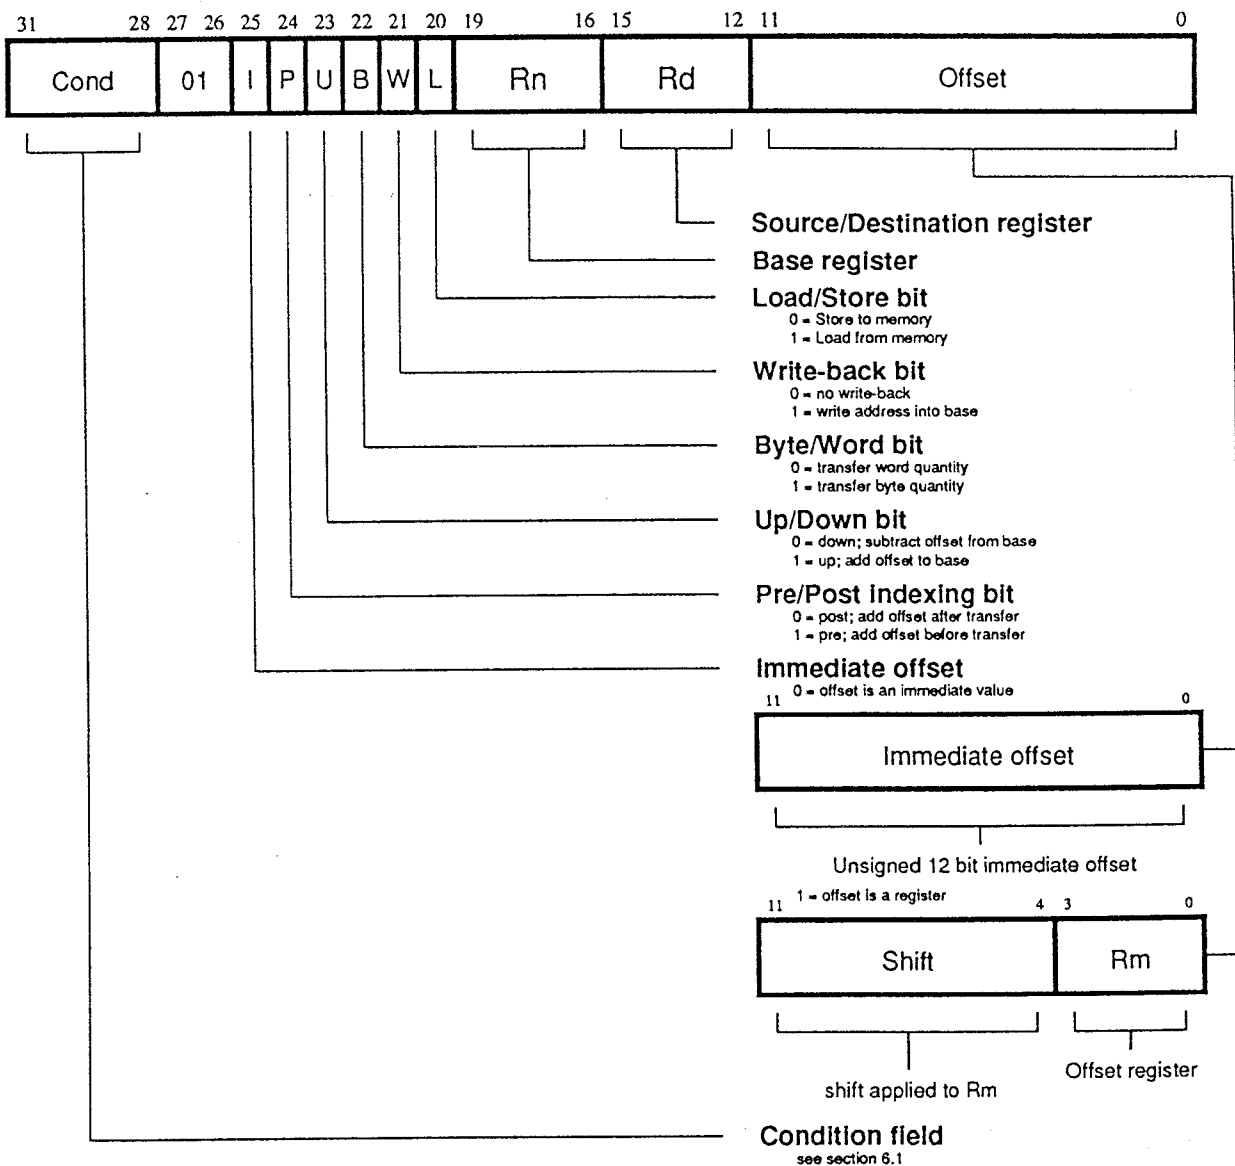
MLAEQS R1,R2,R3,R4 ; conditionally R1:=R2*R3+R4,
                   ; setting condition codes
```

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
ADDS  A,B,A        ; add middle sections
                        ; (couldn't use MLA as we need C correct)
ADDCS D,D,#410000  ; carry from above add
ADDS  C,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)

6.5 Single data transfer (LDR, STR)



The instruction is only executed if the condition is true. The various conditions are defined in section 6.1.

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.

6.5.1 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 Rn. 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, 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 non-user mode code, where setting the W bit forces the **TRANS** pin 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 pin.

6.5.2 Shifted register offset

The 8 shift control bits are described in the data processing instructions (section 6.2.3), but the register specified shift amounts are not available in this instruction class.

6.5.3 Bytes and words

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

A byte load (LDRB) expects the data on bits 0 to 7 if the supplied address is on a word boundary, on bits 8 to 15 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 zeroes.

A byte store (STRB) repeats the bottom 8 bits of the source register four times across the data bus. The external memory system should activate the appropriate byte subsystem to store the data (see chapter 7).

A word load (LDR) should generate a word aligned address. 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. External hardware could perform a double access to memory to allow non-aligned word loads, but existing systems do not support this.

A word store (STR) should generate a word aligned address. The data presented to the data bus are not affected if the address is not word aligned, so if support were required for non-aligned stores external hardware would have to switch bytes around on the bus.

6.5.4 Use of R15

These instructions will never cause the PSR to be modified, even when Rd or Rn is R15.

If R15 is specified as the base register (Rn), the PC is used without the PSR flags. When using the PC as the base register one must remember that it contains an address 8 bytes on from the address of the current instruction.

If R15 is specified as the register offset (Rm), the value presented will be the PC together with the PSR.

When R15 is the source register (Rd) of a register store (STR) instruction, the value stored will be the PC together with the PSR. The stored value of the PC will be 12 bytes on from the address of the instruction. A load register (LDR) with R15 as Rd will change only the PC, and the PSR will be unchanged.

6.5.5 Address exceptions

If the address used for the transfer (ie the unmodified contents of the base register for post-indexed addressing, or the base modified by the offset for pre-indexed addressing) has a logic one in any of the bits 26 to 31, the transfer will not take place and the address exception trap will be taken.

Note that it is only the address actually used for the transfer which is checked. A base containing an address outside the legal range may be used in a pre-indexed transfer if the offset brings the address within the legal range, and likewise a base within the legal range may be modified by post-indexing to outside the legal range without causing an address exception.

6.5.6 Data Aborts

A transfer to or from a legal address may still 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** pin 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.

6.5.7 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 section 6.1.

{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 **TRANS** pin 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:

- * 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.

- * 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>.

- * 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 ARM pipelining.

<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.

6.5.8 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 writeback
                             ; BASE+INDEX to BASE

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

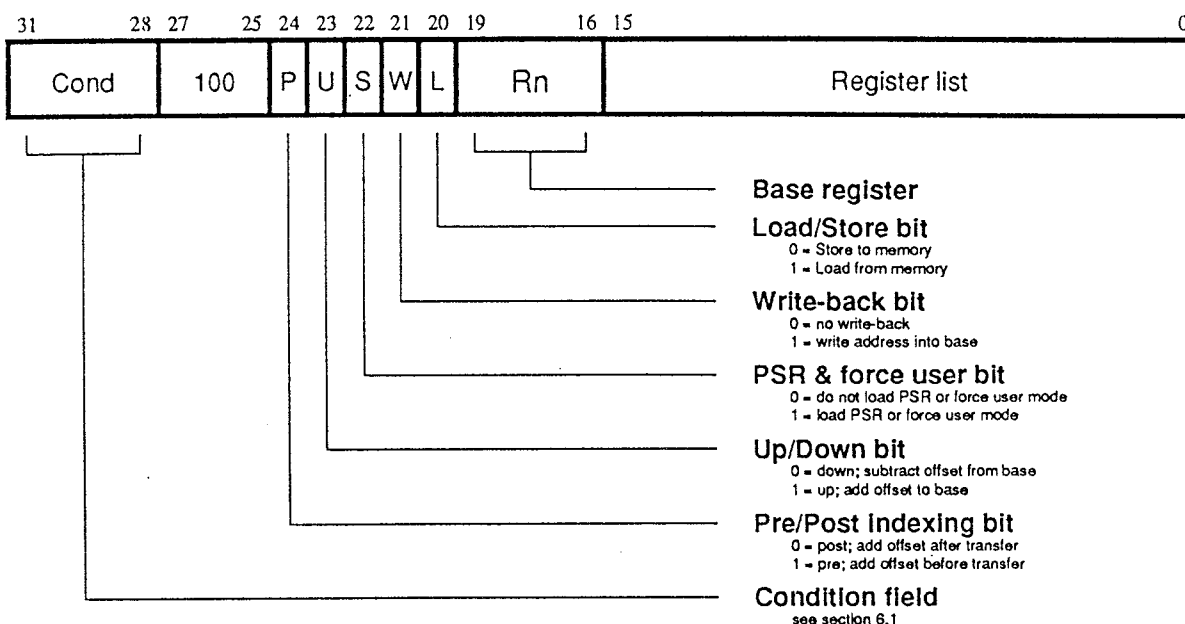
LDR R1, [BASE, INDEX, LSL #2] ; load R1 from contents of
                             ; BASE+INDEX*4

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

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

PLACE
```

6.6 Block data transfer (LDM, STM)



The instruction is only executed if the condition is true. The various conditions are defined in section 6.1.

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.

6.6.1 The register list

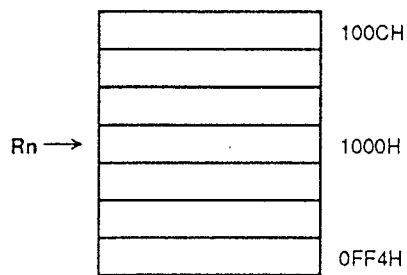
The instruction can cause the transfer of any registers in the current bank (and non-user mode programs can also transfer to and from the user bank, see below). The register list is a 16 bit field in the instruction, with each bit corresponding to a register. A 1 in bit 0 of the register field will cause 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.

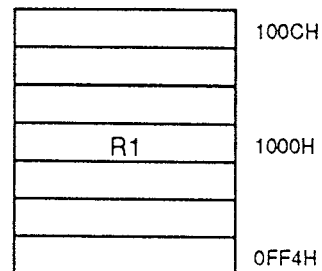
6.6.2 Addressing modes

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). The following figures show the sequence of register transfers, the addresses used, and the value of Rn after the instruction has completed.

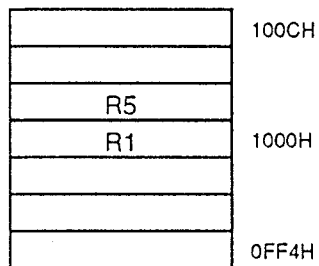
(In all cases, had write back of the modified base not been required (W=0), Rn would have retained its initial value of 1000H unless it was also in the transfer list of a load multiple register instruction, when it would have been overwritten with the loaded value.)

Post-increment addressing

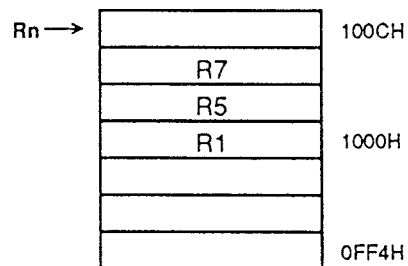
(1)



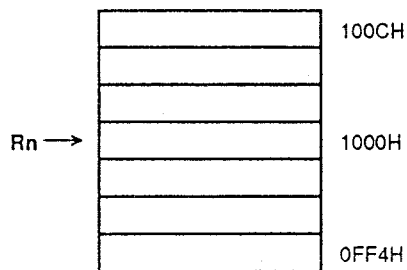
(2)



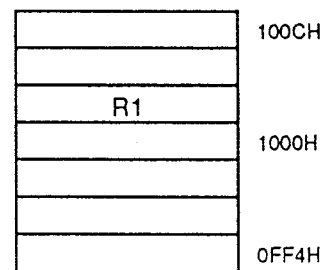
(3)



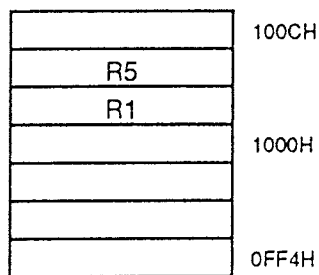
(4)

Pre-increment addressing

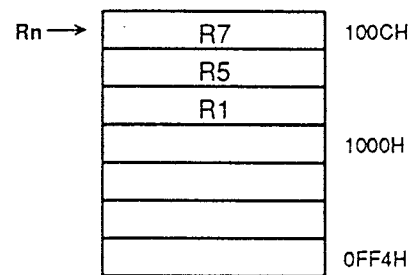
(1)



(2)

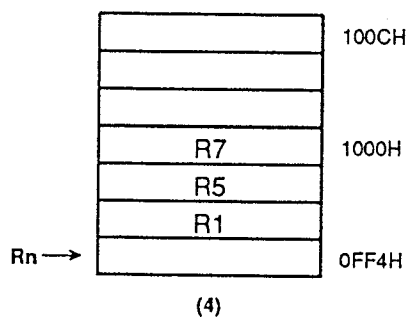
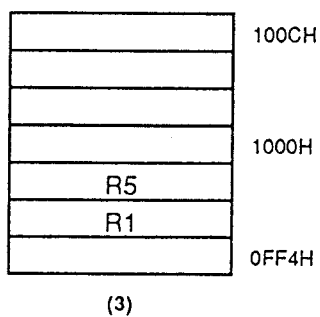
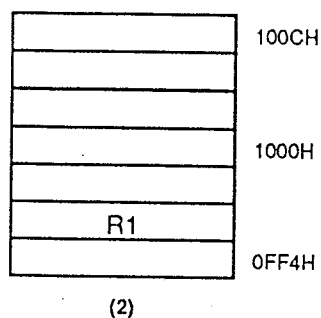
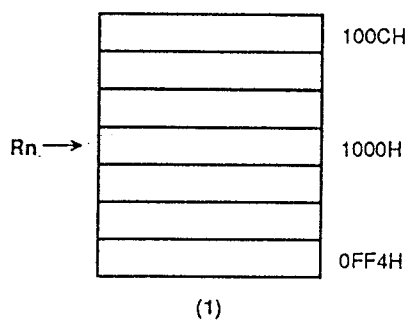


(3)

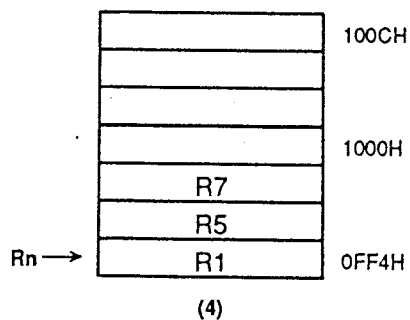
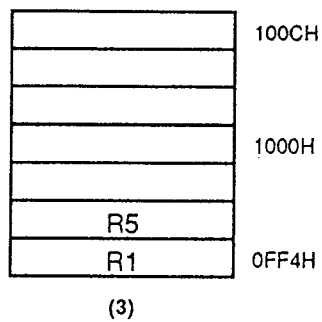
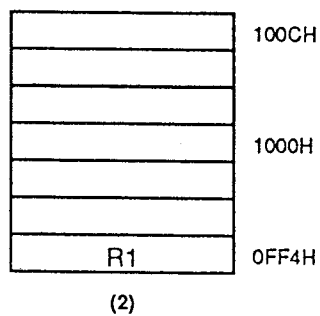
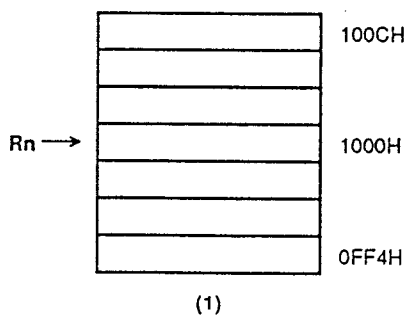


(4)

Post-decrement addressing



Pre-decrement addressing



6.6.3 Transfer of R15

Whenever R15 is stored to memory, the value transferred is the PC together with the PSR flags. The stored value of the PC will be 12 bytes on from the address of the STM instruction.

If R15 is in the transfer list of a load multiple (LDM) instruction the PC is overwritten, and the effect on the PSR is controlled by the S bit. If the S bit is 0 the PSR is preserved unchanged, but if the S bit is 1 the PSR will be overwritten by the corresponding bits of the loaded value. In user mode, however, the I, F, M0 and M1 bits are protected from change whatever the value of the S bit. The mode at the start of the instruction determines whether these bits are protected, and the supervisor may return to the user program, reenabling interrupts and restoring user mode with one LDM instruction.

6.6.4 Forcing transfer of the user bank

For STM instructions the S bit is redundant as the PSR is always stored with the PC whenever R15 is in the transfer list. In user mode programs the S bit is ignored, but in other modes it has a second interpretation. S=1 used to force transfers to take values from the user register bank instead of from the current register bank. This is useful for saving the user state on process switches. Note that when it is so used, write back of the base will also be to the user bank, though the base will be fetched from the current bank. Therefore don't use write back when forcing user bank.

In LDM instructions the S bit is redundant if R15 is not in the transfer list, and again in user mode programs it is ignored in this case. In non-user mode programs where R15 is not in the transfer list, S=1 is used to force loaded values to go to the user registers instead of the current register bank. When so used, care must be taken not to read from a banked register during the following cycle - if in doubt insert a no-op. Again don't use write back when forcing user bank transfer.

6.6.5 Use of R15 as the base

When the base is the PC, the PSR bits will be used to form the address as well, so unless all interrupts are enabled and all flags are zero an address exception will occur. Also, write back is never allowed when the base is the PC (setting the W bit will have no effect).

6.6.6 Inclusion of the base in the register list

When writeback 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. An LDM will always overwrite the updated base if the base is in the list.

6.6.7 Address exceptions

When the address of the first transfer falls outside the legal address space (ie has a 1 somewhere in bits 26 to 31), an address exception trap will be taken. The instruction will first complete in the usual number of cycles, though an STM will be prevented from writing to memory. The processor state will be the same as if a data abort had occurred on the first transfer cycle (see next section).

Only the address of the first transfer is checked in this way; if subsequent addresses over- or under-flow into illegal address space they will be truncated to 26 bits but will not cause an address exception trap.

6.6.8 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 pin HIGH. This can happen on any transfer during a multiple register load or store, and must be recoverable if ARM is to be used in a virtual memory system.

Aborts during STM instructions

If the abort occurs during a store multiple instruction, ARM 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 ARM detects a data abort during a load multiple instruction, it modifies the operation of the instruction to ensure that recovery is possible.

- * Overwriting of registers stops when the abort happens. The aborting load will not take place, nor will the preceding one, but registers two or more positions ahead of the abort (if any) will be loaded. (This guarantees that the PC will be preserved, since it is always the last register to be overwritten.)
- * 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.

6.6.9 Assembler syntax

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

{cond} - two character condition mnemonic, see section 6.1.

Rn is an expression evaluating to a valid register number.

<Rlist> can be either a list of registers and register ranges enclosed in {} (eg {R0,R2-R7,R10}), or an expression evaluating to the 16 bit operand.

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

{^} if present set S bit to load the PSR with the PC, or force transfer of user bank when in non-user 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:

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.

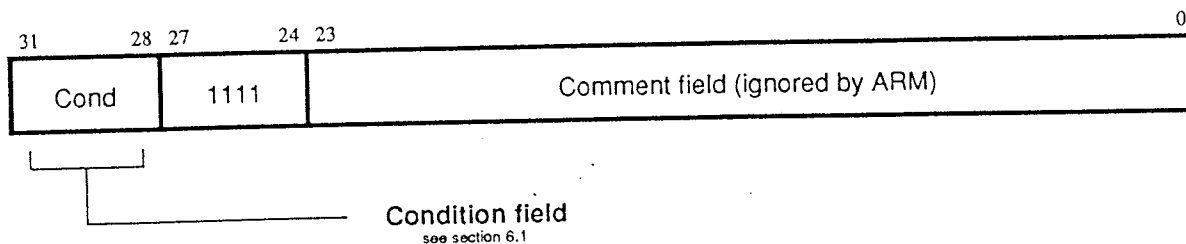
6.6.10 Examples

```
LDMFD SP!, {R0,R1,R2} ; unstack 3 registers
STMIA BASE, {R0-R15}  ; save all registers
```

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
                        ; (also restoring PSR flags)
```

6.7 Software interrupt



The instruction is only executed if the condition is true. The various conditions are defined in section 6.1.

The software interrupt instruction is used to enter supervisor mode in a controlled manner. The instruction causes the software interrupt trap to be taken, which effects the mode change but forces the PC to a fixed value (08H). If this address is suitably protected (by external memory management hardware) from modification by the user, a fully protected operating system may be constructed.

6.7.1 Return from the supervisor

The PC and PSR are saved in R14_svc upon entering the software interrupt trap, with the PC adjusted to point to the word after the SWI instruction. MOVS R15,R14_svc will return to the user program, restore the user PSR and return the processor to user mode.

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.

6.7.2 Comment field

The bottom 24 bits of the instruction are ignored by ARM, 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.

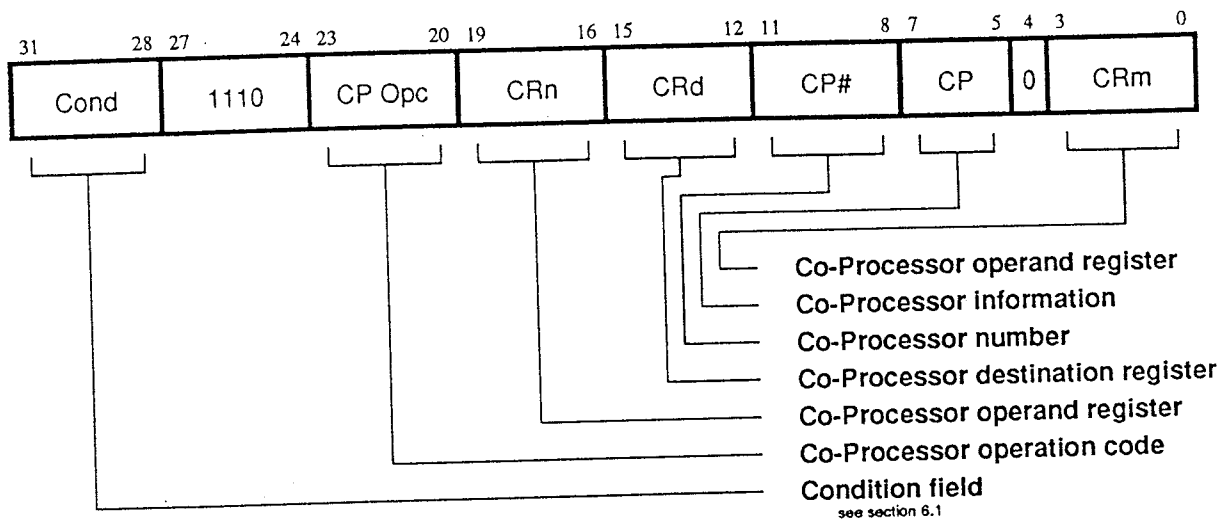
6.7.3 Assembler syntax

SWI{cond} <expression>

{cond} - two character condition mnemonic, see section 6.1.

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

6.8 Co-Processor data operations



The instruction is only executed if the condition is true. The various conditions are defined in section 6.1.

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

6.8.1 The Co-Processor fields

Only bit 4 and bits 24 to 31 are significant to ARM; the remaining bits are used by Co-Processors. The above field names are used by convention, and particular Co-Processors may redefine the use of all fields except CP# as appropriate. The CP# field is used to contain an identifying number (in the range 0 to 15) for each Co-Processor, and a Co-Processor will ignore any instruction which does not contain its number in the CP# field.

The conventional interpretation of the instruction is that the Co-Processor should perform an operation specified in the CP Opc field (and possibly in the CP field) on the contents of CRn and CRm, and place the result in CRd.

6.8.2 Assembler syntax

`CDP{cond} CP#,<expression1>,CRd,CRn,CRm{,<expression2>}`

{cond} - two character condition mnemonic, see section 6.1.

CP# - the unique number of the required Co-Processor.

<expression1> - evaluated to a constant and placed in the CP Opc field.

CRd, CRn and CRm are expressions evaluating to a valid Co-Processor register number.

<expression2> - where present is evaluated to a constant and placed in the CP field.

Chapter 6

6.8.3 Examples

```
CDP 1,10,CR1,CR2,CR3    ; request Co-Proc 1 to do operation 10
                        ; on CR2 and CR3, and put the result in CR1

CDPEQ 2,5,CR1,CR2,CR3,2 ; if Z flag is set request Co-Proc 2 to do
                        ; operation 5 (type 2) on CR2 and CR3,
                        ; and put the result in CR1
```