

CPC

CPC 464/664/6128

Firmware

Soft 968

CPC464/664/6128

FIRMWARE

ROM routines and explanations

Bruce Godden and Paul Overell, Locomotive Software

David Radisic, Amstrad

CAS CATALOG CAS CHEGK CAS IN ABANDON CAS IN CHAR CAS IN CLOSE CAS IN DIRECT CAS IN OPEN CAS INITIALISE CAS NOISY CAS OUT ABANDON CAS OUT CHAR CAS OUT CLOSE. CAS READ CAS RESTORE MOTOR CAS RETURN CAS SET SPEED CAS START MOTOR CAS STOP MOTOR CAS TEST EOF CAS WRITE GRA ASK CURSOR GRA GLEAR WINDOW GRA GET ORIGIN GRA GET PAPER GRA GET PEN GRA GET W HEIGHT GRA GET W WIDTH GRA INITIALISE GRA LINE ABSOLUTE GRA LINE RELATIVE GRA MOVE ABSOLUTE GRA MOVE RELATIVE GRA PLOT ABSOLUTE GRA PLOT RELATIVE GRA RESET GRA SET ORIGIN GRA SET PAPER GRA SET PEN GRA TEST ABSOLUTE GRA TEST RELATIVE GRA WIN HEIGHT GRA WIN WIDTH GRA WR CHAR KL CURR SELECTION HI KL L ROM DISABLE HI KL L ROM ENBLE HI KL LDDR HI KL LDIR HI KL POLL SYNCHRONOUS HI KL PROBE ROM HI KL ROM DESELECT HI KL ROM RESTORE HI KL ROM SELECT HI KL U ROM DISABLE HI KL U ROM ENABLE IND GRA LINE IND GRA PLOT IND GRA TEST IND KM TEST BREAK IND MC WAIT PRINTER IND SCR READ IND SCR WRITE IND TXT DRAW CURSOR IND TXT OUT ACTION IND TXT UNDRAW CURSOR IND TXT UNWRITE IND TXT WRITE CHAR KL ADD FAST TICKER KL ADD FRAME FLY KL ADD TICKER KL CHOKE OFF KL DEL FAST TICKER KL DEL FRAME FLY KL DEL SYNCHRONOUS KL DEL TICKER KL DISARM EVENT KL DO SYNC KL DONE SYNC KL EVENT DISABLE KL EVENT ENABLE KL FIND COMMAND KL INIT BACK KL INIT EVENT KL LOG EXT KL NEW FAST TICKER KL NEW FRAME FLY KL NEXT SYNC KL ROM WALK KL SYNC RESET KL TIME PLEASE KL TIME CONTROL KM GET DELAY KM GET EXPAND KM GET JOYSTICK KM GET REPEAT KM GET SHIFT KM GET STATE KM GET TRANSLATE KM INITIALISE KM READ CHAR KM READ KEY KM RESET KM SET CONTROL KM SET DELAY KM SET EXPAND KM SET REPEAT KM SET SHIFT KM SET TRANSLATE KM TEST KEY KM WAIT CHAR KM WAIT KEY EXT INTERRUPT LOW FAR GALL LOW FRM JUMP LOW INTERRUPT ENTRY LOW KL FAR CALL LOW KL FAR PCHL LOW KL LOW PCHL LOW KL SIDE PCHL LOW LOW JUMP LOW PCDE INSTRUCTION LOW PCHL INSTRUCTION LOW RAM LAM LOW RESET ENTRY LOW SIDE CALL LOW USER RESTART MC BOOT PROGRAM MC BUSY PRINTER MC CLEAR INKS MC JUMP RESTORE MC PRINT CHAR MC RESET PRINTER MC SCREEN OFFSET MC SEND PRINTER MC SET INKS MC SET MODE MC SOUND REGISTER MC START PROGRAM MC WAIT FLYBACK SCR ACCESS SCR CHAR INVERT SCR CHAR LIMITS SCR CHAR POSITION SCR CLEAR SCR DOT POSITION SCR FILL BOX SCR FLOOD BOX SCR GET BORDER SCR GET FLASHING SCR GET INK SCR GET LOGATION SCR HORIZONTAL SCR INITIALISE SCR INK DECODE SCR INK ENCODE SCR NEXT BYTE SCR NEXT LINE SCR PIXELS SCR PREV BYTE SCR PREV LINE SCR REPACK SCR RESET SCR SET BASE SCR SET BORDER SCR SET FLASHING SCR SET INK SCR SET MODE SCR SET OFFSET SCR SW ROLL SCR UNPACK SCR VERTICAL SOUND A ADDRESS SOUND AMPL ENVELOPE SOUND ARM EVENT SOUND CHECK SOUND CONTINUE SOUND HOLD SOUND QUEUE SOUND RELEASE SOUND RESET SOUND T ADDRESS SOUND TONE ENVELOPE TXT CLEAR WINDOW TXT CUR DISABLE TXT CUR ENABLE TXT CUR OFF TXT CUR ON TXT GET BACK TXT GET CONTROLS TXT GET CURSOR TXT GET M TABLE TXT GET MATRIX TXT GET PAPER TXT GET PEN TXT GET WINDOW TXT INITIALISE TXT INVERSE TXT OUTPUT TXT PLACE CURSOR TXT RD CHAR TXT REMOVE CURSOR TXT RESET TXT SET BACK TXT SET COLUMN TXT SET CURSOR TXT SET GRAPHIC TXT SET M TABLE TXT SET MATRIX TXT SET PAPER TXT SET PEN TXT SET ROW TXT STR SELECT TXT SWAP STREAMS TXT VALIDATE TXT VDU DISABLE TXT VDU ENABLE TXT WIN ENABLE TXT WR CHAR KM SET LOCKS KM FLUSH TXT ASK STATE GRA DEFAULT GRA SET BACK GRA SET FIRST GRA SET LINE MASK GRA FROM USER GRA FILL SCR SET POSITION MC PRINT TRANSLATION KM SCAN KEYS

Published by
Amstrad Consumer Electronics plc
Brentwood House
169 Kings Road
Brentwood
Essex

All rights reserved
First edition 1986

Reproduction or translation of any part of this publication without the written permission of the copyright owner is unlawful, Amstrad and Locomotive Software reserve the right to amend or alter the specification without notice. While every effort has been made to verify that this complex software works as described, it is not possible to test any program of this complexity under all possible conditions Therefore the program and this manual are provided "as is" without warranty of any kind, either express or implied.

IBM is a trademark of International Business Machines Corporation.
CP/M, CP/M Plus and Dr Logo are trademarks of Digital Research Inc.

SOFT 968 Copyright © 1986 Locomotive Software Ltd and Amstrad Consumer Electronics plc

The Contents.

1 The Firmware.

- 1.1 The Hardware.
- 1.2 The Division of the Firmware.
- 1.3 Controlling the Firmware.
- 1.4 Jumpblocks.
- 1.5 Conventions.
- 1.6 Routine Documentation.
- 1.7 Example of Patching a Jumpblock

2 ROMs, RAM and the Restart Instructions.

- 2.1 Memory Map.
- 2.2 ROM Selection.
- 2.3 The Restart Instructions.
- 2.4 RAM and the Firmware.
- 2.5 Bank Switching.

3 The Keyboard.

- 3.1 Keyboard Scanning.
- 3.2 Key Translation.
- 3.3 Characters from the Keyboard.
- 3.4 Shift and Caps Lock.
- 3.5 Repeating keys.
- 3.6 Breaks.
- 3.7 Function Keys and Expansion Tokens.
- 3.8 Joysticks.

4 The Text VDU.

- 4.1 Text VDU Coordinate Systems.
- 4.2 Streams.
- 4.3 Text Pen and Paper Inks.
- 4.4 Text Windows.
- 4.5 The Current Position and the Cursor.
- 4.6 Characters and Matrices.
- 4.7 Character Output and Control Codes.

5 The Graphics VDU

- 5.1 Graphics VDU Coordinate Systems.
- 5.2 The Current Graphics Position.
- 5.3 Graphics Pen and Paper Inks.
- 5.4 Graphics Write Mode.
- 5.5 Graphics Window.
- 5.6 Writing Characters.
- 5.7 Drawing Lines.
- 5.8 Filling Areas.

6 The Screen.

- 6.1 Screen Modes.
- 6.2 Inks and Colours.
- 6.3 Screen Addresses.
- 6.4 Screen Memory Map.

7 The Sound Manager.

- 7.1 The Sound Chip.
- 7.2 Tone Periods and Amplitudes.
- 7.3 Enveloping.
- 7.4 Sound Commands.
- 7.5 Sound Queues.
- 7.6 Synchronisation.
- 7.7 Holding Sounds.

8 The Cassette Manager.

- 8.1 File Format.
- 8.2 Record Format.
- 8.3 Bit Format.
- 8.4 The Header Record.
- 8.5 Read and Write Speeds.
- 8.6 Cataloguing.
- 8.7 Reading Files.
- 8.8 Writing Files.
- 8.9 Reading and Writing Simultaneously.
- 8.10 Filenames.
- 8.11 Cassette Manager Messages.
- 8.12 Escape Key.
- 8.13 Low Level Cassette Driving.

9 AMSDOS

- 9.1 Features
- 9.2 Filenames
- 9.3 File Headers
- 9.4 Disc Organisation
- 9.5 Boot Sector
- 9.6 AMSDOS Messages
- 9.7 BIOS Facilities Available to AMSDOS
- 9.8 Store requirements
- 9.9 Extended Disc Parameter Blocks

10 Expansion ROMs, Resident System Extensions and RAM Programs.

- 10.1 ROM Addressing.
- 10.2 The Format of an Expansion RUM.
- 10.3 Foreground ROMs and RAM Programs.
- 10.4 Background RUMs.
- 10.5 Resident System Extensions.
- 10.6 External Commands.
- 10.7 Examples.

11 Interrupts.

- 11.1 The Time Interrupt.
- 11.2 External Interrupts.
- 11.3 Nonmaskable Interrupts.
- 11.4 Interrupts and Events.
- 11.5 Interrupt Queues.

12 Events.

- 12.1 EventClass.
- 12.2 Event Count.
- 12.3 Event Routine.
- 12.4 Disarming and Reinitializing Events.

13 The Machine Pack.

- 13.1 Hardware Interfaces.
- 13.2 The Printer.
- 13.3 Loading and Running Programs.

14 Firmware Jumpblocks.

- 14.1 The Main Jumpblock.
 - 14.1.1 Entries to the Key Manager.
 - 14.1.2 Entries to the TextVDU.
 - 14.1.3 Entries to the Graphics VDU.
 - 14.1.4 Entries to the Screen Pack.
 - 14.1.5 Entries to the Cassette Manager.
 - 14.1.6 Entries to the Sound Manager.
 - 14.1.7 Entries to the Kernel.
 - 14.1.8 Entries to the Machine Pack.
 - 14.1.9 Entries to Jumper.
 - 14.1.10 FurtherEntries.

14.2 Firmware Indirections.

- 14.2.1 Text VDU Indirections.
- 14.2.2 Graphics VDU Indirections.
- 14.2.3 Screen Pack Indirections.
- 14.2.4 Keyboard Manager Indirections.
- 14.2.5 Machine Pack Indirections.
- 14.2.6 Further Indirections.

14.3 The High Kernel Jumpblock.

14.4 The Low Kernel Jumpblock.

15 The Main Firmware Jumpblock.

16 The Firmware Indirections.

17 Kernel High Entries.

18 Low Entries to the Kernel.

19 AMSDOS 'BIOS' Facilities

20 AMSDOS External Commands

Appendices

- I Key Numbering.**
- II Key Translation Tables.**
- III Repeating Keys.**
- IV Function Keys and Expansion Strings.**
- V Inks and Colours.**
- VI Displayed Character Set.**
- VII Text VDU Control Codes.**
- VIII Notes and Tone Periods.**
- IX The Programmable Sound Generator.**
- X Kernel Block Layouts.**
- XI The Alternate Register Set.**
- XII Hardware and Hardware Variants.**
- XIII Hints, Tips, and Workarounds.**
- XIV Printer Translation Tables.**

Preface.

The computers in the CPC range contain ROMs which hold the BASIC interpreter and the firmware'. The firmware consists of low level routines responsible for driving the hardware, handling the screen, handling real-time events and other similar functions. This manual describes the firmware.

This manual is applicable to the full range of CPC machines; 464, 664, 6128 and 464+ DDI-1. The firmware in these machines is not identical. All 464s contain V1.0 of the firmware, all 664s contain V1.1 and all 6128s contain V1.2. All firmware calls are upwards compatible, that is, any firmware call available in V1.0 is available in V1.1 and V1.2 and any firmware call available in V1.1 is available in V1.2. Upwards compatibility will be maintained in any future versions of the firmware.

As stated above this manual relates only to the firmware section of the ROMs. Other AMSOFT publications describe the BASIC interpreter though not at a system level and the various implementations of CP/M. However, the areas in which CP/M and the firmware interact are covered in this manual. Also, areas of the firmware that are affected by the addition of a disc interface are covered in this manual.

1 The Firmware.

This manual describes the firmware of the Amstrad CPC 464/664/6128 microcomputers. It also describes the disc operating systems (CP/M and AMSDOS). It does not describe either the BASIC language supplied with the system or CP/M. The manual does describe certain aspects of the BASIC where these affect other programs and it uses BASIC in certain example programs when describing some features of the firmware. It also describes how to call the firmware from CP/M.

Three versions of the firmware are described. V1.0 (on CPC464), V1.1 (on CPC664) and V1.2 (on CPC6128). Apart from support for bank switching V1.1 and V1.2 are identical and are referred to as V1.1 throughout this manual. It may be necessary for a program to deduce which firmware is fitted in a computer, and this can be achieved by inspecting the on-board ROM's version number (as described in section 10.2) using KL ROM PROBE. This will return a 0, 1 or 2 depending on the version of firmware.

The firmware is the program that resides in the lower ROM and the disc controlling ROM (see section 2). Its function is to control the hardware of the computer and to provide useful facilities for other programs to use. This avoids every program written having to provide its own facilities.

This manual is expected to be of interest to anyone who would like to know how the system works. It is indispensable for programmers writing machine code programs, particularly system programs (e.g. other languages) and games.

The information presented can be extremely detailed. It covers the operation of the firmware from the lowest level (e.g. driving the sound chip) to the highest level (e.g. running a queue of sounds). It is not necessary to understand all the information given to be able to use the firmware, however, a good grasp of how the system works will aid the programmer in selecting the most appropriate method for performing a particular task.

Two disc operating systems are provided: AMSDOS, which enables BASIC programs to use disc files in much the same way as cassette files; and CP/M 2.2 the industry standard operating system (CP/M Plus in the CPC6128, but we do not discuss the differences between the two in this manual). Both AMSDOS and CP/M use the same file structure and may read and write each others file's. CP/M is invoked from BASIC by typing |CPM. Part of CP/M (the CCP and BDOS) is loaded from the disc in drive A:. The CP/M BIOS resides in the disc ROM.

AMSDOS is enabled whenever BASIC is first used. This intercepts most of the cassette firmware routines and redirects them to disc. Thus existing BASIC programs which use cassette files can use disc files with little or no modification. AMSDOS also provides a number of external commands for erasing and renaming files and redirecting the cassette firmware routines.

Provided with the disc system are a number of utility programs for formatting and copying discs and for changing various system parameters. These all run under CP/M.

1.1 The Hardware.

The diagram on the following page gives an indication of the different pieces of hardware in the system and how they connect to each other. For more information on how the hardware works see Appendix XII and the relevant manufacturer's data sheets.

The system centres around the CPU (Central Processing Unit) which is a Z80A microprocessor with a 4MHz clock. Next in importance is the gate array which contains miscellaneous logic to control much of the system. In particular, it controls ink colours, screen mode and ROM enabling (see section 10 and Appendix XII). In conjunction with the CRTC (Cathode Ray Tube Controller), which is a 6845 chip, the gate array generates the video signals for the monitor.

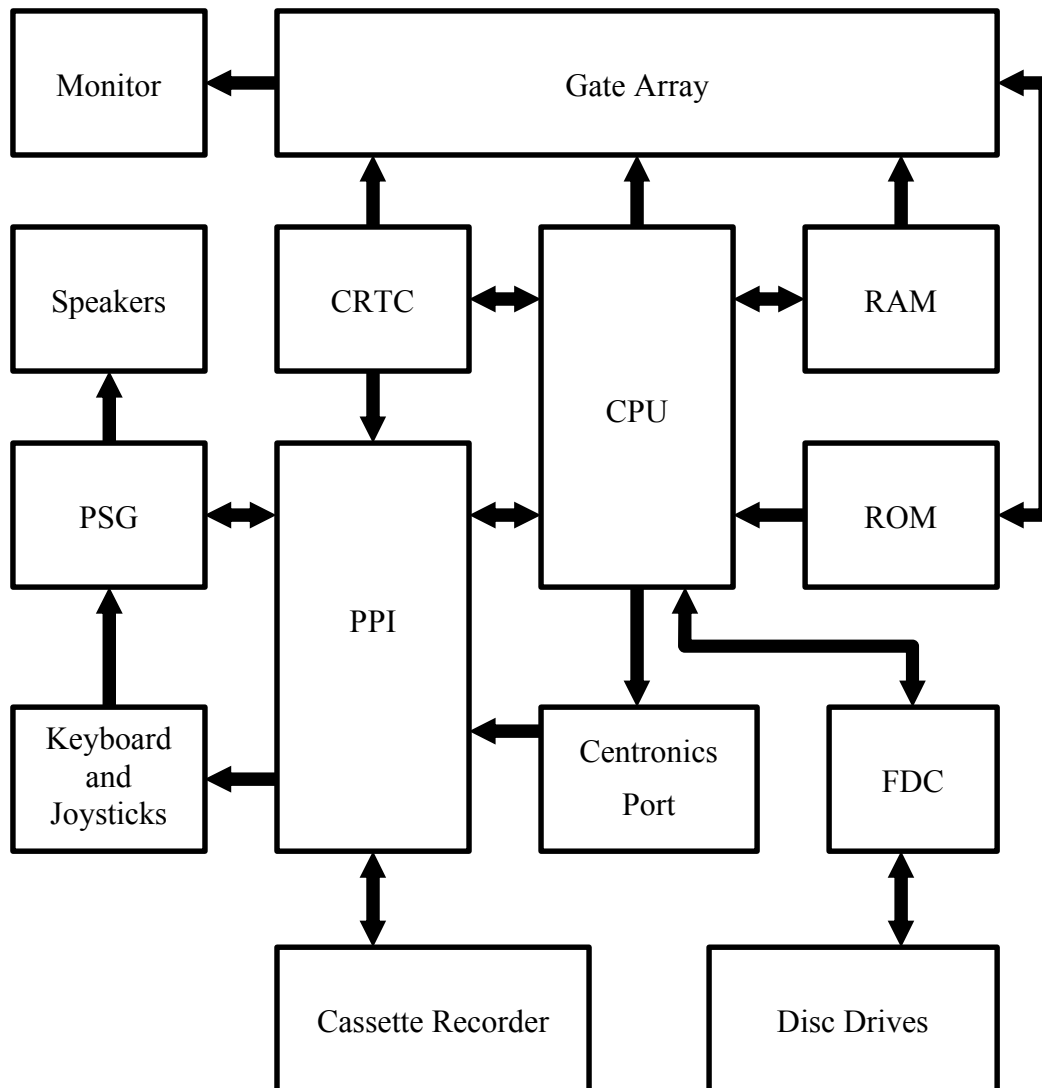
The PSG (Programmable Sound Generator) is an AY-3-8912. This chip has three channels of sound generator, a noise generator, envelope control for each channel and an I/O port. The way the sound generating hardware is used is described in section 7. The I/O port is used in input mode to sense the state of the keyboard and joystick switches.

The FDC (Floppy Disc Controller) is an NEC PD765A chip. Only two disc drives are supported, since the US1 line from the PD765A is ignored. This results in the two disc drives being accessed as drives 0 and 1 and again as 2 and 3. The FDC supports both single and double sided and single and double density mini-floppy disc drives. Note that the clock frequency supplied to the PD765A CLK pin is 4MHz rather than the 8 MHz used with larger disc drives.

Each disc drive takes a single 3" floppy disc. Either side of the disc may be used, depending on which way up the disc is inserted into the drive. The disc interface contains a 16K expansion ROM, 8K of which contains the disc driving software, the remainder being used by DR LOGO.

The PPI (Parallel Peripheral Interface), which is an 8255 chip, is used to control the remainder of the system. It has three ports. Port C is used as an output port to control the cassette recorder motor, to write data to the cassette, to strobe data in or out of the PSG and to select rows of the keyboard. Port B is used as an input port to sense the frame flyback signal, the Centronics port busy signal and various option links and to read data from the cassette. Port A is used to communicate with the PSG and is set into input or output mode as required.

Accesses to memory are synchronised with the video logic - they are constrained to occur on microsecond boundaries. This has the effect of stretching each Z80 M cycle (machine cycle) to be a multiple of 4 T states (clock cycles). In practice this alters the instruction timing so that the effective clock rate is approximately 3.3 MHz.



1.2 The Division of the Firmware.

The firmware is split into 'packs' each dealing with a particular part of the system, usually a hardware device. Each pack has a section of this manual devoted to it where its operation is explained in detail. The system components and their associated packs are:

Keyboard:	Key Manager.
Screen:	Text VDU, Graphics VDU, Screen Pack.
Cassette/Disc:	Cassette Manager/AMSDOS.
Sound:	Sound Manager.
Operating System:	Kernel, Machine Pack, Jumper.

a. Key Manager

The Key Manager is more fully described in section 3. It deals with scanning the keyboard, generating characters, function keys, testing for break and scanning the joysticks.

b. Text VDU

The Text VDU is more fully outlined in section 4. It deals with putting characters on the screen, the cursor and obeying control codes.

c. Graphics VDU

The Graphics VDU is more fully presented in section 5. It deals with plotting points, testing points drawing lines and filling areas on the screen.

d. Screen Pack

The Screen Pack is more fully detailed in section 6. It interfaces the Text and Graphics VDUs with the screen hardware and deals with aspects of the screen that affect both of these packs, such as screen mode or ink colours.

e. Sound Manager

The Sound Manager is more fully discussed in section 7. It deals with queueing, enveloping, synchronising and generating sounds.

f. Cassette Manager/AMSDOS

The Cassette manager is more fully explained in section 8. It deals with reading from tape, writing to tape and cassette motor control.

AMSDOS is explained more fully in section 9. It deals with reading from disc, writing to disc and the disc motor control.

g. Kernel

The Kernel is more fully described in sections 2, 10, 11 and 12. It is the heart of the operating system and deals with interrupts, events, selecting ROMs and running programs.

h. Machine Pack

The Machine Pack is more fully documented in section 13. It deals with the printer and the low level driving of the hardware.

i. Jumper

Jumper, or rather, the main firmware jumpblock is listed in section 14. The entries in the jumpblock are described in detail in section 15. Jumper sets up the firmware jumpblock.

1.3 Controlling the Firmware.

The firmware is controlled by the user calling published routines rather than by the user setting the values of system variables. This will allow the firmware's variable layout to be changed in major ways without the user being affected.

The addresses of the routines the user is to call need to remain constant if the firmware is altered. This is achieved by using jumpblocks (see below).

The advantage of a routine interface is that it allows a number of different system variables to be altered by the firmware in a consistent way in one operation. If the system variables had to be set by the user then the firmware could be left in an indeterminate state if some variables had been set but not others. Also, the routine type of interface ensures that all the required side effects of a change are taken care of automatically without the user being troubled with all the details. An example of this is changing the screen mode (see section 6.1)- changing the size of the screen requires a number of other people to be informed of the change so that illegal screen positions and inks are not used.

1.4 Jumpblocks.

A jumpblock is a series of jump instructions placed in memory at well-known locations. The jumps are to the various routines in the firmware that the user might want to call. Programs that need to use the facilities provided by the routines in the jumpblock should call the appropriate jumpblock entries.

If the firmware is altered then it is quite likely that the addresses of some of the routines available to the user will change. By keeping the address of the jumpblock constant but altering the entries in the jumpblock so that they jump to the new addresses of the routines, the change is hidden from the user (providing that the user is only calling routines via the jumpblock and is not accessing the firmware directly).

To make the change to the firmware completely hidden from the user it is also necessary to keep the entry and exit conditions of the routines accessed via the jumpblock constant. The greater part of this manual is taken up with the detailed entry and exit requirements of the jumpblock entries.

The jumpblock is placed in RAM so that the user can alter the entries in it. This allows the user to trap particular entries and to substitute a new routine that will replace the standard firmware routine. Provided that the new routine obeys the entry and exit requirements of the firmware routine, the substitution will not upset programs unaware of the change.

There are four jumpblocks. These are all listed in section 14. The first and largest jumpblock is the main firmware jumpblock (see sections 14.1 and 15). This allows the user to call most firmware routines. The second jumpblock is the indirections jumpblock (see sections 14.2 and 16). The entries in this jumpblock are used by the firmware at key moments in order to allow the user to alter the action of the firmware. The last two jumpblocks are rather special. They are to do with the Kernel and allow ROMs to be enabled and routines in ROMs to be called. (See sections 14.3, 14.4, 17 and 18).

Section 1.7 below gives an example of how a jumpblock entry might be changed to alter the action of the firmware.

1.5 Conventions.

a. Notation

Processor instructions are generally referred to by their standard Z80 mnemonics. The exceptions that prove the rule are the restart instructions. The mnemonics RST 0 .. RST 7 are used rather than the more usual Z80 mnemonics RST #00 .. RST #38.

The registers are also referred to by their standard Z80 names. The flag register as a whole is referred to as F but the individual flags are called by their full name, e.g. carry. The flags are said to be true when they are set and false when they are clear. Thus a JP NC instruction would jump if carry was false and not if carry was true.

Hexadecimal numbers are indicated by prefixing the number with #, thus # 7F is the number 127 in hex. All numbers not prefixed by # are in decimal.

Large numbers are often abbreviated by writing them as a multiple of 1024. For example, 32K bytes means 32 times 1024 (i.e. 32768) bytes.

b. Usage

Routines, where possible, take and return values in registers. Where more information than may be held in registers is to be passed to a routine, the address of a data area is given. The location in memory of these data areas is sometimes critical, see section 2.4.

Where a routine can succeed or fail this condition is normally passed back in the carry flag. Carry true normally implies success, whilst carry false normally implies failure.

The alternate register set, AF' BC' DE' HL', is reserved for use by the system. The user should not execute either an EX AF,AF' or an EXX instruction as these will have unfortunate consequences. (See Appendix XI for a full description.)

c. General

The logical values true and false are generally represented by # FF and #00 respectively. Often, however, any non-zero value is taken to mean true.

The bits in a byte are numbered 0..7, with bit 0 being the least significant bit and bit 7 being the most significant bit.

Where two byte (word) values are stored (in tables etc) they are always stored with the less significant byte first and the more significant byte second, unless a specific indication to the contrary is given. This is in accordance with the standard way the Z80 stores words.

Tables and the like are always laid out with byte 0 being the first byte of the table. When the address of such a table is given this is the address of byte 0 of the table unless otherwise indicated.

When the computer is turned on (or when it is reset) it completely initializes itself before running any program. This initialization is known as early morning startup, abbreviated to EMS from now on.

1.6 Routine Documentation.

Each routine described in this manual has entry and exit conditions associated with it. Where there are other points of interest about the routine these are normally given in a section after the entry and exit conditions. Such points include whether interrupts are enabled and a fuller description of the parameters and side effects of the routine.

There are two reasons for providing this information. Firstly it tells the user what will happen when the routine is called. Secondly it tells the user what a replacement routine is expected to do.

The entry conditions tell the caller of the routine what the routine expects to be passed to it. When calling a routine all values specified must 'be supplied. Values may only be left out where the routine documents that they are optional. When providing a replacement routine to fit this interface only information that is specified may be used, although not all of it need be used.

The exit conditions tell the caller what values the routine passes back and which processor registers are preserved. Registers that are documented as being corrupted may be changed by the routine or may not. The user should not rely on their contents. When providing a routine to fit this interface it is extremely important that registers documented as being preserved are indeed preserved and that the values returned are compatible with the original routine. Corrupting a register or omitting a result will usually cause the system to fail, often, in subtle and unexpected ways.

Often a routine will have different exit conditions depending on some condition or other (usually whether it worked or not). In these cases the specific differences in the exit conditions are given for each case and all conditions that remain the same irrespective of the case are given in a separate section (marked 'always').

There are abundant examples of routine interfaces in sections 15 to 18.

1.7 Example of Patching a Jumpblock.

The following is an example of how the jumpblocks may be used. At this stage many of the concepts introduced may be unfamiliar to the reader. However, since altering jumpblocks is an important technique for tailoring the system to a particular purpose the example is given here. Later sections will explain the actions taken here.

Suppose an assembler program is being written that is intended to use the printer when it is finished. While this program is being written it would save time and paper if the program could be made to use the screen instead of the printer. However, changing the program itself to use the screen could introduce bugs when it is changed back to using the printer. What is needed is a way of altering the action of the firmware that drives the printer - and this is what a RAM jumpblock is for.

The technique that will be used is to 'connect' the printer to a particular text window. This can be achieved by writing a short routine to send the character to the screen and patching the entry in the jumpblock for sending characters to the printer, MC PRINT CHAR, so that it jumps to this routine instead of its normal routine.

The substitute routine will have to obey the entry/exit conditions for MC PRINT CHAR. These can be found in the full description of this entry in section 15. Briefly they are as follows:

MC PRINT CHAR:

Entry conditions:

A contains character to print.

Exit conditions:

If the character was sent OK:
Carry true.

If the printer timed out:
Carry false.

Always:
A and other flags corrupt.
All other registers preserved.

The action of the substitute routine will be to select the screen stream that the printer output is to appear on, to print the character on the stream and then to restore the stream that was originally selected. To do this the substitute routine will need to call the routines TXT STR SELECT and TXT OUTPUT. Once again the full descriptions of these jumpblock entries can be found in section 15. The entry/exit conditions are as follows:

TXT STR SELECT:

Entry conditions:

A contains stream number to select.

Exit conditions:

A contains previously selected stream number.
HL and flags corrupt.
All other registers preserved.

TXT OUTPUT:

Entry conditions:

A contains character to print.

Exit conditions:

All registers and flags preserved.

The code for the substitute routine could be written as follows (stream 7 has been chosen as the stream on which printer output is to appear):

```
PUSH HL
PUSH BC
LD B, A                ;Save the character to print
LD A, 7                ;Printer stream number
CALL TXT_STR.SELECT   ;Select the printer stream
LD C, A                ;Save the original stream number
LD A, B                ;Get the character again
CALL TXT_OUTPUT        ;Send it to the screen
LD A, C                ;Get the original stream number
CALL TXT_STRSELECT    ;Reselect the original stream

POP BC
POP HL
SCF                    ;The character was sent OK
RET
```

Note the following points:

- 1/ MC PRINT CHAR preserves HL and BC. The routine above uses B and C for temporary storage and HL is corrupted by TXT STR SELECT. HL and BC are therefore pushed and popped to preserve them through the substitute routine.
- 2/ MC PRINT CHAR returns a success/fail indication in the carry flag. Since the routine above can never fail it always sets the carry flag to indicate success.
- 3/ The routine above does not change which text stream is selected. It selects the stream it is going to print on and restores the previously selected stream when it has printed the character. The firmware is written in such a way as to allow routines to restore the original state when they finish if required.

To use the substitute routine it is necessary to patch it into memory and to change the jumpblock entry for MC PRINT CHAR to jump to it. Assume that some memory at #ABOO has been reserved for the substitute routine and that the routine has been patched into memory. The MC PRINT CHAR entry in the jumpblock is at location #BD2B (as can be seen by inspecting section 13.1.8). The three bytes of the entry should be set to the instruction JP # ABOO by patching as follows:

#BD2B	#C3
#BD2C	#00
#BD2D	#AB

From now on all text sent to the printer will appear on the screen on stream 7. Of course, stream 7 should have its window set so that it does not interfere with any other streams using the screen.

This redirection will remain in force until the jumpblock entry is restored. This can be achieved by patching the jumpblock back again or by calling JUMP RESTORE or by causing an EMS initialization to take place by resetting the system.

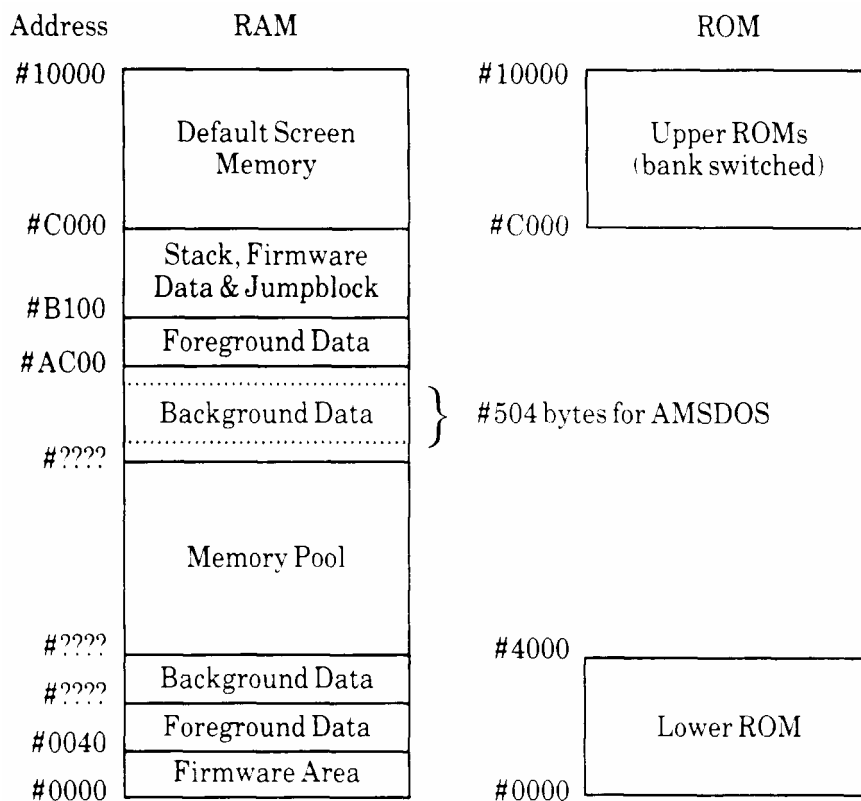
2 ROMs, RAM and the Restart Instructions.

The system has 32K of ROM and 64K of RAM in the Z80's 64K address space. To allow this the ROM can be enabled or disabled as required. Additional expansion ROMs can be selected giving up to 4128K of program area.

All the Z80 restart instructions, except for one, have been reserved for system use. RST 1 to RST 5 are used to extend the instruction set by implementing special call and jump instructions that enable and disable ROMs. RST 6 is available to the user.

2.1 Memory Map.

The memory map is complicated by the fact that into the Z80's address space of 64K bytes has been squeezed 64K bytes of RAM, 32K bytes of ROM and provision for ROM expansion of up to 252*16K (nearly 4M) bytes. The address space is divided as follows:



The sizes of the two background areas depend on the background ROMs fitted to the machine (see section 10).

The upper foreground data area need not have its lower bound at # AC00 but this is the default setting (as used by BASIC). The lower foreground data area

need only be reserved if it is needed (this area is not used by BASIC and is set to zero length). The memory pool left between the background data areas is also for the foreground program to use (see section 10).

The 32K of on-board ROM is split into two sections which are handled separately. Henceforth these will be discussed as if they were separate ROMs. The firmware resides in the lower ROM. The BASIC resides in the upper ROM. This upper ROM is bank switched so that up to 252 expansion ROMs (see section 10) can replace it in the memory map.

2.2 ROM Selection.

There are two mechanisms for switching ROMs in and out of the address space:

a. ROM State.

The upper and lower ROMs may be enabled and disabled separately. When the upper ROM is enabled data read from addresses between #C000 and #FFFF is fetched from the ROM. Similarly, when the lower ROM is enabled data read from addresses between #0000 and #3FFF is fetched from the ROM. When the ROMs are disabled data is fetched from RAM.

Note that the ROM state does not affect writing which always changes the contents of RAM.

b. ROM Select.

Expansion ROMs are supported by switching the upper ROM area between ROMs. Expansion ROMs are addressed by a separate ROM select address byte implemented in I/O space. ROM select addresses are in the range 0...251, providing for up to 252 expansion ROMs.

When the machine is first turned on it selects ROM zero. This will usually select the on board ROM, but an expansion ROM may be fitted at this address, which will pre-empt the on-board ROM.

See section 10 for a description of the use of expansion ROMs.

2.3 The Restart Instructions.

The Kernel supports the store map in a number of ways. In particular a variety of facilities are provided to handle subroutine addresses extended to include ROM select and/or ROM state information. Some of the restart instructions are used to augment the existing Z80 instruction set. The other restarts are reserved.

The firmware area between #0000 and #003F is set up so that the restarts operate whatever the current ROM state is. The user should not alter the contents of this area except as indicated in section 18.

The restarts are as follows. A fuller description of their operation can be found in section 18.

a. The Extended Instruction Set.

LOW JUMP (RST 1)

RST 1 jumps to a routine in the lower 16K of memory. The two bytes following

the restart are assumed to be a 'low address' - so RST 1 can be considered to be a three byte instruction, rather like a JP instruction.

The top 2 bits of the 'low address' define the ROM enable/disable state required; the bottom 14 bits give the actual address (in the range #0000 to #3FFF) to jump to once the ROM state is set up. When the routine returns the ROM state is restored to its original setting.

The firmware jumpblock, through which firmware routines should be called, makes extensive use of LOW JUMPS. These LOW JUMPS request the lower ROM to be enabled, so that the lower ROM may be disabled except when the firmware is active.

SIDE CALL (RST 2)

RST 2 calls a routine in an associated ROM. It has a very specialised use. A foreground program (see section 10) may require more than 16K of ROM. The side call mechanism allows for calls between two, three or four associated ROMs without reference to their actual ROM select addresses, provided that the ROMs are installed next to each other and in order.

The two bytes following the restart instruction give the 'side address' of the routine to call - so the RST 2 can be considered to be a three byte instruction, rather like a CALL instruction. The top 2 bits of the 'side address' specify which of the four ROMs to select; the bottom 14 bits, when added to #C000, give the actual routine address. The upper ROM is enabled, the lower ROM is disabled. Both the ROM state and the ROM select are restored to their original settings when the routine returns.

FAR CALL (RST 3)

RST 3 calls a routine anywhere in memory, in RAM or in any ROM. The two bytes following the restart are assumed to be the address of a 'far address'. The 'far address' is a three byte object, which takes the form:

Bytes 0..1: Actual address of routine to call.

Byte 2: ROM select/state required.

The ROM select/state byte may take the following values:

0. .251: Select the upper ROM at this ROM select address. Enable the upper ROM, disable the lower ROM.

252. .255: No change of ROM select, enable/disable ROMs as follows:

- 252: Enable upper ROM, enable lower ROM.
- 253: Enable upper ROM, disable lower ROM.
- 254: Disable upper ROM, enable lower ROM.
- 255: Disable upper ROM, disable lower ROM.

Note that the 'far address' is not itself contained in the 'instruction', but is pointed at. This is because the ROM select address will depend on the particular order in which the user has chosen to install expansion ROMs and must be established at run time.

Both the ROM state and the ROM select are restored to~ their original settings when the routine returns.

RAM LAM (RST 4)

RST 4 reads the byte from RAM at the address given by HL. It disables both ROMs before reading and restores the state afterwards. This 'instruction' avoids the user having to put a read routine into the central 32K of RAM to access RAM hidden under a ROM.

Writing to a memory location always changes the contents of RAM whatever the ROM enable state.

FIRM JUMP (RST 5)

RST 5 turns on the lower ROM and jumps to a routine. The two bytes following the restart are assumed to be the address to jump to - so RST 5 can be considered to be a three byte instruction, rather like a JP instruction. The lower ROM is enabled before jumping to the routine and is disabled when the routine returns. The state of the upper ROM is left unchanged throughout.

b. The Other Restarts.

RESET (RST 0)

RST 0 resets the system as if the machine has just been turned on.

USER RESTART (RST 6)

RST 6 is available for the user. It could be used to extend the instruction set in the same way that other restarts have been used, or it could be used for another purpose such as a breakpoint instruction in a debugger.

Locations #0030 to #0037 inclusive in RAM may be patched in order to gain control of the restart. If the lower ROM is enabled when the restart is executed then the code in ROM at this address causes the current ROM state to be saved in location #002B. Then the lower ROM is disabled and the firmware jumps to location #0030 in RAM. If the lower ROM is disabled then the restart calls #0030 as normal for this Z80 restart instruction.

INTERRUPT (RST 7)

RST 7 is reserved for interrupts (see section 11), it must not be executed by a program.

2.4 RAM and the Firmware.

The ROM state should be transparent to the user. If the current foreground program (see section 10) is in ROM then the normal ROM state is to have the upper ROM enabled and the lower ROM disabled. If the current foreground program is in RAM then the normal state is to have both ROMs disabled. These states allow the foreground program free access to the memory pool. When a firmware routine is called the lower ROM is enabled and the upper ROM is usually disabled. This allows the firmware free access to the default screen memory (but not to all the memory pool). When the firmware routine returns the ROM state is automatically restored to what it was.

The cases where the ROM state is important are:

a. Stack

The hardware stack should never be below #4000, otherwise serious confusion will occur when the lower ROM is enabled and the stack is used -for example, when interrupts occur or the firmware is called.

Similarly, it is inadvisable to set the stack above #C000 unless it is certain that the upper ROM is never enabled when the stack is in use.

The system provides a stack area immediately below #C000 which is over 256 bytes long. This should be adequate for most purposes.

b. Communication with the firmware.

Most firmware routines take their parameters in registers. However, some use data areas in memory to pass information. Most firmware routines that use data areas in memory read these directly without using RAM LAMs (see above) or the equivalent. These routines are affected by the ROM state and the RUM select. They will read data from a RUM if the RUM is enabled and the routine is given a suitable address. (Note that the jumpblock disables the upper ROM when the firmware is called). Other firmware routines that use data areas in memory always read from RAM. They are unaffected by the ROM state and the ROM select.

Routines that always access RAM will mention this in the description of the routine. Other routines may be assumed to be affected by the ROM state. In particular the various data blocks used by the Kernel must lie in the central 32K of RAM for the Kernel to be able to use them.

c. Communication between upper ROMs.

Programs in upper ROMs may call routines in other ROMs, using the various Kernel facilities. There is no provision in the firmware, however, for a program in one ROM to access constants in another.

The majority of firmware routines are called via the firmware jumpblock, which starts at location #BB00, in the firmware RAM area. The Kernel routines associated with the memory map are called via one of two other jumpblock areas: the LOW area between #0000 and #003F, and the HIGH area starting at #B900. All of these routines and jumpblocks are copied out of the lower ROM into the firmware RAM area when the Kernel is initialized. Thus they work independently of the ROM state.

2.5 Bank Switching

The ULA in the CPC6128 includes circuitry for bank switching 128K of RAM into the 64K memory map described in section 2.1. The bank switched RAM replaces the RAM in the memory map and behaves exactly like it; in particular, it is hidden when a ROM is enabled.

The 128K of bank switched RAM is split into 8 16K blocks, numbered 0..7. There are 8 memory organizations, also numbered 0..7, each of which switches a different set of four blocks into the memory map at #0000.. #3FFF, #4000.. #7FFF, #8000.. #BFFF and #C000.. #FFFF. The user can select an organization by calling KL BANK SELECT.

The blocks available in each organization are as follows:

Organization	Block accessed at memory address			
	#0000	#4000	#8000	#C000
0	0	1	2	3
1	0	1	2	7
2	4	5	6	7
3	0	3	2	7
4	0	4	2	3
5	0	5	2	3
6	0	6	2	3
7	0	7	2	3

During EMS the CPC6128 selects organization 0 and this is the organization normally associated with the firmware. Note that blocks 0 and 2 contain firmware variables, firmware jumpblocks and the stack. All these need to be in their correct places for the firmware to run.

The documentation for a number of firmware routines specifies that data blocks passed to them should be in the central 32K of memory. In most cases it does not matter which blocks are switched into the memory map at these places, however, the Kernel accesses data blocks passed to it (e.g. ticker blocks or RSX command tables) at various times (e.g. during interrupts or event processing) and it has no control over the bank switching at such times. It is up to the user to ensure that the Kernel is only passed data blocks that remain accessible. The simplest solution to this problem is to ensure that all Kernel data blocks are located in block 2 (between #8000 and #BFFF).

Organizations 4.. 7 are the firmware organization with a new block switched into the memory map at #4000. These organizations can be used to access programs or data stored in blocks 4..7.

Organizations 1.. 2 are used by CP/M Plus and are not really suitable for general use. In particular, if organization 2 is selected it is necessary to patch a program into RAM at #0038 to catch interrupts and to bank switch back to a more normal organization (e.g. organization 1) to run the standard interrupt code.

Organization 3 is also used by CP/M Plus but it is of interest since it has the RAM usually used for the screen located at #4000 where it can be accessed without disabling the upper ROM.

Bank switching has no effect on the CRTC. Base addresses #0000, #4000, #8000 and #C000 correspond to the screen being in blocks 0, 1, 2 and 3 respectively. It is not possible to locate the screen in blocks 4..7. However, the firmware routines for accessing the screen memory are affected by bank switching. For example, if a base address of #C000 is set in organization 3 then the firmware will have to be told (using SCR SET POSITION) that the screen memory can be accessed at #4000; if a base address of #4000 is set in organization 3 then the firmware will be unable to access the screen memory since block 1 is not in the memory map. (See section 6.4 for a full description of the screen memory map).

Organizations 4...7 can be used to set up a complete screen in one go by using SCR SET POSITION to make the firmware write to the memory at #4000 without sending a new base address to the screen hardware. Then, when the screen has been finished, the contents of this block can be quickly copied into the block actually being used by the CRTC (using KL LDIR perhaps). For example, a title screen could be set up and bank switched out of the way and then switched back in and copied at a later date when it is wanted.

3 The Keyboard.

The Key Manager is the pack associated with the keyboard. All the attributes of the keyboard are generated and controlled by the Key Manager. These attributes include repeat speed, shift and control keys, function keys and key translation. The joysticks are also scanned by the Key Manager.

The Key Manager has three levels of operation. The lowest level scans the keyboard, the middle level converts the key pressings into key values and the top level converts the key values into characters. The user may access the Key Manager at whichever level is most appropriate for a given program. It is usually unwise, however, for a program to mix accesses at different levels.

3.1 Keyboard Scanning.

The keyboard is completely software scanned. This scan occurs automatically every fiftieth of a second (see KM SCAN KEYS). The keyboard hardware is read and a bit map noting which keys are pressed is constructed. This bit map is available for testing if specific keys are pressed (see KM TEST KEY). As the bit map is constructed keys that are newly pressed are noted and markers are stored in a buffer until needed. If no newly pressed keys are found then the last key pressed may be allowed to repeat if it is still down (see section 3.5). The keyboard is 'debounced' in that a key must be released for two consecutive scans before it is marked as released in the bit map. This 'debouncing' hides multiple operations of the key switch as it opens or closes.

At this stage only four keys are treated specially. The two shift keys and the control key are not stored in the key buffer themselves. Instead, when any other marker is stored the states of the shift and control keys are noted and put into the buffer as well. The escape key generates a marker as normal but may also have other effects depending on whether the break mechanism is armed (see section 3.6).

There is a problem with scanning the keyboard. If three keys at the corners of a rectangle in the key matrix are all pressed at the same time then the key at the fourth corner appears to be pressed as well. There is no way to avoid this problem as it is a feature of the keyboard hardware. All key combinations used by the firmware (and the BASIC) have been especially designed to avoid this effect.

3.2 Key Translation.

When the user asks for a key (KM WAIT KEY or KM READ KEY) the next key pressed marker is read from the key buffer. The marker is converted to a key number and this is looked up in one of three translation tables.

Which table is used depends on whether the shift and control keys were pressed when the key was pressed. One table is used if the control key was pressed, another is used if either shift key was pressed but control was not, the third is used if neither shift nor control keys were pressed. The contents of these tables can be altered by the user as required (by calling KM SET CONTROL, KM SET SHIFT and KM SET TRANSLATE respectively).

The value extracted from the table may be a system token, an expansion token or a character. Expansion tokens and characters are used by the top level of the Key Manager (see 3.3 below) and are passed up from the middle level when they are found in a table. There are three system tokens, which are obeyed immediately they are found in a table. After obeying the token the next marker is read from the buffer and translated.

The default translation tables are described in Appendix II. The immediately obeyed System tokens are:

a. Ignore (# FF)

The key pressed is to be ignored.

b. Shift lock (# FE)

The shift lock is to be toggled (turned on if it is currently off and turned off if it is on).

c. Caps lock(#FD)

The caps lock is to be toggled (turned on if it is off and off if it is on).

3.3 Characters from the Keyboard.

When the user asks the top level for a character (KM WAIT CHAR or KM READ CHAR) a key is fetched from the middle level. If this is a character (#00.. #7F or #A0.. #FC) then it is passed to the user. If it is one of the 32 expansion tokens (#80..#9F) then the string associated with the token is looked up. The characters in this string are passed to the user one at a time with each request for a character until the end of the string is reached.

There is only one character with a special meaning at this level. This is character #EF which is produced when pressing the escape key generates a break event (see section 3.6). It has no effects, it is merely a marker for the place in the buffer where a break event was generated. It is intended to be used to allow all characters before the break to be discarded. This character is not generated by the translation tables and thus cannot be changed by altering them.

A single 'put back' character is supported. When the user puts back a character this character will be returned by the next call to the top level of the Key Manager. This is intended for use by programs that need to test the next character to be read from the keyboard without losing it (when processing breaks perhaps).

In V1.1 firmware it is possible to call KM FLUSH to discard any unused or unwanted characters so that subsequent calls to KM READ CHAR or KM READ KEY will not return values from a previous input. The same effect can be achieved in V1.0 Firmware by repeatedly calling KM READ CHAR until it returns with carry false to indicate that there are no more characters available.

3.4 Shift and Caps Lock.

a. Shift lock

When shift lock is engaged then the keys pressed are translated as if a shift key is pressed.

The shift lock is toggled by a system token (see 3.2 above) which is normally generated by pressing CTRL and CAPS LOCK.

b. Caps lock

When caps lock is engaged then alphabetic characters read from the keyboard are converted to their upper case equivalents. This case conversion is applied before expansion tokens are expanded and so expansions are not capitalised.

The caps lock is toggled by a system token (see 3.2 above) which is normally generated by pressing CAPS LOCK (without control).

In V1.1 firmware it is possible to set the state of the locks as if the SHIFT or CAPS LOCK keys had been pressed by calling KM SET LOCKS.

3.5 Repeating keys.

There is a table, which the user can alter as desired, that specifies which keys are allowed to repeat when held down (see KM SET REPEAT). The default setting for this table is described in Appendix III. Briefly, the default is to allow all keys to repeat except the ESC, TAB, CAPS LOCK, SHIFT, ENTER and CTRL keys and the 12 keys in the numeric keypad (the function keys). -

The speed at which keys repeat and the delay before the first repeat can be set by the user (see KM SET DELAY), The default speed produces up to 25 characters a second with a 0.6 second start up delay.

A key is allowed to repeat if the following conditions are satisfied:

- 1/ The appropriate time has passed since the key was first pressed or it last repeated.
- 2/ The key is still pressed.
- 3/ No other key has been pressed since the key was first pressed.
- 4/ The key is marked as allowed to repeat in the repeat table.
- 5/ There are no keys stored in the key buffer.

Condition 5 above means that the repeat speed and start up delay set the maximum speed at which a key is allowed to repeat. If a program is slow about removing keys from the buffer then the generation of keys will adjust itself to this. Thus it is impossible to get a large number of keys stored in the buffer simply by holding a key pressed.

3.6 Breaks.

Breaks can occur when the keyboard scanner detects that the ESC key is pressed. When the escape key is found to be pressed the indirection KM TEST BREAK is called to deal with the break. The default setting for this routine tests whether the SHIFT, CTRL and ESC keys and no others are pressed. If so then the system is reset (by executing an RST 0), otherwise the break mechanism is invoked.

If the break mechanism is disarmed then no action is taken other than the normal insertion of the marker for the escape key into the key buffer. If the break mechanism is armed then two additional operations take place. Firstly, a special marker is placed into the key buffer that will generate character #EF when it is found (irrespective of the translation tables). This is intended to be used to allow the characters which were in the buffer before the break occurred to be discarded. Secondly, the synchronous break event is 'kicked'.

The break mechanism can be armed or disarmed at any time (by calling KM ARM BREAK or KM DISARM BREAK). The default state is disarmed. When a break is detected the mechanism is disarmed automatically which prevents multiple breaks from occurring.

The method BASIC uses to handle breaks should serve as a model for other programs. BASIC's actions are as follows:

The break mechanism is armed. After each BASIC instruction the synchronous event queue is polled and if a break event is found (because it has been kicked as explained above) the break event routine is run.

The break event routine stops sound generation (SOUND HOLD) and then it discards all characters typed before the break occurred by reading characters from the keyboard (KM READ CHAR) until either the buffer is empty or the break event marker (character #EF) is found. BASIC then turns the cursor on (TXT CUR ON) and waits for the next character to be typed (KM WAIT CHAR).

If the next character is the escape token (character # FC - the default value generated by the ESC key) then a flag is set to make BASIC abandon execution (or run the user's ON BREAK GOSUB subroutine) and the break event routine returns.

If the next character is any character other than escape then the break will be ignored. If it is any character other than space then this is 'put back' (KM CHAR RETURN). Before the event routine returns the cursor is turned off (TXT CUR OFF), sound generation is restarted (SOUND CONTINUE) and the break mechanism is rearmed. BASIC then continues as if nothing had happened.

When reading or writing from the cassette the ESC key is handled in a different manner which is described in section 8.12.

3.7 Function Keys and Expansion Tokens.

The Key Manager allows for 32 expansion tokens (values #80.. #9F) which may be placed in the key translation tables. Each token is associated with a string which is stored in the expansion buffer.

When the user asks the top level for a character a key is fetched from the middle level. If this key is a character it is passed straight back. However, if it is an expansion token then the string associated with the token is looked up. The characters in this string are passed out one at a time with each request for a character until the end of the string is reached. Values #80.. #9F and #EF, #FD.. #FF in the expansion string are treated as characters and are not expanded or obeyed.

The user may set the string associated with an expansion token (see KM SET EXPAND) and may cause any key on the keyboard to generate an expansion token. The default settings for the expansion tokens and the keys with which they are normally associated are given in Appendix IV. The user may also set the size and location of the expansion buffer (see KM EXP BUFFER); the default buffer is at least 100 bytes long.

3.8 Joysticks.

There may be two joysticks connected to the system. These are both scanned in the same way as keys on the keyboard. Indeed, the second joystick occupies the same locations in the key matrix as certain other keys and is indistinguishable from them. The state of the joysticks can be determined by calling the routine KM GET JOYSTICK.

Because the joysticks are scanned like keys the pressing of joystick buttons can be detected like any other key. Firstly, individual direction or buttons can be tested in the key bit map (see section 3.1) by calling KM TEST KEY. Secondly, the joystick buttons generate characters when they are pressed (providing the translation tables are set suitably) and these characters can be detected. The major problem with this latter method is that the rate of generation of characters depends on how fast the keyboard is set to repeat. If the repeat speed is increased to make the joystick more responsive then the keyboard may become impossible to use.

See Appendix I for the numbering of the keys and joystick buttons and see Appendix II for the default translation tables.

4 The Text VDU.

The Text VDU is a character based screen driver. It controls 8 different streams each of which can have an area of screen allocated to it (a window). The Text VDU allows characters to be written to the screen and read from the screen. It also treats certain 'characters' as 'control codes' which can have various effects, from moving the cursor to setting the colour of an ink.

4.1 Text VDU Coordinate Systems.

The Text VDU uses two coordinate systems - logical and physical. Generally the user specifies positions to the Text VDU in logical coordinates. Physical coordinates are used internally and occasionally by the user to specify positions to the Text VDU. Both systems use signed 8 bit numbers and work in character positions. Each character position is 8 pixels (dots) wide and 8 pixels high. This means that the position of a coordinate on the screen depends upon the screen mode.

Physical coordinates have columns running left to right and rows running top to bottom. The character position at the top left corner of the screen is row 0, column 0.

Logical coordinates are similar to physical coordinates except that the character position at the top left corner of the text window is row 1, column 1.

4.2 Streams.

The Text VDU has facilities for handling up to 8 streams at once. Each stream has an independent state (although some facilities are shared and thus affect all streams when altered). The features that are stream dependent are:

- VDU enable.
- Cursor enable (enable or disable, on or off).
- Cursor position.
- Window size.
- Pen and paper inks.
- Character write mode (opaque or transparent).
- Graphics character write mode.

The features that affect all streams include:

- Character matrices.
- Control code buffer.
- Text VDU indirections.
- Screen mode.

All these features are explained in detail in the sections below.

At anytime, the stream which is currently selected may be changed without adverse effects provided that the control code buffer is not in use (see section 4.7 for further explanation). A stream will remain selected until another stream is selected. This means that a program need not know which stream it is using.

The default stream, selected at EMS, is stream 0.

BASIC extends the stream concept to include the printer and cassette/disc files. This extension is not part of the firmware.

4.3 Text Pen and Paper Inks.

Each stream has a pen and a paper ink associated with it. The text pen ink is used to set the foreground pixels in characters (see section 4.6). The text paper is used to set the background pixels in characters and to clear the text window. The pens and papers can be set to any ink that is valid in the current screen mode (see section 6.1). The default setting for a stream has the paper set to ink 0 and the pen set to ink 1. Changing a pen or paper ink does not change the screen; it merely alters how characters will be written in the future.

4.4 Text Windows.

Each stream has a text window associated with it. This window specifies the area of the screen where the stream is permitted to write characters. This allows different streams to use different portions of the screen without interfering with each other.

Windows are trimmed so that they fit within the current screen (whose size varies with the screen mode, see section 6.1). The smallest size window allowed is 1 character wide and 1 character high.

Before writing to the screen the position to write at is forced to lie inside the window (see 4.5 below). This may cause the window to roll. Other operations, such as obeying certain control codes also cause the write position to be forced inside the window.

A text window which does not cover the whole screen is rolled by the firmware copying areas of screen memory around. There is no alternate method available. This makes rolling large windows a fairly time consuming process.

A text window which covers the whole screen is rolled by using the hardware rather than by copying areas of memory. The offset of the start of the screen in the screen memory can be set (see section 6.4). By changing this offset by +80 or -80 the whole screen can be rolled up or down by a line of characters.

It is obviously a good idea to prevent windows that are being used from overlapping. If they are allowed to overlap then the portion in multiple use will merely contain whatever was written to it last. There is no precedence of windows one over another. A window occupying the whole screen will overlap the other windows and so if this window is rolled it will move the contents of the other windows.

The default windows, set up at EMS and after changing screen mode, cover the whole of the screen. All eight windows overlap.

4.5 The Current Position and the Cursor.

Each stream has a current position associated with it. This is where the next character to be printed on the screen is expected to be placed. However, if, when a character is to be printed, the current position is found to lie outside the text window then it is forced inside. The following steps are applied in turn to force the current position inside the window:

- 1/ If the current position is left of the left edge of the window then it is moved to the right edge of the window and up one line.
- 2/ If the current position is right of the right edge of the window then it is moved to the left edge of the window and down one line.
- 3/ If the current position is now above the top line of the window then it is moved to the top line of the window and the contents of the window are rolled down one line.
- 4/ If the current position is now below the bottom line of the window then it is moved to the bottom line of the window and the contents of the window are rolled up one line.

When the cursor is enabled, the current position is marked by the cursor blob. However, before placing the cursor blob on the screen, the current position is forced to lie inside the current window just as it is before a character is placed on the screen. This may cause the current position to move.

If the cursor is disabled then the current position may lie outside the window and it will not be forced inside the window until, for example, a character is printed.

The current position can be changed directly (by calling `TXT SET CURSOR`, `TXT SET ROW` or `TXT SET COLUMN`) or by sending control codes to the Text VDU. The location the current position is moved to is not forced inside the window immediately, but only when the window is to be written to, as described above. This allows the current position to be changed by moving via a position outside the window, if required.

There are two ways to disable the cursor and prevent the cursor blob from appearing on the screen. The first, `cursor on off`, is intended for use by system programs. This is used by BASIC, for example, to hide the cursor unless input is expected. The second, `curs& enable disable`, is intended for use by the user. The cursor blob will only be placed on the screen if it is both on and enabled.

In V1.1 firmware it is possible to interrogate the current enable disable states of the VDU and cursor for the current stream using `TXT ASK STATE`.

The cursor blob is normally an inverse patch. The character at the cursor position is displayed with the text pen and paper inks reversed. This makes it easy to restore the original form of the character position if the cursor is moved. It is possible for the user to alter the form of the cursor blob, if required, by changing the indirections `TXT DRAW CURSOR` and `TXT UNDRAW CURSOR`.

4.6 Characters and Matrices.

A character is displayed on the screen in an area 8 pixels (dots on the monitor) wide and 8 pixels high. Thus the maximum number of characters on the screen depends upon the screen mode, (see section 6.1). Each character has a matrix which is an 8 byte vector that specifies the shape of the character. The first byte of the vector refers to the top line of the character and the last byte to the bottom line of the character. The most significant bit of a byte in the vector refers to the leftmost pixel on a line of the character and the least significant bit refers to the rightmost pixel on a line of the character. If a bit in the matrix is set then the pixel is in the foreground. If a bit is clear then the pixel is in the background.

A foreground pixel in the character is always set to the pen ink. The treatment of a background pixel depends on the character write mode of the VDU. In the default mode, opaque mode, background pixels are set to the paper ink. There is another mode, transparent mode, in which the background pixels are not altered. Thus, in transparent mode, the character is written over the top of the current contents of the screen. This is useful for annotating pictures or generating composite characters.

The Text VDU is capable of printing 256 different characters, although special effort is required to print the first 32 characters which are usually interpreted as control codes. The matrices for the characters are normally stored in the ROM but the user may arrange for any number of the characters to have matrices stored in RAM where they may then be altered. The default setting, at EMS, is to have all the matrices in ROM. (BASIC takes special action during its own initialization to create 16 'user defined' matrices.) The default character set is described in Appendix VI.

When the user sets up a table of user defined matrices, by calling `TXT SET M TABLE`, it is initialized with the current settings of the matrices from ROM or RAM. This means that extending the table does not alter the current matrices. Contracting the table will make the characters lost revert to their default matrices in ROM.

When characters are read from the screen (by calling `TXT RD CHAR`) the pixels on the screen are converted to the form of a matrix. This is compared with the current character matrices to find which character it is. This means that changing the character matrices or altering the screen may make a character unrecognisable, in particular, changing the pen or paper ink can cause confusion. Usually these problems result in the character appearing to be a space (character #20) and so special precautions are taken to avoid generating spaces. After some ink changes real spaces may be read as block graphic characters #80 or #8F.

To allow the user to change how characters are written to and read from the screen, the indirections `TXT WRITE CHAR` and `TXT UNWRITE` are provided.

4.7 Character Output and Control Codes.

The main character output routine for the Text VDU is TXT OUTPUT. This obeys controls codes (characters 0..31) and prints all other characters. Characters sent to TXT OUTPUT pass through various levels of indirection and can be dealt with by various output routines.

TXT OUTPUT uses the TXT OUT ACTION indirection to sort out whether the character is a printing character, is a control code to be obeyed or is the parameter of a control code.

TXT OUT ACTION normally calls TXT WRITE CHAR to print characters on the screen. However, if the graphic character write mode is selected then characters are printed using the Graphics VDU character write routine (see 5.6 below). This mode can be selected on a character by character basis using a control code or on all characters sent (see TXT SET GRAPHIC). When graphic character write mode is selected control codes are not obeyed but are printed by the graphics routine instead.

TXT OUT ACTION deals with a control code in the following manner:

- 1/ The code is stored at the start of the control code buffer.
- 2/ The code is looked up in the control code table to find out how many parameters it requires.
- 3/ If no parameters are required go directly to step 5.
- 4/ If one or more parameters are required then TXT OUT ACTION returns but the next characters sent to it are added to the control code buffer rather than being printed or obeyed. This continues until sufficient parameter characters have been received.
- 5/ The code is looked up in the control code table to get the address of the routine to call to perform the control code and this routine is then executed.
- 6/ The control code buffer is discarded and the next character sent may be printed or may be the start of a new control code sequence.

The user can change the operation of a control code by changing the entry for it in the control code table (see TXT GET CONTROLS). This contains a 3 byte entry for each code and entries are stored in ascending order (i.e. the entry for #00 first, #01 next and so on).

Bits 0...3 of the first byte of each entry specifies the number of parameters required. This must lie in the range 0..9 as the control code buffer is only capable of storing up to 9 parameters.

In V1.1 firmware bit 7 specifies whether the code is affected when the VDU is disabled. If bit 7 is set then the code is to be ignored when the VDU is disabled, otherwise it is to be obeyed.

The second and third bytes are the address of the routine to call to obey the code. This routine should lie in the central 32K of RAM or in the lower ROM (which will be enabled). It should conform to the following entry/exit conditions:

Entry:

A contains the last character added to the buffer.

B contains the number of characters in the buffer (including the control code).

C contains the same as A.

HL contains the address of the control code buffer (points at the control code).

Exit:

AF, BC, DE and HL corrupt.
All other registers preserved.

The control code buffer is shared between all the streams. A control code sequence should be completed before the stream is changed otherwise unexpected effects may occur.

The default control code actions, set at EMS and when TXT RESET is called, are described in Appendix VII.

It is possible to disable a text stream by calling TXT VDU DISABLE. When disabled the stream will not write any characters to the screen and in V1.1 firmware control codes may not be obeyed (as described above). Normal operation can be restored by calling TXT VDU ENABLE. Note, however, that calling these routines will empty the control code buffer. This effect may be used to avoid problems when the state of the control buffer is unknown (when printing an error message perhaps).

5 The Graphics VDU

The Graphics VDU allows individual pixels (dots) on the screen to be set or tested and lines to be drawn. The plotting takes place on an ideal screen that is always 640 points wide and 400 points high. This means that more than one point on the ideal screen will map onto a particular pixel on the real screen. The width of the ideal screen (640 points) is chosen to be the horizontal number of pixels on the screen in the highest resolution mode (mode 2). The height of the ideal screen (400 points) is chosen to be twice the vertical number of pixels on the screen in all modes. This ensures that the aspect ratio of the screen is approximately unity, i.e. a circle looks circular and not elliptical.

5.1 Graphics VDU Coordinate Systems.

The Graphics VDU uses 4 coordinate systems. The user specifies positions in user coordinates or relative coordinates or occasionally in standard coordinates. Internally the Graphics VDU uses base coordinates (or occasionally standard coordinates).

User coordinates, relative coordinates and standard coordinates are all very similar. They all use signed 16 bit numbers and work in points with X-coordinates running left to right and Y-coordinates running bottom to top. The screen is always 400 points high and 640 points wide whatever the screen mode is. This means that a pixel (dot on the screen) is mapped onto by 8 points in mode 0, 4 points in mode 1 and 2 points in mode 2. The origin (coordinate (0,0)) of these systems vary:

In standard coordinates the origin is the point at the bottom left corner of the screen.

The origin of user coordinates can be set by the user. The default origin is at the bottom left corner of the screen. This makes the default user coordinates the same as standard coordinates.

The origin of relative coordinates is the current position (see 5.2 below). This allows plotting to be carried out independently of the position on the screen and is useful if a particular shape is to be repeated on the screen a number of times or if it is inconvenient to keep track of the current location.

Base coordinates are a physical coordinate system which deals with pixels. X-coordinates run left to right and Y-coordinates run bottom to top. Pixel (0,0) is the pixel at the bottom left corner of the screen. Because this coordinate system works in pixels the coordinates of positions on the screen depend upon the screen mode. Base coordinates are unsigned 16 bit numbers and only coordinates that refer to a pixel on the screen are valid.

Graphics routines convert from relative coordinates to user coordinates, if necessary, and then from user coordinates to base coordinates before accessing the physical screen. During the latter conversion there is a loss of accuracy because of the mapping of multiple points onto a single pixel. This could make shapes drawn on the screen appear asymmetrical (particularly circles) but the Graphics VDU avoids this by rounding the coordinates towards the user origin. Thus symmetrical shapes should be drawn symmetrically about the user origin to take advantage of the rounding. If the shape is not centred on the user origin then the asymmetry may reappear.

In V1.1 firmware it is possible to call GRA FROM USER to convert from user to base coordinates - this will make using routines in the lower level screen pack easier (e.g. SCR DOT POSITION, SCR HORIZONTAL, SCR VERTICAL).

5.2 The Current Graphics Position.

The Graphics VDU stores a current position. This is the user coordinate of the last point specified to the Graphics VDU (or the origin after clearing the graphics window). The origin of relative coordinates is specified to be at this point, thus relative coordinates are an offset from the current position.

When drawing a line one end is at the position specified while the other end is at the current graphics position. When drawing a character on the screen using the graphics character write routine the character is placed with the current graphics position being the top left corner of the character.

After plotting or testing a point or drawing a line the current graphics position is moved to the position specified. After writing a character the current graphics position is moved right by the width of a character ready for the next character to be written.

5.3 Graphics Pen and Paper Inks.

The Graphics VDU has a pen (foreground) ink and a paper (background) ink. The graphics pen ink is used to plot pixels and to set foreground pixels when writing characters (see 5.6 below) and drawing lines (see 5.7 below). The graphics paper ink is used to clear the graphics window and to set background pixels when writing characters or drawing lines.

The pen and paper can be set to any ink valid in the current screen mode (see section 6.2). The default has the paper set to ink 0 and the pen set to ink 1. Changing the pen or paper ink does not change the screen it merely alters how pixels will be written in the future.

5.4 Graphics Write Mode.

Pixels plotted by the Graphics VDU are plotted using the current graphics write mode. This specifies how the ink to be plotted interacts with the ink a pixel is currently set to.

There are four write modes:

0: FORCE:	NEW	=	INK
1: EXCLUSIVE-OR:	NEW	=	INK xor OLD
2: AND:	NEW	=	INK and OLD
3: OR:	NEW	=	INK or OLD

NEW is the ink that the pixel will be set to.

OLD is the ink that the pixel is currently set to.

INK is the ink that is to be plotted.

The default Graphics write mode is FORCE mode. The Text VDU always sets pixels as if it is operating in this mode. Also the graphics window is cleared by writing in FORCE mode irrespective of the actual write mode.

Provided that suitable ink settings are chosen, AND mode and OR mode allow particular bits in a pixel to be cleared or set. This allows the Graphics VDU to write in 'bit planes' and by choosing the colours of the inks carefully overlapping shapes can be drawn and automatically hidden behind one another.

If the inks are chosen suitably, EXCLUSIVE-OR mode can be used to plot over the current contents of the screen. It is also useful because a shape can be removed from the screen by redrawing it since exclusive-oring with the same ink twice restores the original setting of a pixel.

The graphics write mode may be set by calling SCR ACCESS or by using a control code (see Appendix VII).

5.5 Graphics Window.

The Graphics VDU allows a single window to be specified. This allows the user to mix text and graphics on the screen without them interfering with each other. If the text windows are allowed to overlap the graphics window then the contents of the graphics window will be moved when the text windows are rolled. The graphics window cannot be rolled.

When plotting points, drawing lines or writing characters no pixel outside the graphics window is ever written. Unlike the text windows no action is taken to force a point inside the window - actions outside the window will be lost. Conversely, when testing points, points outside the window are all deemed to be set to the current graphics paper ink. Points inside the window are written and read as expected.

The graphics window is related to a specific area of the screen and not to the user coordinate system. Thus, changing the origin of the user coordinate system will not move the location of the window on the screen although it does change the user coordinates of points in the window.

The default graphics window, set at EMS and after changing screen mode, covers the whole of the screen.

5.6 Writing Characters.

The Graphics VDU write character routine draws a character with the current graphics position at the top left corner of the character. The current position is moved right by the width of a character in the current screen mode. The distance moved varies; in mode 0 it is 32 points; in mode 1, 16 points; and in mode 2, 8 points. Control codes, characters 0..31, are printed and are not obeyed.

In V1.0 firmware the character is always written opaquely irrespective of what mode the Text VDU is using to write characters, i.e. The character background is set to the graphics paper ink and the foreground is set to the graphics pen ink. However, the current graphics write mode is used to plot the pixels in the character (see 5.4 above).

In V1.1 firmware the foreground to the character is always written in the graphics pen ink using the current graphics write mode. How the background to the character is written depends on the background write mode set by calling GRA SET BACK. If the background mode is opaque then the background to the character is written in the graphics paper ink using the current graphics write mode. If the background mode is transparent then the background pixels are not plotted at all, the current settings of these pixels are left unchanged.

5.7 Drawing Lines (Only applicable to V1.1 firmware)

The Graphics VDU has a number of options that affect the way that lines are drawn on the screen. These include the background write mode, the first point mode, the line mask and the graphics write mode.

The line mask is an 8 bit, bit significant value that specifies whether pixels on the line are foreground or background. If the bit of the mask corresponding to a pixel is one then the pixel is foreground. If the bit is zero then the pixel is background. The mask is used repeatedly along the length of the line. i.e. Bit 7 of the mask corresponds to pixels 1, 9, 17, ... of the line, bit 6 to pixels 2, 10, 18, ... etc.

Foreground pixels of a line are plotted in the graphics pen ink using the graphics write mode. How background pixels are plotted depends upon the background write mode. If the background mode is opaque then the background pixels are plotted in the graphics paper ink using the graphics write mode. If the background mode is transparent then the pixels are not plotted at all, the current settings of the pixels are left unchanged.

The first point mode specifies whether the pixel at the current graphics position is to be included in the line or not. Not plotting the first pixel of a line is particularly useful for drawing lines when the graphics write mode is XOR. For instance, the corner pixels of a box will be plotted twice if the first pixels of lines are plotted and this will result in the pixels not being set in XOR mode.

5.8 Filling Areas (Only applicable to V1.1 firmware)

The Graphics VDU provides a generalised area fill routine. The user draws the outline of the area to be filled using the standard line drawing and pixel plotting routines and then moves the current position to any pixel inside the area and calls the fill routine, GRA FILL. This will set all pixels inside the delimited area to the fill ink irrespective of what the current graphics write mode is set to. (i.e. GRA FILL always works as if FORCE mode was selected).

The fill routine recognises pixels making up the edge of the area by the ink that they are set to. All pixels set to either the current graphics pen ink or the ink that is being used to fill the area with are treated as edge pixels. (These two inks may be the same). Edge pixels need only lie diagonally adjacent to each other, it is not necessary for them to be orthogonally adjacent. Of course, the edges of the graphics window are also treated as edges of the area to fill and so no pixels outside the graphics window will be affected.

GRA FILL uses a buffer supplied by the user for storing information. The more complicated the area the longer the buffer will need be. If the supplied buffer is too short then parts of the area will be ignored and will not be filled. By supplying a long enough buffer any arbitrarily complicated shape may be filled.

6 The Screen Pack.

The Screen Pack is used by the Text and Graphics VDUs to access the hardware of the screen. It also controls the features of the screen that affect both the Text VDU and Graphics VDU, such as what mode the screen is in.

6.1 Screen Modes.

The screen has three modes of operation, numbered 0, 1 and 2. The modes have different resolutions and display different numbers of inks on the screen.

All modes have a vertical resolution of 200 pixels (picture elements or dots on the screen). The horizontal resolution varies from 160 pixels to 640 pixels. As each character is 8 pixels by 8 pixels the number of characters across the screen varies with the mode - from 20 characters to 80 characters. The screen is always 25 characters high.

The number of inks that can be displayed on the screen varies with the screen resolution. When the screen is 640 pixels wide only 2 inks can be displayed, when the screen is 320 pixels wide 4 inks can be displayed and when the screen is 160 pixels wide 16 inks can be displayed.

In summary, the modes are:

Mode	Pixel Size	Character Size	Inks
0	160x200	20x25	16
1	320x200	40x25	4
2	640x200	80x25	2

The default screen mode, set at EMS, is mode 1.

The screen mode is set by calling SCR SET MODE which also has other effects.

Firstly, the screen is cleared to ink 0. If the text and graphics paper inks are not set to ink 0 then this will become apparent on the screen when characters are written or windows are cleared. If the user wishes to alter this screen clearing operation for some reason then it may be intercepted at the SCR MODE CLEAR indirection.

Secondly, the Text and Graphics VDUs are set into standard states. The windows are all set to cover the whole screen. If the pen and paper inks are out of range for the new mode then they are masked (with #01 or #03) to bring them into range. The current text positions are moved to the top left corner of the screen and the text cursors are turned off (see TXT CUR OFF). The current graphics position and the user origin are moved to the bottom left corner of the screen.

6.2 Inks and Colours.

The various screen modes allow pixels (dots on the screen) to be set to different numbers of inks as follows:

Mode 0: 16 inks, 0..15
Mode 1: 4 inks, 0..3
Mode 2: 2 inks, 0..1

How the ink for a pixel is encoded into a byte of screen memory is described in section 6.4. The ink that a pixel is set to determines what colour the pixel is displayed in. However, the colour associated with an ink is not fixed, it can be changed.

There are 27 colours available. Each ink may be set to any of these colours. The border to the screen acts much like an ink and can have its colour specified as well. The display hardware fetches the ink value from the screen memory for each pixel as it is displayed. This ink value is used to access a small area of RAM inside the gate array called the 'palette'. The palette contains the actual colour which is to be displayed by the monitor for that particular ink. Changing entries in the palette thus causes all pixels set to that ink to change colour when they are next displayed (i.e. within 1/50th of a second or so).

In fact the Screen Pack allows two colours to be associated with an ink (or the border). These are loaded into the palette alternately under software control. If the two colours associated with an ink are different then the ink will flash, if the colours are the same then the ink will be steady. The user can change the rate of alternation, from the default of 5 cycles per second, if required (see SCR SET FLASHING).

When specifying colours the Screen Pack uses an ordering that corresponds to a grey scale on a monochrome monitor. This runs from the darkest colour (black), colour 0, to the brightest colour (bright white), colour 26. The colours do not appear to have any particular ordering when viewed on a colour monitor.

The palette uses a different (and apparently nonsensical) numbering scheme for the colours. The Screen Pack automatically translates the grey scale number to the hardware number and vice versa when appropriate. Unless the user is driving the hardware directly the hardware numbers will never be encountered.

The default settings for the colour of each ink and a list of the 27 colours available are given in Appendix V.

6.3 Screen Addresses.

The Screen Pack does not use a coordinate system itself. It uses screen addresses. However, it does work with the physical and base coordinate systems of the Text and Graphics VDUs described in sections 4.1 and 5.1 respectively. In particular, routines are provided to convert positions given in physical or base coordinates to screen addresses.

A screen address is, prosaically enough, the address of a byte within the screen memory. To specify a particular pixel a screen address is often passed to a routine along with a mask that indicates exactly which pixel is required. Routines are provided for stepping a screen address up, down, right and left one byte. (The screen map makes this a non-trivial operation.)

6.4 Screen Memory Map.

The screen is a memory mapped pixel screen. The screen memory fills 16K of RAM in all modes. The default location for the screen, set at EMS, is the 16K of RAM starting at #C000. This lies underneath the upper ROM, when it is enabled, which keeps the screen out of the way of the rest of the system. However, this also means that the upper ROM has to be disabled whenever the screen is read. The firmware jumpblock uses LOW JUMP restarts which turn the upper ROM off to ensure that the screen memory is accessible if required.

It is possible to change the location of the screen memory to any of the 4 16K memory blocks on 16K boundaries (see SCR SET BASE). However, only #C000 and #4000 are useful; #0000 and #8000 both overlap firmware jumpblocks or other system areas. The descriptions below all assume the default screen location at #C000.

In V1.1 firmware it is possible to set the location of the screen that is used by the screen pack routines independently of setting the hardware value. This will then enable text and graphics to be produced in the 'screen' that isn't currently on view - switching to the other possible location (#4000 to #C000) will cause the already prepared graphics etc. to instantly appear - thus avoiding flicker and enabling smooth animation effects.

The screen memory map is not simple. Fortunately it is not necessary to understand it because the Text and Graphics VDUs provide idealised models of the screen. However, to achieve maximum speed for certain applications (such as animated games) it may be necessary to access the screen memory directly.


The screen memory is divided into 8 blocks, each 2K bytes long. Block 0 runs from # C000 to # C7FF, block 1 runs from #C800 to #CFFF, and soon. Each line of pixels on the screen uses 80 consecutive bytes from a block. The top line of the screen comes from block 0, the second line from block 1 and so on until the eighth line which comes from block 7. The sequence starts with block 0 again on the ninth line and repeats in this fashion all the way down the screen. The successive lines in a block are stored consecutively so there are 48 unused bytes at the end of each block.

There is a further complication to this screen map. The description above assumes that the first byte displayed from the block is the first byte of the block. In practice the offset in a block of the first byte to be displayed can be set to any even value (see SCR SET OFFSET). The same offset applies to all eight blocks. A block wraps around from its last byte to its first byte, thus #C7FE, #C7FF and #C000 are consecutive bytes in block 0 and could all be on the same line of the screen. Altering the offset by $\pm 80 \text{ MOD } 2048$ (the length of a line) rolls the screen up or down by one character line (8 pixel lines). This effect is used by the Text VDU when rolling the entire screen.

The meaning of the bytes accessed as described above varies with the screen mode. Each byte stores the inks for 2,4 or 8 pixels. The bits used to encode each pixel are not arranged in an obvious manner. The following table specifies which bits of screen memory are used to encode which pixel in the various modes. The bit numbers given in the table are the bits of the screen byte. They are given in the order of bits in the pixel - the first bit given is most significant bit of the pixel and the last bit is the least significant bit.

	Mode 0	Mode 1	Mode 2
Leftmost pixel	Bits 1,5,3,7	Bits 3,7	Bit 7
		Bits 2,6	Bit 6
		Bits 0,4,2,6	Bit 5
		Bits 1,5	Bit 4
Rightmost pixel	Bits 0,4,2,6	Bits 1,5	Bit 3
		Bits 0,4	Bit 2
		Bits 0,4	Bit 1
		Bits 0,4	Bit 0

The following diagram illustrates the mapping from pixels on the screen to addresses in screen memory for the simple case of a base address of #C000 and an offset of 0.



#C7D0..#C7FF, #CFD0..#CFFF, ..., #FFD0..#FFFF are unused.

On the CPC6128 the base address sets which block will be used for the screen memory. Base addresses of #0000, #4000, #8000 and #C000 correspond to blocks 0, 1, 2, and 3. It is not possible for the screen memory to be located in blocks 4...7. Where the block being used for screen memory actually appears in the memory map depends on the bank switching (see section 2.5).

7 The Sound Manager.

The Sound Manager deals with the sound chip. It allows various envelopes and sounds to be set up and played under the control of the user. Most of the control is achieved using software rather than the sound chip hardware.

7.1 The Sound Chip.

The sound chip used is the General Instruments AY-3-8912. This has three channels and a pseudo-random noise generator that can be connected to any of the channels. The chip has a limited number of amplitude envelopes available (see Appendix IX) but the software enveloping, described below, can achieve all that the hardware is capable of, and more. Tone enveloping is all done by the software, there is no hardware support.

The sound generated by the chip uses square waveforms. There is no way to generate any other waveform.

It is possible to access the sound chip directly should the need arise. However, the routine MC SOUND REGISTER should be used to write to registers of the sound chip. This is because the keyboard is attached to the I/O port of the sound chip and the keyboard scanning routine expects to find the sound chip in a standard state (i.e. not in use). Also, there are timing constraints on accesses to the chip; using MC SOUND REGISTER will avoid consideration of these details.

The sound chip has three independent sound channels. The outputs from these are mixed together to form two stereo channels - sound channels A and B are mixed to form one stereo channel and sound channels B and C are mixed to form the other stereo channel. The stereo sound is available on the output jack on the back of the machine. However, there is only a single internal speaker and so the two stereo channels are mixed together to drive this. The volume of sound from the internal speaker can be controlled by the volume control knob on the side of the machine near the on/off switch. This control overrides the other volume control methods described below.

7.2 Tone Periods and Amplitudes.

The sound chip allows 16 different amplitudes in the range 0.. 15. Amplitude 0 is no sound at all, amplitude 15 is maximum volume.

The pitch of a note to be generated is specified by the period of the note rather than by the frequency. This period is given in 8 microsecond units. Thus, the tone period specified and the frequency of the tone generated are related by the formula:

$$\text{Tone period} = 125\,000 / \text{Frequency}$$

See Appendix VIII for a list of the suggested periods for generating musical notes.


7.3 Enveloping.

Real sounds rarely have a constant volume. Enveloping allows an approximation to the variation in volume of real sounds to be made. The sound is split into a number of sections each of which can increase the volume, decrease the volume, or keep it constant. The length of these sections can be varied, as can the rate of increase or decrease in volume. For example, a note generated by a musical instrument may be considered to have 3 sections as follows:

Attack: The volume of the note rises rapidly to its peak.

Sustain: The volume of the note remains constant while the note is played.

Decay: The volume falls away slowly to zero as the note finishes.



The Sound Manager allows two types of envelopes; amplitude envelopes to control a sound's volume and tone envelopes to control its pitch (the pitch is varied in much the same way as the volume). The user can set up to 15 different envelopes of each type. The exact formats of the data blocks specifying envelopes are given in SOUND AMPL ENVELOPE and SOUND TONE ENVELOPE.

a. Amplitude envelopes.

An amplitude envelope is used to control the volume and length of a sound. It can have up to five sections. Each section can be either a hardware or a software section. Software sections are either absolute or relative.

Hardware sections write values into the sound chip registers 11,12 and 13 to set up a hardware envelope. (See Appendix IX for a description of the sound chip registers). Generally a hardware section will be followed by a software section that does nothing except wait for a time long enough for the hardware envelope to operate.

An absolute software section specifies a volume to set and a time to wait before obeying the next section.

A relative software section specifies an step size, a number of steps and a time to wait. For each step requested, the current volume is changed by the given step size and then the Sound Manager waits for the given time after each step before obeying the next step.

Amplitude envelopes are set by calling SOUND AMPL ENVELOPE.

b. Tone envelopes.

A tone envelope controls the pitch of the sound. It can have up to five sections. Each section can be either an absolute or a relative section. The sections of a tone envelope are not necessarily related to those of an amplitude envelope.

An absolute section specifies a tone period to set and a time to wait before obeying the next section.

A relative section specifies an step size, a number of steps and a time to wait. For each step requested, the current tone period is changed by the given step size and then the Sound Manager waits for the given time after each step before obeying the next step.

If the tone envelope is completed before the sound duration expires (see section 7.40) then the final pitch is held constant. Alternatively, tone envelopes can be set to repeat themselves automatically. This allows tremolo effects to be created.

Tone envelopes are set by calling SOUND TONE ENVELOPE.

7.4 Sound Commands.

When a sound is given to the Sound Manager to be played, by calling SOUND QUEUE, a lot of information needs to be specified. This is described briefly below. The detailed layout of a sound command data block is described in SOUND QUEUE.

a. Initial tone period.

The sound is issued with an initial tone period. The pitch of the sound can be varied from this initial value using a tone envelope. If no tone envelope is specified the pitch remains constant. An initial tone period of zero means no tone is to be generated, presumably the sound is to be pure noise (see (e) below).

b. Initial volume.

The sound is issued with an initial volume. The volume of the sound can be varied from this initial value using an amplitude envelope. If no amplitude envelope is specified then the volume remains constant.

c. Tone envelope.

This specifies which tone envelope to use. If no envelope is specified then the pitch of the sound remains constant.

d. Amplitude envelope.

This specifies which amplitude envelope to use. If no envelope is specified then default system envelope is used. This keeps the volume of the sound constant and lasts for 2 seconds.

e. Noise period.

If the noise period is zero then no noise is to be added to the sound. Any other value sets the period for the pseudo-random noise generator and adds noise to the tone generated. Note that there is only one noise generator and so if two sounds are to use it at the same time they will need to agree on the period.

f. Duration.

The length of a sound can be specified in two ways, either as an absolute time (duration) or as a number of operations of the amplitude envelope. In the latter case the envelope is run one or more times and the sound finishes when the envelope has been executed the specified number of times. In the former case, if the duration finishes before the envelope (if any) then the sound is cut short. If the duration is longer than the envelope then the final amplitude is held until the duration expires.

g. Channels and Synchronisation Bits.

The sound can be issued to one or more channels. If a sound is issued to more than one channel then these channels automatically rendezvous with each other. Rendezvous requirements can be set explicitly as well. Also the sound can be held or the sound queue can be flushed (see section 7.6).

7.5 Sound Queues.

Each channel has a queue associated with it. Each queue has space to store at least three sounds. The sound at the head of each queue may be running and making music on its channel or it may be waiting for various synchronisation requirements (see 7.6 below). When a sound command is issued the sound is placed into the queues for the channels specified by the command. When the sound reaches the head of the queue, and providing its synchronisation requirements are met, it is executed.

If a sound that has the flush bit set is put into a queue then all sounds queued for that channel are discarded and any executing sound is stopped immediately. Thus a sound with the flush bit set will move to the head of the queue immediately and may commence execution.

A routine (SOUND CHECK) is provided to test the status of the sound at the head of a queue and to determine how much free space is in a queue. It is also possible to set up a sound event for each queue (by calling SOUND ARM EVENT). This synchronous event is 'kicked' when the queue has a free space in it. The sound event mechanism allows the generation of sound to be carried on aBa background task whilst some other action is being carried out.

7.6 Synchronisation.

There are two mechanisms to allow sounds on different channels to be synchronised. These are holding sounds and rendezvous. The purpose of synchronisation is to ensure that sounds start simultaneously. For example, a simulation of an instrument might use one channel to generate the fundamental note and another channel to generate the harmonics of the note. The synchronisation mechanism, particularly rendezvous, may be used to ensure that the fundamental and the harmonic sounds start exactly together.

A sound can be specified to be held when it is issued. This means that when it reaches the head of the sound queue it is not executed immediately. Instead it waits until it is explicitly released (by calling SOUND RELEASE) before it starts execution.

A sound can have rendezvous requirements set on it when it is issued. If a sound is issued to more than one channel then these channels all set rendezvous with each other automatically. When a sound with a rendezvous set reaches the head of the sound queue it is not executed immediately. Instead it waits until sounds with matching rendezvous requirements reach the head of their sound queues. Only when all rendezvous sounds are found to be present and ready to run do they start.

For instance, a sound on channel A marked to rendezvous with a sound on channel B will not start until a sound on channel B marked to rendezvous with channel A is ready to start - and vice versa! if a sound is ready to start on channel B that is not marked to rendezvous with channel A then it starts but the sound on channel A continues to wait for its rendezvous.

7.7 Holding Sounds.

It is possible to stop a sound while it is executing by calling SOUND HOLD. This will stop a channel making any sound and will save the state of the sound. The sound can be restarted from where it was held by calling SOUND CONTINUE. However, if a hardware envelope was running when the sound was held then it is impossible to predict the effect of restarting the sound. The hardware envelope may or may not continue from where it was held.

Calling SOUND HOLD is different from setting the hold bit when issuing a sound as described in section 7.6 above. SOUND HOLD stops all sounds being generated at any time whilst the hold bit is a method for synchronising sounds and prevents a particular sound starting when it reaches the head of the queue.

8 The Cassette Manager.

The Cassette Manager deals with reading files from and writing files to tape. These operations can either be performed on a character by character basis or on a whole file at once. There is no hardware support for the cassette, even the timing for reading and writing bits is performed by software.

The format of data on the tape is described in great detail. This will only be of academic interest to most users. More general information can be found in sections 8.4 onwards. In the case of V1.1 machines or the CPC464 with a DDI-1 fitted, a |TAPE command will have to be used to access the Tape Operating System.

8.1 File Format.

A file on tape is split into blocks each with a header record and a data record containing up to 2K (2048) bytes of data. The cassette motor which is under software control is turned off between each file block to allow time to process the data read or to generate the data to be written. The motor start-up gap also serves to separate the blocks from each other.

The general format of a block is as follows:

Motor Start-up	File header record	File data record
----------------	--------------------	------------------

However, the first and last blocks of a file have an extra pause before and after them respectively, to separate files on the tape. Their formats are thus:

First block:

Motor Start-up	Pre-file gap	File header record	File data record
----------------	--------------	--------------------	------------------

Last block:

Motor Start-up	File header record	File data record	Post-file gap
----------------	--------------------	------------------	---------------

There is a strong distinction between the file header record and the file data record. The header record is written using one synchronisation character (#2C) and the data record with another (#16). This means that when the Cassette Manager is searching for a file header it is impossible for it to find a file data record by mistake, and vice versa. See 8.2 below for the use of the synchronisation characters.

8.2 Record Format.

A record can contain any number of data bytes from 1 to 65536. The data is split into segments each of which is 256 bytes long. The last segment is padded out to 256 bytes with zeros when writing if necessary. When reading a record any extra bytes are ignored although they are accumulated into the CRC.

The layout of a record is as follows:

Leader	Segment 1	Segment N	Trailer
--------	-----------	-------	-----------	---------

There are N segments where $256 * N$ is the length of data (plus padding) to be written.

A file header record always contains one segment; a file data record contains from one to eight segments (usually 8 segments).

a. Leader

At the start of all records a leader is written which has the following layout:

Pre-record gap	2048 one bits	Zero bit	Sync byte
----------------	---------------	----------	-----------

The leading gap is there to ensure the failure of any attempt to synchronise on the end of a preceding record or on data that was on the tape and that has been over-recorded.

The long sequence of one bits is used to calculate the speed at which the data was written and hence to calculate the threshold value used to distinguish one bits from zero bits.

The single zero bit is used to mark the impending end of the leader and is also used to determine whether the recording has been inverted (see section 8.3).

The synchronisation byte is there to help prevent spurious synchronisation on sequence of bits such as might be found in a record. If an incorrect value for the sync byte is found then an attempt has been made to synchronise on the middle of a record or on the wrong type of record. This byte is used to distinguish header records from data records in a file block (header records use #2C while data records use #16).

b. Segments

Each segment contains 256 data bytes and has the following format:

Byte 1	Byte 2	Byte 256	CRC 1	CRC 2
--------	--------	-------	----------	-------	-------

‘CRC 1’ is the more significant byte and ‘CRC 2’ the less significant byte of the logical NOT of the CRC calculated for the 256 bytes in the segment. (The CRC polynomial used is $X^{15} + X^{12} + X^5 + 1$ with an initial seed of #FFFF).

c. Trailer

The trailer is simply an extra 32 one bits written to the end of the record.


8.3 Bit Format.

A bit is written to the tape as a period of low level followed by an equal period of high level. A one is written to the tape with these periods twice as long as those for a zero. The length of the period for a zero can be set by the user (see CAS SET SPEED).

The tape circuitry has a tendency to move the positions of edges (transitions from high to low or low to high) so as to balance out the difference between ones and zeros written to tape. Precompensation is used - which adds to the period of a one bit and subtracts from the period of a zero bit to make the waveform closer to the ideal when it is read.

When reading, the speed at which the recording was made is determined by timing the one bits in the record leader. As this is a long sequence of the same bit the edges are not shifted and no precompensation is applied. Since the speed is established independently for each record this automatically takes into account most tape speed variations.

Data is written low-high but may be inverted when read (i.e. high-low). It is important for the firmware to determine whether the waveform being read is inverted or not. If this is not achieved then the bits will not be read properly as the following example shows:



The zero bit in the record leader is used to determine whether the recording has been inverted.

Bytes written to the tape are written with the most significant bit first and the least significant bit last.

8.4 The Header Record.

The header record in a file block contains information about the file and about the data in the following data record. Some of the entries in the header are

used by the system for various purposes. The remaining entries are available for the user to set when writing a file, and to read when reading a file. These entries are the file type (byte 18) and all the user fields (bytes 24..63) including the logical length (bytes 24..25) and the entry address (bytes 26..27). The user fields will all be set to zero if they are not used.

The header is laid out as follows:

System fields

Bytes 0..15	Filename	Padded to 16 bytes with nulls.
Byte 16	Block number	The first block is normally block 1 and block numbers increase by 1 on successive blocks.
Byte 17	Last block	A non-zero value means that this is the last block of a file.
Byte 18	File type	A value recording the type of the file (see below).
Bytes 19..20	Data length	The number of data bytes in the data record.
Bytes 21..22	Data location	Where the data was written from originally.
Byte 23	First block	A non-zero value means that this is the first block of a file.

User fields

Bytes 24..25	Logical length	This is the total length of the file in bytes.
Bytes 26..27	Entry address	The execution address for machine code programs
Bytes 28..63	Unallocated	These are unallocated and may be used as required.

The file type (byte 18) is split into a number of fields:

Bit 0	Protection	If this bit is set the file is protected in some way.
Bits 1.3	File contents	0 = Internal BASIC. 1 = Binary. 2 = Screen image. 3 = ASCII. 4..7 are unallocated.
Bits 4..7	Version	ASCII files should be version 1, all other files should be version 0.

8.5 Read and Write Speeds.

The Cassette Manager is capable of reading and writing data at speeds ranging from 700 baud to 2500 baud. There are two speeds commonly used in this range, 1000 baud (the default speed selected at EMS) and 2000 baud. The default speed is chosen to be near the slowest speed to give maximum reliability. The reliability at 2000 baud is still good, however, particularly when playing back on the same machine that was used to record a tape.

Bits are written to the tape as a single cycle of a tone. The tone for a one always has half the frequency of the tone for a zero. Thus ones are twice as long as zeros on the tape. This means that the baud rates given above are only averages and vary according to the actual data written.

Even with the built in cassette mechanism the Cassette Manager has to precompensate the waveform written to the tape to achieve the speeds quoted. This means that the lengths of bits written are altered (ones lengthened, zeros shortened) to try to make the waveform read closer to the ideal after the edges of the waveform have been shifted by the cassette circuitry.

It is only necessary to set the Cassette Manager's write speed. When reading a record from tape the record leader is used to calculate the speed at which it was written. This also allows for tape speed variations between different machines.

8.6 Cataloguing.

To generate a catalogue from the tape the Cassette Manager reads a sequence of file blocks and prints information from them. The file blocks may come from any file, in any order. Cataloguing continues until the user hits the escape key.

The information is reported as follows:

```
FILENAME block N L Ok
```

FILENAME is either the name of the file or 'Unnamed file' if the filename starts with a null.

The block number, N, indicates which block of the file it is. Normally block 1 is the first block of a file.

L is a character representing the file type and protection status of the file. It is formed by adding #24 (character '\$') to the file type from the header masked with #0F. This gives the following characters:

\$	an unprotected BASIC program.
%	a protected BASIC program.
&	a binary file.
'	a protected binary file
*	an ASCII file.

Other characters are possible but the above are the standard file types that are written by the on board ROM.

The above information is printed when the header record is read correctly.

Ok is printed after the data record has been read correctly.

8.7 Reading Files.

Before a file can be read from it must be opened (by calling CAS IN OPEN). This sets up the filename (see 8.10 below) and reads the first block of the file so that the header can be inspected.

The file may either be opened for character input or for direct input, but not both. The mode of input is set by the first access to the file and not when it is opened. As soon as one mode is selected it becomes impossible to use the other mode to access the file.

Character input (calling CAS IN CHAR) allows the user to read the file sequentially one character at a time. Blocks of the file are read from tape into the buffer as needed. This is intended for reading text files and similar applications.

Direct input (calling CAS IN DIRECT) reads the whole of the file into memory in one go. This is intended for loading machine code programs or screen dumps and similar applications.

Interrupts are disabled whilst reading from tape because this has serious timing constraints. Disabling interrupts will prevent the various timer interrupts (as described in section 10.1) from occurring. In particular this might leave the sound chip making a noise for a long period of time and so the Sound Manager is shut down (see SOUND RESET).

In V1.1 firmware the cassette manager routines for reading files return error codes to indicate the cause of the errors (for compatibility with AMSDOS):

- #00: The user hit escape
- #0E The stream is already/not in use
- #0F: Have reached the end of the file

8.8 Writing Files.

Before a file can be written to it must be opened (by calling CAS OUT OPEN). This sets up the filename (see 8.10 below) and the rest of the header that will be written in each file block.

The file may either be opened for character output or for direct output, but not both. The mode of output is set by the first write to the file and not when it is opened. As soon as one mode is selected it becomes impossible to use the other mode to write to the file.

Character output (calling CAS OUT CHAR) allows the user to write to the file one character at a time. The characters are buffered until a complete block (2048 characters) is ready to be written whereupon a file block is written to the tape.

Direct output (calling CAS OUT DIRECT) writes the whole of the file from memory in one go. The data written is still packaged into 2048 byte blocks.

Whichever output mode is used, it is important to close the output file properly (using CAS OUT CLOSE) otherwise the lastblock of the file will not be written.

Interrupts are disabled whilst writing to tape because this has serious timing constraints. Disabling interrupts will prevent the various timer interrupts (as described in section 10.1) from occurring. In particular this might leave the sound chip making a noise for a long period of time and so the Sound Manager is shut down (see SOUND RESET).

In V1.1 firmware the cassette manager routines for writing files return error codes to indicate the of the errors (for compatibility with AMSDOS):

#00: The user hit escape
#0E The stream is already/not in use

8.9 Reading and Writing Simultaneously.

The Cassette Manager allows two files to be open simultaneously. One must be open for reading and the other for writing. Thus it is possible to read from one file and write to another file at the same time.

When the Cassette Manager is about to read a block it asks the user to press PLAY and this implies that the tape with the file for reading should be loaded. Similarly, when it is about to write a block it asks the user to press REC and PLAY and this implies that the tape to which the file is to be written should be loaded. The Cassette Manager assumes that the tape is not changed and that the appropriate cassette controls remain pressed as requested until a prompt is issued. It also assumes that pressing a key means that the prompt has been obeyed.

It is unwise to attempt to read and write simultaneously with the Cassette Manager messages turned off. The only notification given of which tape should be loaded is in the promptmessages.

8.10 Filenames.

When the user opens a file for reading or writing the name of the file to be read or written is specified. The filename is a string of any 16 characters (#00.. #FF). If the file name specified is longer than 16 characters then it is truncated and if it is shorter than 16 characters it is padded to 16 characters with nulls (character #00).

When opening a file for reading a zero length filename or one that starts with a null has a special meaning - read the next file on the tape. The Cassette Manager searches the tape until it finds the first block of a file and it reads this file. Once the first block of a file has been found the Cassette Manager will only read from that file and no other.

BASIC uses a slightly extended form of the filename. If the first character of a BASIC filename is an exclamation mark (character #21) the BASIC turns the prompt messages off (see 8.11 below) and removes the exclamation mark from the name. This facility is not provided at the Cassette Manager level.

8.11 Cassette Manager Messages.

The Cassette Manager issues a number of messages to prompt and inform the user and to warn when errors have occurred. The messages that prompt or inform the user may be turned on or off as desired (see CAS NOISY). Messages that inform the user of errors cannot be turned off by this mechanism.

a. Prompt messages.

Press PLAY then any key:

This message is issued when the Cassette Manager is about to read the first block of a file from tape or when it is about to read a block after having written to tape (see section 8.9). It indicates that the tape containing the file to be read should be loaded and that the PLAY button on the recorder should be pressed. The Cassette Manager does not issue this message at other times since it assumes that the correct tape is still loaded and that the PLAY button is still pressed.

Press REC and PLAY then any key:

This message is issued when the Cassette Manager is about to write the first block of a file to tape or when it is about to write a block after having read from tape. It indicates that the tape on which the file is to be written should be loaded and that the REC and PLAY buttons on the recorder should be pressed. The Cassette Manager does not issue this message at other times since it assumes that the correct tape is still loaded and that the REC and PLAY buttons are still pressed.

b. Information messages.

Found FILENAME block N

This message is printed when reading from the tape if a header record is found that for any reason does not match the record that was expected. This may indicate that the tape is positioned incorrectly (too early or too late) or that the wrong tape is being played.

Loading FILENAME block N

A block of the file has been found and is being read from tape.

Saving FILENAME block N

A block of the file is being written to tape.

FILENAME in the above messages is the name of the file or 'Unnamed file' if the filename starts with a null.

The block number, N, indicates which block of the file is being read or written. The first block of a file is normally block 1, the second block 2 etc.

c. Error messages

Rewind tape

While searching for a block of the file being read, a higher numbered block than that required has been found. The required block has been missed. This message is often produced after a read error in the required block when the next block is found.

Read error X

An error of some kind occurred whilst reading from the tape. The tape should be rewound and the block played again. The X is a single letter indicating what kind of read error occurred:

'a'	Bit too long	An impossibly long one or zero has been measured. This often indicates reading past the end of the record.
'b'	CRC error	Data was read from tape incorrectly.
'd'	Block too long	The data record contains more than the expected 2048 bytes of data.

Write error a

An error occurred whilst writing to the tape. There is only one possible write error. This indicates that the Cassette Manager was unable to write a bit as fast as was requested. This error will never occur unless the user has set the write speed beyond the maximum possible.

8.12 Escape Key.

The escape key on the keyboard may be used to abandon cassette operations at certain times.

When the Cassette Manager issues one of the prompt messages it calls KM READ CHAR repeatedly to empty the key buffer out. Then it calls KM WAIT KEY to wait until the user presses a key to acknowledge the prompt. If the value generated from the key the user presses is #FC, which is the value normally generated by the escape key, then the Cassette Manager will abandon the read or write and will return an error condition to the caller.

When reading from or writing to the cassette interrupts are disabled and the normal keyscanning mechanism is not active. While reading or writing the record leader the Cassette Manager itself scans the keyboard to test whether key 66, the escape key, is pressed. If this key is found to be pressed then the Cassette Manager abandons the read or write and returns to the caller (with an appropriate error condition). While reading or writing the data in the record there is no way to interrupt the Cassette Manager, thus pressing ESC may not be detected for several seconds.

8.13 Low Level Cassette Driving.

To allow the user to produce a new filing system the record read and write routines, CAS READ and CAS WRITE, are in the firmware jumpblock. There is a third routine at this level, CAS CHECK, whose facilities are not used by the higher levels of the Cassette Manager. It allows the data that has been written to tape to be compared with the data in store. This could be used to perform a read after write check if so desired.

Also available in the firmware jumpblock are routines to turn the cassette motor on and off (CAS START MOTOR and CAS STOP MOTOR). It is not necessary to turn the motor on and off around a call of CAS READ, CAS WRITE or CAS CHECK as these routines automatically turn the motor on and off.

9 AMSDOS

AMSDOS is a disc operating system used with all the CPC range of computers, of course, in the case of the 464 the DDI-1 has to be fitted. AMSDOS enables programs to access disc files in a similar manner to cassette files, indeed existing programs which currently use the cassette should be able to use disc files with little or no modification. The main source of incompatibility will be filenames, AMSDOS filenames must conform to CP/M standards but cassette filenames are far less restricted.

AMSDOS has been designed to complement CP/M, not to compete with it. They share the same file structure and can read and write each other's files.

AMSDOS resides in the same ROM as the CP/M BIOS.

9.1 Features

AMSDOS provides the following facilities:

- Switching the cassette input and output streams to and from disc. So that all the facilities available on the cassette become available on disc.

- Displaying the disc directory.

- Erasing disc files.

- Renaming disc files.

- Selecting the default drive and user.

Whenever AMSDOS creates a new file it is always given a name with a type part of `.$$$` regardless of the given name. When the file is closed any previous version of the file is renamed with a `B A K` type part and the new version is renamed from `.$$$` to its proper name. Any existing `.BAK` version is deleted. This gives an automatic one level file back-up.

For example, if the disc contains the files `FRED.BAS` and `FRED.BAK` and the user opens a file called `FRED.BAS` then AMSDOS will create a new file called `FRED.$$$`. When the file is closed the existing `FRED.BAK` is deleted, `FRED.BAS` is renamed to `FRED.BAK` and `FRED.$$$` is then renamed to `FRED.BAS`.

All AMSDOS facilities are implemented either by intercepting the cassette firmware calls or by external commands.

The intercepted firmware calls are:

CAS IN OPEN
CAS IN CHAR
CAS IN DIRECT
CAS RETURN
CAS TEST EOF
CAS IN CLOSE
CAS IN ABANDON
CAS OUT OPEN
CAS OUTCHAR
CAS OUT DIRECT
CAS OUT CLOSE
CAS OUT ABANDON
CAS CATALOG

The remaining cassette firmware calls are not intercepted and remain unaffected.

Full descriptions of both the tape and disc versions of these routines are given in section 15.

The AMSDOS external commands are:

A	Select default drive A:
B	Select default drive B:
CPM	Coldboot CP/M
DIR	Display disc directory
DISC	Redirect cassette routines to disc
DISC . IN	Redirect cassette input routines to disc
DISC . OUT	Redirect cassette output routines to disc
DRIVE	Select default drive
ERA	Erase files
REN	Rename a file
TAPE	Redirect cassette routines to cassette
TAPE . IN	Redirect cassette input routines to cassette
TAPE . OUT	Redirect cassette output routines to cassette
USER	Select default user

From BASIC all these commands must be preceded by a'|

Some of these commands require parameters.

Full descriptions of these external commands are given in section 20.

9.2 Filenames

AMSDOS filenames are upwards compatible with CP/M filenames. But in addition the user number may also be specified and non-significant spaces are permitted before and after the name and any embedded punctuation.

Examples:

ANAME	Default user, drive, and type
10:WOMBAT.TXT	Default drive, and user number 10
2A:WOMBAT.TXT	User 2, on Drive A:
.	Default drive, user, and all files
5 B : POSSUM .55\$	A name with non-significant spaces
a:aard?ark	Lowercase, AMSDOS will convert to uppercase

If given, the user number must be in the range 0..15, the drive letter must be A or B. If either the user or the drive is given they must be followed by a colon.

The following characters may be used in the name and type parts:

a-z A-Z 0-9 !" # \$ % & ' + - @ - \ } { ~

Any other characters will cause the command to fail with the message:

BAD COMMAND

The characters '?' and '*' are wildcards, that is, when placed within a filename or type it will be interpreted as 'any valid character'. For example if the filename

'G??E?? .B*' was used in the **|DIR** command then the files 'GAME1.BAS' 'GAME1.BAK' 'GAME29.BAS' and 'GREET.BAS', and any other matching combinations, would be displayed in the directory.

When parsing a filename, AMSDOS shifts lower case letters into upper case and removes bit 7.

If the user or drive is omitted then the current default values are assumed. These defaults may be set by the user.

If the type part is omitted then a default type is assumed. This depends on the context in which the name is being used, but usually a default type part of three spaces is assumed.

9.3 File Headers

Cassette files are subdivided into 2K blocks, each of which is preceded by header. CP/M files do not have headers. AMSDOS files may, or may not, have a header depending on the contents of the file. This will not cause problems for programs written in BASIC but is an important difference between cassette and disc files.

Unprotected ASCII files do not have headers. All other AMSDOS files have a single header in the first 128 bytes of the file, the header record. These headers are detected by checksumming the first 67 bytes of the record. If the checksum is as expected then a header is present, if not then there is no header. Thus it is possible, though unlikely, that a file without a header could be mistaken for one with a header.

The format of the header record is as follows:

Bytes	0..63	Cassette/Disc header (see below)
	64..66	Length of the file in bytes, excluding the header record. 24 bit number, least significant byte in lowest address
	67..68	Sixteen bit checksum, sum of bytes 0..66
	69..127	Undefined

The use that the cassette manager makes of the header is described in section 8.4. AMSDOS uses the header as follows:

	Bytes	
Filename	0	User number, #00.. #OF
	1...8	Name part, packed with spaces
	9...11	Type part, padded with spaces
	12...15	#00
Block number	16	Not used, set to 0
Last block	17	Not used, set to 0
File type	18	As per cassette
Data length	19.. .20	As per cassette
Data location	21...22	As per cassette
First block	23	Set to #FF, only used for output files
Logical length	24.. .25	As per cassette
Entry address	26.. .27	As per cassette
Unallocated	28.. .63	As per cassette

When a file without a header is opened for input a fake header is constructed in store as follows:

	Bytes	
Filename	0	User number, #00..#OF
	1..8	Name part, padded with spaces
	9..11	Type part, padded with spaces
	12..15	#00
File type	18	#16, unprotected ASCII version 1
Data location	19..20	Address of 2K buffer
First block	23	#FF

All other fields are set to zero.

9.4 Disc Organization

AMSDOS and the CP/M 2.2 BIOS support three different disc formats: SYSTEM format, DATA ONLY format, and IBM format. The CPIM Plus BIOS supports the SYSTEM and DATA formats but not IBM format.

The BIOS automatically detects the format of a disc. Under CPIM this occurs for drive A at a warm boot and for drive B the first time it is accessed. Under AMSDOS this occurs each time a disc with no open files is accessed. To permit this automatic detection each format has unique sector numbers as detailed below.

3 inch discs are double sided, but only one side may be accessed at a time depending on which way round the user inserts the disc. There may be different formats on the two sides.

Common To All Formats

Single sided (the two sides of a 3 inch disc are treated separately).

512 byte physical sector size.

40 tracks numbered 0 to 39.

1024 byte CP/M block size.

64 directory entries.

System Format

9 sectors per track numbered #41 to #49.

2 reserved tracks.

2 to 1 sector interleave.

The system format is the main format supported, CP/M can only be loaded (Cold Boot) from a system format disc. CP/M 2.2 also requires a system format disc to warm boot. The reserved tracks are used as follows.

Track 0 sector #41:	bootsector.
Track 0 sector #42:	configuration sector
Track 0 sectors #43..#47:	unused
Track 0 sectors #48..#49	and
Track 1 sectors #41..#49:	CCP and BIOS

CP/M Plus only uses Track 0 sector #41 as boot sector
Track 0 sector #42... #49 and Track 1 are unused.

Note: Another format called 'VENDOR' format is a special version of system format which does not contain any software on the two reserved tracks. It is intended for use in software distribution.

Data Only Format

9 sectors per track numbered #C1 to #C9.

0 reserved tracks.

2 to 1 sector interleave.

This format is not recommended for use with CP/M 2.2 since it is not possible to 'warm boot' from it. However, because there is a little more disc space available it is useful for AMSDOS or CP/M Plus.

IBM Format

8 sectors per track numbered 1 to 8

1 reserved track

no sector interleave

This format is logically the same as the single-sided format used by CP/M on the IBM PC. It is intended for specialist use and is not otherwise recommended as it is not possible to 'warm boot' from it.

9.5 Boot Sector

In order that non-CP/M systems may be implemented at a later date the BIOS initialization is performed, in part, by a boot program which is read from the disc before attempting to load CP/M. In the non-CP/M case the boot program would not jump to the warm boot routine but go its own way, using the BIOS and firmware routines as desired.

The boot program is in the boot sector which is the first sector (sector #41) on track 0. During a cold boot the BIOS is initialized into a minimum state before loading and executing the boot program. This state is as follows.

All the routines in the ROM copy of the BIOS jumpblock and all routines in the extended jumpblock are available.

Alternate and IY register saving is enabled.

Interrupts are indirected via the BIOS and run on the BIOS's stack.

Disc messages are enabled.

The initial command buffer is empty.

The IOBYTE at #0003 is initialized to #81 (**LST:=LPT: , PUN:=TTY: , RDR:=TTY: , CON:=CRT:**).

The current drive at #0004 is initialized to #00.

The serial interface is not initialized.

The CCP and BDOS are not in store.

The BIOS RAM jumpblock is not in store.

The CP/M jumps at #0000 and #0005 are not initialized.

The boot sector is read and loaded into store at #0100; the stack pointer is initialized to a value immediately below the BIOS's data (#AD33 is normal) area and the boot program is entered at #0100. The boot program may use store from #0100 upwards until it reaches the stack.

To run CP/M the boot program must, at least, jump to the warm boot entry in the ROM jumpblock.

The boot program for CP/M 2.2 loads and obeys the configuration sector and then warm boots CP/M.

The boot program for CP/M Plus searches for, loads and executes a file with the type part **.EMS**.

The boot program has the following interface:

Entry:

SP= highest address available + 1 (a good place for the stack)
BC = address of ROM copy of BIOS jumpblock (BOOT)

Exit:

To run CP/M the program should jump to the WBOOT entry in the above jumpblock

The ROM copy of the BIOS jumpblock should not be used at any other time (indeed, only the boot program knows where it is).

9.6 AMSDOS Messages

AMSDOS uses the CP/M 2.2 BIOS in order to access the disc. Thus BIOS messages will be displayed in the event of a disc error. This section explains the meaning of the AMSDOS messages.

In the following <drive> means A or B, <filename> means an AMSDOS filename.

Bad command

There is a syntax error in a command or filename.

<filename> already exists

The user is trying to rename a file to a name which is already in use.

<filename> not found

The user is trying to open for input, erase or rename a file that does not exist.

Drive <drive>: directory full

There are no more free directory entries (64 directory entries per disc).

Drive <drive>: disc full

There are no more free disc blocks.

Drive <drive>: disc changed, closing <filename>

The user has changed the disc while files were still open on it.

<filename> is read only

The user is trying to erase or rename a file which is marked read-only. May also be caused by closing a file when existing version of the file is read-only.

9.7 BIOS Facilities Available To AMSDOS

AMSDOS uses the CP/M BIOS 2.2 to access the disc. In order that a program running under AMSDOS may access the disc directly nine of the BIOS extended 'jumpblock routines are available. The routines are accessed as external commands. An example of using these commands is given in section 10.6.

NOTE: The BIOS extended jumpblock itself is not available, indeed it does not exist in an AMSDOS environment.

The BIOS routines available and their command names are as follows:

SET MESSAGE	CtrlA (#01)
SETUP DISC	CtrlB (#02)
SELECT FORMAT	CtrlC (#03)
READ SECTOR	CtrlD (#04)
WRITE SECTOR	CtrlE (#05)
FORMAT TRACK	CtrlF (#06)
MOVE TRACK	CtrlG (#07)
GET DR STATUS	CtrlH (#08)
SET RETRY COUNT	CtrlI (#09)

These routines are described in section 19.

The word at #BE40 contains the address of the disc parameter header vector. Disc parameter headers and extended disc parameter blocks may be patched as required (see section 9.8).

Only the BIOS facilities mentioned here maybe used from a program running under AMSDOS.

9.8 Store requirements

When initialized AMSDOS reserves #500 bytes of memory from the memory pool and the kernel reserves another 4 for its external command chaining information.

When loading a machine code program from disc into store using the AMSDOS routine CAS IN DIRECT it is important that AMSDOS's variables are not overwritten. This presents a problem since in general it is not possible to discover where these variables are! This is because variables for expansion ROMs are allocated dynamically. Note that this problem does not arise when loading from the cassette since the cassette manager's variables are in the firmware variable area.

AMSDOS reserves store from the top of the memory pool so the simplest solution is to always load machine code programs into the bottom of store. The program can then relocate itself to a higher address if required.

Alternatively the machine code program could be loaded in two stages: first load and run a small loader in the bottom of store. The action of MC BOOT PROGRAM will have shut down all RSXs and extension ROMs. The loader program should now initialize AMSDOS using KL INIT BACK thus forcing the AMSDOS variables to be wherever you so wish. The loader can now load the machine code program using the AMSDOS routines CAS IN OPEN, CAS IN DIRECT, and CAS IN CLOSE together with MC START PROGRAM.

In order to initialize AMSDOS using KL INIT BACK, AMSDOS's ROM number is required. To determine AMSDOS's ROM number look at any of the intercepted cassette jumpblock entries with the DISC routines selected. Each entry is a far call, the address part of which points at a three byte far address, the third part of the far address is the ROM number. Obviously this must be done before AMSDOS is shut down.

Existing machine code programs, developed on cassette systems without any expansion ROMs, frequently only use store to #ABFF in order to avoid BASICs variables. These can easily be modified to use AMSDOS. Write some machine code to initialize AMSDOS using KL INIT BACK. AMSDOS will reserve RAM down to #ABFC, almost the same as used by BASIC.

9.9 Extended Disc Parameter Blocks

In order to facilitate reading and writing 'foreign' discs of differing formats, all the parameters concerning a drive are kept in RAM in an extended CP/M disc parameter block (XPB). The knowledgeable user may patch an XPB.

There are two XPBs, one per drive.

XPB structure:

bytes	0...14:	standard CP/M 2.2 DPB (see below).
byte	15:	first sector number.
	16:	number of sectors per track.
	17:	gap length (read/write).
	18:	gap length (format).
	19:	filler byte for formatting.
	20:	$\log_2(\text{sector size}) - 7$, 'N' for PD765A.
	21:	sector size/128.
	22:	reserved : current track (set by BIOS).
	23:	reserved: #00 not aligned, #FF aligned
(setbyBIOS).	24:	#00 autoselect format, #FF don't autoselect format.

The XPB for a drive may be found by accessing the Disc Parameter Header (DPH) vector. The first word of the DPH is the address of the XPB for drive A, the second word is the address of the XPB for drive B. The address of the DPH is stored at location #BE40.

The values stored in the standard CP/M 2.2 DPB (Disc Parameter Block) are often derived from the data block allocation size, BLS, which is the number of bytes in a block and which may be 1024, 2048, 4096, 8192 or 16384. The value of BLS is not stored in the DPB but it may be deduced from values stored there. CP/M plus has a slightly different DPB. The CP/M 2.2 DPB is laid out as follows:

- Bytes 0.. 1 (SPT): Total number of 128 byte records on each track.
- 2 (BSH): $\log_2 \text{BLS} - 7$.
- 3 (BLM): $\text{BLS} / 128 - 1$
- 4 (EXM): If $\text{DSM} < 256$ then $\text{BLS} / 1024 - 1$ else $\text{BLS} / 2048 - 1$.
- 5..6 (DSM): Total size of disc in blocks excluding any reserved tracks.
- 7.. 8 (DRM): Total number of directory entries-1.
- 9.. 10 (ALO/1): Bit significant representation of number of directory blocks (#0080 1, #00C0 2etc).
- 11..12 (CKS): Length of checksum vector. Normally $\text{DRM} / 4 + 1$ but if checksumming is not required then 0.
- 13.. 14 (OFF): Number of reserved tracks. This is also the track on which the directory starts.

The XPBs for the different formats are initialized as follows:

System Format

- 36 records per track
- 3 blockshift
- 7 block mask
- 0 extent mask
- 170 number of blocks-i
- 63 number of directory entries-i
- #00C0 2 directory blocks
- 16 size of checksum vector
- 2 reserved tracks
- #41 first sector number
- 9 sectors per track
- 42 gap length (read/write)
- 82 gap length (format)
- #E9 fillerbyte
- 2 $\log_2 (\text{sector size}) - 7$
- 4 records per track
- 0 current track
- 0 not aligned
- 0 do auto select format

Data Only Format

36	records per track
3	block shift
7	block mask
0	extent mask
179	number of blocks-1
63	number of directory entries-1
#00C0	2 directory blocks
16	size of checksum vector
0	reserved tracks
#C1	first sector number
9	sectors per track
42	gap length (read/write)
82	gap length (format)
#E9	filler byte
2	$\log_2(\text{sector size})-7$
4	records per track
0	current track
0	not aligned
0	do auto select format

IBM Format

32	records per track
3	blockshift
7	block mask
0	extent mask
155	number of blocks-1
63	number of directory entries-1
# 00C0	2 directory blocks
16	size of checksum vector
1	reserved tracks
#01	first sector number
8	sectors per track
42	gap length (read/write)
80	gap length (format)
#E9	fillerbyte
2	$\log_2(\text{sector size})-7$
4	records per track
0	current track
0	not aligned
0	do auto select format

10 Expansion ROMs, Resident System Extensions and RAM Programs.

The system can address up to 252 expansion ROMs, mapped over the top 16K of memory, starting at #C000. The Kernel supports two varieties of expansion ROM, foreground and background. A resident system extension (RSX) is similar in use to a background ROM, but must be loaded into RAM before it can be used.

A foreground ROM contains one or more programs, only one of which may be running at one time. The on-board BASIC is the default foreground program. Other possible foreground programs are:

- other systems, such as FORTH or CP/M.
- applications, such as a Word Processor or Spread Sheet.
- tools, such as an Assembler or Debugger.

A RAM program, once loaded, takes over the machine in much the same way as a foreground ROM program. Games will generally be RAM programs.

There may be up to 16 background ROMs, each of which provides some sort of service independent of the foreground program. It is expected that expansion peripherals will each have an associated background ROM containing suitable support routines. Other background ROMs may augment the existing machine software; for example, by providing further graphics functions.

A resident system extension (RSX), once loaded, provides some sort of service in the same way as a background ROM. An RSX might, for example, provide special support for a given printer - where it is more economic to provide the software on cassette rather than in ROM (or PROM).

10.1 ROM Addressing.

Expansion ROMs have ROM addresses in the range 0..251. To select a given ROM the Kernel sets its ROM address by writing to I/O address #DF00. If a ROM is fitted at the address selected, then all further read accesses to the top 16K of memory will return data from the expansion ROM. If no ROM is fitted at the currently selected ROM address the contents of the on-board ROM are returned.

When the machine is first turned on ROM 0 is selected as the foreground program. If no expansion ROM is fitted at ROM address 0, the on-board ROM is used, and BASIC is entered. If an expansion ROM is fitted at ROM address 0 it takes precedence over the on-board ROM.

In V1.0 firmware background ROMs must be fitted at ROM addresses in the range 1...7. Foreground ROMs must be fitted so that there are contiguous ROMs from address 1. When searching for a foreground ROM the kernel starts at address 0 and works upwards until the first unused address greater than 0 is found.

In V1.1 firmware background ROMs may be fitted at ROM addresses in the range 0.. .15. Foreground ROMs must be fitted contiguously from address 16 or at any background ROM address. When searching for a foreground ROM the kernel starts at address 0 and works upwards until the first unused address greater than 15 is found. In either case if an expansion ROM 0 is fitted the on-board ROM can still be accessed at the first unused ROM address.

The Kernel supports a 'far address' which may be used to call subroutines in expansion ROMs. The 'far address' is a three byte object, the last byte of which is a ROM select number. Since the arrangement of ROMs in an expansion card is quite arbitrary the ROM select part of a 'far address' must be established at run time. The 'sideways' ROM addressing facility allows a foreground program to occupy up to four contiguous ROM select addresses, and supports subroutine calls between the ROMs without requiring the program to know the actual ROM address of any of them.

10.2 The Format of an Expansion ROM.

An expansion ROM may be up to 16K bytes long, the first byte being at address #C000. The first few bytes of the ROM are the 'ROM Prefix' and must take the form:

- Byte 0: ROM Type.
- Byte 1: ROM Mark Number.
- Byte 2: ROM Version Number.
- Byte 3: ROM Modification Level.
- Byte 4: External Command Table.

The ROM type specifies what sort of ROM this is and must take one of the following values:

- 0: Foreground ROM
- 1: Background ROM
- 2: Extension ROM

The on-board ROM must be unique in having bit 7 of the type byte set (thus its type byte is #80). This marker is used to detect the end of foreground ROMs. If a foreground program will not fit into a single ROM then the extra ROMs required should be marked as extension ROMs.

The mark number, version number and modification level may be set to any values required.

The external command table comprises a list of command names and a jumpblock. Each command name is implicitly associated with the same numbered entry in the jumpblock. The table takes the form:

- Bytes 0.. 1 : Address of command name table
- Bytes 2.. 4 : Jumpblock entry 0
- Bytes 5..7 : Jumpblock entry 1
- ...etc : ... etc

The command name table is a list of names, each of which may be up to 16 characters long. The last character of each name must have bit 7 set but no other character may. The table is terminated by a null (character 0) after the last character of the last name. Apart from the fact that all characters must be in the range 0.. 127 and that the first character may not be a null, there are no

restrictions on the characters in command names. However, if unsuitable characters are chosen it may prove impossible for programs such as BASIC to access the commands. BASIC expects alphabetic characters in the command names to be in upper case and will not allow characters such as space or comma in the command name.

The ROM prefix for the on-board ROM is:

```

        ORG    #C000          ; Start of the ROM
        DEFB   #80+0         ; On board ROM, Foreground
        DEFB   1             ; Mark 1
        DEFB   0             ; Version 0
        DEFB   0             ; Modification 0
        DEFW   NAME_TABLE    ; Address of name table.
        JP     START_BASIC.  ; The only entry in the
jumpblock
NAME_TABLE: DEFB 'BASI', 'C'+ #80 ; The only command name
            DEFB  0             ; End of name table marker

```

The ROM prefix for a serial I/O card might be:

```

        ORG    #C000          ; Start of ROM
        DEFB   1             ; BackgroundROM
        DEFB   0             ; Mark0
        DEFB   5             ; Version 5
        DEFB   0             ; Modification 0

        DEFW   NAME_TABLE    ; Address of name table

        JP     EMS_ENTRY     ; 0 Background ROM power-up
entry
        JP     RESET         ;1
        JP     SET_BAUD_RATE ; 2
        JP     GET_CHARACTER ;3
        JP     PUT_CHARACTE  ;4
        ...etc

```

```

NAME_TABLE: DEFB 'SIO DRIVE', 'R'+ #80 ; 0
            DEFB 'SIO.RESE', 'T'+ #80 ; 1
            DEFB 'SIO.SET.BAU', 'D'+ #80 ; 2
            DEFB 'SIO.GET.CHA', 'R'+ #80 ; 3
            DEFB 'SIO.PUT.CHA', 'R'+ #80 ; 4
            ...etc
            DEFB 0             ; End of nametable marker

```

Note that the command name table entry for the power-up entry includes a space.

This is still a legal name but the BASIC will never be able to generate it because of the way it uses spaces. Because BASIC cannot generate the name it is impossible for a BASIC user to call the power-up entry by mistake (see section 10.4).

10.3 Foreground ROMs and RAM Programs.

Each of the entries to a foreground ROM is expected to represent a separate program, whose name is given by the corresponding entry in the name table. The first entry of ROM 0 is the default power-up entry point at the end of EMS.

Once a RAM program has been loaded it is treated much like a foreground ROM, except that it does not have a ROM prefix, and the required entry point is determined separately.

Just before a foreground program is entered the machine is reset to its EMS state; i.e. all the hardware and all the firmware are initialized. The environment and entry conditions are as follows:

Memory:

Section 2 describes the memory layout of the system. Three areas of memory are available to the program:

1. The Static Variables Area.

The area from #AC00 to #B0FF inclusive is reserved for use by the foreground program - although it may use more or less as it requires. It is also possible to reserve a foreground data area starting at #0040 if this is required.

2. The Stack.

The hardware stack is set to an area immediately below #C000 which is at least 256 bytes long.

3. The Main Memory Pool.

Most of the rest of memory will be available to the foreground program, depending on what memory is taken by any background ROMs which the foreground program chooses to initialize.

Registers:

The base and limit of the free memory area are passed to the program in registers.

BC = Address of the highest usable byte in memory. (# B0FF)

DE = Address of the lowest byte in the memory pool. (#0040)

HL = Address of the highest byte in the memory pool. (#ABFF)

Note that the program is free to use any memory between the address given in DE and the address in BC inclusive (i.e. #0040 to #B0FF). The contents of HL reflect the standard allocation for static variables; the program is free to use more, or less, as the mood takes it. Also a foreground data area may be reserved at the bottom of store as well. The program should set HL and DE to reflect the area it is using for variables before initializing any background ROMs (see below).

SP is set to the machine provided area at #C000. The program can depend on at least 256 bytes of stack.

The contents of the other registers is indeterminate. Note that the alternate register set (AF' BC' DE' HL') is not available to the program. (But see Appendix XI).

ROM select and state:

For ROM programs:	The foreground ROM is selected. The upper ROM is enabled. The lower ROM is disabled.
For RAM programs:	No ROM is selected. The upper ROM is disabled. The lower ROM is disabled.

General:

Interrupts are enabled.

All hardware and firmware is in its initial state. In particular any expansion devices fitted have been reset, but not yet initialized.

It is the foreground program's responsibility to initialize any background ROMs required and to load and initialize any RSXs. The Kernel entry 'KL ROM WALK' looks for background ROMs and initializes any that it finds. The Kernel entry 'KL INIT BACK' will initialize a particular background ROM. These entries must be passed the addresses of the lowest and highest bytes in the memory pool which is why the foreground program must reserve its fixed data areas before winding up the background ROMs. The background ROMs may allocate memory for their own use by moving either or both boundaries. If, therefore, the foreground program does allow background ROMs to function it must cope with a memory pool whose bounds are not fixed until after all the background ROMs have been initialized. Note that the location of the foreground program's data areas are fixed whilst a background program must deal with variable data areas.

If background ROMs are not initialized then the memory map is very simple, but since discs, light pens, etc are likely to use background ROMs for support software it is rather limiting not to allow background ROMs even for an apparently 'dedicated' game.

The on-board BASIC initializes all background ROMs at EMS. The user chooses whether to load any RSXs from tape.

10.4 Background ROMs.

Background ROMs lie dormant until initialized by the foreground program. During initialization the background software may allocate itself some memory and initialize any data structures and hardware. Providing the initialization is successful the Kernel places the ROM on the list of possible takers for external commands.

The first entry in a background ROM's jumpblock is its initialization routine. This routine must only be called by the firmware when the ROM is initialized it is not meant for the user to call. Tricks such as including a space in the name makes it impossible for BASIC to generate the correct name and hence impossible for a BASIC user to call the entry. The entry and exit conditions for the initialization routine are:

Entry:

- DE contains the address of the lowest byte in the memory pool.
- HL contains the address of the highest byte in the memory pool.

Exit:

If the initialization was successful:

- Carry true.
- DE contains the new address of the lowest byte in the memory pool.
- HL contains the new address of the highest byte in the memory pool.

If the initialization failed:

- Carry false.
- DE and HL preserved.

Always:

- A, BC and other flags corrupt.
- All other registers preserved.

Notes:

- The upper ROM is enabled and selected.
- The lower ROM is disabled.
- The routine may not use the alternate register set.

The ROM may allocate itself memory (either at the top or the bottom of the memory pool or both), simply by changing the appropriate register and returning the new value. For example, to reserve 256 bytes given an address of #AB7F as the top of the pool the program would subtract 256 from HL giving a new top of pool address of #AA7F. The area reserved would be from #AA80 to #AB7F inclusive.

The carry false return is only recognized in V1.1 firmware. In V1.0 firmware this will be treated as if carry was returned true.

When the initialization routine returns, the Kernel stores the address of the base of the upper area which the ROM has allocated to itself (i.e. HL+1). Whenever an entry in the ROM is called this address is passed in the IY index register. This allows the ROM routines to access its upper variable area easily even though it was allocated dynamically. Access to any lower variable area should be done via pointers in the upper area. Since background ROMs do not use absolute areas of memory, problems of background ROMs clashing with each other or with the foreground program will never arise. Note that a background ROM is very likely to expect that its upper data area lies above #4000 so that it is accessible irrespective of whether the lower ROM is enabled or not.

If the initialization is successful then the Kernel also places the ROM on its list of possible handlers of external commands (see below). Note that when the list is scanned for external commands the latest addition is tried first. The entry KL ROM WALK processes the ROMs in reverse address order (15, 14, ... 0), ignoring any gaps or foreground ROMs, thus the ROMs will be searched in the order 0, 1, .. 15.

10.5 Resident System Extensions.

An RSX is similar to a background ROM. Responsibility for loading an RSX and providing it with memory lies with the foreground program. To fit in with the dynamic allocation of memory to background ROMs it is recommended that RSXs should be position independent or relocated when loaded. An RSX could be relocated by writing a short BASIC 'loader' program which reads the RSX in a format which may be relocated easily and POKEs it into store.

Once an RSX is loaded it may be placed on the list of possible handlers of external commands (see following page) by calling KL LOG EXT, passing it the address of the RSX's external command table and a four byte block of memory (in the central 32K of RAM) for the Kernel's use. The format of the table is exactly the same as for a background ROM see section 10.2). The only difference is in the interpretation of the table - the first entry in the jumpblock is not called automatically by the Kernel and thus need not be the RSX's initialization routine.

For example, the way to add an external command table for a graphics extension for BASIC might be:

```
INITIALIZE:    LD      HL,WORK_SPACE ;RSXpower-up Routine
              LD      BC,RSX_TABLE
              JP      KL_LOG_EXT

WORK_SPACE:   DEFS    4                ;Area for Kernel to use

RSX_TABLE:    DEFW    NAME_TABLE

              JP      DRAW_CIRCLE      ;0
              JP      DRAW_TRIANGLE    ;1
              JP      FILL_AREA        ;2

NAME_TABLE:   DEFB    'CIRCL','E'+#80   ;0
              DEFB    'TRIANG','L'+#80  ;1
              DEFB    'FIL','L'+#80     ;2
              DEFB    #00
```

Note that when the list is scanned for external commands the latest addition is tried first. Since RSX's will, in general, be loaded after background ROMs have been initialized, RSX commands will take precedence over those in background ROMs. The entry and exit conditions for external commands are discussed in the following section (section 10.6).

10.6 External Commands.

Once the foreground program has decided that it has an external command on its hands it should call the Kernel entry KL FIND COMMAND, passing to it a string giving the command name. This routine first attempts to find an RSX or a background ROM whose external command table contains the command. Only those RSXs and ROMs which have been suitably initialized are taken into consideration. If the command is found then the 'far address' of the corresponding jumpblock entry is returned (see section 2.3). If the command is not found the routine starts at ROM 0 and searches for a foreground ROM whose external command table contains the command. If a foreground ROM is found, then the system resets and enters the appropriate foreground program. If no match for the command can be found a failure indication is returned.

Note that the external command mechanism allows both for the finding of background and RSX routines, and for the switching of foreground programs. Note also that the first command name in a background ROM corresponds to the implicit initialization entry, and should not be used as a command.

The first time a background or RSX routine is used the external command mechanism should be used to establish its jumpblock address. This may then be stored and used directly for subsequent calls of the routine. It is foolish to assume that a particular background ROM is always plugged into the same socket or that a relocatable RSX is always located at the same address.

The first time a background or RSX routine is used the external command mechanism should be used to establish its jumpblock address. This may then be stored and used directly for subsequent calls of the routine. It is foolish to assume that a particular background ROM is always plugged into the same socket or that a relocatable RSX is always located at the same address.

It is the foreground program's responsibility to invoke the external command once its address has been found, and to pass it parameters in a suitable form. BASIC in the on-board ROM functions as follows, and should serve as a model for other foreground programs if only to allow common use of commands by other systems:

An external command is identified by a vertical bar ('|') followed by the command name, optionally followed by a list of parameters. The bar does not form part of the command name. The command name must consist of alphabetic characters (which are converted to upper case), numeric characters or dots.

Parameters are passed by value, that is each parameter may be a numeric expression, the calculated value of which is passed, or an address. The number and type of parameters must be agreed between the BASIC program and the command because BASIC performs no checking.

Each parameter passed is a two byte number, whose interpretation depends on its type:

Integer expression: two's complement value of the Integer result.
Real expression: the Real result forced to Unsigned Integer.
Variable reference: address of the value of a variable (for a string this is the address of the descriptor).

A string descriptor is three bytes long. Byte 0 contains the length of the string. Bytes 1 and 2 contain the address where the string is stored. If the string length is 0 then the address of the string is meaningless. String variables may be changed provided that the string descriptor is not altered in any way.

Entry:

A contains the number of parameters.

IX contains the address of the parameters.

IY contains the address of the ROM's upper data area if the command was found in a background ROM. If the command was found in an RSX's external command table then IY is undefined.

Exit:

AF, BC, DE, HL, IX and IY corrupt.

Alternate register set untouched.

Notes:

Index register IX contains the address of the parameters. If there are n parameters then the ith parameter is at offset (n-i)X2 from the index register address. so the 1st parameter is at the largest offset, and the last parameter is pointed to by IX.

The IY register is set by the Kernel and not by BASIC. The A and IX registers

10.7 Examples

a) A simple external command

This example uses the BIOS routine SET MESSAGE that is available as an external command under AMSDOS. SET MESSAGE turns on or off the disc error messages and has the following interface:

SET MESSAGE Command name: Control A

Entry conditions:

A= #00 => Turn disc error messages on.

A = #FF => Turn disc error messages off.

Exit conditions:

A = Previous state.

HL and flags corrupt.

Before it is possible to use the external command it is necessary to establish and store the far address of the routine. This may be performed as follows:

```

name      LD   HL, CMD_NAME          ;Pointer to command
          CALL KL_FIND_COMMAND      ;Ask Kernel where it is
error     JR   NC, ERROR_ROUTINE    ;Command not found
;
          LD   (CMD_FAR_ADDRESS + 0), HL;Store address
          LD   A,C
          LD   (CMD_FAR_ADDRESS + 2), A ;Store ROM number
          ...

CMD_NAME:  DEFB #01+#80             ;Control A = #01
CMD_FAR_ADDRESS:  DEFS 3           ;Area for storing far
address
```

Having found the far address of the routine it can now be called. For example:

```

          LD A, 0                   ;Enable messages
          RST 3                     ;Far CALL
          DEFW CMD_FAR_ADDRESS      ;Pointer to far address
```

b) A complex external command

This example uses the INCHAR external command provided by the Serial interface. INCHAR reads a character from the Serial Interface and has the following interface:

INCHAR Command name: INCHAR

Entry conditions:

A = Number of parameters (should be 2).

IX = Address of parameter block

IX +2 = Address to store status/

IX +0 = Address to store character read.

Exit conditions:

AF, BC, DE, HL, IX AND IY corrupt.

Before it is possible to use the external command it is necessary to establish and store the far address of the routine. This may be performed as follows:

```
name      LD HL, CMD_NAME          ;Pointer to command
          CALL KL_FIND_COMMAND   ;Ask Kernel where it is
error     JR    NC, ERROR_ROUTINE ;Command not found
;
          LD  (CMD_FAR_ADDRESS +0), HL ;Store address
          LD  A,C
          LD  (CMD_FAR_ADDRESS +2),A  ;Store ROM number
          ...

CMD_NAME:      DEFB 'INCHA', 'R'+ #80
CMD_FAR_ADDRESS: DEFS 3          ;Area for storing far
address
```

Having found the far address of the routine it can now be called. For example:

```
          LD  A, 2                ;2 parameters
          LD  IX, PARAM_BLOCK     ;Address of parameter block
          RST 3                   ;FAR CALL
          DEFW CMD_FAR_ADDRESS    ; Pointer to far address
          LD  HL, (STATUS)        ;HL Serial Interface status
          LD  A, (CHAR)           ;A Characterread(ifany)
          ...

PARAM_BLOCK  DEFW STATUS          ;First parameter is status
             DEFW CHAR           ;Second parameter is
character
;
STATUS:      DEFW #0000
CHAR:        DEFW #0000
```


The external command routine that is being called has to pick the parameters out of the parameter block and it might work as follows:

```
LD    L,(IX+0)
LD    H, (IX + 1)           ;HL = Address of string descriptor
;
LD    A,(HL)
INC   HL                   ;A = Length of the string

LD    E,(HL)
INC   HL
LD    D, (HL)
EX    HL, DE               ;HL = Address of string
;
LD    E,(IX+2)
LD    D,(IX +3)           ;DE Index number

...                          ;Look up string

LD    (IX+4),L
LD    (IX +5), H          ;Store resulting reference number
RET
```

11 Interrupts.

There is only one source of interrupts in an unexpanded machine, namely a regular time interrupt. Expansion boards may generate interrupts, but suitable software must be provided to deal with the extra interrupts.

The system runs with interrupts enabled most of the time. It is inadvisable to disable interrupts for a prolonged period if this is avoidable because the time interrupts will be missed.

A number of firmware routines enable interrupts and this is remarked upon in their descriptions. In particular the Kernel routines dealing with ROMs and the restart instructions (e.g. LOW JUMP) enable interrupts.

11.1 The Time Interrupt.

The time interrupt occurs roughly once every 1/300th of a second. On machines with PAL monitors (as in the UK) or SECAM monitors (as in France) the timer is synchronised with frame flyback every sixth tick. On machines using NTSC monitors (as in the US) the timer is synchronised with frame flyback every fifth tick. The time interrupt is processed by the Kernel and presented to the rest of the system in a number of ways:

a. Fast Ticker Interrupts. Period = 1/300th of a second.
For high resolution or very short period timing (not intended for general use).

b. Sound Generation Interrupt. Period = 1/100th of a second.
This interrupt drives the sound generation firmware, but is otherwise not visible to the system.

c. Frame Flyback Interrupt. Period = 1/50th or 1/60th of a second.
For actions which must take place during frame flyback. Ink flashing is performed during a frame flyback interrupt, for example.

d. Ticker Interrupt. Period = 1/50th of a second.
This is the general purpose ticker interrupt. The keyboard is scanned at the start of each ticker interrupt.

e. System Clock.

There is a timer that counts fast ticks i.e. 1/300ths of a second. This can be used to measure elapsed time without setting up a relatively expensive fast tick event. The timer is read by calling KL TIME PLEASE and may be set by calling KL TIME SET.

11.2 External Interrupts.

The Z80 is run in interrupt mode 1. Which is to say that all interrupts cause an RST 7 to be executed by the processor. The interrupt handling code in the Kernel can distinguish between the time interrupt and an external interrupt. It does this by re-enabling interrupts inside the interrupt routine. If the interrupt repeats then it is assumed to be an external interrupt, otherwise it is taken to be a time interrupt. Note that this requires that the source of external interrupts should not clear the interrupt condition until the software resets it. Unless special action is taken in hardware, Z80 peripheral chips will not obey this requirement. In such cases the recommended course of action is described in Appendix XIII.

Before an external interrupt is enabled its interrupt handler must be 'installed'. This is done by copying the 5 bytes at address #003B to a new location and replacing them by suitable code (probably including a jump). When the Kernel detects an external interrupt it calls address #003B in RAM to process the interrupt:

Entry:

No conditions.

Exit:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

Interrupts are disabled and must remain disabled.
The lower ROM is disabled.
The upper ROM select and state are indeterminate.
The alternate register set must not be touched.

The interrupt routine must establish whether it can deal with the interrupt, and if so it must at least clear it. If the interrupt is not the responsibility of the routine then it should jump to the copy of the bytes taken from location #003B which may be competent to deal with the interrupt. This requires the code patched at location #003B to be position independent in case a second external interrupt handler is installed. The code put at #003B at EMS is position independent - it merely returns.

Note that interrupt handling code must be in RAM somewhere between #0040 and #BFFF. Interrupt handlers should be as short as possible. If an interrupt requires a lot of processing beyond that required to clear it, then the interrupt should kick an event to do the work outside the interrupt path.

11.3 Nonmaskable Interrupts.

There is no provision for handling a nonmaskable interrupt (NMI) in the firmware (despite the fact that NMI is available on the external bus connector). Various firmware routines (notably those connected with driving the Centronics port, the PPI to access the sound chip and keyboard, and the cassette) will have timing constraints violated if NMIs occur whilst they are active. It is recommended that NMI should not be used.

11.4 Interrupts and Events.

As a general rule hardware interrupts should be transformed into their software equivalents, 'events', as soon as possible. The handling of events is more flexible than the handling of hardware interrupts - for example there are no restrictions on where event routines may reside, or on interrupt enabling.

Events are described by an event block. This block contains the event class, the event count and an event routine address. When an event occurs the event block is 'kicked' and the Kernel arranges for the event, routine to be called once for each kick (the number of kicks outstanding is kept in the event block). The event routine is not necessarily called immediately. When the event routine is actually run depends on the event class as follows:

a. Express Asynchronous Events.

This is an unusual class of event. The event routine is called immediately during interrupt processing. The routine must be accessible by the interrupt code, it may not enable interrupts, corrupt the IX or IY registers or use the alternate register set. The routine should be as short as possible.

b. Normal Asynchronous Events.

This is the most flexible sort of event. When the event is kicked the event routine is not called, but the event block is placed on the interrupt event pending queue.

Once the current interrupt has been processed, just before the Kernel returns from the interrupt path, any events on the interrupt event pending queue are processed. While the events are being processed the system is running with interrupts enabled and may be regarded as no longer being in the interrupt path. It is using its own stack rather than the main system stack. This private stack is 128 bytes long.

The asynchronous event routine is, therefore, called shortly after the event is kicked and is not restricted in what it may do or where it may be located. The event routine may take as long to run as is needed. Any further kicks received during the time that the event routine is running will be added to the event count and will be processed before returning to the interrupted program.

c. Synchronous Events.

Synchronous events are queued on the synchronous event pending queue. They are not processed until the foreground program allows the queue to be processed. This can be used to control interactions between different parts of programs.

11.5 Interrupt Queues.

The various time interrupts provide three sources of 'kicks' for events. The events to be kicked when each of the interrupts occur are stored on queues, one queue for each source of kicks. The user provides an area of store for the Kernel's use. The size of the area depends on which queue it is for. The last 7 bytes of the area are always an event block which the user should initialize appropriately. Appendix X describes the layout of these blocks in greater detail.

a. Fast Ticker Events.

Events on the fast ticker queue are 'kicked' on each fast ticker interrupt, i.e. every 1/300th of a second. A fast ticker block is 9 bytes long.

b. Ticker Events.

Each event on the ticker queue is associated with a timer. The timer may be a 'one shot', which goes off once, or a repeater, which goes off periodically. The timer counts ticker interrupts, i.e. 1/50ths of a second, and when sufficient have occurred it goes off. Each time the timer associated with an event goes off the event is kicked. A ticker block is 13 bytes long.

c. Frame Flyback Events.

Events on the frame flyback queue are kicked on each frame flyback interrupt, i.e. every 1/50th of a second on PAL or SECAM machines and every 1/60th of a second on NTSC machines. A frame flyback block is 9 bytes long.

12 Events.

The event mechanism is primarily provided by the Kernel to support the handling of interrupts and other external events. However, the mechanism may also be used to handle internal events in complicated programs (such as a simulation, for example). An event is characterised by the following:

a. Event Class (see section 12.1)

Events may be synchronous or asynchronous, express or normal.

b. Event Priority (see section 12.1)

Synchronous events have an associated priority.

c. Event Count (see section 12.2)

Each time an event occurs the count is incremented.
Each time an event is processed the count is decremented.
The event may be disarmed by setting the count negative.

d. Event Routine (see section 12.3)

The address of the routine which is called to process the event.

An event appears to the Kernel as a data block containing the above values (see Appendix X for the exact layout of an event block). The block must be in the central 32K bytes of memory, so that the Kernel can access it without worrying about the ROM enable state.

When an event occurs the associated event block is kicked by calling KL EVENT. If the event count is negative, the 'kick' is ignored, otherwise the event count is incremented (up to a maximum of 127) and the event routine will be called at some time in the future - depending on the event class. When the event routine returns the event count is decremented, unless it has been set to zero or negative in the meantime.

12.1 Event Class.

Events are either synchronous or asynchronous. Asynchronous events are intended for the processing of external events which require almost immediate service. The processing of asynchronous events pre-empts the main program. The processing of synchronous events is under the complete control of the main program, which will, in general, deal with them when it is convenient to do so.

a. Asynchronous Events.

An asynchronous event is processed immediately the event is kicked - or almost immediately if the kick occurs in the interrupt path - see section 11 on interrupts. The Kernel does not provide any interlocks between asynchronous events and the main program or other events, so care must be exercised to avoid interactions. It is most unwise to call routines that are not re-entrant - for example, the firmware screen driving routines.

If the event count is still greater than zero when the event routine returns, it is decremented. If the count remains greater than zero then the process is repeated (the event routine is called again and the event count is decremented) until the count becomes zero or is set negative (see 12.2 below).

b. Synchronous Events.

Synchronous events are not processed when the event is kicked, but are placed on the synchronous event queue, waiting to be processed. Events are queued in descending order of priority - equal priority events after those already on the queue.

The foreground program should poll the synchronous event queue regularly, to see if there are any events outstanding. If there are then it should then process them. The difference between synchronous and asynchronous events is, therefore, that the foreground program decides when synchronous events should be processed, but the event 'kicker' decides when asynchronous events are to be processed. Provided that the foreground program takes suitable care, there should be no difficulty in handling the interactions and resource sharing between synchronous events and the foreground program.

When the foreground program finds the synchronous event queue is not empty it should (but is not constrained to) instruct the Kernel to process the first event on the queue. When a synchronous event routine is run the Kernel remembers the priority of the event. In the event routine the synchronous event queue may be polled, but the Kernel hides any events whose priority is less than or equal to that of the event currently being processed. When the event routine returns the previous event priority is restored - so the processing of events may be nested.

The synchronous event priorities are split into two ranges, express and normal. All express events have higher priorities than all normal events. The Kernel provides a mechanism to disable the processing of normal events, without affecting express events. This may be used to implement 'critical regions' through which 'normal events may interact. The synchronous event 'kicked' by the Key Manager break handling mechanism is an example of an express synchronous event.

12.2 Event Count.

The main purpose of the event count is to keep track of the difference between the number of times the event has been kicked, and the number of times the event has been processed. This ensures that a kick is not missed if it occurs before the previous kick has been processed. The event count is normally incremented when the event is kicked and decremented when the event routine returns. However the exact action depends on the event count as follows:

Increment.

- 128..-2: The count is not changed - the event is ignored.
- 1: This value is illegal.
- 0: The count is incremented and event processing is initiated as required by the event class.
- 1..126: The count is incremented but no further action is taken. The event is waiting for a previous kick to be processed or for processing to complete.
- 127: The count is not changed - the kick is ignored.

Decrement.

- 128: This value is illegal.
- 127..0: The count is not changed - the event has been disarmed.
- 1: The count is decremented and event processing is terminated.
- 2..127: The count is decremented and event processing continues.

Note that the event routine may disarm itself by setting the count negative (by convention to -64) and can discard unwanted kicks by setting its count to one.

12.3 Event Routine.

In general the address of the event routine is given as a 3 byte 'far address' (see section 2 on the memory layout). This allows the routine to be located in any ROM or anywhere in RAM.

A special form of the event class may specify the routine as at a 'near address'. This does not change the ROM state and so the routine must be located either in the lower ROM or in the central 32K of RAM. The ROM select byte of the 'far address' is ignored and the other two bytes taken as the address of the routine. Calling a 'near address' event routine requires a little less work than calling a full 'far address', and is used by the firmware itself.

12.4 Disarming and Reinitializing Events.

Before an event block may be reinitialized the event must be disarmed. This ensures that the event is removed from the various event pending queues and prevents the event queues being corrupted when the event block is initialized. An asynchronous event must not be reinitialized from inside its asynchronous event routine (because in this case disarming the event does not remove the event from the interrupt event pending queue). Synchronous and asynchronous events are disarmed in different manners.

a. Asynchronous Event..

An asynchronous event should be disarmed by calling KL DISARM EVENT. This sets the event count to a negative value (-64) and thus prevents kicks having any effect. If the event is on the interrupt event pending queue then it will be discarded only when an attempt is made to process the event and not immediately that the event is disarmed.

b. Synchronous Events.

A synchronous event should be disarmed by calling KL DEL SYNCRRONOUS. This sets the eventcount to a negative value (-64) and removes the event block from the synchronous event pending queue (if it is on the queue).

The above procedures prevent the event being successfully kicked, they do not prevent attempts being made to kick the event. A fast ticker, frame flyback or ticker event (see section 11.5) will still be on its appropriate queue and will still be receiving regular attempts to kick it. To prevent time being wasted (and the system from being slowed down because of it) the event should be removed from the interrupt queue by calling KL DEL FAST TICKER, KL DEL FRAME FLY or KL DEL TICKER.

13 The Machine Pack.

The Machine Pack deals with the low level driving of the hardware. It also talks to the Centromcs port (and hence the printer) and is in charge of running 'load and go' programs.

13.1 Hardware Interfaces.

The routines provided for driving the hardware are only to be used by those who understand the hardware and how the firmware drives the hardware. The user should not access the hardware directly when a Machine Pack routine is provided for this purpose.

Often there are higher level routines that accomplish the same effects but that also keep the firmware informed of the current settings. Where possible these higher level routines should be used and the Machine Pack routines avoided. Using the Machine Pack routines may cause the firmware to make erroneous assumptions about the current settings and may cause it to go wrong.

The Machine Pack makes certain assumptions about the state of the hardware when it accesses it. In particular, PPI port A is assumed to be in output mode and the sound chip, ULA, CRTC and Centronics port are assumed to be inactive; that is, not halfway through setting a value into a chip register. It is usually essential that interrupts be disabled when accessing the hardware directly.

There are four main areas of the hardware that the Machine Pack deals with:

a. The screen.

There are three aspects of the screen display that can be set using Machine Pack routines. These are the screen mode (set by calling MC SET MODE) and the screen base and offset (set by calling MC SET OFFSET).

The screen mode sets how many pixels are displayed on the screen and how many inks may be used as follows:

Mode	Resolution	Inks
0	160x200	16
1	320x200	4
2	640x200	2

The screen base sets which 16K block of memory is used for the screen memory. Theoretically, any of #0000, #4000, #8000 or #C000 could be used but, in practice, other considerations mean that only #4000 and #C000 are useful.

The screen offset sets which byte in the screen memory is to be displayed first. Changing the offset will move the contents of the screen in one go. This is used for rolling the screen.

A fuller description of the screen layout and its relationship to these aspects can be found in section 6 on the Screen Pack.

If addresses are to be read back from the CRT controller chip, when using a light pen for instance, then careful inspection of the way the screen memory is addressed will be needed to translate the screen address read from the chip to the actual position on the screen.

The Machine Pack also provides a routine (MC WAIT FLYBACK) to wait until frame flyback occurs (the start of the vertical retrace period). This may be used to ensure that operations on the screen are performed with as little disruption as is possible to the picture on the monitor since no picture is generated during this period. As an alternative to waiting for frame flyback explicitly the user should consider setting up a frame flyback event as described in section 11.5.

The vertical retrace period is not very long. Furthermore, approximately 100 microseconds from its start, a time interrupt occurs that will cause the frame flyback events to be processed (see section 11). These may take a significant length of time out of the retrace period.

b. The inks.

The Machine Pack deals with setting the colours of inks. There is a fuller explanation of the relationship between inks and colours in section 6.2. Briefly, the colour for each ink and the border can be specified independently and changed at will. Note, however, that the Machine Pack deals with the hardware representations of colours and not the grey scale colours that the Screen Pack uses and also that an ink may only be set to one colour, the flashing inks are made by the Screen Pack setting two colours alternately.

Two routines are provided for setting the colours of inks. MC SET INKS allows the colours of all 16 inks and the border to be set (although not all of the inks may be visible on the screen in the current mode). MC CLEAR INKS sets the colour of the border ink and sets all 16 inks to the same colour. The latter is used when clearing the screen to make the operation appear instantaneous.

c. The sound chip.

A routine, MC SOUND REGISTER, is provided to write to a register of the sound chip. This is used by the Sound Manager for hardware access.

d. The Centronics port.

Two routines are provided to access the Centronics port. MC BUSY PRINTER tests if it is busy. MC SEND PRINTER strobes data out of it. Data should not be sent while the port is busy.

The Centronics port is used by the printer routines provided in the Machine Pack and described below.

13.2 The Printer.

There is a routine, MC PRINT CHAR, which calls an indirection, MC WAIT PRINTER, for sending characters to the printer, or rather, to the Centronics

port. In V1.1 firmware, before sending the characters MC PRINT CHAR translates them using the printer translation table. The user can set the translations by calling MC PRINT TRANSLATION. The default translation table (see Appendix XIV) only affects characters in the range #AO.. #AF and is designed to make these characters print more reasonably on the DMP-1 printer for various language options.

MC WAIT PRINTER waits until the Centronics port is not busy and then sends the given character to it. If the port remains busy for a long time then the routine times out and returns indicating that it has failed to send the character. This time out can be used to prevent programs 'hanging' because they are waiting for a (possibly non-existent) printer to become ready.

MC WAIT PRINTER allows the user to intercept characters to be sent to the printer. This could allow special escape sequences to be inserted if needed, or it could allow the printer to be disabled or the length of the time out to be changed.

13.3 Loading and Running Programs.

The Machine Pack provides two routines for running programs, MC START PROGRAM and MC BOOT PROGRAM.

MC START PROGRAM is the simpler of the two routines. It completely re-initializes all the firmware and then enters the given program.

MC BOOT PROGRAM is more complex. It is for loading a program into RAM and running it. The user supplies a routine to MC BOOT PROGRAM that will load the program and return its entry point. Before this load routine is called as much of the firmware as is possible is reset so that the area of memory between #0040 and the base of the firmware RAM at #B100 is available for use. If the system were not reset then an active indirection, event or interrupt routine might be overwritten with disastrous consequences.

If the program is loaded successfully by MC BOOT PROGRAM then the firmware is completely initialized and the program is entered. However, if the loading fails then an appropriate message is printed and the previous foreground program is restarted. If the previous program was itself a RAM program then the default ROM is entered instead because it is likely that the previous program was corrupted when the attempt to load the new one was made.

14 Firmware Jumpblocks.

There are a number of jumpblocks provided by the firmware. The largest of these is the main firmware jumpblock. This is intended to be used by programs to access the firmware routines in the lower ROM. BASIC, for instance, uses these jumps. Note, however that the firmware does not use this jumpblock for internal communication with itself. This means that altering the jumpblock will cause BASIC to behave differently but will not cause the firmware to behave differently.

The next most important jumpblock is the indirections jumpblock. The indirections are jumps that are used by the firmware at key points. This allows the user to alter the action of firmware routines. The entries in this jumpblock are not intended for the user to call, only for the firmware to call. Altering an indirection is the method to make the firmware behave differently.

The remaining two jumpblocks are associated with the Kernel. One is a jumpblock to allow the user to call various useful Kernel routines to do with changing ROM states and the like. The other is not a jumpblock as such, just an area where the routines are at published addresses. These are general utility routines and restarts. In general neither of these areas should be altered by the user.

The routines in these jumpblocks are briefly listed below. More complete descriptions of the routines can be found in sections 15, 16, 17, and 18.

AMSDOS provides a number of external commands which allow the user access to the low level disc driving and to high level disc operations.

These commands are accessed using the external command mechanism described in section 10, i.e. The caller passes the command name to KL FIND COMMAND and far calls the resulting routine. More complete descriptions of these commands can be found in sections 19 and 20.

14.1 The Main Jumpblock.

The main firmware jumpblock lies in RAM between addresses #BB00 and #BD5D. Each entry in the jumpblock occupies three bytes and is initialized to use LOW JUMP restarts (RST1) that cause the lower ROM to be enabled, so that the firmware routines can be run, and the upper ROM to be disabled, so that the screen memory is accessible while the firmware is running. Full descriptions of these routines can be found in section 15.

After the jumpblock has been set up at EMS it is patched by the initialization of the AMSDOS ROM to install the disc (rather than cassette) as default but is not otherwise altered by the firmware until the system is reinitialized. If any entries are changed then it is the user's responsibility to undo the alterations. This can be achieved by calling JUMP RESTORE which completely initializes the jumpblock but this will lose any other patches, such as those made by AMSDOS. It is better to copy the original contents of the changed entries back.

14.1.1 Entries to the Key Manager

The Key Manager deals with the keyboard and the joysticks.

	INITIALIZATION		
0	#BB00	KM INITIALISE	Initialize the Key Manager.

1	#BB03	KM RESET	Reset the Key Manager - clear all buffers, restore standard key expansions and indirections.
CHARACTERS			
2	#BB06	KM WAIT CHAR	Wait for next character from the keyboard.
3	#BB09	KM READ CHAR	Test if a character is available from the keyboard.
4	#BB0C	KM CHAR RETURN	Return a single character to the keyboard for next time.
191	#BD3D	KM FLUSH	Discard all pending characters and keys.
5	#BB0F	KM SET EXPAND	Set an expansion string.
6	#BB12	KM GET EXPAND	Get a character from an expansion string.
7	#BB15	KM EXP BUFFER	Allocate a buffer for expansion strings.
KEYS			
8	#BB18	KM WAIT KEY	Wait for next key from the keyboard.
9	#BB1B	KM READ KEY	Test if a key is available from the keyboard.
10	#BB1E	KM TEST KEY	Test if a key is pressed.
190	#BD3A	KM SET LOCKS	Set the Shift Lock and Caps Lock states.
11	#BB21	KM GET STATE	Fetch Caps Lock and Shift Lock states.
12	#BB24	KM GET JOYSTICK	Fetch current state of the joystick(s).
TRANSLATION TABLES			
13	#BB27	KM SET TRANSLATE	Set entry in key translation table without shift or control.
14	#BB2A	KM GET TRANSLATE	Get entry from key translation table without shift or control.
15	#BB2D	KM SET SHIFT	Set entry in key translation table when shift key is pressed.
16	#BB30	KM GET SHIFT	Get entry from key translation table when shift key is pressed.
17	#BB33	KM SET CONTROL	Set entry in key translation table when control key is pressed

18	#BB36	KM GET CONTROL	Get entry from key translation table when control key is pressed.
REPEATING			
19	#BB39	KM SET REPEAT	Set whether a key may repeat.
20	#BB3C	KM GET REPEAT	Ask if a key is allowed to repeat.
21	#BB3F	KM SET DELAY	Set start up delay and repeat speed
22	#BB42	KM GET DELAY	Get start up delay and repeat speed.
BREAKS			
23	#BB45	KM ARM BREAK	Allow break events to be generated.
24	#BB48	KM DISARM BREAK	Prevent break events from being generated.
25	#BB4B	KM BREAK EVENT	Generate a break event (if armed).

14.1.2 Entries to the Text VDU

The Text VDU is a character based screen driver.

INITIALIZATION

26	#BB4E	TXT INITIALISE	Initialize the Text VDU.
27	#BB51	TXT RESET	Reset the Text VDU - restore default indirections and control code functions.
28	#BB54	TXT VDU ENABLE	Allow characters to be placed on the screen.
29	#BB57	TXT VDU DISABLE	Prevent characters from being placed on the screen.
192	#BD40	TXT ASK STATE	Get state of the text VDU.

CHARACTERS

30	#BB5A	TXT OUTPUT	Output a character or control code to the Text VDU.
31	#BB5D	TXT WR CHAR	Write a character onto the screen.
32	#BB60	TXT RD CHAR	Read a character from the screen.

33	#BB63	TXT SET GRAPHIC	Turn on or off the Graphics VDU character writing option.
----	-------	-----------------	---

WINDOWS

34	#BB66	TXT WIN ENABLE	Set the size of the current text window.
35	#BB69	TXT GET WINDOW	Get the size of the current text window.
36	#BB6C	TXT CLEAR WINDOW	Clear current window.

CURSOR

37	#BB6F	TXT SET COLUMN	Set cursor horizontal position.
38	#BB72	TXT SET ROW	Set cursor vertical position.
39	#BB75	TXT SET CURSOR	Set cursor position.
40	#BB78	TXT GET CURSOR	Ask current cursor position.
41	#BB7B	TXT CUR ENABLE	Allow cursor display - user.
42	#BB7E	TXT CUR DISABLE	Disallow cursor display – user.
43	#BB81	TXT CUR ON	Allow cursor display – system.
44	#BB84	TXT CUR OFF	Disallow cursor display - system.
45	#BB87	TXT VALIDATE	Check if a cursor position is within the window.
46	#BB8A	TXT PLACE CURSOR	Put a cursor blob on the screen.
47	#BB8D	TXT REMOVECURSOR	Take a cursor blob off the screen.

INKS

48	#BB90	TXT SET PEN	Set ink for writing characters.
49	#BB93	TXT GET PEN	Get ink for writing characters.
50	#BB96	TXT SET PAPER.	Set ink for writing text background.
51	#BB99	TXT GET PAPER	Get ink for writing text background.
52	#BB9C	TXT INVERSE	Swap current pen and paper inks.
53	#BB9F	TXT SET BACK	Allow or disallow background being written.

54	#BBA2	TXT GET BACK	Ask if background is being written.
----	-------	--------------	-------------------------------------

MATRICES

55	#BBA5	TXT GET MATRIX	Get the address of a character matrix.
56	#BBA8	TXT SET MATRIX	Set a character matrix.
57	#BBAB	TXT SET M TABLE	Set the user defined matrix table address.
58	#BBAE	TXT GET M TABLE	Get user defined matrix table address.

CONTROL CODES

59	#BBB1	TXT GET CONTROLS	Fetch address of control code table.
----	-------	------------------	--------------------------------------

STREAMS

60	#BBB4	TXT STR SELECT	Select a Text VDU stream.
61	#BBB7	TXT SWAP STREAMS	Swap the states of two streams.

14.1.3 Entries to the Graphics VDU

The Graphics VDU deals with individual pixels.

INITIALIZATION

62	#BBBA	GRA INITIALISE	Initialize the Graphics VDU.
63	#BBBD	GRA RESET	Reset the Graphics VDU -restore standard indirections.
193	#BD43	GRA DEFAULT	Set default Graphics VDU modes.

CURRENT POSITION

64	#BBC0	GRA MOVEABSOLUTE	Move to an absolute position.
65	#BBC3	GRA MOVE RELATIVE	Move relative to current position.
66	#BBC6	GRA ASK CURSOR	Get the current position.
67	#BBC9	GRA SET ORIGIN	Set the origin of the user coordinates.
68	#BBCC	GRA GET ORIGIN	Get the origin of the user coordinates.
197	#BD4F	GRA FROM USER	Convert user coordinates to base coordinates.

WINDOW

69	#BBCF	GRA WIN WIDTH	Set left and right edges of the graphics window.
70	#BBD2	GRA WIN HEIGHT	Set the top and bottom edges of the graphics window.
71	#BBD5	GRA GET W WIDTH	Get the left and right edges of the graphics window.
72	#BBD8	GRA GET W HEIGHT	Get the top and bottom edges of the graphics window.
73	#BBDB	GRA CLEAR WINDOW	Clear the graphics window.

INKS

74	#BBDE	GRA SET PEN	Set the graphics plotting ink.
75	#BBE1	GRA GET PEN	Get the current graphics plotting ink.
76	#BBE4	GRA SET PAPER	Set the graphics background ink.
77	#BBE7	GRA GET PAPER	Get the current graphics background ink.
194	#BD46	GRA SET BACK	Set whether background is to be written.

PLOTTING

78	#BBEA	GRA PLOT ABSOLUTE	Plot a point at an absolute position.
79	#BBED	GRA PLOT RELATIVE	Plot a point relative to the current position.

TESTING

80	#BBFO	GRA TEST ABSOLUTE	Test a point at an absolute position.
81	#BBF3	GRA TEST RELATIVE	Test a point relative to the current position.

LINE DRAWING

82	#BBF6	GRA LINE ABSOLUTE	Draw a line to an absolute position.
83	#BBF9	GRA LINE RELATIVE	Draw a line relative to the current position.
195	#BD49	GRA SET FIRST	Set whether first point of a line is to be plotted.
196	#BD4C	GRA SET LINE MASK	Set mask for drawing lines.

AREA FILLING

198 #BD52 GRA FILL Fill an area of the screen.

CHARACTER DRAWING

84 #BBFC GRA WR CHAR Put a character on the screen at the current graphics position.

14.1.4 Entries to the Screen Pack

The Screen Pack interfaces the Text and Graphic VDUs to the screen hardware. Screen functions that affect both text and graphics (e.g. ink colours) are located in the Screen Pack.

INITIALIZATION

85 #BBFF SCRINITIALISE Initialize the Screen Pack.
86 #BC02 SCR RESET Reset the Screen Pack -
restore
standard indirections, ink colours and flash rates.

SCREEN HARDWARE

87 #BC05 SCR SET OFFSET Set the offset of the start of the screen.
88 #BC08 SCR SET BASE Set the area of RAM to use for the screen memory.
199 #BD55 SCR SET POSITION Set the location of the screen memory without moving the screen.
89 #BC0B SCR GET LOCATION Fetch current base and offset settings.

MODE

90 #BC0E SCR SET MODE Set screen into a new mode.
91 #BC11 SCR GET MODE Ask the current screen mode.
92 #BC14 SCR CLEAR Clear the screen (to ink zero).
93 #BC17 SCR CHAR LIMITS Ask size of the screen in characters.

SCREEN ADDRESSES

94 #BC1A SCR CHAR POSITION Convert physical coordinates to a screen position.
95 #BC1D SCR DOT POSITION Convert base coordinates to a screen position.
96 #BC20 SCR NEXT BYTE Step a screen address right one byte.
97 #BC23 SCR PREV BYTE Step a screen address left one byte.
98 #BC26 SCR NEXT LINE Step a screen address down one line.

99 #BC29 SCR PREV LINE Step a screen address up one line.

INKS

100 #BC2C SCR INK ENCODE Encode an ink to cover all pixels in a byte.
101 #BC2F SCR INK DECODE Decode an encoded ink.
102 #BC32 SCR SET INK Set the colours in which to display an ink.
103 #BC35 SCR GET INK Ask the colours an ink is currently displayed in.
104 #BC38 SCR SET BORDER Set the colours in which to display the border.
105 #BC3B SCR GET BORDER Ask the colours the border is currently displayed in.
106 #BC3E SCR SET FLASHING Set the flash periods.
107 #BC41 SCR GET FLASHING Ask the current flash periods.

MISCELLANEOUS

108 #BC44 SCR FILL BOX Fill a character area of the screen with an ink.
109 #BC47 SCR FLOOD BOX Fill a byte area of the screen with an ink.
110 #BC4A SCR CHAR INVERT Invert a character position.
111 #BC4D SCR HW ROLL Move the whole screen up or down eight pixel lines (one character).
112 #BC50 SCR SW ROLL Move an area of the screen up or down eight pixel lines (one character).
113 #BC53 SCR UNPACK Expand a character matrix for the current screen mode.
114 #BC56 SCR REPACK Compress a character matrix to the standard form.

115	#BC59	SCR ACCESS	Set the screen write mode for the Graphics VDU.
116	#BC5C	SCR PIXELS	Write a pixel to the screen ignoring the Graphic VDU write mode.
117	#BC5F	SCR HORIZONTAL	Plot a purely horizontal line
118	#BC62	SCR VERTICAL	Plot a purely vertical line.

14.1.5 Entries to the Cassette Manager/AMSDOS

The Cassette Manager handles reading files from tape and writing files to tape. AMSDOS intercepts the starred entries and redirects them so that they read from and write to disc. The external commands TAPE and DISC can be used to switch between the tape and disc versions of these routines (see section 14.6).

INITIALIZATION

119	#BC65	CAS INITIALISE.	Initialize the Cassette Manager - close all streams, set default speed and enable messages.
120	#BC68	CAS SET SPEED	Set the write speed.
121	#BC6B	CAS NOISY	Enable or disable prompt messages.

MOTOR CONTROL

122	#BC6E	CAS START MOTOR	Start the cassette motor.
123	#BC71	CAS STOP MOTOR	Stop the cassette motor.
124	#BC74	CAS RESTORE MOTOR	Restore previous state of cassette motor.

READING FILES

125	#BC77	*CASINOPEN	Open a file for input.
126	#BC7A	*CAS IN CLOSE	Close the input file properly.
127	#BC7D	*CAS IN ABANDON	Close the input file immediately.
128	#BC80	*CAS IN CHAR	Read a character from the input file.
129	#BC83	*CAS IN DIRECT	Read the input file into store.

130	#BC86	*CAS RETURN	Put the last character read back.
131	#BC89	*CAS TEST EOF	Have we reached the end of the input file yet?

WRITING FILES

132	#BC8C	*CAS OUT OPEN	Open a file for output.
133	#BC8F	*CAS OUT CLOSE	Close the output file properly.
134	#BC92	*CAS OUT ABANDON	Close the output file immediately.
135	#BC95	*CAS OUT CHAR	Write a character to the output file.
136	#BC98	*CAS OUT DIRECT	Write the output file directly from store.

CATALOGUING

137	#BC9B	*CAS CATALOG	Generate a catalogue from the tape.
-----	-------	--------------	-------------------------------------

RECORDS

138	#BC9E	CAS WRITE	Write a record to tape.
139	#BCA1	CAS READ	Read a record from tape.
140	#BCA4	CAS CHECK	Compare a record on tape with the contents of store.

14.1.6 Entries to the Sound Manager

The Sound Manager controls the sound chip.

INITIALIZATION

141	#BCA7	SOUND RESET	Reset the Sound Manager -shut the sound chip up and clear all sound queues.
142	#BCAA	SOUND QUEUE	Add a sound to a sound queue.
143	#BCAD	SOUND CHECK	Ask if there is space in a sound queue.

144 #BCB0	SOUND ARM EVENT	Set up an event to be run when a sound queue becomes not full.
-----------	-----------------	--

SOUNDS

145 #BCB3	SOUND RELEASE	Allow sounds to happen.
146 #BCB6	SOUND HOLD	Stop all sounds in mid flight.
147 #BCB9	SOUND CONTINUE	Restart sounds after they have been stopped.

ENVELOPES

148 #BCBC	SOUND AMPL ENVELOPE	Set up an amplitude envelope
149 #BCBF	SOUND TONE ENVELOPE	Set up a tone envelope.
150 #BCC2	SOUND A ADDRESS	Get the address of an amplitude envelope.
151 #BCC5	SOUND T ADDRESS	Get the address of a tone envelope.

14.1.7 Entries to the Kernel

The Kernel handles synchronous and asynchronous events. It is also in charge of the store map and switching ROMs on and off. Apart from the entries listed below, the Kernel has its own jumpblock and a number of routines whose addresses are published. These extra entries are listed in sections 14.3 and 14.4 below.

INITIALIZATION

152 #BCC8	KL CHOKE OFF	Reset the Kernel - clears all event queues etc.
153 #BCCB	KL ROM WALK	Find and initialize all background ROMs.
154 #BCCE	KL INIT BACK	Initialize a particular background ROM.
155 #BCD1	KL LOG EXT	Introduce an RSX to the firmware.
156 #BCD4	KL FIND COMMAND	Search for an RSX or background ROM or foreground ROM to process a command.

FRAME FLYBACK LIST

157	#BCD7	KL NEW FRAME FLY	Initialize and put a block onto the frame flyback list.
158	#BCDA	KL ADD FRAME FLY	Put a block onto the frame flyback list.
159	#BCDD	KL DEL FRAME FLY	Remove a block from the frame flyback list.

FAST TICK LIST

160	#BCE0	KL NEW FAST TICKER	Initialize and put a block onto the fast tick list.
161	#BCE3	KL ADD FAST TICKER	Put a block onto the fast tick list.
162	#BCE6	KL DEL FAST TICKER	Remove a block from the fast tick list.

TICK LIST

163	#BCE9	KL ADD TICKER	Put a block onto the tick list.
164	#BCEC	KL DEL TICKER	Remove a block from the tick list.

EVENTS

165	#BCEF	KL INIT EVENT	Initialize an event block.
166	#BCF2	KL EVENT	'Kick' an event block.
167	#BCF5	KLSYNCRESET	Clear synchronous event queue.
168	#BCF8	KL DEL SYNCHRONOUS	Remove a synchronous event from the event queue.
169	#BCFB	KL NEXT SYNC	Get the next event from the queue.
170	#BCFE	KL DO SYNC	Perform an event routine.
171	#BD01	KL DONE SYNC	Finish processing an event.
172	#BD04	KL EVENT DISABLE	Disable normal synchronous events.
173	#BD07	KL EVENT ENABLE	Enable normal synchronous events.
174	#BD0A	KL DISARM EVENT	Prevent an event from occurring

ELAPSED TIME

175	#BD0D	KL TIME PLEASE	Ask the elapsed time.
176	#BD10	KL TIME SET	Set the elapsed time.

BANK SWITCHING

201	#BD5B	KL BANK SWITCH	Select a memory organization.
-----	-------	----------------	-------------------------------

14.1.8 Entries to the Machine Pack

The Machine Pack provides an interface to the machine hardware. Most packs use Machine to access any hardware they use. The major exception is the Cassette Manager which, for speed reasons, performs its own hardware access.

PROGRAMS

177	#BD13	MC BOOT PROGRAM	Load and run a foreground program.
178	#BD16	MC START PROGRAM	Run a foreground program.

SCREEN

179	#BD19	MC WAIT FLYBACK	Wait for frame flyback.
180	#BD1C	MC SET MODE	Set the screen mode.
181	#BD1F	MC SCREEN OFFSET	Set the screen offset.
182	#BD22	MC CLEAR INKS	Set all inks to one colour.
183	#BD25	MC SET INKS	Set colours of all the inks.

PRINTER

184	#BD28	MC RESET PRINTER	Reset the printer indirection.
200	#BD58	MC PRINT TRANSLATION	Set the printer translation table.
185	#BD2B	MC PRINT CHAR	Translate a character then send it to the Centronics port.
186	#BD2E	MC BUSY PRINTER	Test if the Centronics port is busy.
187	#BD31	MC SEND PRINTER	Send a character to the Centronics port.

SOUND CHIP

#BD34		MC SOUND REGISTER	Send data to a sound chip
-------	--	-------------------	---------------------------

14.1.9 Entries to Jumper

Jumper sets up the main jumpblock.

INITIALIZATION

189	#BD37	JUMP RESTORE	Restore the standard jumpblock.
-----	-------	--------------	---------------------------------

14.2 Firmware Indirections

The firmware indirections listed here are taken at key points in the firmware thus allowing the user to provide substitute routines for many firmware actions, without having to replace a complete firmware package. These indirections are not intended for the user to call - there is usually a higher level routine in the main firmware jumpblock that is more suitable.

The indirections are set up by the pack to whom they apply whenever its reset (or initialize) routine is called and during EMS; they are not otherwise altered by the firmware.

The indirections are all three bytes long and use standard jump instructions (#C3). If a ROM state other than upper ROMs disabled and lower ROM enabled is required then the appropriate restart instruction might be substituted (see section 2.3). The indirections are to be found between #BDCD and #BDF6.

At this level of operation very little validation is carried out. If incorrect parameters are passed or a substitute routine corrupts a register in defiance of the documented interface then the firmware will probably cease to function as expected.

More detailed descriptions of these routines can be found in section 16.

14.2.1 Text VDU Indirections

0	#BDCD	TXT DRAW CURSOR	Place the cursor blob on the screen (if enabled).
1	#BDD0	TXT UNDRAW CURSOR	Remove the cursor blob from the screen (if enabled).
2	#BDD3	TXT WRITE CHAR	Write a character onto the screen.
3	#BDD6	TXT UNWRITE	Read a character from the screen.
4	#BDD9	TXT OUT ACTION	Output a character or control code.

14.2.2 Graphics VDU Indirections

5	#BDDC	GRA PLOT	Plot a point.
6	#BDDF	GRA TEST	Test a point.
7	#BDE2	GRA LINE	Draw a line.

14.2.3 Screen Pack Indirections

8	#BDE5	SCR READ	Read a pixel from the screen.
9	#BDE8	SCR WRITE	Write pixel(s) to the screen using the current graphics write mode.
10	#BDEB	SCR MODE CLEAR	Clear the screen to ink 0.

14.2.4 Keyboard Manager Indirections

11	#BDEE	KM TEST BREAK	Test for break (or reset).
13	#BDF4	KM SCAN KEYS	Scan the keyboard

14.2.5 Machine Pack Indirections

12	#BDF1	MC WAIT PRINTER	Print a character or time out.
----	-------	-----------------	--------------------------------

14.3 The High Kernel Jumpblock

The high Kernel jumpblock is provided to allow the user to turn ROMs on and off and to access memory underneath ROMs while they are enabled. The entries in this jumpblock are not all jump instructions, some entries are the start of routines, thus the user should not alter any of the entries in this jumpblock. The high Kernel jumpblock occupies store from #B900 upwards. More detailed descriptions of the routines in it can be found in section 17.

0	#B900	KL U ROM ENABLE	Turn on the current upper ROM.
1	#B903	KL U ROM DISABLE	Turn off the upper ROM
2	#B906	KL L ROM ENABLE	Turn on the lower ROM.
3	#B909	KL L ROM DISABLE	Turn off the lower ROM.
4	#B90C	KL ROM RESTORE	Restore the previous ROM state.
5	#B90F	KL ROM SELECT	Select particular upper ROM.
6	#B912	KL CURR SELECTION	Ask which upper ROM is currently selected.
7	#B915	KL PROBE ROM	Ask class and version of a ROM.

8	#B918	KL ROM DESELECT	Restore the previous upper ROM selection.
9	#B91B	KL LDIR	Move store (LDIR) with ROMs disabled.
10	#B91E	KL LDDR	Move store (LDDR) with ROMs disabled.
11	#B921	KL POLL SYNCHRONOUS	Check if an event with higher priority than the current event is pending.
14	#B92A	KL SCAN NEEDED	Ensure keyboard is scanned at next opportunity.

(N.B. there are no entries 12 or 13).

14.4 The Low Kernel Jumpblock.

The Kernel provides a number of useful routines in the area of memory between #0000 and #003F. These are available, in some cases, both as a published routine address and as a restart instruction. In general the routines are available both in ROM and in RAM so whether the lower ROM is enabled does not matter. There are also a couple of areas available for the user to patch to trap RST 6s and interrupts from external hardware.

The low Kernel jumpblock is not intended for the user to alter. However, it may be necessary to alter it under certain circumstances. In particular a program may need to intercept the INTERRUPT ENTRY (by patching the jump at #0038) or the RESET ENTRY (by patching the bytes from #0000..#0007). If a program does change any locations in this jumpblock (other than those in the USER RESTART or EXT INTERRUPT areas) then it is the program's responsibility to ensure that the lower ROM is enabled or the original contents are restored when any other program runs. In particular the program must sort out the state when interrupts occur (hence the need to patch the INTERRUPT ENTRY).

More detailed descriptions of the routines in this jumpblock can be found in section 18.

#0000	RST 0	RESET ENTRY	Completely reset the machine as if powered up.
#0008	RST 1	LOW JUMP	Jump to lower ROM or RAM, takes an inline 'low address' to jump to.
#000B		KL LOW PCHL	Jump to lower ROM or RAM, HL contains the 'low address' to jump to.
#000E		PCBC INSTRUCTION	Jump to address in BC.

#0010	RST 2	SIDE CALL	Call to a sideways ROM, takes inline 'side address' to call.
#0013		KL SIDE PCHL	Call to a sideways ROM, HL contains 'side address' to call.
#0016		PCDE INSTRUCTION	Jump to address in DE
#0018	RST 3	FARCALL	Call a routine in any ROM or RAM, takes an inline address of the 'far address' to call.
#001B		KL FAR PCHL	Call a routine in any ROM or RAM, C and HL contain the 'far address' to call.
#001E		PCHL INSTRUCTION	Jump to address in HL.
#0020	RST 4	RAM LAM	LD A,(HL) with all ROMs disabled.
#0023		KL FAR ICALL	Call a routine in any ROM or RAM, HL points at the 'far address' to call.
#0028	RST 5	FIRM JUMP	Jump to lower ROM, takes an inline address to jump to.
#0030	RST 6	USER RESTART	ROM version saves current ROM state in #002B, turns the lower ROM off and jumps to the RAM version. RAM version may be patched by the user between #0030 and #0037 inclusively.
#0038	RST 7	INTERRUPT ENTRY	This restart is not available as it is used for interrupts (Z80 interrupt mode 1).
#003B		EXT INTERRUPT	When an interrupt occurs on the expansion port the firmware calls location #003B in RAM. The user may patch between #003B and #003F inclusive to trap this occurrence.

15 The Main Firmware Jumpblock.

This section describes in detail the entry and exit conditions and the effects of all the routines in the main firmware jumpblock. The main firmware jumpblock is described in section 14.1.

The user is advised to read the sections on each pack before attempting to understand the jumpblock entries. The relevant sections are:

Key Manager	(KM)	Section 3.
Text VDU	(TXT)	Section 4.
Graphics VDU	(GRA)	Section 5.
Screen Pack	(SCR)	Section 6.
Sound Manager	(SOUND)	Section 7.
Cassette Manager	(CAS)	Section 8.
AMSDOS		Section 9
Kernel	(KL)	Sections 2, 10, 11 and 12.
Machine Pack	(MC)	Section 13.

The top line of each description has the following layout:

<Entry number>: <Entry name> <Entry address>

Entries in the jumpblock are numbered starting from zero. The entry address is the address to call to invoke the firmware routine or the address of the three bytes to patch to intercept the routine. The entry address can be calculated as:

Entry address = Start of jumpblock + 3 * Entry number

Each entry is named and is referred to by name throughout this manual.

The last section of each description is a list of related routines. The user is advised to look at these as the list may include routines as the list may include routines more suited for the application being considered. Conversely the routines may shed further light on how the original routine should be used.

The descriptions of the routines are for the default routine that the entry jumps to. The user may change the entry and this may alter the action of the routine. The user is advised to stick to the entry/exit conditions described otherwise programs that call the routine (BASIC for example) may cease to operate correctly.

0: KM INITIALISE

BB00

Initialize the Key Manager

Action:

Full initialization of the Key Manager (as used during EMS). All Key Manager variables, buffers and indirections are initialized. The previous state of the Key Manager is lost.

Entry Conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.

All other registers preserved.

Notes:

The Key Manager indirection (KM TEST BREAK) is set to its default routine.

The key buffer is set up (to be empty).

The expansion buffer is set up and the expansions are set to their default strings.

The key translation tables are initialized to their default translations.

The repeating key map is initialized to its default state.

The repeat speeds are set to their default values.

Shift and caps lock are turned off.

The break event is disarmed.

See Appendices II, III and IV for the default translation tables, repeating key table and expansion strings.

This routine enables interrupts.

Related entries:

KM RESET

1: KM RESET

#BB03

Reset the Key Manager.

Action:

Reinitializes the Key Manager indirections and buffers.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The Key Manager indirection (KM TEST BREAK) is set to its default routine.
The key buffer is set up(to be empty).
The expansion buffer is set up and the expansions are set to their default strings (see Appendix IV).
The break event is disarmed.

All pending keys and characters are discarded.

This routine enables interrupts.

Related entries:

KM DISARM BREAK

KM EXP BUFFER

KM INITIALISE

2: KM WAIT CHAR

#BB06

Wait for next character from the keyboard.

Action:

Try to get a character from the key buffer or the current expansion string. This routine waits until a character is available if no character is immediately available.

Entry conditions:

No conditions.

Exit conditions:

Carry true.

A contains the character.

Other flags corrupt.

All other registers preserved.

Notes:

The possible sources for generating the next character are, in the order that they are tested:

The 'put back' character.

The next character of an expansion string.

The first character of an expansion string.

A character from a key translation table.

Expansion tokens found in the key translation table are expanded to their associated strings. Expansion tokens found in expansion strings are not expanded but are treated as characters.

Related entries:

KM CHAR RETURN

KMREADCHAR

KM WAIT KEY

3: KM READ CHAR

#BB09

Test if a character is available from the keyboard.

Action:

Try to get a character from the key buffer or the current expansion string. This routine does not wait for a character to become available if there is no character available immediately.

Entry conditions:

No conditions.

Exit conditions:

If there was a character available:

- Carry true.
- A contains the character.

If there was no character available:

- Carry false.
- A corrupt.

Always:

- Other flags corrupt.
- All other registers preserved.

Notes:

The possible sources for generating the next character are, in the order that they are tested:

- The 'put back' character.
- The next character of an expansion string.
- The first character of an expansion string.
- A character from a key translation table.

Expansion tokens in the key translation tables will be expanded to their associated strings. Expansion tokens found in expansion strings are not expanded but are treated as characters.

This routine will always return a character if one is available. It is therefore possible to flush out the Key Manager buffers by calling KM READ CHAR repeatedly until it reports that no character is available.

Related entries:

KM CHAR RETURN

KM FLUSH

KM READ KEY

KM WAIT CHAR

4: KM CHAR RETURN

#BB0C

Return a single character to the keyboard for next time.

Action:

Save a character for the next call of KM READ CHAR or KM WAIT CHAR.

Entry conditions:

A contains the character to put back.

Exit conditions:

All registers and flags preserved.

Notes:

The 'put back' character will be returned before any other character is generated by the keyboard. It will not be expanded (or otherwise dealt with) but will be returned as it is. The 'put back' character need not have been read from the keyboard, it could be inserted by the user for some purpose.

It is only possible to have one 'put back' character. If this routine is called twice without reading a character between these then the first 'put back' will be lost. Furthermore, it is not possible to return character 255 (because this is used as the marker for no 'put back' character).

Related entries:

KM READ CHAR

KM WAIT CHAR

5: KM SET EXPAND

#BB0F

Set an expansion string.

Action:

Set the expansion string associated with an expansion token.

Entry conditions:

B contains the expansion token for the expansion to set.

C contains the length of the string.

HL contains the address of the string.

Exit conditions:

If the expansion is OK:

Carry true.

If the string was too long or the token was invalid:

Carry false.

Always:

A, BC, DE, HL and other flags corrupt.

All other registers preserved.

Notes:

The string to be set may lie anywhere in RAM. Expansion strings cannot be set directly from ROM.

The characters in the string are not expanded (or otherwise dealt with). It is therefore possible to put any character into an expansion string.

If there is insufficient room in the expansion buffer for the new string then no change is made to the expansions.

If the string set is currently being used to generate characters (by KM READ CHAR or KM WAIT CHAR) then the unread portion of the string is discarded. The next character will be read from the key buffer.

This routine enables interrupts.

Related entries:

KM GET EXPAND

KM READ CHAR

KM WAIT CHAR

6: KM GET EXPAND

#BB12

Get a character from an expansion string.

Action:

Read a character from an expansion string. The characters in the string are numbered starting from 0.

Entry conditions:

A contains an expansion token.
L contains the character number.

Exit conditions:

If the character was found:

Carry true.
A contains the character.

If the token was invalid or the string was not long enough:

Carry false.
A corrupt.

Always:

DE and other flags corrupt.
All other registers preserved.

Notes:

The characters in the expansion string are not expanded (or otherwise dealt with). It is therefore possible to put any character into an expansion string.

Related entries:

KM READ CHAR
KM SET EXPAND

7: KM EXP BUFFER

#BB15

Allocate a buffer for expansion strings.

Action:

Set the address and length of the expansion buffer. Initialize the buffer with the default expansion strings.

Entry conditions:

DE contains the address of the buffer.
HL contains the length of the buffer.

Exit conditions:

If the buffer is OK:

Carry true.

If the buffer is too short:

Carry false.

Always:

A, BC, DE, HL and other flags corrupt.

All other registers preserved.

Notes:

The buffer must not be located underneath a ROM and it must be at least 49 bytes long (i.e. have sufficient space for the default expansion strings). If the new buffer is too short then the old buffer is left unchanged.

The default expansion strings are given in Appendix IV.

Any expansion string currently being read is discarded.

This routine enables interrupts.

Related entries:

KM GET EXPAND

KM SET EXPAND

8: KM WAIT KEY

#BB18

Wait for next key from the keyboard.

Action:

Try to get a key from the key buffer. This routine waits until a key is found if no key is immediately available.

Entry conditions:

No conditions.

Exit conditions:

Carry true.

A contains the character or expansion token.

Other flags corrupt.

All other registers preserved.

Notes:

The next key is read from the key buffer and translated using the appropriate key translation table. Expansion tokens are not expanded but are passed out for the user to deal with, as are normal characters. Other Key Manager tokens (shift lock, caps lock and ignore) are obeyed but are not passed out.

Related entries:

KM READ KEY

KM WAIT CHAR

9: KM READ KEY

#BB1B

Test if a key is available from the keyboard.

Action

Try to get a key from the key buffer. This routine does not wait if no key is available immediately.

Entry Conditions:

No conditions.

Exit conditions

If a key was available:

Carry true.
A contains the character or expansion token.

If no key was available:

Carry false.
A corrupt.

Always:

Other flags corrupt.
All other registers preserved.

Notes:

The next key is read from the key buffer and translated using the appropriate key translation table. Expansion tokens are not expanded but are passed out for the user to deal with, as are normal characters. Other Key Manager tokens (shift lock, caps lock and ignore) are obeyed but are not passed out.

This routine will always return a key if one is available. It is therefore possible to flush out the key buffer by calling KM READ KEY repeatedly until it claims no key is available. Note, however, that the 'put back' character or a partially read expansion string is ignored. It is advisable to use KM READ CHAR to flush these out when emptying the Key Manager buffers, or, in V1.1 firmware, to call KM FLUSH.

Related Entries:

KM FLUSH

KM READ CHAR

KM WAIT KEY

10: KM TEST KEY

#BB1E

Test if a key is pressed.

Action:

Test if a particular key or joystick button is pressed. This is done using the key state map rather than by accessing the keyboard hardware.

Entry conditions:

A contains a key number.

Exit conditions:

If the key is pressed:

Zero false.

If the key is not pressed:

Zero true.

Always:

Carry false.

C contains the current shift and control state.

A, HL and other flags corrupt.

All other registers preserved.

Notes:

The shift and control states are automatically read when a key is scanned. If bit 7 is set then the control key is pressed and if bit 5 is set then one of the shift keys is pressed.

The key number is not checked. An invalid key number will generate the correct shift and control states but the state of the key tested will be meaningless.

The key state map which this routine tests is updated by the keyboard scanning routine. Normally this is run every fiftieth of a second and so the state may be out of date by that much. The key debouncing requires that a key should be released for two scans of the keyboard before it is marked as released in the key state map; the pressing of a key is detected immediately.

Related entries:

KM GET JOYSTICK

KM GET STATE

KM READ KEY

11: KM GET STATE

#BB21

Fetch Caps Lock and Shift Lock states.

Action:

Ask if the keyboard is currently shift locked or caps locked.

Entry conditions:

No conditions.

Exit conditions:

L contains the shift lock state.

H contains the caps lock state.

AF corrupt.

All other registers preserved.

Notes:

The lock states are:

#00 means the lock is off

#FF means the lock is on

The default lock states are off.

Related entries:

KM SET LOCKS

KM TEST KEY

12: KM GET JOYSTICK

#BB24

Fetch current state of the joystick(s).

Action:

Ask what the current states of the joysticks are. These are read from the key state map rather than by accessing the keyboard hardware.

Entry conditions:

No conditions.

Exit conditions:

H contains the state of joystick 0.

L contains the state of joystick 1.

A contains the state of joystick 0.

Flags corrupt.

All other registers preserved.

Notes:

In normal operation the key state map is updated by the key scanning routine every fiftieth of a second so the state returned may be slightly out of date.

The joystick states are bit significant as follows:

Bit 0	Up.
Bit 1	Down.
Bit 2	Left.
Bit 3	Right.
Bit 4	Fire 2.
Bit 5	Fire 1.
Bit 6	Spare joystick button (usually unconnected).
Bit 7	Always zero.

If a bit is set then the appropriate button is pressed.

Joystick 1 is indistinguishable from certain keys on the keyboard (see Appendix I).

Related entries:

KM TEST KEY

13: KM SET TRANSLATE

#BB27

Set entry in normal key translation table.

Action:

Set what character or token a key will be translated to when neither shift nor control is pressed.

Entry conditions:

A contains a key number.

B contains the new translation.

Exit conditions:

AF and HL corrupt.

All other registers preserved.

Notes:

If the key number is invalid (greater than 79) then no action is taken.

Most values in the table are treated as characters and are passed back to the user.

However, there are certain special values:

#80.. #9F	are the expansion tokens and are expanded to character strings when KM READ CHAR or KM WAIT CHAR is called although they are passed back like any other character when KM READ KEY or KM WAIT KEY is called.
#FD	is the caps lock token and causes the caps lock to toggle (turn on if off and vice versa).
#FE	is the shift lock token and causes the shift lock to toggle (turn on if off and vice versa).
#FF	is the ignore token and means the key should be thrown away.

Characters #E0.. #FC have special meanings to the BASIC to do with editing, cursoring and breaks.

See Appendix II for a full listing of the default translation tables.

Related entries:

KM GET TRANSLATE

KM SET CONTROL

KM SET SHIFF

14: KM GET TRANSLATE

#BB2A

Get entry from normal key translation table.

Action:

Ask what character or token a key will be translated to when neither shift nor control is pressed.

Entry conditions:

A contains a key number.

Exit conditions:

A contains the current translation.

HL and flags corrupt.

All other registers preserved.

Notes:

The key number is not checked. If it is invalid (greater than 79) then the translation returned is meaningless.

Most values in the table are treated as characters and are passed back to the user. However, there are certain special values:

#80..#9F	are the expansion tokens and are expanded to character strings when KM READ CHAR or KM WAIT CHAR is called although they are passed back like any other character when KM READ KEY or KM WAIT KEY is called.
#FD	is the caps lock token and causes the caps lock to toggle (turn on if off and vice versa).
#FE	is the shift lock token and causes the shift lock to toggle (turn on if off and vice versa).
#FF	is the ignore token and means the key should be thrown away.

Characters #E0..#FC have special meanings to the BASIC to do with editing, cursoring and breaks.

See Appendix II for a full listing of the default translation tables.

Related entries:

KM GET CONTROL

KM GET SHIFT

KM SET TRANSLATE.

15: KM SET SHIFT

#BB2D

Set entry in shifted key translation table.

Action:

Set what character or token a key will be translated to when control is not pressed but shift is pressed or shift lock is on.

Entry conditions:

A contains a key number.

B contains the new translation.

Exit conditions:

AF and HL corrupt.

All other registers preserved.

Notes:

If the key number is invalid (greater than 79) then no action is taken.

Most values in the table are treated as characters and are passed back to the user.

However, there are certain special values:

#80.. #9F	are the expansion tokens and are expanded to character strings when KM READ CHAR or KM WAIT CHAR is called although they are passed back like any other character when KM READ KEY or KM WAIT KEY is called.
#FD	is the caps lock token and causes the caps lock to toggle (turn on if off and vice versa).
#FE	is the shift lock token and causes the shift lock to toggle (turn on if off and vice versa).
#FF	is the ignore token and means the key should be thrown away.

Characters #E0.. #FC have special meanings to the BASIC to do with editing, cursoring and breaks.

See Appendix II for a full listing of the default translation tables.

Related entries:

KM GET SHIFT

KM SET CONTROL

KM SET TRANSLATE

16: KM GET SHIFT

#BB30

Get entry from shifted key translation table.

Action:

Ask what character or token a key will be translated to when control is not pressed but shift is pressed or shift lock is on.

Entry conditions:

A contains a key number.

Exit conditions:

A contains the current translation.

HL and flags corrupt.

All other registers preserved.

Notes:

The key number is not checked. If it is invalid (greater than 79) then the translation returned is meaningless.

Most values in the table are treated as characters and are passed back to the user.

However, there are certain special values:

#80.. #9F	are the expansion tokens and are expanded to character strings when KM READ CHAR or KM WAIT CHAR is called although they are passed back like any other character when KM READ KEY or KM WAIT KEY is called.
#FD	is the caps lock token and causes the caps lock to toggle (turn on if off and vice versa).
#FE	is the shift lock token and causes the shift lock to toggle (turn on if off and vice versa).
#FF	is the ignore token and means the key should be thrown away.

Characters #E0.. #FC have special meanings to the BASIC to do with editing, cursoring and breaks.

See Appendix II for a full listing of the default translation tables.

Related entries:

KM GET CONTROL
KM GET TRANSLATE
KM SET SHIFT

17: KM SET CONTROL

#BB33

Set entry in control key translation table.

Action:

Set what character or token a key will be translated to when control is pressed.

Entry conditions:

A contains a key number.

B contains the new translation.

Exit conditions:

AF and HL corrupt.

All other registers preserved.

Notes:

If the key number is invalid (greater than 79) then no action is taken.

Most values in the table are treated as characters and are passed back to the user. However, there are certain special values:

#80..#9F are the expansion tokens and are expanded to character strings when KM READ CHAR or KM WAIT CHAR is called although they are passed back like any other character when KM READ KEY or KM WAIT KEY is called.

#FD is the caps lock token and causes the caps lock to toggle (turn on if off and vice versa).

#FE is the shift lock token and causes the shift lock to toggle (turn on if off and vice versa).

#FF is the ignore token and means the key should be thrown away.

Characters #E0..#FC have special meanings to the BASIC to do with editing, cursoring and breaks.

See Appendix II for a full listing of the default translation tables.

Related entries:

KM GET CONTROL

KM SET SHIFT

KM SET TRANSLATE

18: KM GET CONTROL

#BB36

Get entry from control key translation table.

Action:

Ask what character or token a key will be translated to when control is pressed.

Entry conditions:

A contains a key number.

Exit conditions:

A contains the current translation.

HL and flags corrupt.

All other registers preserved.

Notes:

The key number is not checked. If it is invalid (greater than 79) then the translation returned is meaningless.

Most values in the table are treated as characters and are passed back to the user. However, there are certain special values:

#80.. #9F are the expansion tokens and are expanded to character strings when KM READ CHAR or KM WAIT CHAR is called although they are passed back like any other character when KM READ KEY or KM WAIT KEY is called.

#FD is the caps lock token and causes the caps lock to toggle (turn on if off and vice versa).

#FE is the shift lock token and causes the shift lock to toggle (turn on if off and vice versa).

#FF is the ignore token and means the key should be thrown away.

Characters #E0.. #FC have special meanings to the BASIC to do with editing, cursoring and breaks.

See Appendix II for a full listing of the default translation tables.

Related entries:

KM GET SHIFT

KM GET TRANSLATE

KM SET CONTROL

19: KM SET REPEAT

#BB39

Set whether a key may repeat.

Action:

Set the entry in the repeating key map that determines whether a key is allowed to repeat or not.

Entry conditions:

If the key is to be allowed to repeat:

B contains #FF.

If the key is not to be allowed to repeat:

B contains #00.

Always:

A contains the key number.

Exit conditions:

AF, BC and HL corrupt.
All other registers preserved.

Notes:

If the key number is invalid (greater than 79) then no action is taken.

The default repeating keys are listed in Appendix III.

Related entries:

KM GET REPEAT

KM SET DELAY

20: KM GET REPEAT

#BB3C

Ask if a key is allowed to repeat.

Action:

Test the entry in the repeating key map that says whether a key is allowed to repeat or not.

Entry conditions:

A contains a key number.

Exit conditions:

If the key is allowed to repeat:

Zero false.

If the key is not allowed to repeat:

Zero true.

Always:

Carry false.

A, HL and other flags corrupt.

All other registers preserved.

Notes:

The key number is not checked. If it is invalid (greater than 79) then the repeat state returned is meaningless.

The default repeating keys are listed in Appendix III.

Related entries:

KM SET REPEAT

21: KM SET DELAY

#BB3F

Set start up delay and repeat speed.

Action:

Set the time before keys first repeat (start up delay) and the time between repeats (repeat speed).

Entry conditions:

H contains the new start up delay.
L contains the new repeat speed.

Exit conditions:

AF corrupt.
All other registers preserved.

Notes:

Both delays are given in scans of the keyboard. The keyboard is scanned every fiftieth of a second.

A start up delay or repeat speed of 0 is taken to mean 256.

The default start up delay is 30 scans (0.6 seconds) and the default repeat speed is 2 scans (0.04 seconds or 25 characters a second).

Note that a key is prevented from repeating (by the key scanner) if the key buffer is not empty. Thus the actual repeat speed is the slower of the supplied repeat speed and the rate at which characters are removed from the buffer. This is intended to prevent the user from getting too far ahead of a program that is running sluggishly.

The start up delay and repeat speed apply to all keys on the keyboard that are set to repeat.

Related entries:

KM GET DELAY
KM SET REPEAT

22: KM GET DELAY

#BB42

Get start up delay and repeat speed.

Action:

Ask the time before keys first repeat (start up delay) and the time between repeats (repeat speed).

Entry conditions:

No conditions.

Exit conditions:

H contains the start up delay.
L contains the repeat speed.

AF corrupt.
All other registers preserved.

Notes:

Both delays are given in scans of the keyboard. The keyboard is scanned every fiftieth of a second.

A repeat speed or start up delay of 0 means 256.

Related entries:

KM SET DELAY

23: KM ARM BREAKS

#BB45

Allow break events to be generated.

Action:

Arm the break mechanism. The next call of KM BREAK EVENT will generate a break event.

Entry conditions:

DE contains the address of the break event routine.
C contains the ROM select address for this routine.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The break mechanism can be disarmed by calling KM DISARM BREAK (or KM RESET).

This routine enables interrupts.

Related entries:

KM BREAK EVENT
KM DISARM BREAK

24: KM DISARM BREAK

#BB48

Prevent break events from being generated.

Action:

Disarm the break mechanism. From now on the generation of break events by KM BREAK EVENT will be suppressed.

Entry conditions:

No conditions.

Exit conditions:

AF and HL corrupt.
All other registers preserved.

Notes:

Break events can be rearmed by calling KM ARM BREAK.

The default state of the break mechanism is disarmed, thus calling KM RESET will also disarm breaks.

This routine enables interrupts.

Related entries:

KM ARM BREAK
KM BREAK EVENT

25: KM BREAK EVENT

#BB4B

Generate a break event (if armed).

Action:

Try to generate a break event.

Entry conditions:

No conditions.

Exit conditions:

AF and HL corrupt.
All other registers preserved.

Notes:

If the break mechanism is disarmed then no action is taken. Otherwise a break event is generated and a special marker is placed into the key buffer. This marker generates a break event token (#EF) when read from the buffer. The break mechanism is automatically disarmed after generating a break event so that multiple breaks can be avoided.

This routine may be run from the interrupt path and thus does not and should not enable interrupts. Note, however, that using a LOW JUMP to call the routine (as the firmware jumpblock is set up to do) does enable interrupts and so the jumpblock may not be used directly from interrupt routines.

Related entries:

KM ARM BREAK

KM DISARM BREAK

26: TXT INITIALISE

#BB4E

Initialize the Text VDU.

Action:

Full initialization of the Text VDU (as used during EMS). All Text VDU variables and indirections are initialized, the previous VDU state is lost.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The Text VDU indirections (TXT DRAW CURSOR, TXT UNDRAW CURSOR, TXT WRITE CHAR, TXT UNWRITE and TXT OUT ACTION) are set to their default routines.

The control code table is set up to perform the default control code actions.

The user defined character table is set to be empty.

Stream 0 is selected.

All streams are set to their default states:

- The text paper (background) is set to ink 0.
- The text pen (foreground) is set to ink 1.
- The text window is set to the entire screen.
- The text cursor is enabled but turned off.
- The character writing mode is set to opaque.
- The VDU is enabled.
- The graphic character write mode is turned off.
- The cursor is moved to the top left corner of the window.

The default character set and the default setting for the control code table are described in ~ppendices VI and VII.

Related Entries:

SCR INITIALISE
TXT RESET

27: TXT RESET

#BB51

Reset the Text VDU.

Action:

Reinitializes the Text VDU indirections and the control code table. Does not affect any other aspect of the Text VDU.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The Text VDU indirections TXT DRAW CURSOR, TXT UNDRAW CURSOR, TXT WRITE CHAR, TXT UNWRITE and TXT OUT ACTION are set to their default routines.

The control code table is set up to perform the default control code actions (see Appendix VII).

Related Entries:

TXT INITIALISE

28: TXT VDU ENABLE

#BB54

Allow characters to be placed on the screen.

Action:

Permit characters to be printed when requested (by calling TXT OUTPUT or TXT WR CHAR). Enabling applies to the currently selected stream. The cursor blob is also enabled (by calling TXT CUR ENABLE).

Entry conditions:

No conditions.

Exit conditions:

AF corrupt.
All other registers preserved.

Notes:

The control code buffer used by TXT OUTPUT is emptied, any incomplete control code sequence will be lost.

Related Entries:

TXT ASK STATE
TXT CUR ENABLE
TXT OUTPUT
TXT VDU DISABLE
TXT WR CHAR

29: TXT VDU DISABLE

#BB57

Prevent characters being placed on the screen.

Action:

Prevents characters being printed on the screen (when TXT OUTPUT or TXT WR CHAR is called). Applies to the currently selected stream. The cursor blob is also disabled (by calling TXT CUR DISABLE).

Entry conditions:

No conditions.

Exit conditions:

AF corrupt.
All other registers preserved.

Notes:

The control code buffer used by TXT OUTPUT is emptied, any incomplete control sequence will be lost.

In V1.0 firmware control codes are still obeyed by TXT OUTPUT. In V1.1 firmware only those control codes which are marked in the control code table will be obeyed; other control codes will be ignored (see section 4.7).

Related Entries:

TXT ASK STATE
TXT CUR DISABLE
TXT OUTPUT
TXTVDU ENABLE
TXT WR CHAR

30: TXT OUTPUT

#BB5A

Output a character or control code to the Text VDU.

Action:

Output characters to the screen and obey control codes (characters #00.. #1F). Works on the currently selected stream.

Entry conditions:

A contains the character to send.

Exit conditions:

All registers and flags preserved.

Notes:

This routine calls the TXT OUT ACTION indirection to do the work of printing the character or obeying the control code described below.

Control codes may take up to 9 parameters. These are the characters sent following the initial control code. The characters sent are stored in the control code buffer until sufficient have been received to make up all the required parameters. The control code buffer is only long enough to accept 9 parameter characters.

There is only one control code buffer for all streams. It is therefore possible to get unpredictable results if the output stream is changed midway through sending a control code sequence.

If the VDU is disabled then no characters will be printed on the screen. In Vi .0 firmware all control codes will still be obeyed but in Vi .i firmware only those codes marked in the control code table as to be obeyed when the VDU is disabled will be obeyed (see section 4.7).

If the graphic character write mode is enabled then all characters and control codes are printed using the Graphics VDU routine, GRA WR CHAR, and are not obeyed. Characters are written in the same way that TXT WR CHAR writes characters.

Related Entries:

GRA WR CHAR

TXT OUT ACTION

TXT SET GRAPHIC

TXT VDU DISABLE

TXT VDU ENABLE

TXT WR CHAR

31: TXT WR CHAR

#BB5D

Write a character to the screen.

Action:

Print a character on the screen at the cursor position of the currently selected stream. Control codes (characters #00...#1F) are printed and not obeyed.

Entry conditions:

A contains the character to print.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

If the VDU is disabled then no character will be printed.

Before printing the character the cursor position is forced to lie within the text window (see TXT VALIDATE). After printing the character the cursor is moved right one character.

To put the character on the screen this routine calls the TXT WRITE CHAR indirection.

Related Entries:

GRA WR CHAR

TXT OUTPUT

TXT RD CHAR

TXT WRITE CHAR

32: TXT RD CHAR

#BB60

Read a character from the screen

Action:

Read a character from the screen at the cursor position of the currently selected stream

Entry Conditions:

No conditions.

Exit Conditions:

If a recognisable character was found:

Carry true

A contains the character read.

If no recognisable character found:

Carry false

A contains zero

Always:

Other flags corrupt

All other registers preserved.

Notes:

In V1.1 firmware the cursor position is forced legal (inside the window) before the character is read. This may cause the screen to roll. The same is not true for V1.0 firmware where the cursor position is not forced legal and steps must be taken to avoid reading characters from outside the window.

The read is performed by comparing the matrix found on the screen with the matrices used to generate characters. As a result changing a character matrix, changing the pen or paper inks, or changing the screen (e.g. drawing a line through a character) may make the character unreadable.

To actually read the character from the screen the TXT UNWRITE indirection is called.

Special precautions are taken against generating inverse space (character #8F). Initially the character is read assuming that the background to the character was written in the current paper ink and treating any other ink as foreground. If this fails to generate a recognisable character or it generates inverse space then another try is made by assuming that the foreground to the character was written in the current pen ink and treating any other ink as background.

The characters are scanned starting with #00 and finishing with #FF

Related Entries:

TXT UNWRITE

TXT WR CHAR

33: TXT SET GRAPHIC

#BB63

Turn on or off the Graphics VDU write character option.

Action:

Enable or disable graphic character writing on the currently selected stream.

Entry conditions:

If graphic writing is to be turned on:

A must be non-zero.

If graphic writing is to be turned off:

A must contain zero.

Exit conditions:

AF corrupt.

All other registers preserved.

Notes:

When graphic character writing is enabled then all characters sent to TXT OUTPUT are printed using the Graphics VDU (see GRA WR CHAR) rather than the Text VDU (see TXT WR CHAR). Also all control codes are printed rather than obeyed. Characters sent to TXT WR CHAR will be printed as normal.

Character printing is not prevented by disabling the Text VDU (with TXT VDU DISABLE) if graphic character writing is enabled.

Related Entries:

GRA WR CHAR

TXT OUTPUT

34: TXT WIN ENABLE

#BB66

Set the size of the current text window.

Action:

Set the boundaries of the window on the currently selected stream. The edges are the first and last character columns inside the window and the first and last character rows inside the window.

Entry conditions:

H contains the physical column of one edge.
D contains the physical column of the other edge.
L contains the physical row of one edge.
E contains the physical row of the other edge.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The edge positions are given in physical screen coordinates. i.e. Row 0, column 0 is the top left corner of the screen and the coordinates are signed numbers.

The window is truncated, if necessary, so that it fits on the screen.

The left column of the window is taken to be the smaller of H and D. The top row of the window is taken to be the smaller of L and E.

The cursor is moved to the top left corner of the window.

The window is not cleared.

If the window covers the whole screen then when the window is rolled the hardware roll routine (see 5CR HW ROLL) will be used. If the window covers less than the whole screen the software roll routine (see SCR SW ROLL) will be used.

The default text window covers the whole screen and is set up when TXT INITIALISE or SCR SET MODE is called.

Related Entries:

TXT GET WINDOW
TXT VALIDATE

35: TXT GET WINDOW

#BB69

Get the size of the current window.

Action:

Get the boundaries of the window on the currently selected stream and whether it covers the whole screen.

Entry conditions:

No conditions.

Exit conditions:

If the window covers the whole screen:
Carry false.

If the window covers less than the whole screen:
Carry true.

Always:

H contains the leftmost column in the window.
D contains the rightmost column in the window.
L contains the topmost row in the window.
E contains the bottommost row in the window.

A corrupt.
All other registers preserved.

Notes:

The boundaries of the window are given in physical coordinates. i.e. Row 0, column 0 is the top left corner of the screen.

The boundaries returned by this routine may not be the same as those set when TXT WIN ENABLE was called because the window is truncated to fit the screen.

Related Entries:

TXT VALIDATE
TXT WIN ENABLE

36: TXT CLEAR WINDOW

#BB6C

Clear current window.

Action:

Clear the text window of the currently selected stream to the paper ink of the currently selected stream.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The cursor is moved to the top left corner of the window.

Related Entries:

GRA CLEAR WINDOW
SCR CLEAR
TXT SET PAPER
TXT WIN ENABLE

37: TXT SET COLUMN

#BB6F

Set cursor horizontal position.

Action:

Move the current position of the currently selected stream to a new column. The cursor blob will be removed from the current position and redrawn at the new position (if the cursor is enabled and turned on).

Entry conditions:

A contains the required logical column for the cursor.

Exit conditions:

AF and HL corrupt.
All other registers preserved.

Notes:

The required column is given in logical coordinates. i.e. Column 1 is the leftmost column of the window.

The cursor may be moved outside the window. However, it will be forced to lie inside the window before any character is written by the Text VDU (see TXT VALIDATE) or the cursor blob is drawn.

Related Entries:

TXT GET CURSOR

TXT SET CURSOR

TXT SET ROW

38: TXT SET ROW

#BB72

Set cursor vertical position.

Action:

Move the current position of the currently selected stream to a new row. The cursor blob will be removed from the current position and redrawn at the new position (if the cursor is enabled and turned on).

Entry conditions:

A contains the required logical row for the cursor.

Exit conditions:

AF and HL corrupt.
All other registers preserved.

Notes:

The required row is given in logical coordinates. i.e. Row 1 is the topmost row of the window.

The cursor may be moved outside the window. However, it will be forced to lie inside the window before any character is written by the Text VDU (see TXT VALIDATE) or the cursor blob is drawn.

Related Entries:

TXT GET CURSOR
TXT SET COLUMN
TXT SET CURSOR

39: TXT SET CURSOR

#BB75

Set cursor position.

Action:

Move the current position of the currently selected stream to a new row and column. The cursor blob will be removed from the current position and redrawn at the new position (if the cursor is enabled and turned on).

Entry conditions:

H contains the required logical column.

L contains the required logical row.

Exit conditions:

AF and HL corrupt.

All other registers preserved.

Notes:

The required position is given in logical coordinates. i.e. Row 1, column 1 is the top left corner of the window.

The cursor position may be moved outside the window. However, it will be forced to lie inside the window before any character is written by the Text VDU (see TXT VALIDATE) or the cursor blob is drawn.

Related Entries:

TXT GET CURSOR

TXT SET COLUMN

TXT SET ROW

40: TXT GET CURSOR

#BB78

Ask current cursor position.

Action:

Get the current location of the cursor and a count of the number of times the window of the currently selected stream has rolled.

Entry conditions:

No conditions.

Exit conditions:

H contains the logical cursor column.

L contains the logical cursor row.

A contains the current roll count.

Flags corrupt.

All other registers are preserved.

Notes:

The cursor position is given in logical coordinates. i.e. Row 1, column 1 is the top left corner of the window.

The roll count passed out has no absolute meaning. It is decremented when the window is rolled up and is incremented when the window is rolled down. It may be used to determine whether the window has rolled by comparing it with a previous value.

The position reported may not be inside the window and is, therefore, not necessarily the position at which the next character will be printed. Use TXT VALIDATE to check this.

Related Entries:

TXT SET COLUMN

TXT SET CURSOR

TXT SET ROW

TXT VALIDATE

41: TXT CUR ENABLE

#BB7B

Allow cursor display - user.

Action:

Allow the cursor blob for the currently selected stream to be placed on the screen. The cursor blob will be placed on the screen immediately unless the cursor is turned off (see TXT CUR OFF).

Entry conditions:

No conditions.

Exit conditions:

AF corrupt.
All other registers preserved.

Notes:

Cursor enabling and disabling is intended for use by the user. It is also used when the VDU is disabled (see TXT VDU ENABLE and TXT VDU DISABLE).

Related Entries:

TXT ASK STATE
TXT CUR DISABLE
TXT CUR ON
TXT DRAW CURSOR
TXT UNDRAW CURSOR

42: TXT CUR DISABLE

#BB7E

Disallow cursor display - user.

Action:

Prevent the cursor blob for the currently selected stream from being placed on the screen. The cursor blob will be removed from the screen immediately if it is currently there.

Entry conditions:

No conditions.

Exit conditions:

AF corrupt.
All other registers preserved.

Notes:

Cursor enabling and disabling is intended for use by the user. It is also used when the VDU is disabled (see TXT VDU ENABLE and TXT VDU DISABLE).

Related Entries:

TXT ASK STATE
TXT CUR ENABLE
TXT CUR OFF
TXT DRAW CURSOR
TXT UNDRAW CURSOR

43: TXT CUR ON

#BB81

Allow cursor display - system.

Action:

Allow the cursor blob for the currently selected stream to be placed on the screen. The cursor blob will be placed on the screen immediately unless the cursor is disabled (see TXT CUR DISABLE).

Entry conditions:

No conditions.

Exit conditions:

All registers and flags preserved.

Notes:

Turning the cursor on and off is intended for use by system ROMs.

Related Entries:

TXT ASK STATE

TXT CUR ENABLE

TXT CUR OFF

TXT DRAW CURSOR

TXT UNDRAW CURSOR

44: TXT CUR OFF

#BB84

Disallow cursor display - system.

Action:

Prevent the cursor blob for the currently selected stream from being placed on the screen. The cursor blob will be removed from the screen immediately if it is currently there.

Entry conditions:

No conditions.

Exit conditions:

All registers and flags preserved.

Notes:

Turning the cursor on and off is intended for use by system ROMs.

Related Entries:

TXT ASK STATE

TXT CUR DISABLE

TXT CUR ON

TXT DRAW CURSOR

TXT UNDRAW CURSOR

45: TXT VALIDATE

#BB87

Check if a cursor position is within the window.

Action:

Check a screen position to see if it lies within the current window. If it does not then determine the position where a character would be printed after applying the rules for forcing the screen position inside the window.

Entry conditions:

H contains the logical column of the position to check.

L contains the logical row of the position to check.

Exit conditions:

If printing at the position would not cause the window to roll:

Carry true.

B corrupt.

If printing at the position would cause the window to roll up:

Carry false.

B contains #FF.

If printing at the position would cause the window to roll down:

Carry false.

B contains #00.

Always:

H contains the logical column at which a character would be printed.

L contains the logical row at which a character would be printed.

A and other flags corrupt.

All other registers preserved.

Notes:

The positions on the screen are given in logical coordinates. i.e. Row 1, column 1 is the top left corner of the window.

Before writing a character or putting the cursor blob on the screen the Text VDU validates the current position, performs any required roll then writes at the appropriate position.

The algorithm to work out the position to print at, from the position to check, is as follows:

- 1/ If the position is right of the right edge of the window it is moved to the left edge of the window on the next line.
- 2/ If the position is left of the left edge of the window it is moved to the right edge of the window on the previous line.
- 3/ If the position is now above the top edge of the window then it is moved to the top edge of the window and the window needs rolling downwards.
- 4/ If the position is now below the bottom edge of the window it is moved to the bottom edge of the window and the window needs rolling upwards.

Related Entries:

SCR HW ROLL

SCR SW ROLL

TXT GET CURSOR

46: TXT PLACE CURSOR

#BB8A

Put a cursor blob on the screen.

Action:

Put a cursor blob on the screen at the cursor position for the currently selected stream.

Entry conditions:

No conditions.

Exit conditions:

AFcorrupt.
All other registers preserved.

Notes:

TXT PLACE CURSOR is provided to allow the user to run multiple cursors in a window. The indirection TXT DRAW CURSOR should be called for merely placing the normal cursor blob on the screen. Higher level routines, such as TXT OUTPUT and TXT SET CURSOR, automatically remove and place the normal cursor when appropriate, the user must deal with any other cursors.

It is not safe to call TXT PLACE CURSOR twice at a particular screen position without calling TXT REMOVE CURSOR in between because this may leave a spurious cursor blob on the screen when the cursor position is moved.

The cursor position is forced to be inside the window before the cursor blob is drawn.

The cursor blob is an inverse patch formed by exclusive-oring the contents of the screen at the cursor position with the exclusive-or of the current pen and paper inks.

Related Entries:

TXT DRAW CURSOR

TXT REMOVE CURSOR

47: TXT REMOVE CURSOR

#BB8D

Take a cursor blob off the screen.

Action:

Take a cursor blob off the screen at the cursor position of the currently selected stream.

Entry conditions:

No conditions.

Exit conditions:

AF corrupt.
All other registers preserved.

Notes:

TXT REMOVE CURSOR is provided to allow the user to run multiple cursors in a window. The indirection TXT UNDRAW CURSOR should be called for merely removing the normal cursor from the screen. Higher level routines, such as TXT OUTPUT and TXT SET CURSOR, automatically remove and place the normal cursor when appropriate, the user must deal with any other cursors.

TXT REMOVE CURSOR should only be used to remove a cursor placed on the screen by calling TXT PLACE CURSOR. The cursor should be removed when the cursor position is to be changed (rolling the window implicitly changes the cursor position) or the screen is to be read or written. Incorrect use of this routine may result in a spurious cursor blob being generated.

The cursor position is forced to be inside the window before the cursor blob is removed (this should not matter as TXT PLACE CURSOR has already done this).

The cursor blob is an inverse patch formed by exclusive-oring the contents of the screen at the cursor position with the exclusive-or of the current pen and paper inks.

Related Entries:

TXT PLACE CURSOR

TXT UNDRAW CURSOR

48: TXT SET PEN

#BB90

Set ink for writing characters.

Action:

Set the text pen ink for the currently selected stream. This is the ink that is used for writing characters (the foreground ink).

Entry conditions:

A contains ink to use.

Exit conditions:

AF and HL corrupt.
All other registers preserved.

Notes:

The ink is masked to bring it within the range of legal inks for the current screen mode. That is with #0F in mode 0, #03 in mode land #01 in mode 2.

The cursor blob will be redrawn using the new ink (if enabled).

Related Entries:

GRA SET PEN

SCR SET INK

TXT GET PEN

TXT SET PAPER

49: TXT GET PEN

#BB93

Get ink for writing characters.

Action:

Ask what the pen ink is set to for the currently selected stream. This is the ink used for writing characters (foreground ink).

Entry conditions:

No conditions.

Exit conditions:

A contains the ink.
Flags corrupt.
All other registers preserved.

Notes:

This routine has no other effects.

Related Entries:

GRA GET PEN
SCR GET INK
TXT GET PAPER
TXT SET PEN

50: TXT SET PAPER

#BB96

Set ink for writing text background.

Action:

Set the text paper ink for the currently selected stream. This is the ink used for writing the background to characters and for clearing the text window.

Entry conditions:

A contains the ink to use.

Exit conditions:

AF and HL corrupt.
All other registers preserved.

Notes:

The ink is masked to bring it within the range of legal inks for the current screen mode. That is with #0F in mode 0, #03 in mode land #01 in mode 2.

The cursor blob will be redrawn using the new ink (if enabled).

This ink will be used when clearing areas of the text window (by TXT CLEAR WINDOW and certain control codes).

This routine does not clear the text window.

Related Entries:

GRA SET PAPER

SCR SET INK

TXT GET PAPER

TXT SET PEN

51: TXT GET PAPER

#BB99

Get ink for writing background.

Action:

Ask what the paper ink is set to for the currently selected stream. This is the ink used for writing the background to characters and for clearing the text window.

Entry conditions:

No conditions.

Exit conditions:

A contains the ink.
Flags corrupt.
All other registers preserved.

Notes:

This routine has no other effects.

Related Entries:

GRA GET PAPER
SCR GET INK
TXT GET PEN
TXT SET PAPER

52: TXT INVERSE

#BB9C

Swap current pen and paper inks over.

Action:

Exchange the text pen and paper (foreground and background) inks for the currently selected stream.

Entry conditions:

No conditions.

Exit conditions:

AF and HL corrupt.
All other registers preserved.

Notes:

In V1.1 firmware the cursor blob is removed and replaced and so the current position is forced legal (inside the window) which may cause the screen to roll. In V1.0 firmware the cursor blob is not redrawn and so it should not be on the screen when this routine is called.

Related Entries:

TXT SET PAPER
TXT SET PEN

53: TXT SET BACK

#BB9F

Allow or disallow background being written.

Action:

Set character write mode to opaque or transparent for the currently selected stream. Opaque mode writes background with the character. Transparent mode writes the character on top of the current contents of the screen.

Entry conditions:

If background is to be written (opaque mode):

A must be zero.

If background is not to be written (transparent mode):

A must be non-zero.

Exit conditions:

AF and HL corrupt.
All other registers preserved.

Notes:

Writing in transparent mode is intended for annotating diagrams and similar applications. It can have unfortunate effects if it is used generally because overwriting a character will not remove the character underneath thus creating an incomprehensible jumble on the screen.

Setting the character write mode does not affect the Graphics VDU. In V1.1 firmware the routine GRA SET BACK sets the equivalent graphics background write mode.

Related Entries:

GRA SET BACK
TXT GET BACK
TXT WR CHAR
TXT WRITE CHAR

54: TXT GET BACK

#BBA2

Ask if background is being written.

Action:

Get the character write mode for the currently selected stream.

Entry conditions:

No conditions.

Exit conditions:

If background is to be written (opaque mode):

A contains zero.

If background is not to be written (transparent mode):

A contains non-zero.

Always:

DE, HL and flags corrupt.
All other registers preserved.

Notes:

This only applies to the Text VDU, the Graphics VDU always writes opaque.

Related Entries:

TXT SET BACK

55: TXT GET MATRIX

#BBA5

Get the address of a character matrix.

Action:

Calculate a pointer to the matrix for a character and determine if it is a user defined matrix.

Entry conditions:

A contains the character whose matrix is to be found.

Exit conditions:

If the matrix is in the user defined matrix table:

Carry true.

If the matrix is in the lower ROM:

Carry false.

Always:

HL contains the address of the matrix.

A and other flags corrupt.

All other registers preserved.

Notes:

The matrix may be in RAM or in ROM. The Text VDU assumes that the appropriate ROMs are enabled or disabled when it calls this routine to get the matrix for a character. (The lower ROM is on, the upper ROM is normally off).

The matrix is stored as an 8 byte bit significant vector. The first byte describes the top line of the character and the last byte the bottom line. Bit 7 of a byte refers to the leftmost pixel of a line and bit 0 to the rightmost pixel. If a bit is set in the matrix then the pixel should be written in the pen ink. If a bit is not set then the pixel should either be written in the paper ink or left alone (depending on the character write mode).

Related Entries:

TXT SET MATRIX

56: TXT SET MATRIX

#BBA8

Set a character matrix.

Action:

Set the matrix for a user defined character. If the character is not user defined then no action is taken.

Entry conditions:

A contains the character whose matrix is to be set.
HL contains the address of the matrix to set.

Exit conditions:

If the character is user definable:
Carry true.

If the character is not user definable:
Carry false.

Always:

A, BC, DE, HL and other flags corrupt.
All other registers preserved.

Notes:

The matrix is stored as an 8 byte bit significant vector. The first byte describes the top line of the character and the last byte the bottom line. Bit 7 of a byte refers to the leftmost pixel of a line and bit 0 to the rightmost pixel. If a bit is set in the matrix then the pixel should be written in the pen ink. If a bit is not set then the pixel should either be written in the paper ink or left alone (depending whether the character write mode is opaque or transparent currently).

The matrix is copied from the area given into the character matrix table without using RAM LAMs thus the matrices can be set from ROM providing it is enabled. (Note however that the jumpblock disables the upper ROM.)

Altering a character matrix changes the matrix for all streams. It does not alter any character on the screen; it changes what will be placed on the screen the next time the character is written.

Related Entries:

TXT GET MATRIX
TXT SET M TABLE

57: TXT SET M TABLE

#BBAB

Set the user defined matrix table address.

Action:

Set the user defined matrix table and the number of characters in the table. The table is initialized with the current matrix settings.

Entry conditions:

DE contains the first character in the table.

HL contains the address of the start of the new table.

Exit conditions:

If there was no user defined matrix table before:

Carry false.

A and HL corrupt.

If there was a user defined matrix table before:

Carry true.

A contains the first character in the old table.

HL contains the address of the old table.

Always:

BC, DE and other flags corrupt.

All other registers preserved.

Notes:

If the first character specified is in the range 0..255 then the matrices for all characters between that character and character 255 are to be stored in the user defined table.

If the first character specified is not in the range 0..255 then the user defined matrix table is deemed to contain no matrices (and the table address passed is ignored).

The table must be $(256 - \text{first char}) * 8$ bytes long. The matrices are stored in the table in ascending order. The table is initialized with the current matrix settings, whether they were previously in RAM or in the ROM.

The table should not be located in RAM underneath a ROM.

It is permissible for the new and old matrix tables to overlap (thus allowing the table to be extended or contracted) providing that matrices in the new table

occupy an address earlier or equal to the address that they occupied in the old table.

All streams share the matrix table so any changes to it will be reflected on all streams.

Related Entries:

TXT GET M TABLE

TXT SET MATRIX

58: TXT GET M TABLE

#BBAE

Get user defined matrix table address.

Action:

Get the address of the current user defined matrix table and the first character in the table.

Entry conditions:

No conditions.

Exit conditions:

If there is no user defined matrix table:

- Carry false
- A and HL corrupt

If there is a user defined matrix table:

- Carry true
- A contains the first character in the table
- HL contains the address of the start of the table.

Always:

- Other flags corrupt
- All other registers preserved.

Notes:

The matrices for characters between the first character and 255 are stored in the table in ascending order. Each matrix is 8 bytes long.

Related Entries:

TXT GET MATRIX
TXT SET M TABLE

59: TXT GET CONTROLS

#BBB1

Fetch address of control code table.

Action:

Get the address of the control code table.

Entry conditions:

No conditions.

Exit conditions:

HL contains the address of the control code table.

All other registers and flags preserved.

Notes:

All streams share one control code table so that any changes made to the table will affect all streams.

The control code table has a 3 byte entry for each control code. The entries are stored in ascending order, so the entry for #00 is first and that for # 1F is last. The first byte of each entry is the number of parameters the control code requires, the other two bytes are the address of the routine to call to process the control code when all its parameters have been received. The routine must be located in the central 32K of RAM and it must obey the following interface:

Entry:

A contains the last character added to the buffer.

B contains the length of the buffer (including the control code).

C contains the same as A.

HL contains the address of the control code buffer (points at the control code).

Exit:

AF, BC, DE, HL corrupt.

All other registers preserved.

As the control code buffer only has space to store 9 parameter characters the number of parameters required should be limited to 9 or fewer.

The control code table is reinitialized to its default routines when TXT RESET is called.

In V1.1 firmware the first byte of each entry also specifies whether the control codes is to be disabled when the VDU is disabled or whether it is always to be obeyed. Bit 7 of the byte is set if the code is to be disabled.

Related Entries:

TXT OUTPUT

60: TXT STR SELECT

#BBB4

Select a Text VDU stream.

Action:

Make a given stream the currently selected stream (if it isn't already).

Entry conditions:

A contains the required stream.

Exit conditions:

A contains the previously selected stream.
HL and flags corrupt.
All other registers preserved.

Notes:

The requested stream number is masked (with #07) to make it into a legal stream number.

Many attributes of the Text VDU may be set independently on different streams. It is important to ensure that the correct stream is selected when any of these are altered. These attributes are:

- Pen ink.
- Paper ink.
- Cursor position.
- Window limits.
- Cursor enable/disable.
- Cursor on/off.
- VDU enable/disable.
- Character write mode.
- Graphic character write mode.

If the stream is already selected then this routine returns quickly. It is not unreasonable to repeatedly select a stream (before each character sent, for example).

Related Entries:

TXT OUTPUT

61: TXT SWAP STREAMS

#BBB7

Swap the states of two streams.

Action:

The stream descriptors for two streams are exchanged. The currently selected stream number remains the same (although its descriptor may have been altered).

Entry conditions:

B contains a stream number.
C contains another stream number.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The stream numbers passed are masked (with #07) to ensure that they are legal stream numbers.

The attributes that are exchanged are:

- Pen ink.
- Paper ink.
- Cursor position.
- Window limits.
- Window roll count.
- Cursor enable/disable.
- Cursor on/off.
- VDU enable/disable.
- Character write mode.
- Graphic character write mode.

Related Entries:

TXT STR SELECT

62: GRA INITIALISE

#BBBA

Initialize the Graphics VDU.

Action:

The Graphics VDU is fully initialized (as during EMS). All Graphic VDU variables and indirections are set to their default values.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The full operation is:

- Set the Graphics VDU indirections (GRA PLOT, GRA TEST and GRA LINE) to their default routines.

- Set the graphic paper to ink 0.

- Set the graphic pen to ink 1.

- Set the user origin to the bottom left corner of the screen.

- Move the current position to the user origin.

- Set the graphics window to cover the whole screen.

- The graphics background write mode is set to opaque.

- The line mask is set to #FF and the first pixel of lines are plotted.

- The graphics window is not cleared.

Related entries:

GRA DEFAULT

GRA RESET

SCR INITIALISE

63: GRA RESET

#BBBD

Reset the Graphics VDU.

Action:

Re-initialize the Graphics VDU indirections to their default routines and set default modes.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

Sets the Graphics VDU indirections (GRA PLOT, GRA TEST and GRA LINE) to their default routines. V1.1 firmware also sets the graphics background mode to opaque, sets the line mask to #FF and sets the first pixel of lines to be plotted.

Related entries:

GRA DEFAULT
GRA INITIALISE

64: GRA MOVE ABSOLUTE

#BBC0

Move to an absolute position.

Action:

Move the current position to an absolute position.

Entry conditions:

DE contains the required user X coordinate.
HL contains the required user Y coordinate.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The new position is given in user coordinates. i.e. Relative to the user origin.

The new position can be outside the graphics window.

The Graphic VDU plotting, testing and line drawing routines all move the current graphics position to the point (or endpoint) specified automatically.

Related entries:

GRA ASK CURSOR
GRA MOVE RELATIVE

65: GRA MOVE RELATIVE

#BBC3

Move relative to current position.

Action:

Move the current position to relative to its current position.

Entry conditions:

DE contains a signed X offset.

HL contains a signed Y offset.

Exit conditions:

AF, BC, DE and HL corrupt.

All other registers preserved.

Notes:

The new position can be outside the graphics window.

The Graphic VDU plotting, testing and line drawing routines all move the current graphics position to the point (*or* endpoint) specified automatically.

Related entries:

GRA ASK CURSOR

GRA MOVE ABSOLUTE

66: GRA ASK CURSOR

#BBC6

Get the current position.

Action:

Ask where the current graphics position is.

Entry conditions:

No conditions.

Exit conditions:

DE contains the user X coordinate.

HL contains the user Y coordinate.

AF corrupt.

All other registers preserved.

Notes:

The current position is given in user coordinates. i.e. Relative to the user origin. The Graphic VDU plotting, testing and line drawing routines all move the current graphics position to the point (or endpoint) specified automatically. Thus, the position returned is probably where the last point was plotted or tested.

Related entries:

GRA MOVE ABSOLUTE

GRA MOVE RELATIVE

67: GRA SET ORIGIN

#BBC9

Set the origin of the user coordinates.

Action:

Set the location of the user origin and move the current position there.

Entry conditions:

DE contains the standard X coordinate of the origin.
HL contains the standard Y coordinate of the origin.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The origin position is given in standard coordinates in which (0,0) is the bottom left corner of the screen.

The default origin position is at (0,0). Whenever the screen mode is changed, by calling SCR SET MODE, the origin is restored to its default position.

Related entries:

GRA FROM USER
GRA GET ORIGIN

68: GRA GET ORIGIN

#BBCC

Get the origin of the user coordinates.

Action:

Ask where the user coordinate origin is located.

Entry conditions:

No conditions.

Exit conditions:

DE contains the standard X coordinate of the origin.

HL contains the standard Y coordinate of the origin.

All other registers preserved.

Notes:

The origin position is given in standard coordinates in which (0,0) is the bottom left corner of the screen.

Related entries:

GRA SET ORIGIN

69: GRA WIN WIDTH

#BBCF

Set the right and left edges of the graphics window.

Action:

Set the horizontal position of the graphics window. The left and right edges are respectively the first and last points that lie inside the window horizontally.

Entry conditions:

DE contains the standard X coordinate of one edge.
HL contains the standard X coordinate of the other edge.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The window edges are given in standard coordinates in which (0,0) is the bottom left corner of the screen and coordinates are signed 16 bit numbers.

The left edge of the window is deemed to be the smaller of the two edges supplied. The window will be truncated, if necessary, to make it fit the screen. The edges are moved to screen byte boundaries so that the window only contains whole bytes (the left edge is moved left, the right edge is moved right). This moves the coordinates of the edges as follows in the various modes:

Mode	Left Edge	Right Edge
0	Multiple of 2	Multiple of 2 minus 1
1	Multiple of 4	Multiple of 4 minus 1
2	Multiple of 8	Multiple of 8 minus 1

The default window covers the whole screen. Whenever the screen mode is changed the window is restored to its default size.

All Graphics VDU point plotting and line drawing routines test whether the points they are about to plot lie inside the window; if they are not then the points are not plotted.

Related entries:

GRA GET W WIDTH
GRA WIN HEIGHT

70: GRA WIN HEIGHT

#BBD2

Set the top and bottom edges of the graphics window.

Action:

Set the vertical position of the graphics window. The top and bottom edges are respectively the last and first points that lie inside the window vertically.

Entry conditions:

DE contains the standard Y coordinate of one edge.
HL contains the standard Y coordinate of the other edge.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The window edges are given in standard coordinates in which (0,0) is the bottom left corner of the screen and coordinates are signed 16 bit numbers.

The top edge will be deemed to be the higher of the two edges supplied.

The window will be truncated, if necessary, to make it fit the screen. The edges will be moved to lie on screen line boundaries so that only whole screen lines are included in the window (the top edge will be moved up, the bottom edge will be moved down). This moves the bottom edge to an even coordinate and the top edge to an odd coordinate.

The default window covers the whole screen. Whenever the screen mode is changed the window is restored to its default size.

All Graphics VDU point plotting and line drawing routines test whether the points they are about to plot lie inside the window; if they do not then the points are not plotted.

Related entries:

GRA GET W HEIGHT
GRA WIN WIDTH

71: GRA GET W WIDTH

#BBD5

Get the left and right edges of the graphics window.

Action:

Ask the horizontal position of the graphics window. The left and right edges are respectively the first and last points that lie inside the window horizontally.

Entry conditions:

No conditions.

Exit conditions:

DE contains the standard X coordinate of the left edge of the window.
HL contains the standard X coordinate of the right edge of the window.

AF corrupt.

All other registers preserved.

Notes:

The window edges are given in standard coordinates in which (0,0) is the bottom left corner of the screen.

The edges may not be exactly the same as those that were set using GRA WIN WIDTH as the window is truncated to fit the screen, and the edges are moved to screen byte boundaries so that the window only contains whole bytes.

Related entries

GRA GET W HEIGHT

GRA WIN WIDTH

72: GRA GET W HEIGHT

#BBD8

Get the top and bottom edges of the graphics window.

Action:

Ask the vertical position of the graphics window. The top and bottom edges are respectively the last and first points that lie inside the window vertically.

Entry conditions:

No conditions.

Exit conditions:

DE contains the standard Y coordinate of the top edge of the window.
HL contains the standard Y coordinate of the bottom edge of the window.

AF corrupt.
All other registers preserved.

Notes:

The window edges are given in standard coordinates. i.e. With (0,0) being the bottom left corner of the screen.

The edges may not be exactly the same as those passed to GRA WIN HEIGHT as the window is truncated to fit the screen, and the edges are moved to lie on screen line boundaries so that only whole screen lines are included in the window.

Related entries:

GRA GET W WIDTH
GRA WIN HEIGHT

73: GRA CLEAR WINDOW

#BBDB

Clear the graphic window.

Action:

Clear the graphics window to the graphics paper ink.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.

All other registers preserved.

Notes:

The current graphics position is moved to the origin of the user coordinates.

Related entries:

GRA SET PAPER

GRA WIN HEIGHT

GRA WIN WIDTH

SCR CLEAR

TXT CLEAR WINDOW

74: GRA SET PEN

#BBDE

Set the graphics plotting ink.

Action:

Set the graphics pen ink. This is the ink used by the Graphics VDU for plotting points, drawing lines and writing characters.

Entry conditions:

A contains the required ink.

Exit conditions:

AF corrupt.
All other registers preserved.

Notes:

The ink is masked to bring it in the range of inks for the current screen mode. In mode 0 the mask is #0F, in mode lit is #03 and in mode 2 it is #01.

In V1.1 firmware the graphics pen ink is taken to delimit the edge of the area to fill when flood filling areas of the screen.

Related entries:

GRA GET PEN
GRA SET PAPER
SCR SET INK
TXT SET PEN

75: GRA GET PEN

#BBE1

Get the current graphics plotting ink

Action:

Ask what the current graphics pen ink is set to. This is the ink used by the Graphics VDU for plotting points, drawing lines and writing characters.

Entry conditions:

No conditions.

Exit conditions:

A contains the ink.
Flags corrupt.
All other registers preserved.

Notes:

This routine has no other effects.

Related entries:

GRA GET PAPER
GRA SET PEN
SCR GET INK
TXT GET PEN

76: GRA SET PAPER

#BBE4

Set the graphics background ink.

Action:

Set the graphics paper ink.

Entry conditions:

A contains the required ink.

Exit conditions:

AF corrupt.
All other registers preserved.

Notes:

The ink is masked to bring it in the range of inks for the current screen mode. In mode 0 the mask is #0F, in mode lit is #03 and in mode 2 it is #01.

The paper ink is the ink that is used for clearing the graphics window, and writing the background to characters. It is assumed to cover everywhere outside the graphics window when testing points.

In V1.1 firmware the graphics paper ink is used to plot pixels corresponding to a zero bit in the line mask when drawing lines.

Related entries:

GRA GET PAPER

GRA SET PEN

SCR GET INK

TXT SET PAPER

77: GRA GET PAPER

#BBE7

Get the current graphics background ink.

Action:

Ask what the current graphics paper ink is set to.

Entry conditions:

No conditions.

Exit conditions:

A contains the ink.

Flags corrupt.

All other registers preserved.

Notes:

The paper ink is the ink that is used for clearing the graphics window, and writing the background to characters. It is assumed to cover everywhere outside the graphics window when testing points.

Related entries:

GRA GET PEN

GRA SET PAPER

SCRGETINK

TXT GET PAPER

78: GRA PLOT ABSOLUTE

#BBEA

Plot a point at an absolute position.

Action:

The current graphic position is moved to the position supplied. If this lies inside the graphics window then the point is plotted in the current graphics pen ink using the current graphics write mode. If the point lies outside the graphics window then no action is taken.

Entry conditions:

DE contains the user X coordinate to plot at.
HL contains the user Y coordinate to plot at.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The position to plot at is given in user coordinates. i.e. Relative to the user origin.

This routine calls the GRA PLOT indirection to plot the point. In its turn GRA PLOT calls the SCR WRITE indirection to set the pixel (if it is in the window).

Related entries:

GRA PLOT
GRA PLOT RELATIVE
GRA TEST ABSOLUTE

79: GRA PLOT RELATIVE

#BBED

Plot a point relative to the current position

Action:

The current graphic position is moved to the position supplied. If this lies inside the graphics window then the point is plotted in the current graphics pen ink using the current graphics write mode. If the point lies outside the graphics window then no action is taken.

Entry conditions:

DE contains a signed X offset.
HL contains a signed Y offset.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The position to plot at is given in relative coordinates. i.e. Relative to the current graphics position.

This routine calls the GRA PLOT indirection to plot the point. In its turn GRA PLOT calls the SCR WRITE indirection to set the pixel (if it is in the window).

Related entries:

GRA PLOT
GRA PLOT ABSOLUTE
GRA TEST RELATIVE

80: GRA TEST ABSOLUTE

#BBF0

Test a point at an absolute position.

Action:

The current graphic position is moved to the position supplied. If this lies inside the graphics window then the pixel is read from the screen and the ink it is set to is decoded and returned. If the position lies outside the graphics window then the current paper ink is returned.

Entry conditions:

DE contains the user X coordinate to test at.
HL contains the user Y coordinate to test at.

Exit conditions:

A contains the ink of the specified point (or the graphics paper ink).

BC, DE, HL and flags corrupt.
All other registers preserved.

Notes:

The position to test is given in user coordinates, i.e. Relative to the user origin.

This routine calls the GRA TEST indirection to test the point. In its turn GRA TEST calls the SCR READ indirection to test the pixel (if it is in the window).

Related entries:

GRA PLOT ABSOLUTE
GRA TEST
GRA TEST RELATIVE

81: GRA TEST RELATIVE

#BBF3

Test a point relative to the current position.

Action:

The current graphic position is moved to the position supplied. If this lies inside the graphics window then the pixel is read from the screen and the ink it is set to is decoded and returned. If the position lies outside the graphics window then the current paper ink is returned.

Entry conditions:

DE contains a signed X offset.
HL contains a signed Y offset.

Exit conditions:

A contains the ink of the specified point (or the graphics paper ink).

BC, DE, HL and flags corrupt.
All other registers preserved.

Notes:

The position to test is given in relative coordinates. i.e. Relative to the current graphics position.

This routine calls the GRA TEST indirection to test the point. In its turn GRA TEST calls the SCR READ indirection to test the pixel (if it is in the window).

Related entries:

GRA PLOT RELATIVE
GRA TEST
GRA TEST ABSOLUTE

82: GRA LINE ABSOLUTE

#BBF6

Draw a line to an absolute position.

Action:

Move the current graphics position to the endpoint supplied. All points between this position and the previous graphics position that lie inside the graphics window may be plotted. Points that lie outside the graphics window are ignored.

Entry conditions:

DE contains the user X coordinate of the endpoint.
HL contains the user Y coordinate of the endpoint.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The position of the end of the line is given in user coordinates. i.e. Relative to the user origin.

In V1.0 firmware the points will be plotted in the current graphics pen ink using the current graphics write mode.

In V1.1 firmware the setting of the line mask determines how pixels on the line will be plotted. The line mask is bit significant and is used repeatedly in the order bit 7, bit 6 down to bit 0 for each pixel in the line. If the bit is one then the pixel is plotted in the graphics pen ink using the current graphics write mode. If the bit is zero then the action taken depends on the graphics background write mode. If the background mode is opaque then the pixel is plotted in the graphics paper ink using the current graphics write mode. If the background mode is transparent then the pixel is not plotted.

In V1.1 firmware the first pixel of the line (that at the previous graphics position) is not plotted if the first point plotting mode is set false.

This routine calls the GRA LINE indirection to draw the line. In its turn GRA LINE calls the SCR WRITE indirection to write the pixels (for pixels in the graphics window).

Related entries:

GRA LINE
GRA LINE RELATIVE
GRA SET BACK
GRA SET FIRST
GRA SET LINE MASK

83: GRA LINE RELATIVE

#BBF9

Draw a line relative to the current position.

Action:

Move the current graphics position to the endpoint supplied. All points between this position and the previous graphics position that lie inside the graphics window may be plotted. Points that lie outside the graphics window are ignored.

Entry conditions:

DE contains the signed X offset of the endpoint. HL contains the signed Y offset of the endpoint.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The position of the end of the line is given in relative coordinates. i.e. Relative to the current graphics position.

In V1.0 firmware the points will be plotted in the current graphics pen ink using the current graphics write mode.

In V1.1 firmware the setting of the line mask determines how pixels on the line will be plotted. The line mask is bit significant and is used repeatedly in the order bit 7, bit 6 down to bit 0 for each pixel in the line. If the bit is one then the pixel is plotted in the graphics pen ink using the current graphics write mode. If the bit is zero then the action taken depends on the graphics background write mode. If the background mode is opaque then the pixel is plotted in the graphics paper ink using the current graphics write mode. If the background mode is transparent then the pixel is not plotted.

In V1.1 firmware the first pixel of the line (that at the previous graphics position) is not plotted if the first point plotting mode is set false.

This routine calls the GRA LINE indirection to draw the line. In its turn GRA LINE calls the SCR WRITE indirection to write the pixels (for pixels in the graphics window).

Related entries:

GRA LINE
GRA LINE ABSOLUTE
GRA SET BACK
GRA SET FIRST
GRA SET LINE MASK

84: GRA WR CHAR

#BBFC

Put a character on the screen at the current graphics position.

Action:

Write a character on the screen at the current graphics position.

Entry conditions:

A contains the character to write.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The character is written with its top left corner being the current graphics position.

All characters are printed, even control codes (characters #00.. #1F).

The current position is moved right by the width of the character (ready for another character to be written). In mode 0 this move is 32 points right, in mode 1 the move is 16 points and in mode 0 it is 8 points.

The character will be plotted in the graphic pen ink. In the case of V1.0 firmware, or V1.1 firmware when the background write mode is set to opaque, the background to the character will be plotted in the graphics paper ink. In the case of V1.1 firmware when the background write mode is set transparent the background pixels are not plotted. Pixels in the character that lie outside the graphics window will not be plotted. The pixels are plotted using the SCR WRITE indirection so they are written using the current graphics write mode.

Related entries:

GRA SET BACK

TXT SET GRAPHIC

TXT WR CHAR

85: SCR INITIALISE

#BBFF

Initialize the Screen Pack.

Action:

Full initialization of the Screen Pack (as used during EMS). All Screen Pack variables and indirections are initialized, also the screen mode and the inks are initialized to their default settings.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The screen indirections (SCR READ, SCR WRITE and SCR MODE CLEAR) are set to their default routines.

The inks are set to their default colours (see Appendix V).

The ink flashing periods are set to their default values.

The screen is put into mode 1.

The screen base is set to put the screen memory at #C000.. #FFFF (under the upper ROM).

The screen offset is set to 0.

The screen is cleared to ink 0.

The Graphics VDU write mode is set to FORCE mode.

The ink flashing frame flyback event is setup.

The initialization is performed in an order that attempts to avoid the previous contents of the screen becoming visible (at EMS the contents will be random).

Related entries:

GRA INITIALISE

SCR RESET

TXT INITIALISE

86: SCR RESET

#BC02

Reset the Screen Pack.

Action:

Reinitializes the Screen Pack indirections and the ink colours. Also reinitializes the flash rate and Graphics VDU write mode.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The screen indirections (SCR READ, SCR WRITE and SCR MODE CLEAR) are set to their default routines.

The inks are set to their default colours (see Appendix V).

The ink flashing periods are set to their default values.

The Graphics VDU write mode is set to FORCE mode.

The inks are not passed to the hardware. This will be done when the inks flash next.

Related entries:

SCR INITIALISE

SCR SET ACCESS

SCR SET FLASHING

SCR SET INK

87: SCR SET OFFSET

#BC05

Set the offset of the start of the screen.

Action:

Set the offset of the first character on the screen. By changing this offset the screen can be rolled.

Entry conditions:

HL contains the required offset.

Exit conditions:

AF and HL corrupt.
All other registers preserved.

Notes:

The offset passed is masked with #07FE to make sure it is not too big and to make sure that the offset is even. (The screen is only capable of rolling in two byte increments).

The screen base and screen offset are combined into a single value and sent to the hardware together.

The screen offset is used by SCR CHAR POSITION and SCR DOT POSITION to calculate screen addresses. If the screen offset is changed merely by calling the Machine Pack routine MC SCREEN OFFSET then the Text and Graphics VDUs will use incorrect screen addresses.

The offset is set to zero when the screen mode is set or the screen is cleared by calling SCR CLEAR.

Related entries:

MC SCREEN OFFSET
SCR GET LOCATION
SCR HW ROLL
SCR SET BASE
SCR SET POSITION

88: SCR SET BASE

#BC08

Set the area of RAM to use for the screen memory.

Action:

Sets the base address of the screen memory. This can be used to move the screen out from underneath the upper ROM or to display a prepared screen instantly.

Entry conditions:

A contains the more significant byte of the base address.

Exit conditions:

AF and HL corrupt.
All other registers preserved.

Notes:

The screen memory can only be located on a 16K boundary so the value passed is masked with #C0. The default screen base, set at EMS, is #C0.

The screen offset is combined with the screen base into a single value which is sent to the hardware.

The screen base address is used by SCR CHAR POSITION and SCR DOT POSITION to calculate screen addresses. If the screen base is changed merely by calling the Machine pack routine MC SCREEN OFFSET then the text and graphics VDUs will use incorrect screen addresses.

The screen memory is not cleared when the screen base is set, use SCR CLEAR to do this.

Related entries:

MC SCREEN OFFSET
SCR GET LOCATION
SCR SET OFFSET
SCR SET POSITION

89: SCR GET LOCATION

#BC0B

Fetch current base and offset settings.

Action:

Ask where the screen memory is located and where the start of the screen is.

Entry conditions:

No conditions.

Exit conditions:

A contains the more significant byte of the base address.
HL contains the current offset.

Flags corrupt.
All other registers preserved.

Notes:

The base and offsets returned by this routine may not be the same as those set using SCR SET BASE or SCR SET OFFSET. This is because the values are masked to make them legal and the screen offset is also changed when the hardware screen rolling routine, SCR HW ROLL, is used.

Related entries:

SCR SET BASE
SCR SET OFFSET
SCR SET POSITION

90: SCR SET MODE

#BC0E

Set screen into a new mode.

Action:

Put the screen into a new mode and make sure that the Text and Graphics VDUs are set up correctly.

Entry conditions:

A contains the required mode.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The mode requested is masked with #03. If the resulting value is 3 then no action is taken. Otherwise one of the following screen modes is set up:

Mode 0:	160 x 200 pixels,	20 x 25 characters.
Mode 1:	320 x 200 pixels,	40 x 25 characters.
Mode 2:	640 x 200 pixels,	80 x 25 characters.

At an early stage the screen is cleared to avoid the old contents of the screen being displayed in the wrong mode. The screen is cleared by calling the SCR MODE CLEAR indirection.

All text and graphics windows are set to cover to whole screen and the graphics user origin is set to the bottom left corner of the screen. The cursor blobs for all text streams are turned off. Stream zero is selected.

The current text and graphics pen and paper inks are masked as appropriate for the new mode (see TXT SET PEN et al). When changing mode to a mode that allows fewer inks on the screen this may cause the pen or paper inks to change.

Related entries:

MC SET MODE

SCR GET MODE

91: SCR GET MODE

#BC11

Ask the current screen mode.

Action:

Fetch and test the current screen mode.

Entry conditions:

No conditions.

Exit conditions:

If current mode is mode 0:

Carry true.

Zero false.

A contains 0.

If current mode is mode 1:

Carry false.

Zero true.

A contains 1.

If current mode is mode 2:

Carry false.

Zero false.

A contains 2.

Always:

Other flags corrupt.

All other registers preserved.

Notes:

The modes are:

Mode 0: 160 x 200 pixels, 20 x 25 characters.

Mode 1: 320 x 200 pixels, 40 x 25 characters.

Mode 2: 640 x 200 pixels, 80 x 25 characters.

Related entries:

SCR SET MODE

92: SCR CLEAR

#BC14

Clear the screen (to ink zero).

Action:

Clear the whole of screen memory to zero.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

At an early stage the ink flashing is turned off and the inks are all set to the same colour as ink 0. This makes the screen clearing appear instantaneous. When all the screen memory has been set to 0 the ink flashing is turned back on (an ink flashing event is added to the frame flyback queue) and all inks are set to their proper colours.

If the text paper ink and the graphics paper ink are not set to ink 0 then this will become apparent on the screen when characters are written or windows are cleared.

The screen offset is set to zero.

Related entries:

GRA CLEAR WINDOW

SCR MODE CLEAR

TXT CLEAR WINDOW

93: SCR CHAR LIMITS

#BC17

Ask the size of the screen in characters.

Action:

Get the last character row and column on the screen in the current mode.

Entry conditions:

No conditions.

Exit conditions:

B contains the physical last column on the screen.

C contains the physical last row on the screen.

AF corrupt.

All other registers preserved.

Notes:

The screen edges are given in physical coordinates. i.e. Row 0, column 0 is the top left corner of the screen. This means that the last column on the screen is 19 in mode 0,39 in mode 1 and 79 in mode 2. The last row on the screen is 24 in all modes.

Related entries:

SCR GET MODE

94: SCR CHAR POSITION

#BC1A

Convert physical coordinates to a screen position.

Action:

Calculate the screen address of the top left corner of a character position on the screen. Also return the width of a character in the current mode.

Entry conditions:

H contains the physical character column.
L contains the physical character row.

Exit conditions:

HL contains the screen address of the top left corner of the character.
B contains the width in bytes of a character in screen memory.

AF corrupt.
All other registers preserved.

Notes:

The character position is given in physical coordinates. i.e. Row 0, column 0 is the top left corner of the screen.

The character position given is not checked for being legal. An illegal position (one outside the limits of the screen) will generate a meaningless screen address.

The conversion to screen address uses the following formula:

$$\text{Screen address} = \text{Screen base} + (\text{Block offset} \text{ MOD } \#0800)$$

where:

$$\text{Block offset} = (\text{Row} * 80) + (\text{Column} * \text{Width}) + \text{Screen offset}$$

and:

Screen base is the address of the start of screen memory.

Width is the width of a character in bytes in the current mode (4 in mode 0, 2 in mode 1, 1 in mode 2).

Screen offset is offset of the first byte to be displayed on the screen.

Related entries:

SCR DOT POSITION

SCR NEXT BYTE

8CR NEXT LINE

SCR PREY BYTE

SCR PREY LINE

95: SCR DOT POSITION

#BC1D

Convert base coordinates to a screen position.

Action:

Calculate the screen address and mask for a pixel. Also return an indication of the number of pixels in a screen byte in the current mode.

Entry conditions:

DE contains the base X coordinate of a pixel.
HL contains the base Y coordinate of a pixel.

Exit conditions:

HL contains the screen address of the pixel.
C contains the mask for the pixel.
B contains one less than the number of pixels in a byte.

AF and DE corrupt.
All other registers preserved.

Notes:

The pixel position is given in base coordinates. i.e. (0,0) is the pixel in the bottom left corner of the screen and each coordinate position refers to a single pixel.

The pixel position is not checked for being legal (within the limits of the screen). If it is not then the screen address calculated is meaningless.

The conversion to screen address uses the following formula:

$$\text{Screen address} = \text{Screen base} + (\text{Line in row} * \#0800) + (\text{Row offset MOD } \#0800)$$

where:

Screen base is the start address of screen memory.
Line in row = (199 - Y coordinate) MOD 8
Row offset = (Row number * 80) + Byte in row + Screen offset

and:

Row number = (199 - Y coordinate)/8
Byte in row = X coordinate/Byte width
Screen offset is offset of the first byte to be displayed on the screen.
Byte width is the number of pixels in a byte in the current mode (2 in mode 0, 4 in mode 1, 8 in mode 2).

X coordinate MOD Byte width is used to calculate the mask for the appropriate pixel.

Related entries:

GRA FROM USER

SCR CHAR POSITION

SCR NEXT BYTE

SCR NEXT LINE

SCR PREY BYTE

SCR PREY LINE

96: SCR NEXT BYTE

#BC20

Step a screen address right one byte.

Action:

Calculate the screen address of the byte right of the supplied screen address.

Entry conditions:

HL contains a screen address.

Exit conditions:

HL contains the updated screen address.

AF corrupt.

All other registers preserved.

Notes:

Moving off the end of the screen line is not prevented. It will simply point the screen address at the next byte in the screen block. Normally this will be the first byte on a screen line 8 screen lines down from the old line (i.e. down one character row). However, moving right off the end of the last screen line in a block will point the screen address at the start of the 48 bytes in the block that are not displayed on the screen.

This routine is intended to be used for moving the screen address when putting characters or drawing lines on the screen.

Related entries:

SCR CHAR POSITION

SCR DOT POSITION

SCR NEXT LINE

5CR PREY BYTE

SCR PREY LINE

97: SCR PREV BYTE

#BC23

Step a screen address left one byte.

Action:

Calculate the screen address of the byte left of the supplied screen address.

Entry conditions:

HL contains a screen address.

Exit conditions:

HL contains the updated screen address.

AF corrupt.

All other registers preserved.

Notes:

Moving off the start of the screen line is not prevented. It will simply point the screen address at the previous byte in the screen block. Normally this will be the last byte on a screen line 8 screen lines up from the old line (i.e. up one character row). However, moving left off the start of the top screen line in a block will point the screen address at the last of the 48 bytes in the block that are not displayed on the screen.

This routine is intended to be used for moving the screen address when putting characters or drawing lines on the screen.

Related entries:

SCR CHAR POSITION

SCR DOT POSITION

SCR NEXT BYTE

SCR NEXT LINE

SCR PREY LINE

98: SCR NEXT LINE

#BC26

Step a screen address down one line.

Action:

Calculate the screen address of the byte below the supplied screen address.

Entry conditions:

HL contains a screen address.

Exit conditions:

HL contains the updated screen address.

AF corrupt.

All other registers preserved.

Notes:

Moving off the bottom of the screen is not prevented (and not recommended).

After moving off the bottom the screen address is not useful.

This routine is intended to be used for moving the screen address when putting characters or drawing lines on the screen.

Related entries:

SCR CHAR POSITION

SCR DOT POSITION

SCR NEXT BYTE

SCR PREY BYTE

SCR PREY LINE

99: SCR PREV LINE

#BC29

Step a screen address up one line.

Action:

Calculate the screen address of the byte above the supplied screen address.

Entry conditions:

HL contains a screen address.

Exit conditions:

HL contains the updated screen address.

AF corrupt.

All other registers preserved.

Notes:

Moving off the top of the screen is not prevented (and not recommended).

After moving off the top the screen address is not useful.

This routine is intended to be used for moving the screen address when putting characters or drawing lines on the screen.

Related entries:

SCR CHAR POSITION

SCR DOT POSITION

SCR NEXT BYTE

SCR NEXT LINE

SCR PREY BYTE

100: SCR INK ENCODE

#BC2C

Encode an ink to cover all pixels in a byte.

Action:

Convert an ink to the encoded form **that** will set all pixels in a byte to the ink. This encoded ink can then be masked to generate the appropriate value to set a single pixel to the ink.

Entry conditions:

A contains an ink number.

Exit conditions:

A contains the encoded ink.

Flags corrupt.

All other registers preserved.

Notes:

The encoding is not trivial as the pixels in a byte are interleaved and also the bits in a pixel are not in the obvious order. The pixel bits are (most significant to least significant):

	Mode 0	Mode 1	Mode 2
Leftmost pixel:	Bits 1,5,3,7	Bits 3,7	Bit 7
		Bits 2,6	Bit 6
		Bits 0,4,2,6	Bit 5
Rightmost pixel:	Bits 0,4,2,6	Bits 1,5	Bit 4
		Bits 0,4	Bit 3
			Bit 2
			Bit 1
			Bit 0

The Text and Graphic VDUs store their pen and paper inks in this encoded form for ease of use internally. This saves time converting the ink for each pixel plotted.

The encoding is different in different modes and so all inks have to be re-encoded when the screen mode is changed. SCR SET MODE does this automatically for the Text VDU and Graphics VDU pen and paper inks.

Related entries:

SCR INK DECODE

101: SCR INK DECODE

#BC2F

Decode an encoded ink.

Action:

Convert an encoded ink to the appropriate ink number.

Entry conditions:

A contains an encoded ink.

Exit conditions:

A contains the ink number.

Flags corrupt.

All other registers preserved.

Notes:

The decoding is performed by decoding the ink of the leftmost pixel in the encoded ink. The ink for this pixel is encoded in the following bits (most significant to least significant) in the various screen modes:

Mode 0: Bits 1,5,3,7

Mode 1: Bits 3,7

Mode 2: Bit 7

Related entries:

SCR INK ENCODE

102: SCR SET INK

#BC32

Set the colours in which to display an ink.

Action:

Set which two colours will be used to display an ink. If the two colours are the same then the ink will remain a steady colour. If the colours are different then the ink will alternate between these two colours.

Entry conditions:

A contains an ink number.
B contains the first colour.
C contains the second colour.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The ink number is masked with #0F to make sure it is legal, and the colours are masked with #1F. Colours 27..31 are not intended for use; they are merely duplicates of other colours available.

The new colours for an ink are not sent to the hardware immediately. They are stored and will appear on the screen when the next frame flyback occurs.

The length of time for which each colour is displayed on the screen can be set by calling SCR SET FLASHING.

The inks are set to their default colours at EMS and when SCR RESET is called.

The various colours available and the default ink colours set are described in Appendix V.

Related entries:

GRA SET PAPER
GRA SET PEN
SCR GET INK
SCR SET BORDER
SCR SET FLASHING
TXT SET PAPER
TXT SET PEN

103: SCR GET INK

#BC35

Ask the colours an ink is currently displayed in.

Action:

Get the two colours that are used to display an ink on the screen.

Entry conditions:

A contains an ink number.

Exit conditions:

B contains the first colour.

C contains the second colour.

AF, DE and HL corrupt.

All other registers preserved.

Notes:

The ink number is masked with #0F to make sure it is legal. The colours returned may not be the same as those supplied to the Screen Pack as the colours are masked when they are set.

The new colours for an ink are not sent to the hardware immediately when they are set. They are stored and appear on the screen when the next frame flyback occurs. This means that the colours returned may not actually be visible to the user yet.

The default settings for the inks and the various colours available are described in Appendix V.

Related entries:

GRA GET PAPER

GRA GET PEN

SCR GET BORDER

SCR SET INK

TXT GET PAPER

TXT GET PEN

104: SCR SET BORDER

#BC38

Set the colours in which to display the border.

Action:

Set which two colours will be used to display the border. If the two colours are the same then the border will remain a steady colour. If the colours are different then the border will alternate between these two colours.

Entry conditions:

B contains the first colour.
C contains the second colour.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The colours are masked with # 1F to ensure that they are legal. Colours 27..31 are not intended for use; they are merely duplicates of other colours available.

The new colours for the border are not sent to the hardware immediately. They are stored and will appear on the screen when the next frame flyback occurs.

The length of time for which each colour is displayed on the screen can be set by calling SCR SET FLASHING.

The border is set to its default colour at EMS and when SCR RESET is called. The default colour and the colours available are described in Appendix V.

Related entries:

SCR GET BORDER
SCR SET FLASHING
SCR SET INK

105: SCR GET BORDER

#BC3B

Ask the colours the border is currently displayed in.

Action:

Get the two colours used to display the border on the screen.

Entry conditions:

No conditions.

Exit conditions:

B contains the first colour.

C contains the second colour.

AF, DE and HL corrupt.

All other registers preserved.

Notes:

The colours returned may not be the same as those supplied to the Screen Pack as they are masked when they are set.

The new colours for the border are not sent to the hardware immediately when they are set. They are stored and appear on the screen when the next frame flyback occurs. This means that the colours returned may not actually be visible to the user yet. The default border colour and the colours available are described in Appendix V.

Related entries:

SCR GET INK

SCR SET BORDER

106: SCR SET FLASHING

#BC3E

Set the flash periods.

Action:

Set for how long each of the two colours for the inks and the border are to be displayed on the screen. These settings apply to all inks and the border.

Entry conditions:

H contains the period for the first colour.

L contains the period for the second colour.

Exit conditions:

AF and HL corrupt.

All other registers preserved.

Notes:

The flash periods are given in frame flybacks (1/50 or 1/60 of a second). A period of 0 is taken to mean a period of 256.

The default setting for the flash periods is 10 frame flybacks (1/5 or 1/6 of a second). This is set at EMS and when SCR RESET is called.

The new flash periods are not used immediately but when the inks next flash.

Related entries:

SCR GET FLASHING

SCR SET BORDER

SCR SET INK

107: SCR GET FLASHING

#BC41

Ask the current flash periods.

Action:

Get the time for which each of the two colours associated with an ink or the border is displayed.

Entry conditions:

No conditions.

Exit conditions:

H contains the period for the first colour.
L contains the period for the second colour.

AF corrupt.
All other registers preserved.

Notes:

The flash periods are given in frame flybacks (1/50 or 1/60 of a second).

A period of 0 means 256.

Related entries:

SCR SET FLASHING

108: SCR FILL BOX

#BC44

Fill a character area of the screen with an ink.

Action:

Fill a rectangular area of the screen with an ink. The boundaries of this area are given in character positions.

Entry conditions:

A contains the encoded ink to fill the area with.

H contains the physical left column of the area to fill.

D contains the physical right column of the area to fill.

L contains the physical top row of the area to fill.

E contains the physical bottom row of the area to fill.

Exit conditions:

AF, BC, DE and HL corrupt.

All other registers preserved.

Notes:

The area boundaries are given in physical coordinates. i.e. Row 0, column 0 is the top left corner of the screen. They are not checked for legality. If illegal boundaries are passed (edges off the screen) then unpredictable effects may occur.

The screen is written directly without using any other write routine. The current Graphics VDU write mode is therefore ignored.

Related entries:

SCR CLEAR

SCR FLOOD BOX

TXT CLEAR WINDOW

109: SCR FLOOD BOX

#BC47

Fill a byte area of the screen.

Action:

Fill a rectangular area of the screen with an ink. The boundaries of the area must lie on byte boundaries. This routine will not fill an arbitrary area of the screen to a pixel boundary.

Entry conditions:

C contains the encoded ink to fill the area with.

HL contains the screen address of the top left corner of the area to fill.

D contains the (unsigned) width of the area to fill in bytes.

E contains the (unsigned) height of the area to fill in screen lines.

Exit conditions:

AF, BC, DE and HL corrupt.

All other registers preserved.

Notes:

The whole of the rectangle being cleared must lie on the screen. If any of it lies off the screen then unpredictable effects may occur.

A height or width of 0 is taken to mean 256 (which is too large to fit on the screen).

The screen is written directly without using any other write routine. The current

Graphics VDU write mode is therefore ignored.

Related entries:

GRA CLEAR WINDOW

SCR CLEAR

SCR FILL BOX

110: SCR CHAR INVERT

#BC4A

Invert a character position.

Action:

All pixels at a character position that are written in one ink are rewritten in a second ink, and vice versa. This gives an inverse effect to the character position. Inverting the character a second time will restore the original inks. This effect is used to draw the Text VDU cursors.

Entry conditions:

B contains an encoded ink.
C contains another encoded ink.
H contains a physical character column.
L contains a physical character row.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The character position is given in physical coordinates. i.e. Row 0, column 0 is the top left corner of the screen.

The character position given is not checked for being legal. An illegal position (one outside the limits of the screen) will have unpredictable effects.

All pixels at the character position are exclusive-ored with the exclusive-or of the two inks supplied. Pixels at the character position that are set to one of the two inks supplied will therefore be set to the other supplied ink. Pixels set to other inks will also be altered.

Related entries:

TXT PLACE CURSOR

TXT REMOVE CURSOR

111: SCR HW ROLL

#BC4D

Move the whole screen up or down eight pixel lines (one character).

Action:

Roll the screen using the hardware. The new line appearing on the screen is cleared.

Entry conditions:

If the screen is to roll down:
B must be zero.

If the screen is to roll up:
B must be non-zero.

Always:
A contains the encoded ink to clear the new line to.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The screen is rolled by changing the screen offset (see SCR SET OFFSET).

Rolling the screen upwards moves the screen contents up and clears the new bottom line. The screen offset is therefore increased by 80 (MOD #0800).

Rolling the screen downwards moves the screen contents down and clears the new top line. The screen offset is therefore decreased by 80 (MOD #0800).

The new line is cleared by writing to it directly thus the Graphics VDU write mode is ignored.

The Text VDU roll count is not changed by this routine (see TXT GET WINDOW).

Special precautions are taken to make sure that the screen is kept looking presentable during the rolling and in particular during the clearing of the new line. Principally this consists of clearing the new line in two parts. First the part that is not visible on the screen (by virtue of the screen addressing) is cleared. Then, after waiting for frame flyback and changing the screen offset, the second half of the the line that was part of the line that just rolled off the screen is cleared.

Related entries:

SCR SET OFFSET
SCR SW ROLL

112: SCR SW ROLL

#BC50

Move an area of the screen up or down eight pixel lines (one character).

Action:

Roll an area of the screen by copying. The area to be rolled is specified in character positions.

Entry conditions:

If the screen is to roll down:

B must be zero.

If the screen is to roll up:

B must be non-zero.

Always:

A contains the encoded ink to clear the new line to.

H contains the physical left column of the area to roll.

D contains the physical right column of the area to roll.

L contains the physical top row of the area to roll.

E contains the physical bottom row of the area to roll.

Exit conditions:

AF, BC, DE and HL corrupt.

All other registers preserved.

Notes:

The area boundaries are given in physical coordinates. i.e. Row 0, column 0 is the top left corner of the screen. The boundaries are not checked for legality. If illegal boundaries are passed (edges off the screen) then unpredictable effects may occur.

Rolling the area upwards moves the area contents up and clears the new bottom line. Rolling the area downwards moves the area contents down and clears the new top line.

The line is cleared by writing to it directly; the Graphics VDU write mode is ignored.

The Text VDU roll count is not changed by this routine (see TXT GET WINDOW).

Special precautions are taken to make sure that the screen is kept looking presentable during the rolling. Principally this consists of waiting for frame flyback before performing the copy.

Related entries:

SCR HW ROLL

113: SCR UNPACK

#BC53

Expand a character matrix for the current screen mode.

Action:

Convert a matrix from its standard form to a set of pixel masks as appropriate for the current screen mode.

Entry conditions:

HL contains the address of a matrix.

DE contains the address of an area to unpack into.

Exit conditions:

AF, BC, DE and HL corrupt.

All other registers preserved.

Notes:

The matrix is converted into a series of masks which cover all the screen bytes in the character. This means that each byte of the matrix is converted to 4 bytes in mode 0, 2 bytes in mode 1 and 1 byte in mode 2. Thus the unpacking area must be 32, 16 or 8 bytes long.

If a bit in the matrix is set then the appropriate pixel mask is included in the unpacked version (the bits are set to one). Otherwise the pixel mask is not included in the unpacked version (the bits are set to zero).

Related entries:

SCR REPACK

114: SCR REPACK

#BC56

Compress a character matrix to the standard form.

Action:

A character on the screen is converted to a matrix by comparing each pixel with an ink. If the pixel is set to that ink then the appropriate bit in the character matrix is set, otherwise the bit is cleared.

Entry conditions:

A contains the encoded ink to match against.
H contains the physical character column to read from.
L contains the physical character row to read from.
DE contains the address of the area to construct the matrix in.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The character position is given in physical coordinates in which row 0, column 0 is the top left corner of the screen.

The character position given is not checked for legality. An illegal position (one outside the limits of the screen) will have unpredictable effects.

The matrix produced has the normal layout. It is 8 bytes long, stored top line first and bottom line last, the most significant bit of a byte refers to the leftmost pixel of a line and the least significant bit to the rightmost pixel.

Because the pixels are tested for being set to only one ink the matrix produced is not an exact representation of what is on the screen. It may be necessary, when trying to read characters from the screen, to repack using various different inks.

Related entries:

SCR UNPACK
TXT RD CHAR

115: SCR ACCESS

#BC59

Set the screen write mode for the Graphics VDU.

Action:

Set the Graphics VDU write mode so that the Graphics VDU plots pixels by writing, anding, oring or exclusive-oring.

Entry conditions:

A contains the required write mode.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The write mode is masked with #03 to make it legal. The write modes are:

0: FORCE mode:	NEW = INK
1: XOR mode:	NEW = INK exclusive-or OLD
2: AND mode:	NEW = INK and OLD
3: OR mode:	NEW = INK or OLD

NEW is the final setting of the pixel.
OLD is the current setting of the pixel.
INK is the ink being plotted.

The default mode is FORCE mode (mode 0) and this is set at EMS and when SCR RESET is called.

Setting the write mode affects how the indirection routine SCR WRITE sets pixels. Graphics VDU plotting routines call this indirection to set pixels and so the write mode affects the Graphics VDU. No Text VDU routines call this indirection (they set pixels on the screen directly) and so the write mode does not affect the Text VDU. The routines that clear areas of the screen (e.g. GRA CLEAR WINDOW) act like the Text VDU and are unaffected by the write mode.

Related entries:

GRA DEFAULT
SCR INITIALISE
SCR RESET
SCR WRITE

116: SCR PIXELS

#BC5C

Write a pixel to the screen ignoring the Graphics VDU write mode.

Action:

Write a pixel or pixels to the screen. The position to write at is given by a screen address and pixel mask. The pixel is always set to the ink supplied whatever mode of writing the Graphics VDU is using.

Entry conditions:

B contains the encoded ink to write.
C contains the mask for the pixel(s).
HL contains the screen address of the pixel(s).

Exit conditions:

AF corrupt.
All other registers preserved.

Notes:

The screen address is not checked and so passing an invalid screen address will have unpredictable results.

The pixel mask may be a combined mask for more than one pixel (thus speeding up plotting in certain cases).

To plot a pixel using the Graphics VDU write mode SCR WRITE should be called. SCR PIXELS is equivalent to calling SCR WRITE when the default mode (FORCE mode) is selected. The Text VDU sets the pixels in characters using FORCE mode.

Related entries:

SCR WRITE

117: SCR HORIZONTAL

#BC5F

Plot a purely horizontal line.

Action:

Draw a line on the screen that runs horizontally. The pixels on the line are plotted using the SCR WRITE indirection and thus use the current Graphics VDU write mode.

Entry conditions:

A contains the encoded ink to draw in.
DE contains the base X coordinate of the start of the line.
BC contains the base X coordinate of the end of the line.
HL contains the base Y coordinate of the line.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The endpoints of the line are given in base coordinates. i.e. (0,0) is the pixel in the bottom left corner of the screen and each coordinate position refers to a single pixel.

The endpoints are not checked for being legal (within the limits of the screen). If they are not legal then unpredictable effects may occur.

The start X coordinate must be less than or equal to the end X coordinate.

This routine may be used to duplicate the method that the Graphics VDU uses for plotting lines - it splits a line that is more horizontal than vertical into a number of segments that are purely horizontal and plots these separately.

Related entries:

GRA FROM USER
GRA LINE ABSOLUTE
GRA LINE RELATIVE
SCR VERTICAL

118: SCR VERTICAL

#BC62

Plot a purely vertical line.

Action:

Draw a line on the screen that runs vertically. The SCR WRITE indirection is used to plot pixels in the line thus the current Graphics VDU write mode is used.

Entry conditions:

A contains the encoded ink to draw in.

DE contains the base X coordinate of the line.

HL contains the base Y coordinate of the start of the line.

BC contains the base Y coordinate of the end of the line.

Exit conditions:

AF, BC, DE and HL corrupt.

All other registers preserved.

Notes:

The endpoints of the line are given in base coordinates. i.e. (0,0) is the pixel in the bottom left corner of the screen and each coordinate position refers to a single pixel.

The endpoints are not checked for being legal (within the limits of the screen). If they are not legal then unpredictable effects may occur.

The start Y coordinate must be less than or equal to the end Y coordinate.

This routine may be used to duplicate the method that the Graphics VDU uses for plotting lines - it splits a line that is more vertical than horizontal into a number of segments that are purely vertical and plots these separately.

Related entries:

GRA FROM USER

GRA LINE ABSOLUTE

GRA LINE RELATIVE

SCR HORIZONTAL

119: CAS INITIALISE

#BC65

Initialize the Cassette Manager.

Action:

Full initialization of the Cassette Manager (as used during EMS).

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

Operations carried out are:

- All streams are marked closed.
- The default write speed is set up.
- The prompt messages are turned on.
- The cassette motor is turned off (except on V1.0 firmware).

Related entries:

CAS IN ABANDON
CAS NOISY
CAS OUT ABANDON
CAS SET SPEED
CAS STOP MOTOR

120: CAS SET SPEED

#BC68

Set the write speed.

Action:

Set the length to write bits and the amount of write precompensation to apply.

Entry conditions:

HL contains the length of half a zero bit.

A contains the precompensation to apply.

Exit conditions:

AF and HL corrupt.

All other registers preserved.

Notes:

The speed supplied is the length of half a zero bit in microseconds. A one bit is written as twice the length of a zero bit. The speed supplied can be related to the average baud rate (assuming equal numbers of ones and zeros) by the following equation:

$$\begin{aligned} \text{Average baud rate} \\ &= 1\,000\,000 / (3 * \text{Half zero length}) \\ &= 333\,333 / \text{Half zero length} \end{aligned}$$

The half zero length must lie between 130 and 480 microseconds. Values outside this range will cause read and write errors.

The precompensation supplied is the extra length, in microseconds, to add to half a one bit and to subtract from half a zero bit under certain conditions. The amount of precompensation required varies with the speed (more is required at higher baud rates).

The precompensation may lie between 0 and 255 microseconds although the higher settings are not useful as they will cause read and write errors.

The default half zero length and precompensation settings are 333 microseconds (1000 baud) and 25 microseconds respectively. The commonly used faster setting is 167 microseconds (2000 baud) with 50 microseconds of precompensation. These values have been determined after extensive testing and the user is advised to stick to them.

Related entries:

CAS INITIALISE

121: CAS NOISY

#BC6B

Enable or disable prompt messages.

Action:

Disabling messages will prevent the prompt and information messages from being printed. It will not prevent error messages from being printed. Enabling messages allows all messages to be printed.

Entry conditions:

If messages are to be enabled:

A must be zero.

If messages are to be disabled:

A must be non-zero.

Exit conditions:

AF corrupt.

All other registers preserved.

Notes:

The prompt and information messages which are turned off are:

Press PLAY then any key:

Press REC and PLAY then any key:

Found <FILENAME> block <N>

Loading <FILENAME> block <N>

Saving <FILENAME> block <N>

The error messages which are not turned off are:

Read error <x>

Write error a

Rewind tape

Related entries:

CAS INITIALISE

122: CAS START MOTOR

#BC6E

Start the cassette motor.

Action:

Turn the cassette motor on and wait for it to pick up speed if it was previously off.

Entry conditions:

No conditions.

Exit conditions:

If the motor turned on OK:

 Carry true.

If the user hit escape:

 Carry false.

Always:

 A contains the previous motor state.

 Other flags corrupt.

 All other registers preserved.

Notes:

If the motor is not already on then the routine waits for approximately two seconds to allow the tape to reach full speed.

The motor is always turned on by this routine. If the user hits the escape key then the time spent waiting for the motor to pick up speed is truncated.

The previous motor state may be passed to CAS RESTORE MOTOR.

Related entries:

CAS RESTORE MOTOR

CAS STOP MOTOR

123: CAS STOP MOTOR

#BC71

Stop the cassette motor.

Action:

Turn the cassette motor off and return its previous state.

Entry conditions:

No conditions.

Exit conditions:

If the motor was turned off OK:

 Carry true.

If the user hit escape:

 Carry false.

Always:

 A contains the previous motor state.

 Other flags corrupt.

 All other registers preserved.

Notes:

The motor is always turned off by this routine. There is no delay to allow the motor to slow down.

The previous motor state may be passed to CAS RESTORE MOTOR.

Related entries:

CAS RESTORE MOTOR

CAS START MOTOR

124: CAS RESTORE MOTOR

#BC74

Restore previous state of cassette motor.

Action:

Turn the cassette motor on or off again. Wait for motor to pick up speed when turning the motor on if it is currently off.

Entry conditions:

A contains the previous motor state.

Exit conditions:

If the motor was turned on or off OK:

 Carry true.

If the user hit escape:

 Carry false.

Always:

 A and other flags corrupt.

 All other registers preserved.

Notes:

This routine uses the previous motor state as returned by CAS START MOTOR or CAS STOP MOTOR.

If calling this routine results in the motor being turned on when it is currently off then the routine waits for approximately two seconds to allow the tape to reach full speed.

The motor is always turned on or off (as appropriate) by this routine. If the user hits the escape key then this merely truncates the time spent waiting for the motor to pick up speed.

Related entries:

CAS START MOTOR

CAS STOP MOTOR

125: CAS IN OPEN

#BC77

Open a file for input.

Action:

Set up the read stream for reading a file and read the first block.

Entry conditions:

B contains the length of the filename.

HL contains the address of the filename.

DE contains the address of a 2K buffer to use.

Exit conditions:

If the file was opened OK:

Carry true.

Zero false.

HL contains the address of a buffer containing the file header.

DE contains the data location (from the header).

BC contains the logical file length (from the header).

A contains the file type (from the header).

If the stream is in use:

Carry false.

Zero false.

In V1.1: A contains an error number (#0E).

In V1.0: A corrupt.

BC, DE and HL corrupt.

If the user hit escape:

Carry false.

Zero true.

In V1.1: A contains an error number(#00).

In V1.0: A corrupt

BC, DE and HL corrupt.

Always:

IX and other flags corrupt. All other registers preserved

Notes:

This routine can return two error numbers:

#00: The user hit escape.

#0E: The stream is already in use.

The 2K buffer (2048 bytes) supplied is used to store the contents of a block of the file when it is read from tape. It will remain in use until the file is closed by calling either CAS IN CLOSE or CAS IN ABANDON. The buffer may lie anywhere in memory, even underneath a ROM.

The filename passed is copied into the read stream descriptor. If it is longer than 16 characters then it is truncated to 16 characters. If it is shorter than 16 characters then it is padded with nulls (#00) to 16 characters. While the filename may contain any character, it is best to avoid nulls. Lower case ASCII letters (characters #61.. #7A) are converted to their upper case equivalents (characters #41.. #5A). The filename may lie anywhere in RAM, even underneath a ROM.

The filename is normally the name of the file that is to be read. However, a zero length filename (or one starting with a null) is treated specially. It is taken to mean read the next file on the tape.

When the file is opened for reading the first block of the file is read immediately. The address of the area where the header from this block is stored is passed back to the user so that information can be extracted from it. This area will lie in the central 32K of RAM. The user is not allowed to write to the header, only to read from it. The Cassette Manager uses some fields in the header for its own purposes and so these may differ from those read from the tape. The file type, logical length, entry point and all user fields will remain unchanged. (See section 8 for a description of the header.)

Related entries:

CAS IN ABANDON

CASINCHAR

CAS IN CLOSE

CAS INDIRECT

CAS IN OPEN (DISC)

CAS OUT OPEN

125 CAS IN OPEN (DISC)

#BC77

Open a file for input.

Action:

Set up the read stream for reading a file and read the header if there is one, otherwise create a fake header in store.

Entry conditions:

B contains the length of the filename.
HL contains the address of the filename.
DE contains the address of a 2K buffer to use.

Exit conditions:

If the file was opened OK:

Carry true.
Zero false.
HL contains the address of a buffer containing the file header.
DE contains the data location (from the header).
BC contains the logical file length (from the header).
A contains the file type from the header).

If the stream is already open:

Carry false.
Zero false.
A contains an error number (#0E).
BC, DE and HL corrupt.

If the open failed for any other reason:

Carry false.
Zero true.
A contains an error number.
BC, DE and HL corrupt.

Always:

IX and other flags corrupt. All other registers preserved.

Notes:

The 2K buffer (2048 bytes) supplied is used to store the contents of the file when it is read from disc. It will remain in use until the file is closed by calling either CAS IN CLOSE or CAS IN ABANDON. The buffer may lie anywhere in memory, even underneath a RUM.

The filename must conform to the AMSDOS conventions with no wild cards. The filename may lie anywhere in RAM, even underneath a RUM.

If the type part of the filename is omitted AMSDOS will attempt to open, in turn, a file with the following type parts ' . ', ' . **BAS**', ' . **BIN**', If none of these exist then the open will fail.

When the file is opened the first record of the file is read immediately. If this record contains a header then it is copied into store, otherwise a fake header is constructed in store. The address of the area where the header is stored is passed back to the user so that information can be extracted from it, This area will lie in the central 32K of RAM. The user is not allowed to write to the header, only to read from it. AMSDOS uses some fields in the header for its own purposes and so these may differ from those read from the disc. The file type, logical length, entry point and all user fields will remain unchanged.

Related entries:

CAS IN ABANDON (DISC)

CAS IN CHAR (DISC)

CAS IN CLOSE (DISC)

CAS IN DIRECT (DISC)

CAS IN OPEN

CAS OUT OPEN (DISC)

126: CAS IN CLOSE

#BC7A

Close the input file properly.

Action:

Mark the read stream as closed.

Entry conditions:

No conditions.

Exit conditions:

If the stream was closed OK:

Carry true.

A corrupt.

If the stream was not open:

Carry false.

In V1.1: A contains a error number(#0E).

In V1.0: A corrupt.

Always:

BC, DE, HL and other flags corrupt.

All other registers preserved.

Notes:

This routine can only return one error number:

#0E: The stream is not open.

This routine should be called to close a file after reading from it using either CAS IN CHAR or CAS IN DIRECT.

The user may reclaim the buffer passed to CAS IN OPEN after calling this routine.

Related entries:

CAS IN ABANDON

CAS IN CLOSE (DISC)

CAS IN OPEN

CAS OUT CLOSE

126 CAS IN CLOSE (DISC)

#BC7A

Close the input file properly.

Action:

Mark the read stream as closed.

Entry conditions:

No conditions.

Exit conditions:

If the stream was closed OK:

Carry true.

Zero false.

A corrupt.

If the stream is not open:

Carry false.

Zero false.

A contains an error number(#0E). If the close failed for any other reason:

Carry false.

Zero true.

A contains an error number.

Always:

BC, DE, HL and other flags corrupt.

All other registers preserved.

Notes:

This routine should be called to close a file after reading from it using either CAS IN CHAR or CAS IN DIRECT.

The user may reclaim the buffer passed to CAS IN OPEN after calling this routine.

The drive motor is turned off immediately after the input file is closed. This is done so that a loaded program which takes over the machine is not left with the motor running indefinitely.

Related entries:

CAS IN ABANDON (DISC)

CAS IN CLOSE

CAS IN OPEN (DISC)

CAS OUT CLOSE (DISC)

127: CAS IN ABANDON

#BC7D

Close the input file immediately.

Action:

Abandon reading from the read stream and close it.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

This routine is intended for use after an error or in similar circumstances.

The user may reclaim the buffer passed to CAS IN OPEN after calling this routine.

Related entries:

CAS IN ABANDON (DISC)
CAS IN CLOSE
CAS IN OPEN
CAS OUT ABANDON

127 CAS IN ABANDON (DISC)

#BC7D

Close the input file immediately.

Action:

Abandon reading from the read stream and close it.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

This routine is intended for use after an error or in similar circumstances.

The user may reclaim the buffer passed to CAS IN OPEN after calling this routine.

Related entries:

CAS IN ABANDON
CAS IN CLOSE (DISC)
CAS IN OPEN (DISC)
CAS OUT ABANDON (DISC)

128: CAS IN CHAR

#BC80

Read a character from the input file.

Action:

Read a character from the input stream. Fetches blocks from tape as required.

Entry conditions:

No conditions.

Exit conditions:

If the character was read OK:

Carry true.

Zero false.

A contains the character read from the file.

If the end of the file was found:

Carry false.

Zero false.

In V1.1: A contains an error number(#0E or #0F).

In V1.0: A corrupt If the user hit escape:

Carry false.

Zero true.

In V1.1: A contains an error number(#00).

In V1.0: A corrupt

Always:

IX and other flags corrupt.

All other registers preserved.

Notes:

This routine can return three error numbers:

#00: The user hit escape.

#0E: The stream is not open for reading or the user hit escape previously.

#0F: Have reached the end of the file.

Once the first character has been read from a file it can only be used for character by character access. It is not possible to switch to direct reading (by CAS IN DIRECT).

Related entries:

CAS IN CHAR (DISC)

CAS IN CLOSE

CAS IN DIRECT

CAS IN OPEN

CAS OUT CHAR

CAS RETURN

CAS TEST EOF

128 CAS IN CHAR (DISC)

#BC80

Read a character from the input file.

Action:

Read a character from the input stream.

Entry conditions:

No conditions.

Exit conditions:

If the character was read OK:

Carry true.

Zero false.

A contains the character read from the file.

If the end of the file was found, or stream not open as expected:

Carry false.

Zero false.

A contains an error number(#0E, #0F or #1A).

If failed for any other reason:

Carry false.

Zero true.

A contains an error number.

Always:

IX and other flags corrupt.

All other registers preserved.

Notes:

Once the first character has been read from a file the rest of the file may only be read character by character (using CAS IN CHAR), It is not possible to switch to direct reading (by CAS IN DIRECT).

The CP/M end of file character (#1A) is treated as end of file (carry false, zero false). However, it is possible to continue reading characters until the hard end of file. The error number returned is set to #1A for soft (CP/M) end of file and #0F for hard end of file. The action of spotting soft end of file is not performed by the equivalent cassette version of the routine and it will never return #1A when carry is false.

If a file containing binary data is read using this routine then it will be necessary to spot soft EOF and ignore it.

Related entries

CAS IN CHAR

CAS IN CLOSE (DISC)

CAS IN DIRECT (DISC)

CAS IN OPEN (DISC)

CAS OUT CHAR (DISC)

CAS RETURN (DISC)

CAS TEST EOF (DISC)

129: CAS IN DIRECT

#BC83

Read the input file into store.

Action:

Read the input file directly into store in one go rather than one character at a time.

Entry conditions:

HL contains the address to put the file (anywhere in RAM).

Exit conditions:

If the file was read OK:

Carry true.

Zero false.

HL contains the entry address (from the header).

A corrupt.

If the file was not open as expected:

Carry false.

Zero false.

In V1.1: A contains an error number(#OE).

In V1.0: A corrupt.

HL corrupt.

If the user hit escape:

Carry false.

Zero true.

In V1.1: A contains an error number(#00).

In V1.0: A corrupt.

HL corrupt.

Always:

BC, DE, IX and other flags corrupt.

All other registers preserved.

Notes:

This routine can return two error numbers:

#00: The user hit escape.

#0E: The stream is not open for reading directly or escape hit previously.

The read stream must be newly opened (by CAS IN OPEN). If the stream has been used for character access (by calling CAS IN CHAR) then it is not possible to directly read the file. Neither is it possible to directly read from the file more than once. This will merely corrupt the copy of the file read.

The buffer of data read when the stream was opened is copied to its correct position and the remainder of the file (if any) is also read into store.

Related entries:

CAS IN DIRECT (DISC)

CAS IN CHAR

CAS IN CLOSE

CAS IN OPEN

CAS OUT DIRECT

129 CAS IN DIRECT (DISC)

#BC83

Read the input file into store.

Action:

Read the input file directly into store in one go rather than one character at a time.

Entry conditions:

HL contains the address to put the file (anywhere in RAM).

Exit conditions:

If the file was read OK:

- Carry true.
- Zero false.
- HL contains the entry address (from the header).
- A corrupt.

If the stream is not open as expected:

- Carry false.
- Zero false.
- A contains an error number (#OE).
- HL corrupt.

If the read failed for any other reason:

- Carry false.
- Zero true.
- A contains an error number.
- HL corrupt.

Always:

- BC, DE, IX and other flags corrupt.
- All other registers preserved.

Notes:

The read stream must be newly opened (by CAS IN OPEN). If the stream has been used for character access (by calling CAS IN CHAR or CAS TEST EOF) then it is not possible to directly read the file. Neither is it possible to directly read from the file more than once. Any attempt to do so will corrupt the copy of the file read.)

If the file has a header then the number of bytes read is that recorded in the 24 bit file length field (bytes 64..66 of the disc file header. If there is no header the file is read until hard end of file.

The CP/M end of file character, #1A, is not treated as end of file.

Related entries:

CAS IN CHAR (DISC)
CAS IN CLOSE (DISC)
CAS IN DIRECT

CAS IN OPEN (DISC)
CAS OUT DIRECT (DISC)

130: CAS RETURN

#BC86

Put the last character read back.

Action:

Put the last character read by CAS IN CHAR back into the read buffer. The character will here-read next time CAS IN CHAR is called.

Entry conditions:

No conditions.

Exit conditions:

All registers and flags preserved.

Notes:

It is only possible to use this routine to return the last character that has been read by CAS IN CHAR. At least one character must have been read since:

- or the stream was opened
- or the last character was returned
- or the last test for the end of file was made.

Related entries:

CAS IN CHAR

CAS RETURN (DISC)

130 CAS RETURN (DISC)

#BC86

Put the last character read back.

Action:

Put the last character read by CAS IN CHAR back into the read buffer. The character will here-read next time CAS IN CHAR is called.

Entry conditions:

No conditions.

Exit conditions:

All registers and flags preserved.

Notes:

It is only possible to use this routine to return the last character that has been read by CAS IN CHAR. At least one character must have been read since:

- the stream was opened
- or the last character was returned
- or the last test for the end of file was made.

Related entries:

CAS IN CHAR (DISC)
CAS RETURN

131: CAS TEST EOF

#BC89

Have we reached the end of the input file yet?

Action:

Test if the end of the input file has been reached.

Entry conditions:

No conditions.

Exit conditions:

If the end of the file was not found:

- Carry true.
- Zero false.
- A corrupt.

If the end of the file was found:

- Carry false.
- Zero false.
- In V1.1: A contains an error number(#0E or #0F).
- In V1.0: A corrupt.

If the user hit escape:

- Carry false.
- Zero true.
- In V1.1: A contains a error number(#00).
- In V1.0: A corrupt.

Always:

- IX and other flags corrupt.
- All other registers preserved.

Notes:

This routine can return three error numbers:

- #00: The user hit escape.
- #0E: The stream is not open for reading characters or the user hit escape previously.
- #0F: Have reached the end of the file.

Calling this routine puts the stream into character input mode. It is not possible to use direct reading after calling this routine.

It is not possible to call CAS RETURN after this routine has been called. A character must be read first.

Related entries:

- CAS IN CHAR
- CAS TEST EOF (DISC)

131: CAS TEST EOF (DISC)

#BC89

Have we reached the end of the input file yet?

Action:

Test if the end of the input file has been reached.

Entry conditions:

No conditions.

Exit conditions:

If the end of the file was not found:

Carry true.

Zero false.

A corrupt.

If the end of the file was found or stream not open as expected:

Carry false.

Zero false.

A contains an error number (#0E, #0F or #1A).

If failed for any other reason:

Carry false.

Zero true.

A contains an error number.

Always:

IX and other flags corrupt. All other registers preserved.

Notes:

This routine will report end of file if either there are no more characters in the file or if the next character to be read is the CP/M end of file character, #1A.

Calling this routine puts the stream into character input mode. It is not possible to use direct reading after calling this routine.

It is not possible to call CAS RETURN after this routine has been called. A character must be read first.

Related entries:

CAS IN CHAR (DISC)

CAS TEST EOF

132: CAS OUT OPEN

#BC8C

Open a file for output.

Action:

Set up the write stream for output.

Entry conditions:

B contains the length of the filename.

HL contains the address of the filename.

DE contains the address of a 2K buffer to use.

Exit conditions:

If the user hit escape:

Carry false.

Zero true.

In V1.1: A contains an error number(#00).

In V1.0: A corrupt.

HL corrupt.

If the stream is in use already:

Carry false.

Zero false.

In V1.i: A contains an error number(#OE),

In Vi.0: A corrupt.

HL corrupt.

If the file was opened OK:

Carry true.

Zero false.

A corrupt.

HL contains the address of a buffer containing the header that will be written to each file block.

Always:

BC, DE, IX and other flags corrupt.

All other registers preserved.

Notes:

This routine can only return two error numbers.

#00: The user hit escape.

#0E: The stream is already open.

When writing files character by character the 2K buffer (2048 bytes) supplied is used to store the contents of a block of the file before it is written to tape. It will remain in use until the file is closed by calling either CAS OUT CLOSE or CAS OUT ABANDON. The buffer may reside anywhere in memory - even underneath a ROM

When the stream is opened for writing, a header is set up which will be written at the start of each block of the file. Many of the fields in the header are set by the Cassette Manager but the remainder are available for use by the user. The address of this header is passed to the user so that information can be stored in it. The user may write to the file type, logical length, entry point and all user fields. The user is not allowed to write to any other field in the header. The user settable fields are all zeroized initially, with the exception of the file type which is set to unprotected ASCII version 1. (See section 8.4 for a description of the header.)

The filename passed is copied into the write stream descriptor. If it is longer than 16 characters then it is truncated to 16 characters. If it is shorter than 16 characters then it is padded with nulls (#00) to 16 characters. While the filename may contain any character, it is best to avoid nulls. Lower case ASCII letters (characters #61..#7A) are converted to their upper case equivalents (characters #41..#5A). The filename may lie anywhere in RAM, even underneath a ROM.

Related entries:

CAS IN OPEN
CAS OUT ABANDON
CAS OUT CHAR
CAS OUT CLOSE
CAS OUT DIRECT
CAS OUT OPEN (DISC)

132 CAS OUT OPEN (DISC)

#BC8C

Open a file for output.

Action:

Set up the write stream for output.

Entry conditions:

B contains the length of the filename.

HL contains the address of the filename.

DE contains the address of a 2K buffer to use.

Exit conditions:

If the file was opened OK:

Carry true.

Zero false.

HL contains the address of a buffer containing the header.

A corrupt.

If the stream is open already:

Carry false.

Zero false.

A contains an error number (#0E).

HL corrupt.

If the open failed for any other reason:

Carry false.

Zero true.

A contains an error number.

HL corrupt.

Always:

BC, DE, IX and other flags corrupt.

All other registers preserved.

Notes:

When characters are output to the file using CAS OUT CHAR the 2K buffer supplied is used by AMSDOS to buffer the output. It will remain in use until the file is closed by calling either CAS OUT CLOSE or CAS OUT ABANDON. The buffer may reside anywhere in memory - even underneath a ROM.

The filename passed must conform to AMSDOS conventions with no wild cards. It is copied into the write stream header. The filename may lie anywhere in RAM - even underneath a ROM.

The file is opened with a type part of ' . \$\$\$' regardless of the type part supplied. Any existing file with the same name and a type part of ' . \$\$\$' is deleted. The file is renamed to its supplied name when CAS OUT CLOSE is called.

When the stream is opened a header is set up. Many of the fields in the header are set by AMSDOS but the remainder are available for use by the user. The address of this header is passed to the user so that information can be stored in it. The user may write to the file type, logical length, entry point and all user fields. The user is not allowed to write to any other field in the header. The user settable fields are all zeroized initially, with the exception of the file type which is set to unprotected ASCII version 1.

The header type field must be written to before CAS OUT CHAR or CAS OUT DIRECT is called. The type field must not be altered after calling either of these routines. If the file type is set to any type other than unprotected ASCII then space will be reserved for the header which will be written when the file is closed.

Related entries:

CAS IN OPEN (DISC)
CAS OUT ABANDON (DISC)
CAS OUT CHAR (DISC)
CAS OUT CLOSE (DISC)
CAS OUT DIRECT (DISC)
CAS OUT OPEN

133: CAS OUT CLOSE

#BC8F

Close the output file properly.

Action:

Mark the write stream as closed and write the last buffer of data to tape.

Entry conditions:

No conditions.

Exit conditions:

If the stream was closed OK:

Carry true.

Zero false.

A corrupt.

If the stream was not open:

Carry false.

Zero false.

In V1.1: A contains an error number(#OE).

In V1.0: A corrupt. If the user hit escape:

Carry false.

Zero true.

In V1.1: A contains an error number(#00).

In V1.0: A corrupt.

Always:

BC, DE, HL, IX and other flags corrupt.

All other registers preserved.

Notes:

This routine can return two error numbers:

#00: The user hit escape.

#0E: The stream is not open.

It is necessary to call this routine after using CAS OUT CHAR or CAS OUT DIRECT to cause the last block of data to be written to the tape. If the block is zero bytes long (nothing has been written to the file) then nothing is written to tape.

If the writing is to be abandoned then CAS OUT ABANDON should be called as this does not write the last block of data to the tape.

If the user hits escape during the writing of the last block then the file is left open and is not closed.

The user may reclaim the buffer passed to CAS OUT OPEN after calling this routine.

Related entries:

CAS IN CLOSE

CAS OUT ABANDON

CAS OUT CLOSE (DISC)

CAS OUT OPEN

133 CAS OUT CLOSE (DISC)

#BC8F

Close the output file properly.

Action:

Mark the write stream as closed and give it its correct name.

Entry conditions:

No conditions.

Exit conditions:

If the stream was closed OK:

Carry true.
Zero false.
A corrupt.

If the stream is not open:

Carry false.
Zero false.
A contains an error number (# OE).

If the close failed for any other reason:

Carry false.
Zero true.
A contains an error number.

Always:

BC, DE, HL, IX and other flags corrupt.
All other registers preserved.

Notes:

It is necessary to call this routine after using CAS OUT CHAR or CAS OUT DIRECT to ensure that all the data is written to the disc, to write the header to the start of the file and to give the file its true name.

If no data has been written to the file then it is abandoned and nothing is written to disc. This is for compatibility with cassette routines.

When the file was opened it was given the type part of ' . \$\$\$'. This routine will rename the file to its true name and rename any existing version to have a ' . **BAK**' type part. This ensures that any previous version of the file is automatically kept as a backup. Any existing ' . **BAK**' version is deleted. If, when the file was opened, the caller did not specify a type part then AMSDOS will use the type part ' . **BAS**' for BASIC files, ' . **BIN**' for binary files and ' . ' for all other files, as specified by the file type field in the header.

If the actual length of the file is not a multiple of 128 bytes (a CP/M record) then a CP/M end of file character, #1A, is added to the file. This additional character is not recorded in the length of the file.

If writing is to be abandoned then CAS OUT ABANDON should be called as this does not write any more data to disc.

The user may reclaim the buffer passed to CAS OUT OPEN after calling this routine.

Related entries:

CAS IN CLOSE (DISC)
CAS OUT ABANDON (DISC)
CAS OUT CLOSE
CAS OUT OPEN (DISC)

134: CAS OUT ABANDON

#BC92

Close the output file immediately.

Action:

Abandon the output file and mark the write stream closed. Any unwritten data is discarded and not written to tape.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

This routine is intended for use after an error or in similar circumstances.

Related entries:

CAS IN ABANDON
CAS OUT ABANDON (DISC)
CAS OUT CLOSE
CAS OUT OPEN

134 CAS OUT ABANDON (DISC)

#BC92

Close the output file immediately.

Action:

Abandon the output file and mark the write stream closed. Any unwritten data is discarded and not written to disc.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.

All other registers preserved.

Notes:

This routine is intended for use after an error or in similar circumstances.

If more than one 16K physical extent has already been written to disc then the file will appear in the disc directory with a type part of '• \$\$\$'. Otherwise the file will disappear. This is because each 16K of a file requires a directory entry, A directory entry is not written to disc until the 16K has been written or the file is closed (CAS OUT CLOSE).

Related entries:

CAS IN ABANDON (DISC)

CAS OUT ABANDON

CAS OUT CLOSE (DISC)

CAS OUT OPEN (DISC)

135: CAS OUT CHAR

#BC95

Write a character to the output file.

Action:

Add a character to the buffer for the write stream. If the buffer is already full then it is written to tape before the new character is inserted.

Entry conditions:

A contains the character to write.

Exit conditions:

If the character was written OK:

Carry true.

Zero false.

A corrupt.

If the file was not open as expected:

Carry false.

Zero false.

In V1.1: A contains an error number(#OE).

In V1.0: A corrupt. If the user hit escape:

Carry false.

Zero true.

In V1.1: A contains an error number(#00).

In V1.0: A corrupt.

Always:

A, IX and other flags corrupt.

All other registers preserved.

Notes:

This routine can return two error numbers:

#00: The user hit escape.

#0E: The stream is not open for writing characters or the user hit escape previously.

If this routine returns the file not open as expected condition then either the user has hit escape previously or the file has been written using CAS OUT DIRECT. In either case, or if escape is pressed, the character sent will be discarded.

It is necessary to call CAS OUT CLOSE after sending all the characters to the file to ensure that the last block of the file is written to the tape.

Once this routine has been called it is not possible to switch to directly writing the file.

Related entries:

CAS IN CHAR

CAS OUT CHAR (DISC)

CAS OUT CLOSE

CAS OUT DIRECT

CAS OUT OPEN

135 CAS OUT CHAR (DISC)

#BC95

Write a character to the output file.

Action:

Add a character to the buffer for the write stream. If the buffer is already full then it is written to disc before the new character is inserted.

Entry conditions:

A contains the character to write.

Exit conditions:

If the character was written OK:

Carry true.

Zero false.

A corrupt.

If the stream is not open as expected:

Carry false.

Zero false.

A contains an error number (#0E).

If failed for any other reason:

Carry false.

Zero true.

A contains a error number.

Always:

IX and other flags corrupt.

All other registers preserved.

Notes:

It is necessary to call CAS OUT CLOSE after sending all the characters to the file to ensure that the file is correctly written to disc.

Once this routine has been called it is not possible to switch to directly writing the file (CAS OUT DIRECT).

Related entries:

CAS IN CHAR (DISC)

CAS OUT CHAR

CAS OUT CLOSE

CAS OUT DIRECT (DISC)

CAS OUT OPEN (DISC)

136: CAS OUT DIRECT

#BC98

Write the output file directly from store.

Action:

Write the contents of store directly out to the output file.

Entry conditions:

HL contains the address of the data to write.

DE contains the length of the data to write.

BC contains the entry address (to go into the header).

A contains the file type (to go into the header).

Exit conditions:

If the file was written OK:

Carry true.

Zero false.

A corrupt.

If the file was not open as expected:

Carry false.

Zero false.

In V1.1: A contains in error number(#OE).

In V1.0: A corrupt. If the user hit escape:

Carry false.

Zero true.

In V1.1: A contains a error number(#00).

In V1.0: A corrupt.

Always:

BC, DE, HL, IX and other flags corrupt.

All other registers preserved.

Notes:

This routine can return two error numbers:

#00: The user hit escape.

#0E: The stream is not newly opened.

After writing the file it must be closed using CAS OUT CLOSE to ensure that the last block of the file is written to tape.

It is not possible to change the method for writing files from character output (using CAS OUT CHAR) to direct output (using CAS OUT DIRECT) or vice versa once the method has been chosen. Nor is it possible to directly write a file in two or more parts by calling CAS OUT DIRECT more than once - this will write corrupt data. Attempting to break these rules will result in a file not open as expected error.

Related entries:

CAS IN DIRECT
CAS OUT CLOSE

CAS OUT DIRECT (DISC)
CAS OUT OPEN

136 CAS OUT DIRECT (DISC)

#BC98

Write the output file directly from store.

Action:

Write the contents of store directly out to the output file.

Entry conditions:

HL contains the address of the data to write (to go into the header).

DE contains the length of the data to write (to go into the header).

BC contains the entry address (to go into the header).

A contains the file type (to go into the header).

Exit conditions:

If the file was written OK:

Carry true.

Zero false.

A corrupt.

If the stream is not open as expected:

Carry false.

Zero false.

A contains an error number (#0E).

If failed for any other reason:

Carry false.

Zero true.

A contains an error number.

Always:

BC, DE, HL, IX and other flags corrupt.

All other registers preserved.

Notes:

After writing the file it must be closed using CAS OUT CLOSE to ensure that the file is written to disc.

It is not possible to change the method for writing files from character output (using CAS OUT CHAR) to direct output (using CAS OUT DIRECT) or vice versa once the method has been chosen. Nor is it possible to directly write a file in two or more parts by calling CAS OUT DIRECT more than once - this will write corrupt data.

Related entries:

CAS IN DIRECT (DISC)

CAS OUT CLOSE (DISC)

CAS OUT DIRECT

CAS OUT OPEN (DISC)

137: CAS CATALOG

#BC9B

Generate a catalogue from the tape.

Action:

Read file blocks to check their validity and print information about them on the screen.

Entry conditions:

DE contains the address of a 2K buffer to use.

Exit conditions:

If the cataloguing went OK:

- Carry true.
- Zero false.
- A corrupt.

If the read stream was in use:

- Carry false.
- Zero false.
- In V1.1: A contains an error number (#0E).
- In V1.0: A corrupt.

Always:

- BC, DE, HL, IX and other flags corrupt.
- All other registers preserved.

Notes:

This routine can only return one error number:

- #0E: The stream is already in use.

This routine uses the read stream and so the stream must be closed when it is called. The read stream remains closed when this routine exits. The write stream is unaffected by this routine.

The prompt messages are turned on (see CAS NOISY) by this routine.

When cataloguing the Cassette Manager reads a header record, prints information from it and then reads the data record. This cycle repeats until the user hits the escape key. The information printed is as follows:

FILENAME block N T Ok

FILENAME is the name of the file on the tape, or 'Unnamed file' if the filename starts with a null (character #00).

N is the number of the block. Block 1 is normally the first block in a file.

T is a representation of the file type of the file. It is formed by adding #24 (the character '\$') to the file type byte masked with #0F (to remove the version number field). The standard file types are thus:

\$	a BASIC program file
%	a protected BASIC program file
*	an ASCII text file (default file type)
&	a binary file
'	a protected binary file

Other file types are possible but will not have been written by the BASIC in the on-board ROM. See section 8.4 for a description of the file type byte.

Ok is printed after the end of the data record. This shows that the data was read without errors and also serves to indicate the end of the data on tape (to help avoid over-recording a tape file).

Related entries:
CAS CATALOG (DISC)
CAS NOISY

137 CAS CATALOG (DISC)

#BC9B

Display the disc directory.

Action:

Display the disc directory for the current drive and current user. The directory is sorted into alphabetical order and displayed in as many columns as will fit in the current text window (stream #0). The size in Kbytes is shown along side each file. The total amount of free space on the disc is also shown.

Entry conditions:

DE contains the address of a 2K buffer to use.

Exit conditions:

If the cataloging went OK:

Carry true.
Zero false.
A corrupt.

If failed for any reason:

Carry false.
Zero true.
A contains an error number.

Always:

BC, DE, HL, IX and other flags corrupt.
All other registers preserved.

Notes:

Files marked SYS are not shown.

Files marked R/O are shown with a '*' after the file name.

Unlike the cassette version of this routine, the disc input stream is not required. (Note: BASIC abandons both the input and output streams when generating the catalogue.)

Related entries:

CAS CATALOG
| DIR

138: CAS WRITE

#BC9E

Write a record to tape.

Action:

Write a record to the cassette. This routine is used by the higher level routines (CAS OUT CHAR, CAS OUT DIRECT and CAS OUT CLOSE) to write the header and data records that make up a tape file.

Entry conditions:

HL contains the address of the data to write.

DE contains the length of the data to write.

A contains the sync character to write at the end of the leader.

Exit conditions:

If the record was written OK:

Carry true.

A corrupt.

If an error occurred or the user hit escape:

Carry false.

A contains an error code.

Always:

BC, DE, HL, IX corrupt.

All other registers preserved.

Notes:

A data length of 0 passed to this routine is taken to mean 65536 bytes and all of memory will be written to tape. (This is unlikely to be useful).

The data to be written may lie anywhere in RAM, even underneath a ROM.

The sync character is used to distinguish header records (sync is #2C) from data records (sync is #16). Other sync characters could be used but the resulting record would require special action to be taken to read it.

The error codes returned by this routine are:

- | | | |
|---|---------|---|
| 0 | Break | The user hit the escape key. |
| 1 | Overrun | The Cassette Manager was unable to get back to writing a bit fast enough. |

Because reading and writing the tape requires stringent timing considerations interrupts are disabled whilst the tape is being written (potentially a period of over 5 minutes). It would be unpleasant to have the sound chip making a noise for all this time so the Sound Manager is shut down (SOUND RESET). When writing to the tape has finished interrupts are re-enabled.

The cassette motor is started by this routine (in case it is not already on) and restored to its previous state when writing is completed.

Related entries:

CAS CHECK

CAS READ

139: CAS READ

#BCA1

Read a record from tape.

Action:

Read a whole or part record from the cassette. This routine is used by the higher level routines (CAS IN CHAR, CAS IN DIRECT and CAS CATALOG amongst others) to read the header and data records that make up a file.

Entry conditions:

HL contains the address to put the data read.

DE contains the length of the data to read.

A contains the sync character expected at the end of the leader.

Exit conditions:

If record was read OK:

Carry true.

A corrupt.

If an error occurred or the user hit escape:

Carry false.

A contains an error code.

Always:

BC, DE, HL, IX and other flags corrupt.

All other registers preserved.

Notes:

A data length of 0 passed to this routine is taken to mean 65536 bytes. (This is not useful).

It is not necessary to read the whole of a record from tape. If the length passed is less than the actual length of the record then only that number of bytes will be read. Trying to read more bytes from a record than were written will produce an error, usually an overflow error (see below).

The sync character is used to distinguish header records (sync is #2C) from data records (sync is #16). Other sync characters could be used if the record was written that way.

The error codes returned by this routine are:

- 0 Break The user hit the escape key.
- 1 Overflow The Cassette Manager found a bit that was too long to read.
- 2 CRC A CRC failure was detected.

The cassette motor is started by this routine (in case it is not already on) and restored to its previous state when reading is completed.

Because reading the tape requires stringent timing considerations, interrupts are disabled whilst the tape is being read (potentially a period of over 5 minutes). It would be unpleasant to have the sound chip making a noise for all this time so the Sound Manager is shut down (SOUND RESET). When reading from the tape has finished interrupts are re-enabled.

Related entries:

CAS CHECK
CAS WRITE

140: CAS CHECK

#BCA4

Compare a record on tape with the contents of store.

Action:

Check that a tape record contains a correct version of the data supplied. This routine is intended to be used after writing records to check that they were written correctly.

Entry conditions:

HL contains the address of the data to check.

DE contains the length of the data to check.

A contains the sync character expected at the end of the leader.

Exit conditions:

If the record checked OK:

Carry true.

A corrupt.

If an error occurred or the user hit escape:

Carry false.

A contains an error code.

Always:

BC, DE, HL, IX and other flags corrupt.

All other registers preserved.

Notes:

A data length of 0 passed to this routine is taken to mean 65536 bytes. (This is bound to produce a check failure).

It is not necessary to check the whole of a record on tape. If the length passed is less than the actual length of the record then only that number of bytes will be checked. Trying to check more bytes in a record than were written will produce an error of some sort (see below).

The data to be checked may lie anywhere in RAM, even underneath a ROM.

The sync character is used to distinguish header records (sync is #2C) from data records (sync is #16). Other sync characters could be used.

The error codes returned by this routine are:

- 0 Break The user hit the escape key.
- 1 Overrun The Cassette Manager found a bit that was too long to read.
- 2 CRC A CRC failure was detected.
- 3 Different The data read from tape did not agree with that in memory.

The cassette motor is started by this routine (in case it is not already on) and restored to its previous state when checking is completed.

Because reading from the tape requires stringent timing considerations, interrupts are disabled whilst the tape is being checked (potentially a period of over 5 minutes). It would be unpleasant to have the sound chip making a noise for all this time so the Sound Manager is shut down (SOUND RESET). When checking has finished interrupts are re-enabled.

Related entries:

CAS READ

CAS WRITE

141: SOUND RESET

#BCA7

Reset the Sound Manager.

Action:

Re-initialize the Sound Manager - shut the sound chip up and clear all queues.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The sound queues are cleared.
Any current sound is stopped.
The sound generator chip is silenced.

This routine enables interrupts.

Related entries:

SOUND HOLD

142: SOUND QUEUE

#BCAA

Add a sound to a sound queue.

Action:

Try to add a sound to the sound queue of one or more channels. If the sound queue of any of the channels is full then no sound will be issued to any channel.

Entry conditions:

HL contains the address of a sound program which must lie in the central 32K of RAM.

Exit conditions:

If the sound was added to the queue(s):

Carry true.
HL corrupt.

If at least one queue was full:

Carry false.
HL preserved.

Always:

A, BC, DE, IX and other flags corrupt.
All other registers preserved.

Notes:

The sound program is laid out as follows:

Byte 0:	Channels to use and rendezvous requirements.
Byte 1:	Amplitude envelope to use.
Byte 2:	Tone envelope to use.
Bytes 3...4:	Tone period.
Byte 5:	Noise period.
Byte 6:	Initial amplitude.
Bytes 7...8:	Duration or envelope repeat count.

All values in the sound program are masked into the appropriate range before being used.

The channels to issue the sound on are encoded into byte 0 as follows:

Bit 0:	Issue on channel A.
Bit 1:	Issue on channel B.
Bit 2:	Issue on channel C.

The rendezvous requirements are encoded into byte 0 as follows:

- Bit 3: Rendezvous with channel A.
- Bit 4: Rendezvous with channel B.
- Bit 5: Rendezvous with channel C.
- Bit 6: Hold until released.
- Bit 7: Flush queue

A channel will ignore an order to rendezvous with itself. Sounds issued on multiple channels implicitly rendezvous with each other. Sounds that are ordered to rendezvous will be issued to the sound generator starting at the same time.

Setting the hold bit prevents the sound from running until it is released by calling SOUND RELEASE (or a routine having a similar effect). Setting the flush bit will empty the queue and abandon any currently active sound thus allowing the new sound to start immediately.

The amplitude envelope is in the range 0..15. Envelopes 1..15 are the amplitude envelopes that can be set using SOUND AMPL ENVELOPE. Envelope 0 means use no amplitude envelope, simply hold the initial amplitude for 2 seconds or the duration specified.

The tone envelope is in the range 0..15. Envelopes 1..15 are the tone envelopes that can be set using SOUND TONE ENVELOPE. Envelope 0 means use no tone envelope, simply hold the initial tone.

A tone period of 0 means do not generate any tone. Tone periods in the range 1..4095 specify the period of the tone in 8 microsecond units.

The noise period is in the range 0..31. Noise periods 1..31 specify the period of the noise component of a sound. A noise period of 0 means use no noise.

The initial amplitude is in the range 0..15. Amplitude 0 being no initial sound, amplitude 15 being maximum volume.

Bytes 7 and 8 store the sound time. If this is zero then the amplitude envelope is obeyed once. If the sound time is negative then the amplitude envelope is obeyed minus the sound time number of times (i.e. 1..32768 times). If the sound time is positive but not zero then it is taken to be the duration of the sound in 1/100s of a second.

If a duration is specified when an amplitude envelope is in use then the duration given sets the length of the sound. If the duration is shorter than the envelope then the envelope is truncated. If the duration is longer than the envelope then the final amplitude of the envelope is sustained until the duration expires. Tone envelopes are treated in much the same way as amplitude envelopes except that they never specify the length of the sound.

The sound event that is run when a sound queue has a free slot is disarmed on the channels specified in this command.

All sounds currently held by SOUND HOLD are automatically released when this routine is called. Also, the sound queue event is disarmed (see SOUND ARM EVENT). SOUND QUEUE may enable interrupts.

Related entries:
SOUND ARM EVENT
SOUND CHECK

143: SOUND CHECK

#BCAD

Ask if there is space in a sound queue.

Action:

Ask the status of a sound channel. The status includes the number of free spaces in the sound queue and whether the channel is held.

Entry conditions:

A contains the bit for the channel to test.

Exit conditions:

A contains the channel status.
BC, DE, HL and flags corrupt.
All other registers preserved.

Notes:

The channel to ask the status of is encoded as follows:

- Bit 0: Ask about channel A.
- Bit 1: Ask about channel B.
- Bit 2: Ask about channel C.

If more than one bit is set then the status of only one channel is returned. The channels are tested in the order given above.

The status returned is encoded as follows:

- Bits 0..2: Contain the number of free slots in the channel's sound queue.
- Bit 3: The channel is awaiting a rendezvous with channel A.
- Bit 4: The channel is awaiting a rendezvous with channel B.
- Bit 5: The channel is awaiting a rendezvous with channel C.
- Bit 6: The channel is held.
- Bit 7: The channel is active (producing a sound).

Calling this routine disarms the sound queue event that occurs when the queue has a free slot for the channel returned (see SOUND ARM EVENT).
This routine may enable interrupts.

Related entries:

SOUND ARM EVENT
SOUND QUEUE

144: SOUND ARM EVENT

#BCB0

Set up an event to be run when a sound queue becomes empty.

Action:

Arm the sound event to be run when a free slot occurs in a channel's sound queue.

Entry conditions:

A contains the bit for the channel to arm.
HL contains the address of an event block.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The channel for which to arm the event is encoded as follows:

Bit 0: Arm channel A.
Bit 1: Arm channel B.
Bit 2: Arm channel C.

If more than one bit is set then only one channel is armed. The channels are armed in the order given above.

The event block passed must be initialized (by KL INIT EVENT).

The event will be 'kicked' when a free slot occurs in the queue. If there is a free slot in the queue when this routine is called then the event will be 'kicked' immediately.

The sound event is disarmed automatically when SOUND QUEUE or SOUND CHECK is called. It is also disarmed when the event is run. Thus, the event routine will need to rearm the sound event to keep it running continuously.

This routine may enable interrupts.

Related entries:

KL INIT EVENT
SOUND CHECK
SOUND QUEUE

145: SOUND RELEASE

#BCB3

Allow sounds which are individually held to start.

Action:

Release held sounds on a number of channels. This allows sounds that were marked with a hold bit when they were set up by SOUND QUEUE to start (other factors willing).

Entry conditions:

A contains bits for the channels to release.

Exit conditions:

AF, BC, DE, HL and IX corrupt.
All other registers preserved.

Notes:

The channels to release are encoded as follows:

- Bit 0: Release channel A.
- Bit 1: Release channel B.
- Bit 2: Release channel C.

All channels that are specified are released.

All sounds currently held by SOUND HOLD are automatically released.

This routine may enable interrupts.

Related entries:

SOUND QUEUE

146: SOUND HOLD

#BCB6

Stop all sounds in midflight.

Action:

This stops all sounds immediately. The sounds can be started again by calling SOUND CONTINUE.

Entry conditions:

No conditions.

Exit conditions:

If a sound was active:

Carry true.

If no sound was active:

Carry false.

Always:

A, BC, HL and other flags corrupt.
All other registers preserved.

Notes:

Sounds that are held by this routine are automatically restarted when SOUND QUEUE or SOUND RELEASE are called as well as when SOUND CONTINUE itself is called.

The sound is stopped by halting the execution of sound and tone envelopes and setting the sound chip volume to zero for all channels. When the sound is restarted it will continue from as near where it was stopped as is possible.

This routine enables interrupts.

Related entries:

SOUND CONTINUE
SOUND RESET

147: SOUND CONTINUE

#BCB9

Restart sounds after they have all been held.

Action:

Allow sounds that have been held by calling SOUND HOLD to continue.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and IX corrupt.
All other registers preserved.

Notes:

If no sounds are held then no action is taken.

This routine may enable interrupts.

Related entries:

SOUND HOLD
SOUND RELEASE

148: SOUND AMPL ENVELOPE

#BCBC

Set up an amplitude envelope.

Action:

Set up one of the 15 user programmable amplitude (volume) envelopes.

Entry conditions:

A contains an envelope number.

HL contains the address of an amplitude data block.

Exit conditions:

If envelope has been setup OK:

Carry true.

HL contains the address of the data block + 16.

A and BC corrupt.

If envelope number is invalid:

Carry false.

A, B, and HL preserved.

Always:

DE and other flags corrupt.

All other registers preserved.

Notes:

The envelope to set up is specified by a number in the range 1.. 15. No envelope is set up if a number outside this range is passed.

The amplitude data block is copied into the amplitude envelope. The data block may lie in ROM or in RAM. It may not lie in RAM hidden underneath a ROM.

The amplitude data block has the following layout:

Byte 0: Count of sections in the envelope.

Bytes 1..3: First section of the envelope.

Bytes 4..6: Second section of the envelope.

Bytes 7..9: Third section of the envelope.

Bytes 10..12: Fourth section of the envelope.

Bytes 13..15: Fifth section of the envelope.

The first byte of the amplitude data block specifies the number of sections used in the envelope. Sections not used need not be set up. An envelope using no sections has a special meaning - hold a constant volume lasting for 2 seconds.

The number of sections to use is not checked, if a number outside the range 0..5 is supplied then this will have unpredictable effects. This should be avoided.

Each section of the amplitude data block can specify either a hardware or a software envelope. This is indicated by the first byte of the section.

A software envelope section is laid out as follows:

Byte 0: Step count.
Byte 1: Step size.
Byte 2: Pause time.

The fact that this is a software envelope section rather than a hardware envelope section is indicated by byte 0 not having bit 7 set.

If the step count is in the range 1..127 then the step size is added to the volume that number of times with a wait equal to the pause time in 1/100s of a second after each addition.

If the step count is 0 the the step size is taken to be an absolute volume setting. A single wait of the pause time in 1/100s of a second is made.

After calculating the new volume this is masked with #0F to make sure it is legal. Thus, all arithmetic on the volume is carried out modulo 16.

A pause time of 0 is taken to mean 256 1/100s of a second.

A hardware envelope section is laid out as follows:

Byte 0: Envelope shape.
Byte 1..2: Envelope period.

The fact that this is a hardware envelope section rather than a software envelope section is indicated by byte 0 having bit 7 set.

The envelope shape (masked with #7F) is sent to register 13 of the sound generator. This sets the shape of the hardware envelope and whether it repeats (see Appendix IX for details).

The envelope period is sent to registers 11 and 12 of the sound generator. These set the the length of the hardware envelope (see Appendix IX for details).

The section after a hardware section should be a pause long enough to allow the hardware envelope to operate. A pause can be constructed using a software envelope with a step size of 0 and with the repeat count and pause time juggled to give the right total time.

There is no protection against changing an envelope whilst it is in use. This could have unpredictable effects and should be avoided.

The length of the sound can either be determined by the duration supplied when the sound is queued or by the envelope terminating (see SOUND QUEUE). If a duration is set that is shorter than the envelope then the envelope is truncated. If the duration is longer than the envelope then the final volume is sustained until the duration expires.

Related entries:

SOUND A ADDRESS

SOUND TONE ENVELOPE

149: SOUND TONE ENVELOPE

#BCBF

Set up a tone envelope.

Action:

Set up one of the 15 user programmable tone envelopes.

Entry conditions:

A contains an envelope number.

HL contains the address of a tone data block.

Exit conditions:

If the envelope has been set up OK:

Carry true.

HL contains the address of the data block + 16.

A and BC corrupt.

If the envelope number is invalid:

Carry false.

A, BC and HL preserved.

Always:

DE and other flags corrupt.

All other registers preserved.

Notes:

The envelope to set up is specified by a number in the range 1.15. No envelope is set up if a number outside this range is passed.

The tone data block is copied into the tone envelope. The data block may lie in ROM or in RAM. It may not lie in RAM hidden underneath a ROM.

The tone data block has the following layout:

Byte 0: Count of sections in the envelope.

Bytes 1..3: First section of the envelope.

Bytes 4..6: Second section of the envelope.

Bytes 7..9: Third section of the envelope.

Bytes 10.. 12: Fourth section of the envelope.

Bytes 13.. 15: Fifth section of the envelope.

The first byte of the tone data block (masked with #7F) specifies the number of sections used in the envelope. Sections not used need not be set up. An envelope using no sections will not alter the tone (i.e. no enveloping). The number of sections to use is not checked, if a number outside the range 0.5 is supplied then this will have unpredictable effects. This should be avoided.

The top bit, bit 7, of the count is used to indicate a repeating envelope. If this bit is set then when the last section of the envelope finishes the first will be used again.

Each section of the tone data block is laid out as follows:

Byte 0: Step count.
Byte 1: Step size.
Byte 2: Pause time.

If the step count lies in the range #00.. #EF then the section is a relative section. The step size is sign extended (bit 7 is copied to bits 8.. 15) and is added to the current tone period the number of times specified by the step count. After each addition a wait of the pause time in 1/100s of a second is made. The sound chip only uses the lower 12 bits of the tone period so all arithmetic is carried out modulo #1000.

A step count of 0 is taken to mean 1 step whilst a pause time of 0 is taken to mean 256 1/100s of a second.

If the step count lies in the range #F0.. #FF then the section is an absolute section. The least significant four bits of the step count are taken to be the most significant byte of the tone period and the step size is taken to be the least significant byte. This tone period is set immediately and is followed by a pause whose length is set by the pause time in 1/100s of a second.

There is no protection against changing an envelope whilst it is in use. This could have unpredictable effects and should be avoided.

If the tone envelope finishes before the end of the sound (as set when the sound was queued) then the final tone is held. i.e. The tone envelope does not affect the length of the sound.

Related entries:

SOUND AMPL ENVELOPE

SOUND T ADDRESS

150: SOUND A ADDRESS

#BCC2

Get the address of an amplitude envelope.

Action:

Ask where the data area for an amplitude envelope is stored.

Entry conditions:

A contains an envelope number.

Exit conditions:

If the envelope was found OK:

Carry true.

HL contains the address of the amplitude envelope.

BC contains the length of an envelope (16 bytes).

If the envelope number was invalid:

Carry false.

HL corrupt.

BC preserved.

Always:

A and other flags corrupt.

All other registers preserved.

Notes:

The envelope number must lie in the range 1.. 15.

The amplitude envelope is laid out as described in SOUND AMPL ENVELOPE.

Related entries:

SOUND AMPL ENVELOPE

SOUND T ADDRESS

151: SOUND T ADDRESS #BCC5

Get the address of a tone envelope.

Action:

Ask where the data area for a tone envelope is stored.

Entry conditions:

A contains an envelope number.

Exit conditions:

If the envelope was found OK:

Carry true.

HL contains the address of the tone envelope.

BC contains the length of an envelope (16 bytes).

If the envelope number was invalid:

Carry false.

HL corrupt.

BC preserved.

Always:

A and other flags corrupt.

All other registers preserved.

Notes:

The envelope number must lie in the range 1.. 15.

The tone envelope is laid out as described in SOUND TONE ENVELOPE.

Related entries:

SOUND A ADDRESS

SOUND TONE ENVELOPE

The values returned by KL CHOKE OFF are used by MC BOOT PROGRAM if the program load fails.

This information is included for the reader's interest. MC BOOT PROGRAM is the recommended means of loading and entering a RAM foreground program. MC START PROGRAM is the recommended means of entering a ROM foreground program, or a RAM foreground program which has already been loaded.

KL CHOKE OFF disables interrupts.

Related entries:

MC BOOT PROGRAM

MC START PROGRAM

153: KL ROM WALK

#BCCB

Find and initialize all background ROMs.

Action:

Background ROMs provide support for expansion hardware or augment the software facilities of the machine. If the facilities provided by the background ROMs are to be available, the foreground program must initialize them. This routine finds and initializes all background ROMs.

Entry conditions:

DE contains address of the first usable byte of memory (lowest address).
HL contains address of the last usable byte of memory (highest address).

Exit conditions:

DE contains the address of the new first usable byte of memory.
HL contains the address of the new last usable byte of memory.
AF and BC corrupt.
All other registers preserved

Notes:

When a foreground program is entered it is passed the addresses of the first and last bytes in memory which it may use. The area of memory outside this is used to store firmware variables, the stack, the jumpblocks and the screen memory. From the area available for a foreground program to use, the areas for background programs to use must be allocated.

The foreground program should initialize background ROMs at an early stage, before it uses the memory it is given. It may choose whether to enable background ROMs or not. KL INIT BACK may be used to initialize a particular background ROM or this routine may be used to initialize all available background ROMs.

KL ROM WALK inspects the ROMs at ROM select addresses in the range 1..7 in V1.0 firmware and 0..15 in V1.1 firmware. The power-up initialization entry of each background ROM found is called (unless it is the current foreground ROM in V1.1 firmware). This entry may allocate some memory for the background ROM's use by adjusting the values in DE and HL before returning. Once the ROM has been initialized the Kernel adds it to the list of external command servers, and notes the base of the area which the ROM has allocated to itself at the top of memory (if any). Subsequent FAR CALLs to entries in the ROM will automatically set the IY index register to point at the ROM's upper memory area.

See section 10.4 for a full description of background ROMs.

Related entries:

KL FIND COMMAND
KL INIT BACK
KL LOG EXT

154: KL INIT BACK

#BCCE

Initialize a particular background ROM.

Action:

Background ROMs provide support for expansion hardware or augment the software facilities of the machine. If the facilities provided by the background ROMs are to be available the foreground program must initialize them. This routine selects and initializes a particular background ROM.

Entry conditions:

C contains the ROM select address of the ROM to initialize.

DE contains address of the first usable byte of memory (lowest address).

HL contains address of the last usable byte of memory (highest address).

Exit conditions:

DE contains the address of the new first usable byte of memory.

HL contains the address of the new last usable byte of memory.

AF and B corrupt.

All other registers preserved.

Notes:

The ROM select address must be in the range 1..7 in V1.0 and 0...15 in V1.1 firmware and the ROM at this address must be a background ROM or the request will be ignored. In V1.1 firmware the request will be ignored if the ROM is the current foreground ROM.

When a foreground program is entered it is passed the addresses of the first and last bytes in memory which it may use. The area of memory outside this is used to store firmware variables, the stack, the jumpblocks and the screen memory. From the area available for a foreground program to use, the areas for background programs to use must be allocated.

The foreground program should initialize background ROMs at an early stage, before it uses the memory it is given. It may choose whether to enable background ROMs or not. KL ROM WALK maybe used to initialize all available ROMs or this routine may be used to initialize particular ROMs.

This routine causes the background ROM's power-up initialization entry to be called. This entry may allocate some memory for the background ROM's use by adjusting the values in DE and HL before returning. Once the ROM has been initialized the Kernel adds it to the list of external command servers, and notes the base of the area which the ROM has allocated to itself at the top of memory (if any). Subsequent FAR CALLs to entries in the ROM will automatically set the IY index register to point at the ROM's upper memory area.

See section 10.4 for a full description of background ROMs.

Related entries:

KL FIND COMMAND

KL LOG EXT

KL ROM WALK

155: KL LOG EXT

#BCD1

Introduce an RSX to the Firmware.

Action:

RSXs (Resident System Extensions) are similar to background ROMs, but are loaded into RAM. This routine must be called to include the RSX on the Kernel's list of external command servers.

Entry conditions:

BC contains the address of the RSX's command table.
HL contains the address of a 4 byte area of RAM for the Kernel's use.

Exit conditions:

DE corrupt.
All other registers preserved.

Notes:

Both the RSX's command table and the Kernel's storage area must lie in the central 32K of memory, i.e. not under a ROM.

The format of a command table is described in section 10.2 and RSXs are discussed in section 10.5.

Related entries:

KL FIND COMMAND
KL INIT BACK

156: KL FIND COMMAND

#BCD4

Search for an RSX, background ROM or foreground ROM to process a command.

Action:

All expansion ROMs and RSXs have command tables of the same form. This routine searches all RSXs and background ROMs on the Kernel's list of external command servers looking for a match for the given command name. If the name is found, then the 'far address' of the associated routine is returned. If the command is not a background or RSX command then all the foreground ROMs that can be found are searched for a foreground program with the given name. If a foreground program is found then the system immediately enters it.

Entry conditions:

HL contains the address of the command name to search for.

Exit conditions:

If an RSX or background ROM command was found:

Carry true.

C contains the ROM select address.

HL contains the address of the routine.

If the command was not found:

Carry false.

C and HL corrupt.

Always

A, B and DE corrupt.

All other registers preserved.

Notes:

The command name passed must be in RAM but may lie underneath a ROM. The name may be any number of characters long but only the first 16 characters are significant. All alphabetic characters in the name should be in upper case and the last character of the name should have bit 7 set.

The ROM select and routine addresses returned are suitable for calling KL FAR PCHL.

The list of external command servers is generated as background ROMs and RSXs are initialized (see KL ROM WALK, KL INIT BACK and KL LOG EXT). The command tables are scanned in the opposite order to that in which the command servers were introduced. Thus, RSXs will tend to take precedence over background ROMs, since RSX's are, in general, initialized after background ROMs. Background ROMs are normally initialized in reverse order of ROM select address, so lower numbered ROMS will take precedence over higher.

See section 10.2 for a full description of the format of expansion ROM command tables.

The first entry in a background ROM's command name table (the one associated with the power-up entry) may be used as the ROM's name. KL FIND COMMAND may be used, therefore, to find out whether a particular background ROM has been initialized.

When searching for a foreground program, ROMs are inspected starting with ROM 0 and working upwards. The search ceases when the first unused ROM address greater than 0 on V1.0 firmware and greater than 15 on V1. 1 firmware is found.

The on-board BASIC may be entered by searching for and invoking the command 'BASIC'.

If a foreground ROM command is found the ROM is entered unconditionally and this routine never returns.

Related entries:

KL INIT BACK

KL LOGEXT

KL ROM WALK

MC START PROGRAM

157: KL NEW FRAME FLY

#BCD7

Initialize and put a block onto the frame flyback list.

Action:

The Kernel maintains a list of events to be kicked each time frame flyback occurs. This routine initializes a block and adds it to the list.

Entry conditions:

HL contains the address of the frame flyback block.

B contains the event class.

C contains the ROM select address of the event routine.

DE contains the address of the event routine.

Exit conditions:

AF, DE and HL corrupt.

All other registers preserved.

Notes:

The frame flyback block is 9 bytes long and must lie in the central 32K of RAM. The last 7 bytes of the frame flyback block are an event block which is initialized to reflect the parameters passed in B, C and DE (see KL INIT EVENT). The exact layout of a frame flyback block is described in Appendix X.

The frame flyback block is appended to the frame flyback list if it is not already on it. This routine enables interrupts.

Related entries:

KL ADD FRAME FLY

KL DEL FRAME FLY

KL INIT EVENT

158: KL ADD FRAME FLY

#BCDA

Put a block onto the frame flyback list.

Action:

The Kernel maintains a list of events to be kicked each time frame flyback occurs. This routine adds a block to the list.

Entry conditions:

HL contains the address of the frame flyback block.

Exit conditions:

AF, DE and HL corrupt.
All other registers preserved.

Notes:

The frame flyback block is 9 bytes long and it must lie in the central 32K of RAM. The last 7 bytes of the frame flyback block are an event block which must be initialized separately before calling this routine. The exact layout of a frame flyback block is described in Appendix X.

The block is appended to the frame flyback list if it is not already on it. This routine enables interrupts.

Related entries:

KL DEL FRAME FLY
KL INIT EVENT
KL NEW FRAME FLY

159: KL DEL FRAME FLY

#BCDD

Remove a block from the frame flyback list.

Action:

The Kernel maintains a list of events to be kicked each time frame flyback occurs. This routine removes a block from the list.

Entry conditions:

HL contains the address of the frame flyback block.

Exit conditions:

AF, DE and HL corrupt.
All other registers preserved.

Notes:

This routine does nothing if the block is not on the list.

Removing a block from the list only prevents the event being kicked again. It does not affect any outstanding frame flyback events.

This routine enables interrupts.

Related entries:

KL ADD FRAME FLY

KL NEW FRAME FLY

160: KL NEW FAST TICKER

#BCEO

Initialize and put a block onto the fast ticker list.

Action:

The Kernel maintains a list of events to be kicked each time the 1/300th of a second timer interrupt occurs. This is known as the fast ticker list. This routine initializes a block and adds it to the list.

Entry conditions:

HL contains the address of the fast ticker block.

B contains the event class.

C contains the ROM select address of the event routine.

DE contains the address of the event routine.

Exit conditions:

AF, DE and HL corrupt.

All other registers preserved.

Notes:

The fast ticker block is 9 bytes long and must lie in the central 32K of RAM. The last 7 bytes of the fast ticker block are an event block which is initialized to reflect the parameters passed in B, C and DE (see KL INIT EVENT). The exact layout of a fast ticker block is described in Appendix X.

The fast ticker block is appended to the fast ticker list if it is not already on it.

The fast ticker facility is not intended for general use. However, it does allow relatively short times to be measured giving greater resolution than the general ticker facilities.

This routine enables interrupts.

Related entries:

KL ADD FAST TICKER

KL ADD TICKER

KL DEL FAST TICKER

KL INIT EVENT

KL TIME PLEASE

161: KL ADD FAST TICKER

#BCE3

Put a block onto the fast ticker list.

Action:

The Kernel maintains a list of events to be kicked each time the 1/300th of a second timer interrupt occurs. This is known as the fast ticker list. This routine adds a block to the list.

Entry conditions:

HL contains the address of the fast ticker block.

Exit conditions:

AF, DE and HL corrupt.

All other registers preserved.

Notes:

The fast ticker block is 9 bytes long and must lie in the central 32K of RAM. The last 7 bytes of the fast ticker block are an event block which must be initialized before calling this routine. The exact layout of a fast ticker block is described in Appendix X.

The fast ticker block is appended to the fast ticker list if it is not already on it.

The fast ticker facility is not intended for general use. However, it does allow relatively short times to be measured giving greater resolution than the general ticker facilities.

This routine enables interrupts.

Related entries:

KL ADD TICKER

KL DEL FAST TICKER

KL INIT EVENT

KL NEW FAST TICKER

KL TIME PLEASE

162: KL DEL FAST TICKER

#BCE6

Remove a block from the fast ticker list.

Action:

The Kernel maintains a list of events to be kicked each time the 1/300th of a second timer interrupt occurs. This is known as the fast ticker list. This routine removes a block from the list.

Entry conditions:

HL contains the address of the fast ticker block.

Exit conditions:

AF, DE and iL corrupt.
All other registers preserved.

Notes:

This routine does nothing if the block is not on the list.

Removing a block from the list only prevents the event from being kicked again. It does not affect any outstanding fast ticker events.

This routine enables interrupts.

Related entries:

KL ADD FAST TICKER
KL DEL TICKER
KL NEW FAST TICKER

163: KL ADD TICKER

#BCE9

Put a block onto the tick list.

Action:

The general purpose timing facility measures time in 1/50th of a second units. The Kernel maintains a list of tick blocks each of which contains a count and a recharge value. Every 1/50th of a second the Kernel processes all the tick blocks, decrementing the count entry of each. If the count entry of a block becomes zero the event contained in the block is 'kicked', and the count is set to the recharge value.

Entry conditions:

HL contains the address of the tick block.
DE contains the initial value for the count entry.
BC contains the value for the recharge entry.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The tick block is 13 bytes long and must lie in the central 32K of memory. The last 7 bytes of the tick block are an event block which must be initialized before this routine is called. The exact layout of a tick block is described in Appendix X.

The count and recharge entries in the block are set. The block is then appended to the tick list if it is not already on the list. This routine may be used, therefore, to change the count and recharge entries of an existing block.

Blocks with a count entry of zero are ignored when the list is processed. Setting a recharge value of zero, therefore, sets up the block as a 'one shot timer'. Since it takes the Kernel time to ignore a tick block, any redundant blocks should be removed from the list as soon as possible.

It is not possible to predict, particularly with synchronous events, how long it will be after the 'kick' before the event routine is actually called. Notwithstanding these delays, the ticker may be used to obtain an exact number of 'kicks' in a given period since the recharge mechanism immediately resets the count. The event counting mechanism will ensure that 'kicks' are not missed, provided that there are never more than 127 outstanding at once.

This routine enables interrupts.

Related entries:

KL ADD FAST TICKER
KL DEL TICKER
KLINITEVE

164: KL DEL TICKER

#BCEC

Remove block from the tick list.

Action:

If the given block is on the tick list it is removed. The contents of the block are not affected.

Entry conditions:

HL contains the address of the tick block.

Exit conditions:

If the tick block was found on the tick list:

Carry true.

DE contains the count remaining before the next event.

If the tick block was not found on the tick list:

Carry false.

DE corrupt.

Always:

A, HL and other flags corrupt.

All other registers preserved.

Notes:

The contents of the block are not affected by removing it from the list. In particular the continued processing of outstanding events is not affected. The block could be put back on the list at a later date and it could continue counting where it left off.

This routine enables interrupts.

Related entries:

KL ADD TICKER

KL DEL FAST TICKER

165: KL INIT EVENT

#BCEF

Initialize an event block.

Action:

Initialize all entries in an event block.

Entry conditions:

HL contains the address of the event block.

B contains the event class.

C contains the RUM select address of the event routine.

DE contains the address of the event routine

Exit conditions:

HL contains the address of the event block + 7

All other registers preserved.

Notes:

The event block is 7 bytes long and must lie in the central 32K of RAM. The layout of an event block is described in Appendix X. See section 12 for a general discussion of events.

The ROM select and address of the routine are the 'far address' of the event routine (see section 2).

The event class is bit significant as follows:

- Bit 0: Near address.
- Bits 1..4: Synchronous event priority.
- Bit 5: Must be zero.
- Bit 6: Express event.
- Bit 7: Asynchronous event.

If the asynchronous bit is set then the event is an asynchronous event, otherwise it is a synchronous event. Asynchronous events do not have priorities and so the priority field is ignored.

If the express event bit is set then the event is an express event. The meaning of this depends on whether the event is synchronous or asynchronous.

All express synchronous events have higher priorities than any normal synchronous event. The priority of a synchronous event is encoded in bits 1.4 of the class, the higher the number the greater the priority. No event may have priority

0. The processing of normal synchronous events may be disabled (by calling KL EVENT DISABLE), while that of express synchronous events may not.

An express asynchronous event will have its event routine called directly from the interrupt path. A normal asynchronous event is processed just before returning from the interrupt (with interrupts enabled).

If the near address bit is set then the event routine is located either in the lower ROM or in the central 32K of RAM. The ROM select address is ignored and the routine is called directly, rather than through the FAR CALL mechanism, thus reducing the event processing overhead. Where possible, asynchronous events should be at 'near addresses'. Express asynchronous events must always be at 'near addresses'.

Event blocks appear in various other blocks handled by the Kernel, including frame flyback, fast ticker and tick blocks. This routine is used to initialize the event block parts of these.

The bytes after the last byte of the event block, even where the block forms part of another block, are not used by the Kernel. When the event routine is called the address of the block is passed to it, so the user may append further information about the event to the block. This allows several similar events to share the same event routine, each event having its 'own' variables appended to its event block.

The event routine has the following entry and exit conditions:

Entry:

If the event routine is at a 'far address':

HL contains the address of byte 5 of the event block
(so any appended data can start at address HL + 2).

If the event routine is at a 'near address':

DE contains the address of byte 6 of the event block
(so any appended data can start at address DE + 1)

Exit:

AF, BC, DE and HL corrupt.
All other registers preserved.

The event routine may use the IX and IY registers but must preserve them. It may not use the second register set. Express asynchronous events may not enable interrupts.

KL INIT EVENT enables interrupts.

Related entries:

KL DEL SYNCHRONOUS

KL DISARM EVENT

KL EVENT

KL NEW FAST TICKER

KL NEW FRAME FLY

KL NEW TICKER

KL SYNC RESET

166: KL EVENT

#BCF2

‘Kick’ an event block.

Action:

The event mechanism arranges that an event routine be called in response to each ‘kick’ of an event block. KL EVENT performs the ‘kick’.

Entry conditions:

HL contains the address of the event block.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

Unlike the vast majority of Kernel routines this routine may be called from the interrupt path. Because the LOW JUMP instruction in the main firmwarejumpblock enables interrupts the user must pick the address part of the ‘low address’ out of the jumpblock and mask off the top two bits to extract the address in the lower ROM of KL EVENT. The following code does this:

```
LD    DE,(KL_EVENT+ 1) ;extract address part of LOW JUMP
RES   7,D                ;remove upper ROM state from ‘low address’
RES   6,D                ;remove lower ROM state from ‘low address’
CALL  PCDE_INSTRUCTION ; CALL KL EVENT
```

(If the user is going to perform this operation repeatedly it is recommended that the address should be extracted once and should be stored somewhere).

The effect of the ‘kick’ depends on the event count in the event block:

Count < 0: The event is disarmed, and kicking it has no effect.

Count > 0: There are other kicks outstanding and the event is being processed. This kick simply increments the count (unless it has already reached the maximum of 127). Once event processing has begun it continues until the count becomes zero or the event is disarmed.

Count = 0: The event is armed but event processing is not active.
The count is incremented and event processing initiated.

How event processing is initiated depends on the event class.

Synchronous Events.

Synchronous events are added to the synchronous event queue in priority order. It is the responsibility of the foreground program to process the synchronous event queue regularly.

Synchronous event routines are called when the foreground program calls KL DO SYNC, the event count is then dealt with when KL DONE SYNC is called.

Asynchronous Events.

a. Not in the Interrupt Path

The event routine is called immediately. When the routine returns, if the event count greater than zero it is decremented. If the count is still greater than zero then the procedure is repeated.

b. In the Interrupt Path - Normal Asynchronous Event

The event is placed on the interrupt event pending queue. On exit from the interrupt path the Kernel processes all events on the interrupt pending queue as described in (a) above. This means that normal asynchronous event routines are called in an extension of normal (non-interrupt) processing between interrupt return and the main program. The routine is, therefore, not subject to the restrictions imposed on interrupt path routines.

c. In the Interrupt Path - Express Asynchronous Event

The event routine is called immediately, in the interrupt path. The routine must be at a 'near address' (see KL INIT EVENT). Under no circumstances may the routine enable interrupts.

KL EVENT enables interrupts unless it is called from the interrupt path.

Related entries:

KL INIT EVENT

KL NEXT SYNC

KL POLL SYNCHRONOUS

KL SYNC RESET

167: KL SYNC RESET

#BCF5

Clear synchronous event queue.

Action:

The synchronous event queue is set empty - any outstanding events are simply discarded. The current event priority, used by KL POLL SYNCHRONOUS and KL NEXT SYNC to mask out lower priority events, is reset.

Entry conditions:

No conditions.

Exit conditions:

AF and HL corrupt.

All other registers preserved.

Notes:

It is the user's responsibility to ensure that the discarded events and any currently active events are reset. The event count of discarded events will be greater than zero, so any further 'kicks' will simply increment the count, but not add the event to the synchronous event queue - the events are, therefore, effectively disarmed.

Related entries:

KL DEL SYNCHRONOUS

KL NEXT SYNC

KL POLL SYNCHRONOUS

168: KL DEL SYNCHRONOUS

#BCF8

Remove a synchronous event from the event queue.

Action:

The event is disarmed. If it is on the synchronous event queue then it is removed.

Entry conditions:

HL contains the address of the event block.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

Deleting an event from the queue prevents the outstanding 'kicks' from being processed.

Before a synchronous event block is reset or reinitialized this routine should be used to ensure that it is not currently pending.

This routine enables interrupts.

Related entries:

KL DISARM EVENT
KL INIT EVENT
KL SYNC RESET

169: KL NEXT SYNC

#BCFB

Get next event from the queue.

Action:

If there is an event on the synchronous event queue whose priority is greater than the current event priority (if any), then remove the event from the queue, set the current event priority to that of the event removed and return the previous event priority.

Entry conditions:

No conditions.

Exit conditions:

If there is an event to be processed:

Carry true.

HL contains the address of the event block.

A contains the previous event priority (if any). If there is no event to be processed:

Carry false.

A and HL corrupt.

Always:

DE corrupt.

All other registers preserved.

Notes:

KL NEXT SYNC returns the address of the next event to be processed, if any, which it has taken off the synchronous event queue and whose priority has now been set as the event priority mask.

The foreground program should call KL POLL SYNCHRONOUS regularly to check for outstanding events. KL POLL SYNCHRONOUS is a short routine in RAM, so calling it imposes little overhead. If there is an event outstanding then the above procedure should be invoked, and should be repeated until the event queue is empty.

The current event priority mechanism allows event routines to poll for, and process, events of higher priority. The priority returned by this routine must be preserved until it is passed to KL DONE SYNC.

KL NEXT SYNC enables interrupts.

The procedure for processing synchronous events is as follows:

TRY.AGAIN:

```
CALL KL_NEXTSYNC ; return next event, if any
JR    NC,?????   ; jump if no event to process
;
PUSH  HL          ; save address of event
PUSH  AF          ; save previous event priority
CALL  KL_DO_SYNC ; call the event routine
POP   AF
POP   HL
;
CALL  KL_DONE_SYNC ; reset the event priority mask, deal with the
; event count and put the event back on the
; queue if the count is still greater
; than zero
JR    TRYAGAIN    ; see if any events are still awaiting
; processing
```

Related entries:

KL DONE SYNC

KLDO SYNC

KLEVENT

KL INIT EVENT

KL POLL SYNCHRONOUS

170: KL DO SYNC

#BCFE

Perform an event routine.

Action:

Call the event routine for a given event.

Entry conditions:

HL contains the address of the event block.

Exit conditions:

AF, BC, DE and HL corrupt.

All other registers preserved.

Notes:

This routine is intended to be called to process an event after KL NEXT SYNC has found it to be pending. Use of this entry at any other time is not recommended.

See KL NEXT SYNC above for the general scheme for processing synchronous events.

KL DO SYNC does not itself affect the event count.

Related entries:

KL DONE SYNC

KL NEXT SYNC

171: KL DONE SYNC

#BDO1

Finish processing an event.

Action:

Once a synchronous event has been processed, by invoking its event routine via KL DO SYNC, this entry must be called to restore the current event priority and to deal with the event count. If the count remains greater than zero the event block is placed back on the synchronous event queue.

Entry conditions:

A contains the previous event priority.
HL contains the address of the event block.

Exit conditions:

AF, BC, DE and HL corrupt
All other registers preserved.

Notes:

This routine is intended to be called after calling KL NEXT SYNC, to find a pending event, and KL DO' SYNC, to run the event routine. It uses the previous event priority and the event block address returned by KL NEXT SYNC. Other uses of this entry are not recommended.

See KL NEXT SYNC above for the general scheme for processing synchronous events.

Restoring the current event priority is an essential step in maintaining the synchronous event priority scheme.

If the event count is greater than zero then it is decremented. If the count is still greater than zero then there are further events outstanding and the event is placed back on the synchronous event queue. The event may be disarmed between KL NEXT SYNC and KL DONE SYNC. Setting the event count to one before calling KL DONE SYNC forces multiple events to be treated as a single event.

KL DONE SYNC may enable interrupts.

Related entries:

KL DO SYNC
KL NEXT SYNC

172: KL EVENT DISABLE

#BDO4

Disable normal synchronous events.

Action:

Prevent normal synchronous events from being processed but allow express synchronous events to be processed. This is achieved by setting the current event priority higher than any possible normal synchronous event priority.

Entry conditions:

No conditions.

Exit conditions:

HL corrupt.
All other registers preserved.

Notes:

KL EVENT DISABLE does not prevent events from being kicked. The effect is to 'mask off all pending normal synchronous events so that they are hidden from the foreground program (when KL POLL SYNCHRONOUS or KL NEXT SYNC are called) and hence are not processed.

KL EVENT ENABLE reverses the effect of KL EVENT DISABLE.

It is not possible to disable synchronous events permanently from inside a synchronous event routine as the previous current event priority is restored when the event routine returns.

Related entries:

KL DISARM EVENT
KL EVENT ENABLE
KL NEXT SYNC
KL POLL SYNCHRONOUS

173: KL EVENT ENABLE

#BDO7

Enable normal synchronous events.

Action:

Allows normal and express synchronous events to be processed.

Entry conditions:

No conditions.

Exit conditions:

HL corrupt.

All other registers preserved.

Notes:

Events are enabled by default. KL EVENT ENABLE reverses the effect of KL EVENT DISABLE.

It is not possible to enable synchronous events permanently from inside a synchronous event routine as the current event priority which is used to disable events is restored when the event routine returns.

Related entries:

KL EVENT DISABLE

KL NEXT SYNC

KL POLL SYNCHRONOUS

174: KL DISARM EVENT

#BDOA

Prevent an event from occurring.

Action:

Disarms the event by setting the event count to a negative value. Any further 'kicks' (calls of KL EVENT) for the event will be ignored, any outstanding events are discarded.

Entry conditions:

HL contains the address of the event block.

Exit conditions:

AF corrupt.
All other registers preserved.

Notes:

KL DISARM EVENT should only be used with asynchronous events. Synchronous events may be disarmed by calling KL DEL SYNCHRONOUS, which also ensures that the event is not on the synchronous event queue.

The event may be rearmed by reinitializing it (KL INIT EVENT) or by setting the event count (byte 2) of the event block to zero.

Related entries:

KL DEL SYNCHRONOUS
KL INIT EVENT

175: KL TIME PLEASE

#BDOD

Ask the elapsed time.

Action:

The Kernel maintains a count which it increments on each time interrupt. The count, therefore, measures time in 1/300th of a second units. This routine returns the current count.

Entry conditions:

No conditions.

Exit conditions:

DEHL contains the four byte count (D contains the most significant byte and L the least significant byte).

All other registers preserved.

Notes:

The count is zeroized when the machine is turned on or reset. The count may be set to another starting value by KL TIME SET.

The count is not kept up to date if interrupts are disabled for long periods, such as while reading and writing the cassette.

The four byte count overflows after approximately:

- 14,316,558 Seconds
- = 238,609 Minutes
- = 3,977 Hours
- = 166 Days

This routine enables interrupts.

Related entries:

KL TIME SET

176: KL TIME SET

#BD10

Set the elapsed time.

Action:

The Kernel maintains a count which it increments on each time interrupt. The count, therefore, measures time in 1/300th of a second units. This routine sets the count to a given value.

Entry conditions:

DEHL contains the four byte count to set (D contains the most significant byte and L the least significant byte).

Exit conditions:

AF corrupt.

All other registers preserved.

Notes:

The four byte count overflows after approximately:

	14,316,558	Seconds
	238,609	Minutes
=	3,977	Hours
=	166	Days

KL TIME SET may be used to set the count to the actual time of day, so that the Kernel then maintains a real clock rather than a simple measure of the time elapsed since the last reset.

The count is not kept up to date if interrupts are disabled for long periods, such as while reading and writing the cassette.

This routine enables interrupts.

Related entries:

KLTIME PLEASE

177: MC BOOT PROGRAM

#BD13

Load and run a program.

Action:

Shut down as much of the system as possible then load a program into RAM and run it. If the load fails then the previous foreground program is restarted.

Entry conditions:

HL contains the address of the routine to call to load the program.

Exit conditions:

Does not exit!

Notes:

The system is partially reset before attempting to load the program. External interrupts are disabled, as are all timer, frame flyback and keyboard break events. Sound generation is turned off, indirections are set to their default routines and the stack is reset to the default system stack. This process ensures that no memory outside the firmware variables area is in use when loading the program. Overwriting an active event block or indirection routine could otherwise have unfortunate consequences.

The partial system reset does not change the ROM state or ROM selection. The routine run to load the program must be in accessible RAM or an enabled ROM. Note that the firmware jumpblock normally enables the lower ROM and disables the upper ROM and so the routine must normally be in RAM above #4000 or in the lower ROM.

The routine run to load the program is free to use any store from #0040 up to the base of the firmware variables area (#B100) and may alter indirections and arm external device interrupts as required. It should obey the following exit conditions:

If the program loaded successfully:

Carry true.

HL contains the program entry point.

If the program failed to load:

Carry false.

HL corrupt.

Always:

A, BC, DE, IX, IY and other flags corrupt.

After a successful load the firmware is completely initialized (as at EMS) and the program is entered at the entry address returned by the load routine. Returning from the program will reset the system (perform RST 0).

After an unsuccessful load an appropriate error message is printed and the previous foreground program is restarted. If the previous foreground was itself a RAM program then the default ROM is entered instead as the program may have been corrupted during the failed loading.

Related entries:

CAS IN DIRECT

KL CHOKE OFF

MC START PROGRAM

178: MC START PROGRAM

#BD16

Run a foreground program.

Action:

Fully initialize the system and enter a program.

Entry conditions:

HL contains the entry point address.

C contains the required ROM selection.

Exit conditions:

Never exits!

Notes:

HL and C comprise the ‘far address’ of the entry point of the foreground program (see section 2).

When entering a foreground program in ROM the ROM selection should be that required to select the appropriate ROM. When entering a foreground program in RAM the ROM selection should be used to enable or disable ROMs as the RAM program requires (ROM select addresses # FC.. # FF).

This routine carries out a full EMS initialization of the firmware before entering the program. Returning from the program will reset the system (perform RST 0). MC START PROGRAM is intended for running programs in ROM or programs that have already been loaded into RAM. To load and run a RAM program use MC BOOT PROGRAM.

Related entries:

MC BOOT PROGRAM

RESET ENTRY (RST0)

179: MC WAIT FLYBACK

#BD19

Wait for frame flyback.

Action:

Wait until frame flyback occurs.

Entry conditions:

No conditions.

Exit conditions:

All registers and flags preserved.

Notes:

Frame flyback is a signal generated by the CRT controller to signal the start of the vertical retrace period. During this period the screen is not being written and so major operations can be performed on the screen without producing unsightly effects. A prime example is rolling the screen.

The frame flyback signal only lasts for a couple of hundred microseconds but the vertical retrace period is much longer than this. However, there will be a ticker interrupt in the middle of frame flyback which may cause the foreground processing to be suspended for a significant length of time. It is important, therefore, to perform any critical processing as soon after frame flyback is detected as is possible.

This routine returns immediately if frame flyback is occurring when it is called. It does not wait for the start of frame flyback (use a frame flyback event to do this).

Related entries:

KL ADD FRAME FLY

180: MC SET MODE

#BD1C

Set the screen mode.

Action:

Load the hardware with the required screen mode.

Entry conditions:

A contains the required mode.

Exit conditions:

AF corrupt.

All other registers preserved.

Notes:

The required mode is checked and no action is taken if it is invalid. If it is valid then the new value is sent to the hardware.

The screen modes are:

- 0: 160 x 200 pixels, 20 x 25 characters.
- 1: 320 x 200 pixels, 40 x 25 characters.
- 2: 640 x 200 pixels, 80 x 25 characters.

Altering the screen mode without notifying the Screen Pack will produce peculiar effects on the screen. In general SCR SET MODE should be called to change screen mode. This, in its turn, sets the new mode into the hardware.

Related entries:

SCR SET MODE

181: MC SCREEN OFFSET

#BD1F

Set the screen offset.

Action:

Load the hardware with the offset of the first byte on the screen inside a 2K screen block and which 16K block the screen memory is located in.

Entry conditions:

A contains the new screen base.

HL contains the new screen offset.

Exit conditions:

AF corrupt.

All other registers preserved.

Notes:

The screen base address is masked with #C0 to make sure it refers to a valid 16K memory area. The default screen base is #C0 (the screen is underneath the upper ROM).

The screen offset is masked with #07FE to make it legal. Note that bit 0 is ignored as the hardware only uses even offsets.

If the screen base or offset is changed without notifying the Screen Pack then unexpected effects may occur on the screen. In general SCR SET BASE or SCR SET OFFSET should be called. These, in their turn, send the values to the hardware.

Related entries:

SCR SET BASE

SCR SET OFFSET

182: MC CLEAR INKS

#BD22

Set all inks to one colour.

Action:

Set the colour of the border and set the colour of all the inks. All the inks are set to the same colour thus giving the impression that the screen has been cleared instantly.

Entry conditions:

DE contains the address of an ink vector.

Exit conditions:

AF corrupt.

All other registers preserved.

Notes:

The ink vector has the form:

Byte 0: Colour for the border.

Byte 1: Colour for all inks.

The colours supplied are the numbers used by the hardware rather than the grey scale numbers supplied to SCR SET INK (see Appendix V).

After the screen has been cleared (or whatever) the correct ink colours can be set by calling MC SET INKS.

This routine sets the colours for all 16 inks whether they can be displayed on the screen in the current mode or not.

This ink clearing technique is used by the Screen Pack when clearing the screen or changing mode (by SCR CLEAR and SCR SET MODE)

Related entries:

MC SET INKS

183: MC SET INKS

#BD25

Set colours of all the inks.

Action:

Set the colours of all the inks and the border.

Entry conditions:

DE contains the address of an ink vector.

Exit conditions:

AF corrupt.

All other registers preserved.

Notes:

The ink vector passed has the following layout:

Byte 0: Colour for the border.

Byte 1: Colour for ink 0.

Byte 2: Colour for ink 1.

.....

Byte 16: Colour for ink 15.

The colours supplied are the numbers required by the hardware rather than the grey scale numbers supplied to SCR SET INK (see Appendix V).

This routine sets the colours for all inks including those that cannot be visible in the current screen mode. However, it is only necessary to supply sensible colours for the visible inks.

The Screen Pack sets the colours for all the inks each time the inks flash and after an ink colour has been changed (by calling SCR SET INK or SCR SET BORDER).

Related entries:

MC CLEAR INKS

SCR SET BORDER

SCR SET INK

184: MC RESET PRINTER

#BD28

Reset the printer indirection.

Action:

Set the printer indirection, MC WAIT PRINTER, to its default routine and, in V1.1 firmware, set up the default printer translation table.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.

All other registers preserved.

Notes:

The default printer translation table is described in Appendix XIV. This is designed to drive the DMP-1 printer. It only translates the additional characters in the character set (#A0.. #AF); it does not translate any of the standard ASCII characters or the graphic characters.

Related entries:

MC WAIT PRINTER

MC PRINT CHAR

185: MC PRINT CHAR

#BD2B

Try to send a character to the Centronics port.

Action:

Send a character to the printer (Centronics port) or time out if the printer is busy for too long.

Entry conditions:

A contains the character to send (bit 7 ignored).

Exit conditions:

If the character was sent OK:

Carry true.

If the printer timed out:

Carry false.

All ways:

A and other flags corrupt.

All other registers preserved.

Notes:

In V1.1 firmware, the character to be sent is translated using the printer translation table as set by MC PRINT TRANSLATION. If the supplied character is not found in the table then it is sent as supplied without translation. However, if the character is found in the translation table then the corresponding translation is sent instead; unless the translation is #FF in which case the character is ignored and nothing is sent.

This routine calls the Machine Pack indirection MC WAIT PRINTER to send the character. The default indirection routine waits for the Centronics port to become non-busy then sends the character. If the port remains busy for too long (approximately 0.4 seconds) then the routine times out and the character is not sent. This time out is provided so that the caller can test for break whilst driving the printer.

Related entries:

MC PRINT TRANSLATION

MC RESET PRINTER

MC WAIT PRINTER

186: MC BUSY PRINTER

#BD2E

Test if the Centronics port is busy.

Action:

Test if the printer (Centronics port) is busy.

Entry conditions:

No conditions.

Exit conditions:

If Centronics port is busy:

 Carry true.

If Centronics port is idle:

 Carry false.

Always:

 Other flags corrupt.

 All other registers preserved.

Notes:

This routine has no other effects.

Related entries:

MC SEND PRINTER

187: MC SEND PRINTER

#BD31

Send a character to the Centronics port.

Action:

Send a character to the printer (Centronics port) which must not be busy.

Entry conditions:

A contains the character to send (bit 7 ignored).

Exit conditions:

Carry true.

A and other flags corrupt.

All other registers preserved.

Notes:

The printer must not be busy when a character is sent. The higher level routine MC PRINT CHAR will automatically wait for the printer to become non-busy and should be used in preference.

Related entries:

MC BUSY PRINTER

MC PRINT CHAR

188: MC SOUND REGISTER

#BD34

Send data to a sound chip register.

Action:

Set a sound chip sound register. This is a rather convoluted action because of the way the hardware has been designed.

Entry conditions:

A contains the sound chip register number.
C contains the data to send,

Exit conditions:

AF and BC corrupt.
All other registers preserved.

Notes:

This routine enables interrupts.

Related entries:

None!

189: JUMP RESTORE

#BD37

Restore the standard jumpblock.

Action:

Set the main firmware jumpblock to its standard state as described in sections 14.1 and 15.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

This routine may be used to restore the jumpblock to its standard routines after the user has changed entries in it. The whole of the jumpblock is set up so care must be taken if other programs, such as AMSDOS, have patched it.

The indirections jumpblock is set up piecemeal by the various packs' initialization and reset routines. JUMP RESTORE does not set up the indirections.

Related entries:

GRA RESET

KM RESET

MC RESET PRINTER

SCR RESET

TXT RESET

190 KM SET LOCKS

#BD3A

Set the shift and caps lock states.

Action:

Turn the shift and caps locks on or off.

Entry conditions:

H contains the required caps lock state.
L contains the required shift lock state.

Exit conditions:

AF corrupt.
All other registers preserved.

Notes:

This routine is not available on V1.0 firmware.

The lock states are:

#00 means that the lock is to be turned off.

#FF means that the lock is to be turned on.

The default lock states are off.

Related entries:

KM GET STATE

191 KM FLUSH

#BD3D

Flush the keyboard buffers.

Action:

Discard all pending keys from the key buffer, the 'put back' character and any current expansion string.

Entry conditions:

No conditions.

Exit conditions:

AF corrupt.
All other registers preserved.

Notes:

This routine is not available on V1.0 firmware.

The next character that will be returned by KM READ CHAR (or a similar routine) after KM FLUSH is called will be the first character that the user types after the call of KM FLUSH since all the pending characters will have been discarded.

On V1.0 firmware the effect of this routine can be simulated by repeatedly calling KM READ CHAR until it comes back with carry false.

Related entries:

KM READ CHAR

KM READ KEY

192 TXT ASK STATE

#BD40

Get the state of the Text VDU.

Action:

Return the VDU enable/disable state and the cursor on/off and cursor enable/disable states of the current selected stream.

Entry conditions:

No conditions.

Exit conditions:

A contains the stream state.

Flags corrupt.

All other registers preserved.

Notes:

This routine is not available on V1.0 firmware.

The stream state is returned as follows:

Bit 0	0	cursor enabled.	1	cursor disabled.
Bit 1	0	cursor on,	1	cursor off.
Bits 2..6		are undefined.		
Bit 7	0	VDU disabled,	1	VDUenabled.

Related entries:

TXT CUR DISABLE

TXT CUR ENABLE

TXT CUR OFF

TXT CUR ON

TXT VDU DISABLE

TXT VDU ENABLE

193 GRA DEFAULT

#BD43

Set default Graphics VDU modes.

Action:

Sets the graphics write mode, background mode, first pixel mode and line mask to their default settings.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and iLL corrupt.
All other registers preserved.

Notes:

This routine is not available on V1.0 firmware.

This routine sets the following modes:

Graphics write mode is set to force.

Graphics background mode is set to opaque.

First point mode is set to plot the first pixel of lines.

Line mask is set to give continuous lines (mask of #FF).

Related entries:

GRA INITIALISE

GRA RESET

GRA SET BACK

GRA SET FIRST

GRA SET LINE MASK

SCR ACCESS

194 GRA SET BACK

#BD46

Set whether background is to be written.

Action:

Set the graphics background write mode to opaque or transparent. This affects how GRA LINE ABSOLUTE, GRA LINE RELATIVE and GRA WR CHAR write 'background' pixels. In opaque mode the pixels are written in the current paper ink using the current graphics write mode. In transparent mode these pixels are not plotted at all.

Entry conditions:

If background is to be written (opaque mode):

A must be zero.

If background is not to be written (transparent mode):

A must be non-zero.

Exit conditions:

All registers and flags preserved.

Notes:

This routine is not available on V1 .0 firmware.

Transparent write mode is useful for annotating diagrams and for similar applications.

The graphics background write mode is similar to but independent of the character write mode of each stream of the Text VDU.

The default setting is opaque mode.

Related entries:

GRA DEFAULT

GRA LINE

GRA LINE ABSOLUTE

GRA LINE RELATIVE

GRA SET LINE MASK

GRA WR CHAR

TXT SETBACK

195 GRA SET FIRST

#BD49

Set whether the first point of a line is to be plotted.

Action:

Turn plotting of the first pixel of lines on or off.

Entry conditions:

If the first pixel is not to be plotted:

A contains zero.

If the first pixel is to be plotted:

A contains non-zero.

Exit conditions:

All registers and flags preserved.

Notes:

This routine is not available on V1.0 firmware.

Turning off the plotting of the first pixel of a line is particularly useful when drawing using XOR graphics write mode. For example, if a box is drawn in XOR mode when the first pixels of lines are being plotted then the corner pixels will be plotted twice and will therefore not be set. By not plotting the first pixel of lines this effect is avoided.

The default setting for this mode is to plot the first pixel.

Related entries:

GRA DEFAULT

GRA LINE

GRA LINE ABSOLUTE

GRA LINE RELATIVE

196 GRA SET LINE MASK

#BD4C

Set the line mask for plotting pixels of lines.

Action:

Set the line mask that specifies how pixels on lines are to be plotted. Where a bit in the mask is set the pixel will be plotted in the foreground (in graphics pen ink using the graphics write mode), Where a pixel in the mask is not set the pixel will either be plotted in the graphics paper ink using the graphics write mode or it will not be plotted at all depending on the graphics background write mode.

Entry conditions:

A contains the line mask to use.

Exit conditions:

All registers and flags preserved.

Notes:

This routine is not available on V1.0 Firmware.

The line mask is used starting with bit 7 and running to bit 0 and then starting with bit 7 again. Successive lines will use the mask as it was left when the previous line finished, the mask is not reset between lines.

The line mask specifies how pixels are to be plotted. This means that the same mask will give noticeably different effects in the various screen modes.

The mask is applied to the line running from left to right or from bottom to top, depending of the angle of the line, irrespective of which way round the end points of the line are specified.

If the first pixel of the line is not being plotted then the line mask is applied to the second pixel of the line first. It is not stepped on for the missing first pixel.

The default line mask is #FF which plots the whole line in the foreground.

Related entries:

GRA DEFAULT

GRA LINE

GRA LINE ABSOLUTE

GRA LINE RELATIVE

GRA SET BACK

197 GRA FROM USER

#BD4F

Convert user coordinates to base coordinates.

Action:

Convert the coordinates of a point from user coordinates to base coordinates rounding as appropriate.

Entry conditions:

DE contains the user X coordinate.
HL contains the user Y coordinate.

Exit conditions:

DE contains the base X coordinate.
HL contains the base Y coordinate.

AF corrupt.
All other registers preserved.

Notes:

This routine is not available on V1.0 firmware.

The following formulae are used to convert between the coordinate systems:

$$\begin{aligned}\text{Base X} &= (\text{Origin X} + \text{Rounded X}) / \text{Points per pixel} \\ \text{Rounded X} &= (\text{UserX} + \text{Round factor}) \text{ AND Round mask}\end{aligned}$$

Where:

	Round factor		Round mask	Points per pixel
	+ve userX	-ve userX		
Mode 0:	0	0	#FFFF	1
Mode 1:	0	1	#FFFE	2
Mode 2:	0	3	#FFFC	4

$$\begin{aligned}\text{Base Y} &= (\text{OriginY} + \text{Rounded Y}) / \text{Points per pixel} \\ \text{Rounded Y} &= (\text{UserY} + \text{Round factor}) \text{ AND Round mask}\end{aligned}$$

Where:

Round factor	= 0 for +ve user Y
	= 1 for -ve user Y
Roundmask	= #FFFE
Points per pixel	= 2

This routine is particularly useful when calling Screen pack routines which take the positions of points in base coordinates.

Related entries:

GRA SET ORIGIN

SCR DOT POSITION

198 GRA FILL

#BD52

Fill an area of the screen.

Action:

Fill an area of the screen containing the current graphics position and bounded by the edge of the window and pixels set to the pen ink.

Entry conditions:

A contains the (unencoded) ink to fill the area with
HL contains the address of a buffer.
DE contains the length of the buffer.

Exit conditions:

If the area was filled successfully:
Carry true.

If the area was not filled:
Carry false.

Always:

A, BC, DE, HL and other flags corrupt.
All other registers preserved.

Notes:

This routine is not available on V1.0 firmware.

The filling algorithm treats pixels set to the current pen ink and pixels set to the ink that is being used for filling as delimiters of the edge of the area. The fill ink and the pen ink maybe the same ink.

Pixels that are filled are set to the fill ink. The graphics write mode does not affect the way that pixels are written when filling.

The filling algorithm only moves up, down, right or left. It does not move diagonally and so the algorithm will not 'escape' through a gap between edge pixels that are diagonally adjacent. This means that the edge can be delimited using the normal lines drawn by the Graphics VDU.

The filling algorithm avoids recursing. Instead it stores 'interesting points', places that the algorithm has chosen one route to fill but might have chosen another route, in the buffer supplied by the user. The buffer may lie anywhere in RAM. Each 'interesting point' stored uses 7 bytes of the buffer and there is an overhead of 1 byte used to mark the end of the buffer. Thus a buffer 64 bytes long will allow 9 'interesting points' to be stored which should be sufficient for filling most simple areas.

The area to be filled may be as complicated as required but the more complicated the shape the longer the 'interesting point' buffer needs to be.

The failure return from this routine can occur for three reasons. Firstly, the current graphics position may be outside the window. Secondly, the pixel at current graphics position may be edge (pen or fill ink), In these cases the routine will return without filling anything. Thirdly the algorithm may exhaust the 'interesting point' buffer in which case some portions of the area will not be filled.

Related entries:

GRA SET PEN

199 SCR SET POSITION

#BD55

Set the location of the screen memory.

Action:

Tell the Screen pack the screen base and offset without telling the hardware.

Entry conditions:

A contains the screen base.
HL contains the screen offset.

Exit conditions:

A contains the screen base masked as required.
HL contains the screen offset masked as required.

Flags corrupt.
All other registers preserved.

Notes:

This routine is not available on V1.0 firmware.

This routine changes the location of the screen without notifying the hardware of the change. This effect may be used to construct a second screen of text or graphics without clearing the previous screen. When the new screen has been constructed the hardware may be notified and the picture will appear instantly.

In general the user is advised to set the base using SCR SET BASE and the offset using SCR SET OFFSET.

The screen base is masked with #C0 and the screen offset with #07FE to make the values legal.

Related entries:

SCR GET LOCATION

SCR SET BASE

SCR SET OFFSET

200 MC PRINT TRANSLATION

#BD58

Set the printer translation table.

Action:

Set how characters are to be translated before being sent to the printer.

Entry conditions:

HL contains the address of the table.

Exit conditions:

If the table is too long (more than 20 entries):

Carry false.

If the table is OK:

Carry true.

Always:

A, BC, DE, HL and other flags corrupt.

All other registers preserved.

Notes:

This routine is not available on V1.0 Firmware.

The supplied translation table may lie anywhere in RAM, This routine copies the table and so the memory may be re-used if required.

The format of the table is as follows:

Byte 0:	Number of entries in the table (N).
Bytes 1,2:	Entry 1
....
Bytes 2N-1, 2N:	Entry N

The format of each two byte entry is as follows:

Byte 0:	Character to be translated.
Byte 1:	Character to translate to.

If the character to translate to is #FF then the character is ignored and nothing is sent to the printer.

Translation of characters by the printer driver occurs in MC PRINT CHAR. The default translation table is set up when MC RESET PRINTER is called. The default table is designed to drive the DMP-1 printer (see Appendix XIV).

Related entries:

MC PRINT CHAR

201: KL BANK SWITCH

#BD5B

Select a memory organization.

Action:

Set which RAM banks are switched into the 64k of RAM in the memory map.

Entry conditions:

A contains new organization.

Exit conditions:

A contains old organization.

Flags corrupt.

All other registers preserved.

Notes:

This routine is only available on the CPC6128 (ie. V1.2 firmware). The memory organizations and bank switching are discussed fully in section 2.5.

It is inadvisable to bank switch to a memory organization where the code that is being executed, or stack are inaccessible!

Related entries:

KL L ROM DISABLE

KL L ROM ENABLE

KL ROM SELECT

KL U ROM DISABLE

KL U ROM ENABLE

16 The Firmware Indirections.

This section gives the detailed entry and exit conditions and effects of the routines in the `indirectionsjumpblock`. See section 14.2 for a list of these routines.

The firmware indirections are taken by the firmware at key points. They allow the user to intercept and alter a number of firmware actions without having to provide a complete new firmware package.

The descriptions given are for the default settings of the indirections. Replacement routines need not perform all the actions that the default routine performs although they are advised to do so.

IND: TXT DRAW CURSOR

#BDCD

Place the cursor blob on the screen (if enabled).

Default action:

If the cursor is enabled and turned on then the cursor blob is drawn on the screen. If not then no action is taken. The current text position is forced into the window (see TXT VALIDATE) and the cursor blob is written at the resulting position. The cursor blob is an inverse patch. This routine will only be called twice if TXT UNDRAW CURSOR is called in between.

Entry conditions:

No conditions.

Exit conditions:

AF corrupt.
All other registers preserved.

Notes:

This indirection is provided to allow the user to change the form of the cursor blob. See TXT PLACE CURSOR for a description of how the cursor blob is normally written.

The Text VDU routines call this indirection whenever the cursor is placed on the screen. All the Text VDU routines that read from the screen, write to the screen or change the current position remove the cursor (using TXT UNDRAW CURSOR) before performing their action and place it back on the screen afterwards (using TXT DRAW CURSOR). An example of such a routine is TXT WR CHAR that writes a character on the screen.

This indirection is set up when TXT INITIALISE or TXT RESET is called.

Related entries:

TXT PLACE CURSOR

TXT UNDRAW CURSOR

IND: TXT UNDRAW CURSOR

#BDD0

Remove the cursor blob from the screen (if enabled).

Default action:

If the cursor is enabled and turned on then the cursor blob is removed from the screen. If not then no action is taken. The cursor blob is an inverse patch. This routine will only be called after TXT DRAW CURSOR has been used to place the cursor on the screen.

Entry conditions:

No conditions.

Exit conditions:

AF corrupt.

All other registers preserved.

Notes:

This indirection is provided to allow the user to change the form of the cursor blob. See TXT REMOVE CURSOR for a description of how the cursor blob is normally removed.

The Text VDU routines call this indirection to remove the cursor from the screen. All the Text VDU routines that read from the screen, write to the screen or change the current position remove the cursor (using TXT UNDRAW CURSOR) before performing their action and place it back on the screen afterwards (using TXT DRAW CURSOR). An example of such a routine is TXT WR CHAR that writes a character on the screen.

This indirection is set up when TXT INITIALISE or TXT RESET is called.

Related entries:

TXT DRAW CURSOR

TXT REMOVE CURSOR

IND: TXT WRITE CHAR

#BDD3

Write a character onto the screen.

Default action:

Put a character on the screen at a character position.

Entry conditions:

A contains the character to write.

H contains the physical column to write at.

L contains the physical row to write at.

Exit conditions:

AF, BC, DE and HL corrupt.

All other registers preserved.

Notes:

The character position to write at is given in physical coordinates. i.e. Row 0, column 0 is the top left corner of the screen. The position is not checked for legality.

TXT WRITE CHAR is called by TXT WR CHAR to print a character on the screen. The removing of the cursor blob and the calculation of the new current position are performed by TXT WR CHAR and not by TXT WRITE CHAR.

This indirection is set up when TXT INITIALISE or TXT RESET is called.

Related entries:

TXT OUTPUT

TXT UNWRITE

TXT WR CHAR

IND: TXT UNWRITE

#BDD6

Read a character from the screen.

Default action:

Try to read a character from the screen at a character position.

Entry conditions:

H contains the physical column to read from.

L contains the physical row to read from.

Exit conditions:

If a readable character was found:

Carry true.

A contains the character read.

If no recognisable character was found:

Carry false.

A contains zero.

Always:

BC, DE, HL and other flags corrupt.

All other registers preserved.

Notes:

The character position to read from is given in physical coordinates. i.e. Row 0, column 0 is the top left corner of the screen. The position is not checked for legality.

This indirection is called by TXT RD CHAR to read a character from the screen. TXT RD CHAR removes the cursor from the screen before calling this indirection.

The read is performed by comparing the matrix found on the screen with the matrices used to generate characters. As a result changing a character matrix, changing the pen or paper inks or changing the screen (e.g. drawing a line through a character) may make the character unreadable. In particular the cursor blob will cause confusion and so it should not be on the screen

Special precautions are taken against generating inverse space (character #8F). Initially the character is read assuming that the background to the character was written in the current paper ink. If this fails to generate a recognisable character or it generates inverse space then another try is made by assuming that the character was written in the current pen ink.

The characters are scanned starting with #00 and finishing with #FF. Thus, if there are two possible character matrices that match the screen, the smaller of the two character numbers will be returned.

This indirection is set up when TXT INITIALISE or TXT RESET is called.

Related entries:

TXT RD CHAR

TXT WRITE CHAR

IND: TXT OUT ACTION

#BDD9

Output a character or control code.

Default action:

Print a character on the screen or obey a control code (characters #00..#1F). Works on the currently selected stream (except as noted below).

Entry conditions:

A contains the character or code.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

This indirection is called by TXT OUTPUT to do the work of printing characters or obeying the control codes. It is provided to allow the user to change the method of dealing with characters and control codes or to allow the user to redirect output (to the printer for example). TXT OUTPUT merely preserves the registers around the call of TXT OUT ACTION.

Control codes may take up to 9 parameters. These are the characters sent following the initial control code. The characters sent are stored in a buffer until sufficient have been received to make up all the required parameters. The control code buffer is only long enough to accept 9 parameter characters.

There is only one control code buffer which is shared between all the streams. It is, therefore, possible to get unpredictable results if the output stream is changed part of the way through sending a control code sequence.

If the VDU is disabled then no characters will be printed on the screen. In V1.1 firmware control codes that are specially marked in the control code table will not be obeyed if the VDU is disabled. Other control codes and all control codes in V1.0 firmware will be obeyed.

If the graphics character write mode is enabled then all characters and control codes are printed using the Graphics VDU (see GRA WR CHAR) and are not obeyed. Normally characters are written by the Text VDU (see TXT WR CHAR).

This indirection is set up when TXT INITIALISE or TXT RESET is called.

Related entries:

TXT OUTPUT

TXT WR CHAR

IND: GRA PLOT

#BDDC

Plot a point.

Default action:

Check if the point lies inside the current window and if so write it in the current graphics pen ink and using the current graphics write mode. The current graphics position is always moved to the specified point.

Entry conditions:

DE contains the user X coordinate of the point to plot.
HL contains the user Y coordinate of the point to plot.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The position of the point to plot is given in user coordinates, i.e. relative to the user origin.

This indirection is called by GRA PLOT RELATIVE and GRA PLOT ABSOLUTE to plot the point requested. It is provided to allow the user to change the method for plotting (to output to an X-Y plotter for example). GRA PLOT RELATIVE converts from relative to user coordinates and then calls this indirection; GRA PLOT ABSOLUTE calls this indirection immediately.

To write the point on the screen the SCR WRITE indirection is used. Thus the point is plotted using the current graphics write mode.

This indirection is set up when GRA INITIALISE or GRA RESET is called.

Related entries:

GRA PLOT ABSOLUTE

GRA PLOT RELATIVE

GRA TEST

SCR WRITE

IND: GRA TEST

#BDDF

Test a point.

Default action:

Check if the point is inside the graphics window and return the ink it is currently set to if so. Otherwise, return the current graphic paper ink. The current graphics position is always moved to the specified point.

Entry conditions:

DE contains the user X coordinate of the point to test.
HL contains the user Y coordinate of the point to test.

Exit conditions:

A contains the decoded ink of the specified point.
BC, DE, HL and flags corrupt.
All other registers preserved.

Notes:

The position of the point to test is given in user coordinates, i.e. relative to the user origin.

This indirection is used by GRA TEST RELATIVE and GRA TEST ABSOLUTE to test the point requested. It is provided to allow the user to change the method for testing (comparing with the current pen ink for example). GRA TEST RELATIVE converts from relative to user coordinates and then calls this indirection; GRA TEST ABSOLUTE calls this indirection immediately.

To test the ink of a point inside the window the SCR READ indirection is used. This indirection is set up when GRA INITIALISE or GRA RESET is called.

Related entries:

GRA PLOT
GRA TEST ABSOLUTE
GRA TEST RELATIVE
SCR READ

IND: GRA LINE

#BDE2

Draw a line.

Default action:

Draw a line between the current graphics position and the given endpoint using the current graphics write mode. Points on the line that lie outside the current graphics window will not be plotted. The current graphics position is always moved to the specified endpoint.

Entry conditions:

DE contains the user X coordinate of the endpoint.
HL contains the user Y coordinate of the endpoint.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The position of the endpoint is given in user coordinates, i.e. relative to the user origin.

This indirection is used by GRA LINE RELATIVE and GRA LINE ABSOLUTE to draw the line requested. It is provided to allow the user to change the method for line drawing (to output to an X-Y plotter for example). GRA LINE RELATIVE converts from relative to user coordinates and then calls the indirection; GRA LINE ABSOLUTE calls the indirection immediately.

The line is split up into horizontal or vertical sections that are drawn separately (see SCR HORIZONTAL and SCR VERTICAL). The SCR WRITE indirection is called to write the points in these sections. This means that the line is plotted using the current graphics write mode.

In V1.0 firmware the line is plotted in the current pen ink. But in V1.1 firmware the setting of the line mask determines how pixels on the line will be plotted. The line mask is bit significant and is used repeatedly in the order bit 7, bit 6 down to bit 0 for each pixel in the line. If the bit is one then the pixel is plotted in the graphics pen ink. If the bit is zero then the action taken depends on the graphics background write mode. If the background mode is opaque then the pixel is plotted in the graphics paper ink. If the background mode is transparent then the pixel is not plotted.

In V1.1 firmware the first pixel of the line (that at the current graphics position) is not plotted if the first point plotting mode is set false.

This indirection is set up when GRA INITIALISE or GRA RESET is called.

Related entries:

GRA LINE ABSOLUTE

GRA LINE RELATIVE

GRA SET BACK

GRA SET FIRST

GRA SET LINE MASK

SCR HORIZONTAL

SCR VERTICAL

IND: SCR READ

#BDE5

Read a pixel from the screen.

Default action:

Read a pixel from the screen and decode its ink.

Entry conditions:

HL contains the screen address of the pixel.

C contains the mask for the pixel.

Exit conditions:

A contains the decoded ink that the pixel was set to.

Flags corrupt.

All other, registers preserved.

Notes:

The mask supplied must be a mask for a single pixel otherwise the decoding of the ink read from the screen will not work correctly.

This indirection is set up when SCR INITIALISE or SCR RESET is called. It is called by GRA TEST.

Related entries:

GRA TEST

SCR WRITE

IND: SCR WRITE

#BDE8

Write pixel(s) using the current graphics write mode.

Default action:

Plot a pixel or pixels on the screen using the current graphics write mode.

Entry conditions:

HL contains the screen address of the pixel(s). C contains the mask for the pixel(s).

B contains the encoded ink to plot with.

Exit conditions:

AF corrupt.

All other registers preserved.

Notes:

The pixel mask supplied can be for a single pixel or more than one pixel (or even no pixels). The ink supplied should be encoded to cover the whole of a byte (see SCR INK ENCODE).

The pixel is plotted using the current Graphics VDU write mode. These modes are:

- FORCE Pixel is set to the new ink irrespective of the old ink.
- XOR Pixel is set to the ink formed by exclusive-oring the new ink for the pixel and its current setting.
- AND Pixel is set to the ink formed by anding the new ink for the pixel and its current setting.
- OR pixel is set to the ink formed by oring the new ink for the pixel and its current setting.

The write mode can be set by calling SCR ACCESS appropriately.

This indirection is called by all Graphics VDU write routines, in particular GRA

PLOT RELATIVE, GRA PLOT ABSOLUTE, GRA LINE RELATIVE, GRA LINE ABSOLUTE and GRA WR CHAR, to plot pixels on the screen. It is provided to allow the user to intercept the lowest level of point plotting (perhaps to add yet another plotting mode).

This indirection is set up when SCR INITIALISE or SCR RESET is called.

Related entries:

GRA PLOT

SCR ACCESS

SCR PIXELS

SCR READ

IND: SCR MODE CLEAR

#BDEB

Clear the screen to ink 0.

Default action:

Clear the screen memory to zeros. This indirection is provided to allow the user to prevent the screen being cleared after the mode is changed.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

Normally this indirection performs the actions described in SCR CLEAR.

In V1.0 firmware it is necessary for the user to set up the inks if this indirection is intercepted (see Appendix XIII). In V1.1 firmware the screen pack sets up the inks for the user after SCR MODE CLEAR has been called.

This indirection is set up when SCR INITIALISE or SCR RESET is called. NB. When this indirection is called the text and graphics VDUs are in non-standard states.

Related entries:

SCR CLEAR
SCR SET MODE

IND: KM TEST BREAK

#BDEE

Test for break (or reset).

Default action:

Test if the escape key is pressed, if not then no action is taken. If escape, shift and control are all pressed and no other keys then the system is reset. Otherwise, a break event is reported (see KM BREAK EVENT).

Entry conditions:

Interrupts disabled.
C contains shift and control key states.

Exit conditions:

AF and HL corrupt.
All other registers preserved.

Notes:

This indirection is called by the firmware from the interrupt path. Thus interrupts are disabled and they must remain disabled.

If bit 7 of C is set then the control key is pressed. If bit 5 of C is set then one of the shift keys is pressed.

This indirection is called after the keys have been scanned and the escape key was found to have been pressed. It is provided to allow the user to alter the action of a break (particularly to prevent the system reset, see RESET ENTRY).

This indirection is set up when KM INITIALISE or KM RESET is called.

Related entries:

KM BREAK EVENT

IND: MC WAIT PRINTER

#BDF1

Print a character or time out.

Default action:

Wait for the Centronics port to become not busy and then send a character to it. If the port remains busy for a long time the routine times out and the character is not sent.

Entry conditions:

A contains the character to send.

Exit conditions:

If the character was sent OK:

Carry true.

If the Centronics port timed out:

Carry false.

Always:

A and BC corrupt.

All other registers preserved.

Notes:

This indirection is provided to allow the user to drive the printer in a different way. For example, 'escape sequences' could be handled or the time out could be changed.

This indirection is called by the routine MC PRINT CHAR. It tests whether the printer is busy in the same way as MC BUSY PRINTER and sends the character in the same way as MC SEND PRINTER.

This indirection is set up when MC RESET PRINTER is called.

Related entries:

MC BUSY PRINTER

MC PRINT CHAR

MC SEND PRINTER

IND KM SCAN KEYS

#BDF4

Scan the keyboard.

Default action:

Scans the keyboard and updates the key state map. Newly pressed keys are detected and appropriate markers are inserted into the key buffer.

Entry conditions:

No conditions except that interrupts must be disabled.

Exit conditions:

AF, BC, DE and HL corrupt.

All other registers preserved and interrupts remain disabled.

Notes:

This indirection is called every fiftieth of a second during a ticker interrupt. The repeat speeds and start-up delays of the keys are measured in scans of the keyboard and hence fiftieths of a second.

If the escape key is pressed then the indirection KM TEST BREAK is called to process the break.

Related entries:

KL SCAN NEEDED

KM READ KEY

KM TEST BREAK

KM TEST KEY

17 The High Kernel Jumpblock.

Separate from the main firmware jumpblock is a small jumpblock for Kernel routines associated with ROM state and ROM selection. The routines accessed through this jumpblock are all RAM resident, to avoid confusion while the ROM state and ROM select are changed! The RAM area is copied out of ROM during the power-up initialization. The jumpblock should not be altered by the user.

The entry KL POLL SYNCHRONOUS is the 'odd man out' amongst the routines in this jumpblock. Unlike the other synchronous event handling routines, which are in the lower ROM, this routine is RAM resident. This minimises the overhead involved in polling for synchronous events.

A brief listing of the entries in this jumpblock can be found in section 14.3. A discussion of ROMs and the memory map can be found in section 2, further discussion of ROMs can be found in section 10 and a discussion of events can be found in section 12.

HI: KL U ROM ENABLE

#B900

Enable the upper ROM.

Action:

Enables the currently selected upper ROM. Reading from addresses #C000 and up will now return the contents of the ROM.

Entry conditions:

No conditions.

Exit conditions:

A contains the previous ROM state.
Flags corrupt.

All other registers preserved.

Notes:

The mechanisms provided for calling subroutines in the upper ROM and for selecting upper ROMs automatically enable the ROM as required. This routine is used by the firmware but is otherwise of little use.

The previous ROM state may be passed to KL ROM RESTORE to reset the state to what it was before this routine was called.

This routine enables interrupts.

Related entries:

KL L ROM ENABLE
KL ROM RESTORE
KL ROM SELECT
KL U ROM DISABLE

HI: KL U ROM DISABLE

#B903

Disable the upper ROM.

Action:

Disables the upper ROM. Reading from addresses #C000 and up will now return the contents of the RAM.

Entry conditions:

No conditions.

Exit conditions:

A contains the previous ROM state.

Flags corrupt.

All other registers preserved.

Notes:

Disabling the upper ROM gives read access to the top 16K of RAM, which is usually used as the screen memory. Note that the mapping of a location in screen memory to pixels on the screen depends on the mode and on the screen offset.

It is inadvisable to disable the upper ROM while executing instructions in it!

The previous ROM state may be passed to KL ROM RESTORE to reset the state to what it was before this routine was called.

This routine enables interrupts.

Related entries:

KL L ROM DISABLE

KL ROM RESTORE

KL U ROM ENABLE

HI: KL L ROM ENABLE

#B906

Enable the lower ROM.

Action:

Enables the lower ROM. Reading from addresses below #4000 will now return the contents of the ROM.

Entry conditions:

No conditions.

Exit conditions:

A contains the previous ROM state.

Flags corrupt.

All other registers preserved.

Notes:

In general the lower ROM is disabled except when a firmware routine is called. The firmware jumpblock arranges to enable the lower ROM automatically and to disable it again when the routine returns. This routine is used by the firmware but is otherwise of little use.

The previous ROM state may be passed to KL ROM RESTORE to reset the state to what it was before this routine was called.

This routine enables interrupts.

Related entries:

KL L ROM DISABLE

KL ROM RESTORE.

KL U ROM ENABLE

HI: KL L ROM DISABLE

#B909

Disable the lower ROM.

Action:

Disables the lower ROM. Reading from addresses below #4000 will now return the contents of the RAM.

Entry conditions:

No conditions.

Exit conditions:

A contains the previous ROM state.
Flags corrupt.
All other registers preserved.

Notes:

In general the lower ROM is disabled except when a firmware routine is called. The firmware jumpblock arranges to enable the lower ROM automatically and to disable it again when the routine returns.

The previous ROM state may be passed to KL ROM RESTORE to reset the state to what it was before this routine was called.

This routine enables interrupts.

Related entries:

KL L ROM ENABLE
KL ROM RESTORE
KL U ROM DISABLE

HI: KL ROM RESTORE

#B90C

Restore the previous ROM state.

Action:

The ROM state change routines all return a value giving the previous ROM state. Given that value KL ROM RESTORE will reset the state to what it was before the change.

Entry conditions:

A contains the previous ROM state.

Exit conditions:

AF corrupt.

All other registers preserved.

Notes:

The previous ROM state is the value returned by one of:

- KL U ROM ENABLE
- KL U ROM DISABLE
- KL L ROM ENABLE
- KL L ROM DISABLE
- KL ROM SELECT

It is possible to use KL U ROM DISABLE to reverse the effect of a call of KL U ROM ENABLE (amongst various other combinations). However, calling KL ROM RESTORE is the preferred method since it restores the state to what it was, which might have been enabled anyway.

This routine enables interrupts.

Related entries:

- KL L ROM DISABLE
- KL L ROM ENABLE
- KL ROM SELECT
- KL U ROM DISABLE
- KL U ROM ENABLE

HI: KL ROM SELECT

#B90F

Select a particular upper ROM.

Action:

Select a given upper ROM and enable the upper ROM.

Entry conditions:

C contains the ROM select address of the required ROM.

Exit conditions:

C contains the ROM select address of the previously selected ROM.

B contains the previous ROM state.

AF corrupt.

All other registers preserved.

Notes:

The previous state can be passed to KL ROM RESTORE to reset the ROM enable to what it was. Both the previous state and the previous selection can be passed to KL ROM DESELECT to restore the state to what it was and to select the previously selected ROM again.

The mechanisms provided for calling routines in expansion ROMs automatically perform ROM selection as required see section 21.

It is inadvisable to select another upper ROM whilst executing instructions in the upper ROM.

This routine enables interrupts.

Related entries:

KL CURR SELECTION

KL PROBE ROM

KL ROM DESELECT

KL ROM RESTORE

HI: KL CURR SELECTION

#B912

Ask which upper ROM is currently selected.

Action:

Returns the ROM select address of the currently selected upper ROM.

Entry conditions:

No conditions.

Exit conditions:

A contains the ROM select address of the currently selected ROM.

All other registers and flags preserved.

Notes:

It is not possible to predict the ROM select address at which any particular expansion ROM will be fitted. The 'far address' used to reference subroutines in expansion ROMs includes a ROM select byte which must be set up at run time. This routine returns the ROM select address of the current ROM so that it can set up suitable 'far addresses'.

Related entries:

KL PROBE ROM

KL ROM SELECT

HI: KL PROBE ROM

#B915

Ask class and version of a ROM.

Action:

The first few bytes of all upper ROMs contain information in a standard form about the ROM. This routine extracts the class, mark number and version number bytes from the ROM at the given ROM select address.

Entry conditions:

C contains the ROM select address of the ROM to probe.

Exit conditions:

A contains the ROM's class.

L contains the ROM's mark number.

H contains the ROM's version number.

B and flags corrupt.

All other registers preserved.

Notes:

The ROM class returned may take any of the following values:

- 0: Foreground ROM.
- 1: Background ROM.
- 2: Extension foreground ROM.
- #80: On board ROM (the built in BASIC foreground program).

Selecting a ROM address where no ROM is fitted implicitly selects the on-board ROM and so it will return #80 as its class.

The meaning of the mark and version numbers depends on the ROM.

See section 10 for a description of expansion ROMs.

This routine enables interrupts.

Related entries:

KL ROM SELECT

KL CURR SELECTION

HI: KL ROM DESELECT

#B918

Restore previous upper ROM selection.

Action:

Set the ROM state and upper ROM selection to what they were before KL ROM SELECT was called.

Entry conditions:

C contains the ROM select address of the previously selected ROM.
B contains the previous ROM state.

Exit conditions:

C contains the ROM select address of the currently selected ROM.
B corrupt.
All other registers and flags preserved

Notes:

The previous ROM selection and state are the values returned by KL ROM SELECT. The currently selected ROM returned by this routine is the ROM that was selected by calling KL ROM SELECT (unless further selections have been made).

The mechanisms provided for calling subroutines in expansion ROMs automatically perform ROM selection as required.

It is inadvisable to select another upper ROM whilst executing instructions in the upper ROM.

This routine enables interrupts.

Related entries:

KL CURR SELECTION

KL ROM RESTORE

KL ROM SELECT

HI: KL LDIR

#B91B

Move store (LDIR) with ROMs turned off.

Action:

Performs an LDIR instruction (LoaD Increment and Repeat) with both upper and lower ROMs disabled.

Entry conditions:

BC, DE, HL as required by the LDIR instruction.

Exit conditions:

F, BC, DE, HL as set by the LDIR instruction.
All other registers preserved.

Notes:

This routine may be used to move areas of RAM irrespective of the ROM state.

This routine enables interrupts.

Related entries:

KL LDDR

RAM LAM (RST4)

HI: KL LDDR

#B91E

Move store (LDDR) with ROMs turned off.

Action:

Performs an LDDR instruction (LoaD Decrement and Repeat) with both upper and lower ROMs disabled.

Entry conditions:

BC, DE, HL as required by LDDR instruction.

Exit conditions:

F, BC, DE, HL as set by LDDR instruction.
All other registers preserved.

Notes:

This routine may be used to move areas of RAM irrespective of the ROM state.

This routine enables interrupts.

Related entries:

KL LDIR

RAM LAM (RST4)

HI: KL POLL SYNCHRONOUS

#B921

Check if an event with higher priority than the current event is pending.

Action:

If the synchronous event queue is not empty then the priority of the highest priority pending event is compared with the current event's priority (if any).

Entry conditions:

No conditions.

Exit conditions:

If there is a higher priority event pending:

Carry true.

If there is no higher priority event pending:

Carry false.

Always:

A and other flags corrupt.

All other registers preserved.

Notes:

This routine is in the high jumpblock to minimise the overhead of polling for synchronous events. If the synchronous event queue is empty then the routine takes only a few instructions.

While a synchronous event is being processed the Kernel remembers its priority. The synchronous event routine may itself poll the synchronous event queue, but only events of a higher priority than itself are notified to it.

This routine may enable interrupts.

Related entries:

KL EVENT

KL DONE SYNC

KL DO SYNC

KL NEXT SYNC

HI: KL SCAN NEEDED

#B92A

Ensure keyboard is scanned at next opportunity.

Action:

Force the Key Manager to scan the keyboard when the next ticker interrupt occurs. This may be used to reduce the probability of key pressings being missed while interrupts are disabled.

Entry conditions:

No conditions.

Exit conditions:

AF and HL corrupt.
All other registers preserved.

Notes:

The keyboard is normally scanned on every sixth ticker interrupt (every fiftieth of a second). While interrupts are disabled the ticks are lost and the keyboard will not be scanned. If interrupts are disabled for a significant period (more than three ticks) then this routine should be called just before interrupts are re-enabled. If interrupts are disabled for a long time more than twelve ticks) then the user might consider calling this routine and re-enabling interrupts for a short time every fiftieth of a second.

Related entries:

KM SCAN KEYS

18 The Low Kernel Jumpblock.

The bottom of memory, from #0000 to #003F inclusive, is occupied by the code for the restart (RST) instructions and a number of Kernel entries. Most of these entries are concerned with access to subroutines in ROM and RAM. The RST's are:

RST 0 performs a system reset.

RST instructions 1 to 5 inclusive have been used to extend the Z80 instruction set to provide extra CALL and JUMP instructions, which use addresses extended to include ROM state and ROM select components.

RST 6 is available to the user.

RST 7 is used by interrupts.

Since all the entries supplied must be available whether the lower ROM is enabled or not, the area is copied into RAM from the ROM during power-up initialization.

The user is not intended to alter this jumpblock (except where noted in the USER RESTART and EXT INTERRUPT areas). If the user does change the area then it is the user's responsibility to ensure that the changes do not affect other programs. To some extent this can be achieved by ensuring that the lower ROM is always enabled when other programs are running. However, since the other programs may disable the lower ROM this is insufficient in most cases. Ideally the original jumpblock contents should be restored where there is any doubt.

Section 2 contains a discussion of ROMs and the memory map and section 10 contains a general discussion of external ROMs. A brief list of the routines in this area can be found in section 14.4.

LOW: RESET ENTRY

RST0 #0000

Completely reset the machine as if powered up.

Action:

When the machine is first turned on execution starts here. Calling or jumping to #0000, or executing RST 0, resets the machine to its initial power-up state.

Entry conditions:

No conditions.

Exit conditions:

Does not return!

Notes:

All hardware is reset and the firmware is completely initialized. Once all tables and jumpblocks have been set up, control is passed to the default entry in ROM 0 (see section 10).

Related Entries.

MC START PROGRAM

LOW: LOW JUMP

RST 1 #0008

Jump to lower ROM or RAM, takes inline 'low address' to jump to.

Action:

RST 1 is used to extend the instruction set. It is an expanded form of the jump instruction. It should be followed by a 2 byte 'low address' which specifies the location to jump to and the required ROM state.

Entry conditions:

All registers and flags are passed to the target routine untouched.


Exit conditions:

All registers and flags are as set by the target routine.

Notes:

The 'low address' following the restart instruction is laid out as follows:

Bit: 15 14 13 0



If the 'U' bit is set then the upper ROM is disabled.

If the 'L' bit is set then the lower ROM is disabled.

'Address' is the actual address of the target routine to jump to once the ROM state has been set.

When the target routine returns the ROM state is restored to what it was before the jump. To accomplish this 4 bytes are pushed onto the stack and so care should be taken when indexing up the stack (to find the address of inline parameters, for example).

The LOW JUMP, RST 1, 'instruction' may replace the first byte of a JP (jump) instruction. It is intended for use in jumpblocks. The main firmware jumpblock is made up almost exclusively of LOW JUMP 'instructions'.

It is assumed that the destination of the jump is a routine which will return in the usual way. The restart instruction itself does not return. The value at the top of the stack when a LOW JUMP is executed must, therefore, be a return address.

Executing a LOW JUMP enables interrupts.

Related entries:

FAR CALL (RST3)
FIRM JUMP (RST5)
KL FAR ICALL
KL FAR PCHL
KL LOW PCHL

LOW: KL LOW PCHL

#000B

Jump to lower ROM or RAM.

Register HL contains the 'low address' to jump to.

Action:

Takes a 'low address' as a parameter and jumps to it. The 'low address' specifies both the address to jump to and the ROM state required.

Entry conditions:

HL contains the 'low address' to jump to.


All registers and flags are passed to the target routine untouched.

Exit conditions:

All registers and flags are as set by the target routine.

Notes:

The 'low address' is laid out as follows:



If the 'U' bit is set then the upper ROM is disabled.

If the 'L' bit is set then the lower ROM is disabled.

'Address' is the actual address of the target routine to jump to once the ROM state has been set.

When the target routine returns the ROM state is restored to what it was before the jump. To accomplish this 4 bytes are pushed onto the stack and so care should be taken when indexing up the stack (to find the address of inline parameters, for example).

It is assumed that the destination of the jump is a routine which will return in the usual way. The value at the top of the stack when a LOW PCHL is executed must, therefore, be a return address.

Interrupts are enabled.

Related entries:

KL FAR ICALL

KL FAR PCHL

LOW JUMP (RST1)

PCHL INSTRUCTION

LOW: PCBC INSTRUCTION

#000E

Jump to address in BC.

Action:

Equivalent to the JP (HL) instruction (or PCHL in some assembler dialects), except that the destination is in BC not HL.

Entry conditions:

BC contains the address to jump to.

All registers and flags are passed to the target routine untouched.

Exit conditions:

All registers and flags are as set by the target routine.

Notes:

Calling PCBC INSTRUCTION is a useful way of invoking a routine whose address has been picked out of a table or otherwise established at run time.

Related entries:

KLFAR PCHL
KL LOW PCHL
KL SIDE PCHL
PCDE INSTRUCTION
PCHL INSTRUCTION

LOW: SIDE CALL

RST2 #0010

Call to a sideways ROM, takes inline 'side address' to call.

Action:

RST 2 is used to extend the instruction set. It is an expanded form of the CALL instruction. It should be followed by a 2 byte 'side address' which specifies the location to call and the required ROM selection.

Entry conditions:

All registers and flags are passed to the target routine untouched except for IY (which is set to point at a background ROM's upper data area).

Exit conditions:

IY corrupt.

All other registers and flags are as set by the target routine.

Notes:

The 'side address' following the restart instruction is laid out as follows:

Bit: 15 14 13 0

Off	Address
-----	---------

'Off' gives a value in the range 0.3, which, when added to the ROM select address of the main foreground ROM, gives the ROM select address of the required ROM.

After #C000 has been added to it, 'address' is the address of the routine to call.

The target routine returns to the instruction immediately following the inline 'side address'. The ROM select and ROM state are restored to what they were before the call. To accomplish this 6 bytes are pushed onto the stack and so care should be taken when indexing up the stack (to find the address of inline parameters, for example).

When the target routine is entered the lower ROM is disabled and the appropriate upper ROM is selected and enabled.

SIDE CALLs are provided to support foreground programs split over a number of ROMs (up to four). See section 9 on expansion ROMs.

Interrupts are enabled.

Related entries:

FAR CALL (RST3)

KL SIDE PCHL

LOW: KL SIDE PCHL

#0013

Call to a sideways ROM, HL contains 'side address' to call.

Action:

Takes a 'side address' and calls it. The 'side address' specifies the address of the routine to call and which upper ROM to select.

Entry conditions:

HL contains the 'side address' to call.

All registers and flags are passed to the target routine untouched except for IY (which is set to point at a background ROM's upper data area).

Exit conditions:

IY corrupt:

All other registers and flags are as set by the target routine.

Notes:

The 'side address' is laid out as follows:

Bit: 15 14 13 0

Off	Address
-----	---------

'Off' gives a value in the range 0.3, which, when added to the ROM select address of the main foreground ROM, gives the ROM select address of the required ROM.

After #C000 has been added to it, 'address' is the address of the routine to call.

When the target routine is entered the lower ROM is disabled and the appropriate upper ROM is selected and enabled.

When the target routine returns the ROM select and ROM state are restored to what they were before the call. This is accomplished by pushing 6 bytes onto the stack and so care should be taken when indexing up the stack (to find the address of inline parameters, for example).

Side calls are provided to support foreground programs split over a number of ROMs (up to four). See section 10 on external ROMs.

Interrupts are enabled.

Related entries:

FAR CALL (RST3)

KL FAR ICALL

KL FAR PCHL

LOW: PCDE INSTRUCTION

#0016

Jump to address in DE.

Action:

Equivalent to the JP (HL) instruction (or PCHL in some assembler dialects), except that the destination is in DE not HL.

Entry conditions:

DE contains the address to jump to.
All registers and flags are passed to the target routine untouched.

Exit conditions:

All registers and flags are as set by the target routine.

Notes:

Calling PCDE INSTRUCTION is a useful way of invoking a routine whose address has been picked out of a table or otherwise established at run time.

Related entries:

KL FAR PCHL
KL LOW PCHL
KL SIDE PCHL
PCBC INSTRUCTION
PCHL INSTRUCTION

LOW: FAR CALL

RST3 #0018

Call subroutine in RAM or any ROM, takes inline address of 'far address'.

Action:

RST 3 is used to extend the instruction set. It is an expanded form of the CALL instruction that allows routines to be called anywhere in RAM or in any ROM. The restart is followed by the address of a 3 byte 'far address' which specifies the location to call and the required ROM state and ROM selection.

Entry conditions:

All registers and flags are passed to the target routine untouched except for IY (which is set to point at a background ROM's upper data area).

Exit conditions:

IY preserved.

All other registers and flags are as set by the target routine.

Notes:

The restart instruction takes a 2 byte inline parameter which is the address of a 'far address'. The 'far address' is laid out as follows:

Byte: 0 1 2

Address	ROM
---------	-----

Bytes 0...1 give the address of the routine to call.

Byte 2 is the ROM select byte which takes values as follows:

- #00..#FB: Select the given ROM, enable upper, disable lower.
- #FC: No change of ROM selection. enable upper, enable lower.
- #FD: No change of ROM selection, enable upper, disable lower.
- #FE: No change of ROM selection, disable upper, enable lower.
- #FF: No change of ROM selection, disable upper, disable lower.

The reason that the 'far address' is not contained in the FAR CALL instruction directly is because the ROM select byte for routines in ROM will depend upon the particular configuration of expansion ROMs on the machine and must therefore be established and set at run time.

Registers are passed to the target routine untouched except for the IY register. When entering a background ROM this is set to point at the base of the ROM's upper data area. (See section 10.4 and KL INIT BACK).

The target routine returns to the instruction immediately following the inline parameter. The ROM select and ROM state are restored to what they were before the call. This is accomplished by pushing values on the stack and so care should be taken when indexing up the stack after a FAR CALL instruction. (The stack usage is 4 bytes for ROM select bytes in the range #FC..#FF and 6 bytes for ROM select bytes in the range #00..#FB.)

Interrupts are enabled.

Related entries:

KL FAR ICALL

KL FAR PCHL

LOW JUMP (RST1)

SIDE CALL (RST2)

LOW: KL FAR PCHL

#001B

Call subroutine in RAM or any ROM.
C and HL contain the 'far address' to call.

Action:

The far call mechanism allows subroutines to be called anywhere in RAM or in any ROM. This routine takes a 'far address' and calls the given routine setting the requested ROM state and ROM selection.

Entry conditions:

HL contains the address of the routine to call.
C contains the ROM select byte.

All registers and flags are passed to the target routine untouched except for IY (which is set to point at a background ROM's upper data area).

Exit conditions:

IY preserved.

All other registers and flags are as set by the target routine.

Notes:

The ROM select byte takes values as follows:

- #00.. #FB: Select the given ROM, enable upper, disable lower.
- #FC: No change of ROM selection, enable upper, enable lower.
- #FD: No change of ROM selection. enable upper, disable lower.
- #FE: No change of ROM selection. disable upper, enable lower.
- #FF: No change of ROM selection. disable upper, disable lower.

Registers are passed to the target routine untouched except for the IY index register. When entering a background ROM this is set to point at the base of the ROM's upper data area. (See section 10.4 and KL INIT BACK).

When the target routine returns, the ROM select and ROM state are restored to what they were before the call. This is accomplished by pushing values onto the stack and so care should be taken when indexing up the stack after using this routine. (The stack usage is 4 bytes for ROM select bytes in the range #FC.. #FF and 6 bytes for ROM select bytes in the range #00.. #FB.

Interrupts are enabled.

Related entries:

FAR CALL (RST3)

KL FAR ICALL

KL LOW PCHL

KL SIDE PCHL

LOW: PCHL INSTRUCTION

#001E

Jump to address in HL.

Action:

Entry comprises a JP (HL) instruction (or PCHL in some assembler dialects).

Entry conditions:

HL contains the address to jump to.

All registers and flags are passed to the target routine untouched.

Exit conditions:

All registers and flags are as set by the target routine.

Notes:

Calling PCHL INSTRUCTION is a useful way of invoking a routine whose address has been picked out of a table or otherwise established at run time.

Related entries:

KL FAR PCHL

KL LOW PCHL

KL SIDE PCHL

PCBC INSTRUCTION

PCDE INSTRUCTION

LOW: RAM LAM

RST4 #0020

LD A,(HL) with all ROMs disabled.

Action:

RST 4 is used to extend the instruction set. It is equivalent to a LD A, (HL) instruction except that it always reads from RAM irrespective of whether the ROMs are enabled or not.

Entry conditions:

HL contains the address of the location to read.

Exit conditions:

A contains the value read from the given location.
All other registers and flags preserved.

Notes:

Writing to a location always writes to RAM, even if the location is in one of the ROM areas and the ROM is enabled. The RAM LAM, RST 4, 'instruction' is the read equivalent.

Interrupts are enabled.

Related entries:

KL LDDR

KL LDIR

LOW: KL FAR ICALL

#0023

Call subroutine in RAM or any ROM, HL points at 'far address'.

Action:

The far call mechanism allows subroutines to be called anywhere in RAM or in any ROM. This routine takes the address of a 'far address' and calls the given routine setting the ROM state and ROM selection required.

Entry conditions:

HL contains the address of the 'far address' to call.
All registers and flags are passed to the target routine untouched except for IY (which is set to point at a background ROM's upper data area).

Exit conditions:

IY preserved.
All other registers and flags are as set by the target routine.

Notes:

The parameter passed is the address of a 3 byte 'far address'. This is laid out as follows:

Byte: 0 1 2

Address	ROM
---------	-----

Bytes 0...1 give the address of the routine to call.
Byte 2 is the ROM select byte which takes values as follows:

- #00..#FB: Select the given ROM, enable upper, disable lower.
- #FC: No change of ROM selection. enable upper, enable lower.
- #FD: No change of ROM selection, enable upper, disable lower.
- #FE: No change of ROM selection, disable upper, enable lower.
- #FF: No change of ROM selection, disable upper, disable lower.

Registers are passed to the target routine untouched except for the IY index register. When entering a background ROM this is set to point at the base of the ROM's upper data area. (See section 10.4 and KL INIT BACK).

When the target routine returns, the ROM select and ROM state are restored to what they were before the call. This involves pushing values onto the stack and so care should be taken in indexing up the stack after calling this routine. (The stack usage is 4 bytes for ROM select bytes in the range #FC..#FF and 6 bytes for ROM select bytes in the range #00..#FB.)

Interrupts are enabled.

Related entries:

KL FAR CALL

KL FAR PCHL

LOW: FIRM JUMP

RST5 #0028

Jump to lower ROM, takes inline address to jump to.

Action:

RST 5 is used to extend the instruction set. It is an expanded form of the jump instruction for jumping to routines in the lower ROM or into the central 32K of RAM. The restart is followed by the address of the routine to jump to.

Entry conditions:

All registers and flags are passed to the target routine untouched.

Exit conditions:

All registers and flags are as set by the target routine.

Notes:

The lower ROM is enabled before the jump is taken and is disabled (rather than restored) when the target routine returns. Neither the upper ROM state nor the ROM selection are changed. Two bytes are pushed onto the stack and so care should be taken when indexing up the stack (to find the address of inline parameters, for example).

It is assumed that the destination of the jump is a routine which will return in the usual way. The restart instruction itself does not return, The value at top of stack when a FIRM JUMP is executed must, therefore, be a return address.

The FIRM JUMP, RST 5, 'instruction' may replace the first byte of a JP (jump) instruction, particularly in jumpblocks, much like a LOW JUMP. A FIRM JUMP is slightly faster than a LOW JUMP but a LOW JUMP is more flexible in dealing with ROM states.

Interrupts are enabled.

Related entries:

LOW JUMP (RST1)

LOW: USER RESTART

RST6 #0030

Undedicated RST instruction.

Action:

The eight bytes from #0030 to #0037 inclusive may be patched as required.

Entry conditions:

Unknown.

Exit conditions:

Unknown.

Notes:

If the lower ROM is disabled when an RST 6 instruction is executed then the instructions patched into locations #0030 to #0037 are executed in the normal way.

If the lower ROM is enabled when the RST 6 instruction is executed then the firmware disables the lower ROM and jumps to #0030 to execute the instructions planted by the user.

Generally the lower ROM is disabled except while the firmware is active. Since there are no RST 6s in the firmware there should be no problem about the ROM state when a RST 6 is executed. However, to cope with all eventualities, if the lower ROM is found to be enabled when the restart is executed then the ROM state before the lower ROM is disabled is saved in location #002B. If the lower ROM is found to be disabled then location #002B is left untouched. The value stored is suitable to be passed to KL ROM RESTORE to re-enable the ROM (although KL L ROM ENABLE will have the same effect).

The user can detect whether the lower ROM was enabled when the restart was executed if location #002B is set to zero when the RST 6 area is patched and after processing each restart. If #002B is zero when the RST 6 code is entered then the lower ROM was disabled, and if it is non-zero then the lower ROM was enabled.

The default action for RST 6 as set at power-up is to perform a RST 0, i.e. a system reset.

Related entries:

None.

LOW:INTERRUPT ENTRY RST 7 #0038

Hardware interrupt entry point.

Action:

The Z80 runs in interrupt mode 1, which treats normal interrupts as RST 7 instructions. The firmware interrupt handler looks after the built in regular time interrupt. External interrupts, generated by expansion hardware, are passed on to user supplied software.

Entry conditions:

No conditions.

Exit conditions:

All registers and flags preserved.

Notes:

The user must not use RST 7s as these are dedicated to the processing of interrupts.

If the interrupt is from an external source then the user supplied interrupt routine, EXT INTERRUPT, is called.

See section 10 for a fuller discussion of interrupts.

The user may patch this area (#0038.. #003A inclusive) to intercept interrupts if it is absolutely necessary (see Appendix XI, particularly section c).

Related entries:

EXT INTERRUPT

LOW: EXT INTERRUPT

#003B

External interrupt routine.

Action:

The five bytes from # 003B to # 003F inclusive must be patched by the user if there are going to be any external interrupts. When an external interrupt is detected by the firmware interrupt handler the lower ROM is disabled and the code at #003B is called.

Entry conditions:

No conditions.

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

When the routine is called interrupts are disabled and they must remain disabled. Under no circumstances may the user enable interrupts or use the second register set. Before the routine returns it must clear the interrupt source. See section 11.2 for a discussion of external interrupts.

When an interrupt routine is set up the current contents of #003B.. #003F should be copied elsewhere before they are replaced. If, when the routine is called, it discovers that its hardware is not responsible for the interrupt then it should jump to the copy of the previous external interrupt routine (whose hardware may be responsible).

The purpose of an interrupt routine is to clear the interrupt as quickly as possible, and perhaps perform a minimum of processing. While in the interrupt path no further interrupts are acknowledged. If the interrupt generates a substantial work load, then it should be translated into an event, so that the system is not delayed in the interrupt path for any longer than necessary (see section 11.3).

The interrupt routine must be in RAM at addresses lower than #C000 (as the ROM enable and disable routines cannot be called from the interrupt path).

The default external interrupt routine merely returns. This means that the interrupt will not be cleared and so it will repeat as soon as interrupts are re-enabled. This will cause the machine to 'lock up'.

Related entries:

INTERRUPT ENTRY

KL EVENT

19 AMSDOS BIOS Facilities

AMSDOS uses the CP/M BIOS to access the disc. In order that a program running under AMSDOS may access the disc directly nine of the BIOS extended jumpblock routines are available.

The routines are accessed as external commands. To find the address of the required routines use KL FIND COMMAND. The command names are single control characters (Ctrl A...Ctrl I) as these cannot be typed in from BASIC. An example of how to use these commands can be found in section 10.7.

N.B. The BIOS extended jumpblock itself is not available, indeed it does not exist in an AMSDOS environment.

The BIOS routines available and their command names are as follows:

SET MESSAGE	Ctrl A	(#01)
SET UP DISC	Ctrl B	(#02)
SELECT FORMAT	Ctrl C	(#03)
READ SECTOR	Ctrl D	(#04)
WRITE SECTOR	Ctrl E	(#05)
FORMAT TRACK	Ctrl F	(#06)
MOVE TRACK	Ctrl G	(#07)
GET DR STATUS	Ctrl H	(#08)
SET RETRY COUNT	Ctrl I	(#09)

The word at #BE40 contains the address of the disc parameter header vector. Disc parameter headers and extended disc parameter blocks may be patched as required. For more details on this see section 9.9.

When one of these routines fails (carry false) it returns an error number in the A register. This is referred to in the following pages as the 'error status byte'. If bit 7 of this byte is set then the error has already been reported to the user. Bit 6 is used to distinguish errors generated by the floppy disc controller from those generated by AMSDOS.

If bit 6 is zero then the error numbers are as follows:

#0E	the file is not open as expected.
#0F	hard end of file.
#10	bad command, usually caused by an incorrect filename.
#11	file already exists.
#12	file doesn't exist.
#13	directory is full.
#14	disc is full.
#15	disc has been changed with files open on it.
#16	file is read-only.
# 1A	soft end of file (explained in Appendix XIII).

Also V1.1 machines have error #00

#00	user has hit escape.
-----	----------------------

If bit 6 is one then the error was detected by the floppy disc controller and the other bits are returned as follows:

bit 5	data error	- CRC error on data or ID field
bit 4	overflow error	
bit 3	drive not ready	- there is no disc in the drive
bit 2	no data	- can't find the sector
bit 1	not writable	- disc is write protected
bit 0	address mark missing	

In addition the routine may also return the address of the 'error status buffer'. The first byte of this is a count of the number of bytes received during the results phase of the μ PD765A disc controller, the following bytes are those received. For information on the meaning of these bytes see the manufacturer's documentation.

On the following pages are the interfaces to the intercepted routines:

BIOS: SET MESSAGE

(CTRL/A)

Enable or disable the disc error messages.

Action:

When disc error messages are enabled and an error occurs the BIOS will display error messages on the screen and interact with the user. When disabled no messages are displayed.

Entry conditions:

If messages are to be enabled:

A = #00

If messages are to be disabled:

A = #FF

Exit conditions:

A contains previous state.
HL and flags corrupt.
All other registers preserved.

Notes:

The default state is ENABLED.

Related entries:

SET RETRY COUNT

BIOS: SETUP DISC

(CTRL/B)

Set disc parameters.

Action:

Sets the values for the motor on, motor off, write current off and head settle times. Sends a SPECIFY command to the floppy disc controller.

Entry conditions:

HL contains address of parameter block

Format of the parameter block:

bytes 0, 1	motor on timeout in 20 millisecond units.
bytes 2, 3	motor off timeout in 20 millisecond units.
byte 4	write current off time in 10 microsecond units.
byte 5	head settle time in 1 millisecond units.
byte 6	step rate time in 1 millisecond units.
byte 7	head unload delay (as per μ PD765A SPECIFY command).
byte 8	bits 7...1: head load delay, bit 0: non-DMA mode (as per μ PD765A SPECIFY command).

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

The values given are used for both drives. When using two differing drives use the slower of the two times.

The default values are:

motor on timeout	50
motor off timeout	250
write current off time	175
head settle time	15
step rate time	12
head loadtime	1
head unload time	1
non-DMA mode	1

A motor on time of zero will lock the system up. A motor off time of zero will never turn the motor off.

The standard boot program calls this routine to reset some of the disc parameters as specified in the configuration sector, that is, motor on and off timeouts and the step rate.

Related entries

SELECT FORMAT

SET RETRY COUNT

BIOS: SELECT FORMAT

(CTRL/C)

Select a disc format.

Action:

This routine initializes the extended disc parameter block for the given format. Normally the BIOS automatically detects the format of a disc when SELDSK is called by looking at the sector numbers, but for programs such as formatters it is necessary to pre-set the format.

Entry conditions:

A contains first sector number of required format

#41	system format
#C1	data only format
#01	IBM format

E contains drive number:

#00	A:
#01	B:

Exit conditions:

AF, BC, DE and HL corrupt.
All other registers preserved.

Notes:

Bytes 0...21 of the extended disc parameter block are completely reset, all previous values are lost. Bytes 22...24 (track, align flag, auto-select flag) are not affected. See chapter 2.15.

To set a non-standard format the user may patch the extended disc parameter block directly.

BIOS: READ SECTOR

(CTRL/D)

Read a sector from disc.

Action:

Read the specified sector into store.

Entry conditions:

HL contains address of sector buffer

E contains drive number

#00 A:

#01 B:

D contains track number

C contains sector number

Exit conditions:

If sector read OK:

Carry true.

A contains 0

HL preserved

If failed to read sector correctly:

Carry false.

A contains error status byte (as defined above).

HL contains address of error status buffer.

Always:

Other flags corrupt.

All other registers preserved.

Notes:

The sector buffer may lie anywhere in RAM, even under a ROM.

Related entries:

WRITE SECTOR

BIOS: WRITE SECTOR

(CTRL/E)

Write a sector to disc.

Action:

Write the required sector from store.

Entry conditions:

HL contains address of sector buffer

E contains drive number

#00 A:

#01 B:

D contains track number

C contains sector number

Exit conditions:

If sector written OK:

Carry true.

A contains 0.

HL preserved.

If failed to write sector correctly:

Carry false.

A contains error status byte as defined above.

HL contains address of error status buffer.

Always:

Other flags corrupt.

All other registers preserved.

Notes:

The sector buffer maybe anywhere in RAM, even underneath a ROM.

Related entries:

READ SECTOR

BIOS: FORMAT TRACK

(CTRL/F)

Format an entire track.

Action:

Format a track.

Entry conditions:

HL contains address of header information buffer

E contains drive number

#00 A:

#01 B:

D contains track number

Format of header information:

sector entry for first sector
sector entry for second sector

...

sector entry for last sector

sector entry format:

byte 0: track number
byte 1: head number
byte 2 : sector number
byte 3 : \log_2 (sector size) - 7

Exit conditions:

If track formatted OK:

Carry true.
A contains 0.
HL preserved.

If failed to format track correctly:

Carry false.
A contains error status byte as defined above.
HL contains address of error status buffer.

Always:

Other flags corrupt.
All other registers preserved.

Notes:

The extended DPB must be preset for the required format (see SELECT FORMAT).

Related entries:

SELECT FORMAT

BIOS: MOVE TRACK

(CTRL/G)

Move to specified track.

Action:

Move head to specified track without verifying the move.

Entry conditions:

E contains drive number

#00 A:

#01 B:

D contains track number

Exit conditions:

If moved to the track OK:

Carry true.

A contains 0.

HL preserved.

If failed to move to the track:

Carry false.

A contains error status byte as defined above

HL contains address of error status buffer

Always

Other flags corrupt.

All other registers preserved.

Notes:

This routine is intended as a diagnostic aid and need not normally be used because the read/write/format routines all seek to the correct track automatically.

Related entries:

None.

BIOS: GET DR STATUS

(CTRL/H)

Return status for specified drive.

Action:

This routine returns status register 3 of the floppy disc controller as defined below for the specified drive.

bit 7	undefined	
bit 6	write protect	- The write protect line is true.
bit 5	drive ready	- The ready line is true.
bit 4	track zero	- The track zero line is true.
bit 3	undefined	
bit 2	head address	- always zero.
bit 1	unit select 1	- unit select 1, always zero.
bit 0	unit select 0	- currently selected drive.

Entry conditions:

A contains drive number

#00 A:
#01 B:

Exit conditions:

If carry true

A contains Drive status byte as defined above

HL preserved

If carry false

HL contains address of error status buffer, second byte = Drive status byte
as defined above

A corrupt

Always

Other flags corrupt

All other registers preserved

Notes

This routine returns carry to indicate which set of exit conditions have occurred. No other meaning should be attached to the state of carry.

Related entries:

SELECT FORMAT
READ SECTOR
WRITE SECTOR
FORMAT TRACK
MOVE TRACK

BIOS: SET RETRY COUNT

(CTRL/I)

Set the number of retries for reading/writing/formatting.

Action:

Sets the number of times an operation is retried in the event of an error.

Entry conditions:

A contains new value for retry count

Exit conditions:

A contains old value of retry count.

HL and flags corrupt.

All other registers preserved.

Notes:

The pattern of retries is as follows. Each 'Try' counts one. The retry pattern is repeated until either the operation succeeds or the number of tries has reached the retry count:

- Try
- Try
- Move in one track and back again
- Try
- Move out one track and back again
- Try
- Move to inner track and back again
- Try
- Try
- Move in one track and back again
- Try
- Move out one track and back again
- Try
- Move to outer track and back again
- Repeat

The default value is 16, i.e. twice around the above loop.

Related entries:

READ SECTOR

WRITE SECTOR

FORMAT TRACK

20 AMSDOS External Commands

The disc ROM contained in the 664/6128 and DDI-1 interface contains a number of external commands. These may be found by using KL FIND COMMAND as described in section 10. Facilities available include switching between drive A and B, switching between the tape and disc operating system, cold booting into CPM, displaying a directory of a disc, and erasing and renaming files.

When using the following commands under BASIC they should be preceded by '[' (obtained by pressing shift and @ together) to indicate that they are external commands. When calling KL FIND COMMAND the '[' should not be used.

Since these commands are designed to be called from BASIC they use the parameter block mechanism for passing parameters to them (as described in sections 10.6 and 10.7). Briefly, the parameter block is pointed at by IX and A contains the number of parameters supplied. IX points at the last parameter and preceding parameters are at positive offsets from IX. The value stored in the parameter block for an integer parameter is the integer itself, and the value stored for a string parameter is the address of the string descriptor. The first byte of the string descriptor is the length of the string; the second and third bytes are a pointer to the characters. The routine will corrupt registers AF, BC, DE, HL, IX and IY.

The external commands contained in the disc ROM are:

- A
- B
- CPM
- DIR
- DISC
- DISC.IN
- DISC.OUT
- DRIVE
- ERA
- REN
- TAPE
- TAPE.IN
- TAPE.OUT
- USER

A

Action:

Set the default drive to drive A.

Parameters:

None.

Notes:

This command is equivalent to the DRIVE command with 'A' as a parameter.

This command will fail if AMSDOS is unable to determine the format of the disc in drive A. In which case the default drive will not be changed.

When AMSDOS is initialized the default drive is set to drive A.

Related entries:**B****DRIVE**

B

Action:

Set the default drive to drive B.

Parameters:

None.

Notes:

This command is equivalent to the DRIVE command with 'B' as a parameter.

The command will fail if AMSDOS is unable to determine the format of the disc in drive B. In which case the default drive is not changed.

When AMSDOS is initialized the default drive is set to drive A.

Related entries:**A****DRIVE**

CPM

Action:

This command shuts down BASIC and AMSDOS and cold boots CP/M.

Parameters:

None.

Notes:

MC START PROGRAM is used so all ticker chains etc are lost.

The CP/M utility AMSDOS.COM performs the inverse function and restores AMSDOS and BASIC.

This command can also be used to boot any program that uses the CP/M bootstrap mechanism as described in section 9.5.

Related entries:

None.

DIR

Action:

Display the disc directory.

Parameters:

One optional string parameter.

Notes:

The parameter is a filename, possibly containing wild cards, only those files which match this filename are displayed. If the parameter is omitted then '*.*' is assumed.

The total amount of free space on the disc is shown in Kbytes.

The directory is displayed in as many columns as will fit in the text window of the currently selected stream.

Files marked SYS are not shown.

Files without an extent zero are not shown.

Unlike CAS CATALOG (DISC) the directory is neither sorted nor are the sizes shown. The output is similar to that of the CP/M DIR command.

Related entries:

A

B

DRIVE

USER

DISC

Action:

This command redirects both the tape input and output firmware entries to their disc counterparts.

Parameters:

None

Notes:

The redirected firmware entries are:

CAS IN OPEN
CAS IN CLOSE
CAS IN ABANDON
CAS IN CHAR
CAS IN DIRECT
CAS RETURN
CAS TEST EOF
CAS CATALOG
CAS OUT OPEN
CAS OUT CLOSE
CAS OUT ABANDON
CAS OUT CHAR
CAS OUT DIRECT

DISC is equivalent to the two commands DISC.IN DISC.OUT.

Related entries:

DISC.IN
DISC.OUT
TAPE
TAPE.IN
TAPE.OUT

DISC.IN

Action:

This command redirects the tape input firmware entries to their disc counterparts.

Parameters:

None

Notes:

The redirected firmware entries are:

CAS IN OPEN
CAS IN CLOSE
CAS IN ABANDON
CAS IN CHAR
CAS IN DIRECT
CAS RETURN
CAS TEST EOF
CAS CATALOG

Related entries:

DISC
DISC.OUT
TAPE
TAPE.IN
TAPE.OUT

DISC.OUT

Action:

This command redirects the tape output firmware entries to their disc counterparts.

Parameters:

None.

Notes:

The redirected firmware entries are:

CAS OUT OPEN
CAS OUT CLOSE
CAS OUT ABANDON
CAS OUT CHAR
CAS OUT DIRECT

Related entries

DISC
DISC.IN
TAPE
TAPE.IN
TAPE.OUT

DRIVE

Action:

Set the current default drive.

Parameters:

One string parameter.

Notes:

The string parameter must be a single letter in the range 'A'..'P' or 'a'..'p'.
Drives 'C'..'P' are for future enhancement.

The command will fail if AMSDOS is unable to determine the format of the disc in the requested drive. In which case the default drive will remain unchanged.

When AMSDOS is initialized the default drive is set to drive A.

Related entries:

A

B

ERA

Action:

Erase files.

Parameters:

One string parameter.

Notes:

The string parameter is a filename, possibly containing wild cards. All files which match this filename are erased.

A file which matches the filename but is marked R/O will not be erased. In this event a message is displayed for each 16K (extent) of the file.

If none of the files on the disc match the filename then an error message is displayed.

Related entries:

A

B

DRIVE

USER

REN

Action:

Rename a file.

Parameters:

Two string parameters.

Notes:

The first string parameter is the new name for the file. A file of this name must not already exist. The second parameter is the name of the file to be renamed.

Neither name may contain wild cards.

Both files must be on the same drive.

The files may be in different users.

If the file to be renamed is marked R/O then an error message is displayed and the file is not renamed.

The renamed file will have attributes R/W DIR regardless of the original file's attributes. If the file to be renamed does not exist then an error message is displayed.

Related entries:

A

B

DRIVE

USER

TAPE

Action:

This command restores the tape firmware entries to the state they were before AMSDOS was initialized.

Parameters:

None.

Notes:

The restored firmware entries are:

CAS IN OPEN
CAS IN CLOSE
CAS IN ABANDON
CAS IN CHAR
CAS IN DIRECT
CAS RETURN
CAS TEST EOF
CAS CATALOG
CAS OUT OPEN
CAS OUT CLOSE
CAS OUT ABANDON
CAS OUT CHAR
CAS OUT DIRECT

TAPE is equivalent to the two commands TAPE.IN TAPE.OUT

Note that any patches made to the jumpblock entries for these routines made before a DISC, DISC.IN or DISC.OUT command was executed will be lost. However, any patches made to these entries before AMSDOS was initialized will be restored.

Related entries:

DISC
DISC.IN
DISC.OUT
TAPE.IN
TAPE.OUT

TAPE.IN

Action:

This command restores the tape input firmware entries to the state they were before AMSDOS was initialized.

Parameters:

None.

Notes:

The restored firmware entries are:

CAS IN OPEN
CAS IN CLOSE
CAS IN ABANDON
CAS IN CHAR
CAS IN DIRECT
CAS RETURN
CAS TEST EOF
CAS CATALOG

Note that any patches to the jumpblock entries for these routines made before a DISC or DISC.IN command was executed will be lost. However, any patches made to these entries before AMSDOS was initialized will be restored.

Related entries:

DISC
DISC.IN
DISC.OUT
TAPE
TAPE.OUT

TAPE.OUT

Action:

This command restores the tape output firmware entries to the state they were before AMSDOS was initialized.

Parameters:

None.

Notes:

The restored firmware entries are:

CAS OUT OPEN
CAS OUT CLOSE
CAS OUT ABANDON
CAS OUT CHAR
CAS OUT DIRECT

Note that any patches to the jumpblock entries for these routines made before a DISC or DISC.OUT command was executed will be lost. However, any patches made to these entries before AMSDOS was initialized will be restored.

Related entries:

DISC
DISC.IN
DISC.OUT
TAPE
TAPE.IN

USER

Action:

Set the default user number.

Parameters:

One integer parameter.

Notes:

The user number must be in the range 0.. 15. Any other parameter will cause an error and the default user will remain unchanged. When AMSDOS is initialized the default user number is set to 0.

Related entries:

DIR

ERA

REN

Appendix I

Key Numbering.


The various tables in the Key Manager, such as the translation tables or the repeating key table, are all accessed by key number. The numbering of the keys (and joysticks) is given in the following diagrams:

Note that the physical layout of the 6128 keyboard differs from the 464 and 664 keyboards but that the keys with the same key numbers have the same symbols on the keytops and generate the same values. The keys that have moved are marked with an asterisk.

464/664 Main Keyboard

66	64	65	57	56	49	48	41	40	33	32	25	24	16	79
68	67	59	58	50	51	43	42	35	34	27	26	17	18	
70	69	60	61	53	52	44	45	37	36	29	28	19		
21	71	63	62	55	54	46	38	39	31	30	22	21		
47											23*			

6128 Main Keyboard/Function/Numeric Keypad




66	64	65	57	56	49	48	41	40	33	32	25	24	16	79	10	11	3
68	67	59	58	50	51	43	42	35	34	27	26	17	18	20	12	4	
70	69	60	61	53	52	44	45	37	36	29	28	19			13	14	5
21	71	63	62	55	54	46	38	39	31	30	22	21	15	0	7		
23	9	47							6		8	2	1				


464/664 Function/Numeric Keypad

10	11	3
20	12	4
13	14	5
15	7*	6*

4641664 Cursor Keys



464/664/6128 Joysticks



Note that joystick 1 overlays keys 48..53 on the main keyboard and is indistinguishable from them.

The following table translates key numbers in the opposite direction, from the key number to the inscription on the keytop. Where there is a symbol on the keytop an appropriate word is used, RIGHT for the right cursor key for example. Brackets around the inscription are used to distinguish the various areas of the keyboard as follows:

- {..} Function key (numeric keypad)
- (..) Joystick 0.
- [..] Joystick 1.

	0	1	2	3	4	5	6	7
0	UP	RIGHT	DOWN	{9}	{6}	{3}	{ENTER}	{}
8	LEFT	COPY	{7}	{8}	{5}	{1}	{2}	{0}
16	CLR	[ENTER]	{4}	SHIFT	\	CTRL
24	↑	-	@	P	;	:	/	.
32	0	9	O	I	L	K	M	,
40	8	7	U	Y	H	J	N	SPACE
48	6	5	R	T	G	F	B	V
	[UP]	[DOWN]	[LEFT]	[RIGHT]	[FIRE2]	[FIRE1]	[SPARE]	
56	4	3	E	W	S	D	C	X
64	1	2	ESC	Q	TAB	A	CAPS	Z
72	(UP)	(DOWN)	(LEFT)	(RIGHT)	(FIRE2)	(FIRE1)	(SPARE)	DEL

Appendix II

Key Translation Tables.

See section 3, and section 3.2 in particular, for a description of key translation. Also, Appendix I, which gives the key numbering scheme, maybe of interest.

The diagrams given in this Appendix are for the 464 and 664 keyboards. The 6128 keyboard is similar but some keys have been moved. The new location of the keys can be deduced by inspecting Appendix I.

There are three keyboard translation tables used. These convert a key into its associated character or token. One table is used to translate keys when the control key is pressed, one is used to translate keys when the shift key is pressed or the shift lock is on but the control key is not pressed, the last is used to translate keys when neither shift nor control is pressed.

The diagrams following describe the default translation tables. Where possible the correct character has been placed on the key. The actual value for each of these characters can be found in Appendix VI on the character set. In the cases where the key produces a value which is not a printable ASCII character the abbreviations in the following table will be used. The default settings of the expansion tokens are given in Appendix IV.

Characters and Codes.

NUL	#00	ASCII control code.
SOH	#01	ASCII control code.
STX	#02	ASCII control code.
ETX	#03	ASCII control code.
EOT	#04	ASCII control code.
ENQ	#05	ASCII control code.
ACK	#06	ASCII control code.
BEL	#07	ASCII control code.
BS	#08	ASCII control code.
HT	#09	ASCII control code.
LF	#0A	ASCII control code.
VT	#0B	ASCII control code.
FF	#0C	ASCII control code.
CR	#0D	ASCII control code.
SO	#0E	ASCII control code.
SI	#0F	ASCII control code.

DLE	#10	ASCII control code.
DC1	#11	ASCII control code.
DC2	#12	ASCII control code.
DC3	#13	ASCII control code.
DC4	#14	ASCII control code.
NAK	#15	ASCII control code.
SYN	#16	ASCII control code.
ETB	#17	ASCII control code.
CAN	#18	ASCII control code.
EM	#19	ASCII control code.
SUB	#1A	ASCII control code.
ESC	#1B	ASCII control code.
FS	#1C	ASCII control code.
GS	#1D	ASCII control code.
RS	#1E	ASCII control code.
US	#1F	ASCII control code.
SPACE	#20	ASCII space character.
UP	#5E	Up arrow.
DEL	#7F	ASCII code.
LB	#A3	Pound character.

Expansion Tokens.

F0	#80	Function key 0.
F1	#81	Function key 1.
F2	#82	Function key 2.
F3	#83	Function key 3.
F4	#84	Function key 4.
F5	#85	Function key 5.
F6	#86	Function key 6.
F7	#87	Function key 7.
F8	#88	Function key 8.
F9	#89	Function key 9.

F.	#8A	Function key full stop.
FEN	#8B	Function key enter without control pressed.
FRUN	#8C	Function key enter with control pressed.

Edit and Cursor Codes.

COPY	#E0	Copy key.
INS	#E1	Insert/overwrite toggle key.
WUP	#F0	Write cursor up.
WDN	#F1	Write cursor down.
WLT	#F2	Write cursor left.
WRT	#F3	Write cursor right.
RUP	#F4	Read cursor up.
RDN	#F5	Read cursor down.
RLT	#F6	Read cursor left.
RRT	#F7	Read cursor right.
BEG	#F8	Write cursor to start of text.
END	#F9	Write cursor to end of text.
STA	#FA	Write cursor to start of line.
FIN	#FB	Write cursor to end of line.

System Tokens.

BRK	#FC	Breakkeytoken.
CAPS	#FD	Capslocktoggletoken.
SHIFT	#FE	Shiftlocktoggletoken.
	#FF	Ignore.

Keys that are not marked in the diagrams following generate the system ignore token, if #FF.

Normal Translation Table.

The following diagram describes the translation when neither shift nor control is pressed.

Main Keyboard

BRK	1	2	3	4	5	6	7	8	9	0	-	UP	DLE	DEL
HT	q	w	e	r	t	y	u	i	o	p	@	[CR	
CAPS	a	s	d	f	g	h	j	k	l	:	;]		
	z	x	c	v	b	n	m	,	.	/	\			
S P A C E														

Function/Numeric Keypad

F7	F8	F9
F4	F5	F6
F1	F2	F3
F0	F.	FEN

Cursor Keys

	WUP	
WLT	COPY	WRT
	WDN	

JOYSTICK 0

	VT	
BS	FIRE2 x	FIRE1 z
	LF	HT

JOYSTICK 1

	6	
r	FIRE2 f	FIRE1 g
	5	t

Shift Translation Table.

The following diagram describes the translation when either shift key is pressed, or the shift lock is on, but the control key is not pressed.

Main Keyboard

BRK	!	"	#	\$	%	&	'	()	_	=	LB	DLE	DEL
HT	Q	W	E	R	T	Y	U	I	O	P		{	CR	
CAPS	A	S	D	F	G	H	J	K	L	*	+	}		
	Z	Z	C	V	B	N	M	<	>	?	:			
S P A C E														

Function/Numeric Keypad

F7	F8	F9
F4	F5	F6
F1	F2	F3
F0	F.	FEN

Cursor Keys

	RUP	
RLT	COPY	RRT
	RDN	

JOYSTICK 0

	VT		
BS	FIRE2 X	FIRE1 Z	HT
	LF		

JOYSTICK 1

	%		
R	FIRE2 G	FIRE1 F	T
	&		

Control Translation Table.

The following diagram describes the translation when the control key is pressed.

Main Keyboard

BRK		~									US		RS	DLE	DEL		
INS	DC1	ETB	ENQ	DC2	DC4	EM	NAK	HT	SI	DLE	NUL	ESC	CR				
SHIFT	SOH	DC3	EOT	ACK	BEL	ES	LF	VT	FF			GS					
	SUB	CAN	ETX	SYN	STX	SO	CR					FS					
<table border="1" style="margin: auto;"> <tr> <td></td><td></td> </tr> </table>																	


Function/Numeric Keypad

F7	F8	F9
F4	F5	F6
F1	F2	F3
F0	F.	FEN


Cuxsor Keys

	BEG	
STA	COPY	FIN
	END	

JOYSTICK 0



JOYSTICK 1




Appendix III

Repeating Keys.


Which keys are allowed to repeat may be set by the user. See section 3 (and section 3.5 in particular) for a full description of repeating keys. Also, see Appendix I which gives the key numbering scheme and indicates which keys have moved on the 6128 keyboard.

The default repeating key table is described in the following diagrams. Keys which are not allowed to repeat are marked with an asterisk.


Main Keyboard




Function/Numeric Keypad




Cursor Keys



JOYSTICK 0



JOYSTICK 1



Appendix IV

Function Keys and Expansion Strings.

Function keys are more fully explained in section 3, and in section 3.7 in particular. The following table specifies the default string for each expansion token and which key the token is associated with by default.

Token	Value	Default String	Default Key
0	#80	0	Function Key 0.
1	#81	1	Function Key 1.
2	#82	2	Function Key 2.
3	#83	3	Function Key 3.
4	#84	4	Function Key 4.
5	#85	5	Function Key 5.
6	#86	6	Function Key 6.
7	#87	7	Function Key 7.
8	#88	8	Function Key 8.
9	#89	9	Function Key 9.
10	#8A	.	Function Key full stop.
11	#8B		Function Key enter.
12	#8C	RUN ”	Function Key enter with control.
13..31	#8D..#9F		None.

Tokens 13..31 are all set to empty strings and none of them are defaulted to associate with a key.

stands for carriage return (Character #0D)

Appendix V

Inks and Colours.

A full discussion of inks and colours can be found in section 6.2. This appendix lists the colours that are available and the default settings for the inks.

There are 27 colours available. The Screen Pack refers to these colours by a grey scale number so that colour 0 is the darkest and colour 26 is the brightest. The hardware requires these grey scales to be translated into the hardware code for the colour. It is unlikely that the user will ever need to deal with the hardware numbers, they are merely given for information.

Grey Scale	Colour	HW Number
0	Black	20
1	Blue	4
2	Bright blue	21
3	Red	28
4	Magenta	24
5	Mauve	29
6	Bright red	12
7	Purple	5
8	Bright magenta	13
9	Green	22
10	Cyan	6
11	Skyblue	23
12	Yellow	30
13	White	0
14	Pastel blue	31
15	Orange	14
16	Pink	7
17	Pastel magenta	15
18	Bright green	18
19	Sea green	2
20	Bright cyan	19
21	Lime	26
22	Pastel green	25
23	Pastel cyan	27
24	Bright yellow	10
25	Pastel yellow	3
26	Bright white	11

The user can set the colours in which the 16 inks and the border are displayed. The following table gives the default settings:

Ink	Colour	Colour Numbers
Border	Blue	1/1
0	Blue	1/1
1	Bright yellow	24/24
2	Bright cyan	20/20
3	Bright red	6/6
4	Bright white	26/26
5	Black	0/0
6	Bright blue	2/2
7	Bright magenta	8/8
8	Cyan	10/10
9	Yellow	12/12
10	Pastel blue	14/14
11	Pink	16/16
12	Bright green	18/18
13	Pastel green	22/22
14	Flashing blue / bright yellow	1/24
15	Flashing sky blue / pink	11/16

Appendix VI

















Displayed Character Set.

There are 256 symbols in the displayed character set. All of these can be printed, although it requires special effort to print characters 0..31 which are often interpreted as control codes. The user can set the matrix for any or all characters (see section 4.6). The following lists describe the default character set.

The character set is split into a number of areas for ease of description:

0..31	(#00.. # 1F)	ASCII control codes.
32..127	(#20.. #7F)	ASCII characters.
128..143	(#80.. #8F)	Block graphics.
144..159	(#90..#9F)	Line graphics.
160..191	(#A0.. #BF)	Further characters.
192..255	(#C0.. #FF)	Miscellaneous graphic symbols.

a. ASCII Control Codes.

0	#00	NUL		Square.
1	#01	SOH		Upside down L.
2	#02	STX		Upside down T.
3	#03	ETX		Backwards L.
4	#04	EOT		Lightning flash.
5	#05	ENQ		Square with a diagonal cross.
6	#06	ACK		Tick.
7	#07	BEL		Bell (semi-circle with feet).
8	#08	BS		Left pointing arrow.
9	#09	HT		Right pointing arrow.
10	#0A	LF		Down pointing arrow.
11	#0B	VT		Up pointing arrow.
12	#0C	FF		Christmas tree (down pointing arrow with a tail).
13	#0D	CR		Bent left pointing arrow.
14	#0E	SO		Circle with a diagonal cross.
15	#0F	SI		Circle with a central dot.

16	#10	DLE	☐	Square with a horizontal bar.
17	#11	DC1	⊙	Circle with three o'clock.
18	#12	DC2	⊙	Circle with half past three.
19	#13	DC3	⊙	Circle with half past nine.
20	#14	DC4	⊙	Circle with nine o'clock.
21	#15	NAK	✘	Crossed out tick.
22	#16	SYN	⏏	Square wave.
23	#17	ETB	├	Sideways T.
24	#18	CAN	⌚	Hourglass.
25	#19	EM	⚬	Vertical bar with a central blob.
26	#1A	SUB	❏	Backwards question mark.
27	#1B	ESC	⊙	Circle with a horizontal bar.
28	#1C	FS	⏏	Square with nine o'clock.
29	#1D	GS	⏏	Square with half past nine.
30	#1E	RS	⏏	Square with half past three.
31	#1F	US	⏏	Square with three o'clock.

b. ASCII Characters.


Characters 32-127 (#20.. #7F) are listed in the following table. They make up the standard ASCII character set.

















	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
32 #20		!	"	#	\$	%	&	'	()	*	+	,	-	.	/
48 #30	0	1	2	3	4	5	6	7	8	9	:	;	<	=	>	?
64 #40	@	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
80 #50	P	Q	R	S	T	U	V	W	X	Y	Z	[\]	↑	-
96 #60		a	b	c	d	e	f	g	h	i	j	k	l	m	n	o
112 #70	p	q	r	s	t	u	v	w	x	y	z	{		}	~	

c. Block Graphics.

Characters 128..143 (#80..#8F) are a set of block graphics. Each character is divided into four cells. Bits 0..3 of the character number determine which cells are filled. If the appropriate bit is set then the cell is filled in, otherwise it is left blank. The cells are:

Bit 0	Bit 1
Bit 2	Bit 3


In the following list cells that are filled are marked with an , cells that are blank are marked with an open square.

128	#80	Block graphic	
129	#81	Block graphic	
130	#82	Block graphic	
131	#83	Block graphic	
132	#84	Block graphic	
133	#85	Block graphic	
134	#86	Block graphic	
135	#87	Block graphic	
136	#88	Block graphic	
137	#89	Block graphic	
138	#8A	Block graphic	
139	#8B	Block graphic	
140	#8C	Block graphic	
141	#8D	Block graphic	
142	#8E	Block graphic	
143	#8F	Block graphic	

d. Line Graphics

Characters 144..159 (#90..#9F) are a set of line graphics. The lines join the centre of the character to the centre of an edge. Each of the lines is associated with a bit of the character number. If the bit is set then the line is present, if the bit is not set then the line is not present. The central block of the character is always set.

The lines are associated with bits as follows:



144	#90	Line graphic	
145	#91	Line graphic	
146	#92	Line graphic	
147	#93	Line graphic	
148	#94	Line graphic	
149	#95	Line graphic	
150	#96	Line graphic	
151	#97	Line graphic	
152	#98	Line graphic	
153	#99	Line graphic	
154	#9A	Line graphic	
155	#9B	Line graphic	
156	#9C	Line graphic	
157	#9D	Line graphic	
158	#9E	Line graphic	
159	#9F	Line graphic	

























e. Further Characters.

160	#A0	^	Circumflex
161	#A1	´	Acute accent
162	#A2	¨	Umlaut
163	#A3	£	Pound
164	#A4	©	Copyright
165	#A5	¶	Pilcrow
166	#A6	§	Section
167	#A7	‘	Open single quote (pairs with character 39)
168	#A8	¼	One quarter
169	#A9	½	One half
170	#AA	¾	Three quarters
171	#AB	±	Plus or minus
172	#AC	÷	Division
173	#AD	¬	Not
174	#AE	¿	Inverted question mark
175	#AF	¡	Inverted exclamation mark
176	#B0	α	Lower case alpha
177	#B1	β	Lower case beta
178	#B2	γ	Lower case gamma
179	#B3	δ	Lower case delta






180	#B4	€	Lower case epsilon
181	#B5	θ	Lower case theta
182	#B6	λ	Lower case lambda
183	#B7	μ	Lower case mu
184	#B8	π	Lower case pi
185	#B9	σ	Lower case sigma
186	#BA	φ	Lower case phi
187	#BB	ψ	Lower case psi
188	#BC	χ	Lower case chi
189	#BD	ω	Lower case omega
190	#BE	Σ	Upper case sigma
191	#BF	Ω	Upper case omega

f. Miscellaneous Graphics Symbols.

192	#C0	↙	Diagonal line joining top to left.
193	#C 1	↘	Diagonal line joining top to right.
194	#C2	↗	Diagonal line joining bottom to right.
195	#C3	↖	Diagonal line joining bottom to left.
196	#C4	↖↘	Diagonal lines joining top to left and right.
197	#C5	↘↗	Diagonal lines joining right to top and bottom.
198	#C6	↗↖	Diagonal lines joining bottom to right and left.
199	#C7	↖↘	Diagonal lines joining left to top and bottom.
200	#C8	↖↘↗↖	Diagonal lines joining top to left and bottom to right.
201	# C9	↘↗↖↘	Diagonal lines joining top to right and bottom to left.
202	# CA	◊	Diamond joining all edges.

203	#CB		Both major diagonals (Large X)
204	#CC		Forwards major diagonal (Large slash)
205	#CD		Backwards major diagonal (large backslash)
206	#CE		Chequered pattern
207	#CF		Shading
208	#DO		Line along top edge
209	#D1		Line along right edge
210	#D2		Line along bottom edge
211	#D3		Line along left edge
212	#D4		Triangle filling top left corner
213	#D5		Triangle filling top right corner
214	#D6		Triangle filling bottom right corner
215	#D7		Triangle filling bottom left corner
216	#D8		Top half shaded
217	#D9		Right half shaded
218	#DA		Bottom half shaded
219	#DB		Left half shaded
220	#DC		Shaded triangle filling top left corner
221	#DD		Shaded triangle filling top right corner
222	#DE		Shaded triangle filling bottom right corner
223	#DF		Shaded triangle filling bottom left corner
224	#E0		Happy face
225	#E1		Sad face
226	#E2		Club

227	#E3	◆	Diamond
228	#E4	♥	Heart
229	#E5	♠	Spade
230	#E6	○	Empty circle
231	#E7	●	Filled circle
232	#E8	□	Empty square
233	#E9	■	Filled square
234	#EA	♂	Male (Mars).
235	#EB	♀	Female (Venus).
236	#EC	♣	Crochet
237	#ED	♪	Quaver
238	#EE	✳	Star
239	#EF	🚀	Rocket
240	#F0	↑	Up pointing arrow head
241	#F1	↓	Down pointing arrow head
242	#F2	←	Left pointing arrow head
243	#F3	→	Right pointing arrow head
244	#F4	▲	Up pointing triangle
245	#F5	▼	Down pointing triangle
246	#F6	▶	Right pointing triangle
247	#F7	◀	Left pointing triangle
248	#F8	🕺	Dancing person standing
249	#F9	🕸	Dancing person doing splits
250	#FA	🕺	Dancing person with left leg out

251	#FB		Dancing person with right leg out
252	#FC		Bomb
253	#FD		Mushroom cloud
254	#FE		Up and down arrow
255	#FF		Right and left arrow

Appendix VII

Text VDU Control Codes.

Character values in the range 0-31 sent to the main Text VDU output routine (TXT OUTPUT) do not produce a character on the screen, but are interpreted as control codes. These codes may affect the meaning of one or more of the following characters, which are the code's parameters.

All control codes work on the currently selected stream unless otherwise indicated. For instance, setting the pen, code 15, sets the text pen ink for the currently selected stream whilst setting the colour of an ink, code 28, will affect all streams (and the Graphics VDU).

Certain codes force the current position (the cursor position) to a legal position inside the current window before they are obeyed. This is explained in more detail in section 4.5. The cursor may be left in an illegal position.

The following table specifies the default actions for the control codes. By changing entries in the control code table the action of these codes can be altered as desired. See section 4.7 for a full description.

Code	Name	Params	Action
0	NUL	0	In V1.1 firmware: Force the cursor to a legal position(see TXTVALIDATE). In V1.0 firmware: No effect.
1	SOH	1	Print the character given by the parameter (see TXT WR CHAR). This allows characters 0..31 to be printed.
2	STX	0	Disable the cursor blob (see TXT CUR DISABLE). Reverses the effect of ETX (code 3).
3	ETX	0	Enable the cursor blob (see TXT CUR ENABLE). Reverses the effect of STX (code 2).
4	EOT	1	Set the screen mode given by the parameter (see SCR SET MODE). The parameter is taken MOD 4 and the value 3 is ignored: 0 sets mode 0 (160 x 200). 1 sets mode 1 (320 x 200). 2 sets mode 2 (640 x 200).
5	ENQ	1	Print the character given by the parameter using the Graphics VDU as if the graphic character write mode was active (see TXT SET GRAPHIC and GRA WR CHAR).

6	ACK	0	Enable the VDU (see TXT VDU ENABLE). Reverses the effect of NAK (code 2 1).
7	BEL	0	Makes a short bleep, sound. Note that this flushes the sound queues.
8	BS	0	Make the current position legal then move left one character.
9	TAB	0	Make the current position legal then move right one character.
10	LF	0	Make the current position legal then move down one line.
11	VT	0	Make the current position legal then move up one line.
12	FF	0	Clear the current window and move the current position to the top left corner (see TXT CLEAR WINDOW).
13	CR	0	Make the current position legal and then move it to the left edge of the window on the current line (see TXT SET COLUMN).
14	SO	1	Set the paper ink to the ink given by the parameter (see TXT SET PAPER). Parameter is taken MOD 16.
15	SI	1	Set the pen ink to the ink given by the parameter (see TXT SET PEN). Parameter is taken MOD 16.
16	DLE	0	Make the current position legal then clear it to the current paper ink.
17	DC1	0	Make the current position legal then clear from the left edge of the window to the current position inclusive. The affected cells are set to the current paper ink.
18	DC2	0	Make the current position legal then clear from it to the right edge of the window inclusive. The affected cells are set to the current paper ink.
19	DC3	0	Make the current position legal then clear from the start of the window to the current position inclusive. The affected cells are set to the current paper ink.
20	DC4	0	Make the current position legal then clear from it to the end of the window inclusive. The affected cells are set to the current paper ink.
21	NAK	0	Disable the VDU (see TXT VDU DISABLE). Reverses the effect of ACK (code 6).

22	SYN	1	<p>Set the character write mode from the parameter (see TXT SET BACK). The parameter is taken MOD 2 and:</p> <p style="padding-left: 40px;">0 sets opaque mode (the default mode). 1 sets transparent mode.</p>
23	ETB	1	<p>Set the Graphics VDU write mode from the parameter (see SCR ACCESS). The parameter is taken MOD 4 and:</p> <p style="padding-left: 40px;">0 sets FORCE mode (the default mode). 1 sets XOR mode. 2 sets AND mode. 3 sets OR mode.</p>
24	CAN	0	<p>Exchange the current pen and paper inks (see TXT INVERSE).</p>
25	EM	9	<p>Set the matrix for a character (see TXT SET MATRIX). The first parameter specifies which character is to be set. The next 8 parameters are the matrix for the character (given top to bottom). If the character is not user definable then no action is taken.</p>
26	SUB	4	<p>Set the limits of the text window (see TXT WIN ENABLE). The first two parameters specify the left and right columns of the window (the smaller is the left column); the last two parameters specify the top and bottom rows of the window (the smaller is the top row).</p>
27	ESC	0	<p>No effect - available for user.</p>
28	FS	3	<p>Set the colours in which to display an ink (see SCR SET INK). The first parameter is taken MOD 16 and specifies which ink is to be set, The second and third parameters are taken MOD 32 and specify the two colours for the ink.</p>
29	GS	2	<p>Set the colours with which to display the border (see SCR SET BORDER). The two parameters are taken MOD 32 and specify the two colours for the border.</p>
30	RS	0	<p>Move the current position to the top left corner of the window (see TXT SET CURSOR).</p>
31	US	2	<p>Move the current position to a given position in the current window (see TXT SET CURSOR). The first parameter specifies the column to move to, the second parameter specifies the row to move to (row 1, column 1 is the top left corner of the window).</p>

Appendix VIII

Notes and Tone Periods.

The tables which follow give the recommended tone period settings for notes in the even tempered scale for the full eight octave range. The period is calculated from the note frequency as follows (since the period is given in 8 microsecond units):

$$\text{Period} = 125000/\text{Frequency}$$

The frequency for each note is calculated from International A as follows:

$$\text{Frequency} = 440 * (2 ^ {(\text{Octave} + (\text{N} - 10) / 12)})$$

where:

- Octave is the octave number. 0 is the octave including International A (and middle C), -1 is the octave below, + 1 is the octave above etc.
- N is the note number. 1 is C, 2 is C#, 3 is D etc.

The period is an integer value and so the frequency of the note produced is not exactly the required frequency. The relative error is given in the tables below. This is calculated as follows:

$$\text{Error} = (\text{Required frequency} - \text{Actual frequency}) / \text{Required frequency}$$

Note	Frequency	Period	Error	Octave -3
C	32.703	3822	#0EEE	-0.007%
C#	34.648	3608	#0E18	+0.007%
D	36.708	3405	#0D4D	-0.007%
D#	38.891	3214	#0C8E	-0.004%
E	41.203	3034	#0BDA	+0.009%
F	43.654	2863	#0B2F	-0.016%
F#	46.249	2703	#0A8F	+0.009%
G	48.999	2551	#09F7	-0.002%
G#	51.913	2408	#0968	+0.005%
A	55.000	2273	#08E1	+0.012%
A#	58.270	2145	#0861	-0.008%
B	61.735	2025	#07E9	+0.011%

Note	Frequency	Period		Error	Octave -2
C	65.406	1911	#0777	-0.007%	
C#	69.296	1804	#070C	+0.007%	
D	73.416	1703	#06A7	+0.022%	
D#	77.782	1607	#0647	-0.004%	
E	82.407	1517	#05ED	+0.009%	
F	87.307	1432	#0598	+0.019%	
F#	92.499	1351	#0547	-0.028%	
G	97.999	1276	#04FC	+0.037%	
G#	103.826	1204	#04D4	+0.005%	
A	110.000	1136	#0470	-0.032%	
A#	116.541	1073	#0431	+0.039%	
B	123.471	1012	#03F4	-0.038%	

Note	Frequency	Period		Error	Octave -1
C	130.813	956	#03DC	+0.046%	
C#	138.591	902	#0386	+0.007%	
D	146.832	851	#0353	-0.037%	
D#	155.564	804	#0324	+0.058%	
E	164.814	758	#02F6	-0.057%	
F	174.614	716	#02CC	+0.019%	
F#	184.997	676	#02A4	+0.046%	
G	195.998	638	#027E	+0.037%	
G#	207.652	602	#025A	+0.005%	
A	220.000	568	#0238	-0.032%	
A#	233.082	536	#0218	-0.055%	
B	246.942	506	#01FA	-0.038%	

Note	Frequency	Period		Error	Octave 0
C	261.626	478	#01DE	+0.046%	Middle C
C#	277.183	451	#01C3	+0.007%	
D	293.665	426	#01AA	+0.081%	
D#	311.127	402	#0192	+0.058%	
E	329.628	379	#017B	-0.057%	
F	349.228	358	#0166	+0.019%	
F#	369.994	338	#0152	+0.046%	
G	391.995	319	#013F	+0.037%	
G#	415.305	301	#012D	+0.005%	
A	440.000	284	#011C	-0.032%	
A#	466.164	268	#010C	-0.055%	
B	493.883	253	#00FD	-0.038%	

Note	Frequency	Period		Error	Octave 1
C	523.251	239	#00EF	+0.046%	
C#	554.365	225	#0E1	-0.215%	
D	587.330	213	#00D5	+0.081%	
D#	622.254	201	#00C9	+0.058%	
E	659.255	190	#0BE	+0.206%	
F	698.457	179	#00B3	+0.019%	
F#	739.989	169	#00A9	+0.046%	
G	783.991	159	#009F	-0.277%	
G#	830.609	150	#0096	-0.328%	
A	880.000	142	#008E	-0.032%	
A#	932.328	134	#0086	-0.055%	
B	987.767	127	#007F	+0.356%	

Note	Frequency	Period		Error	Octave 2
C	1046.502	119	#0077	-0.374%	
C#	1108.731	113	#0071	+0.229%	
D	1174.659	106	#006A	-0.390%	
D#	1244.508	100	#0064	-0.441%	
E	1318.510	95	#005F	+0.206%	
F	1396.913	89	#0059	-0.543%	
F#	1479.978	84	#0054	-0.548%	
G	1567.982	80	#0050	+0.350%	
G#	1661.219	75	#004B	-0.328%	
A	1760.000	71	#0047	-0.032%	
A#	1864.655	67	#0043	-0.055%	
B	1975.533	63	#003F	-0.435%	

Note	Frequency	Period		Error	Octave 3
C	2093.004	60	#003C	+0.462%	
C#	2217.461	56	#0038	-0.662%	
D	2349.318	53	#0035	-0.390%	
D#	2489.016	50	#0032	-0.441%	
E	2637.021	47	#002F	-0.855%	
F	2793.826	45	#002D	+0.574%	
F#	2959.955	42	#002A	-0.548%	
G	3135.963	40	#0028	+0.350%	
G#	3322.438	38	#0026	+0.992%	
A	3520.000	36	#2924	+1.357%	
A#	3729.310	34	#0022	+1.417%	
B	3951.066	32	#0020	+1.134%	

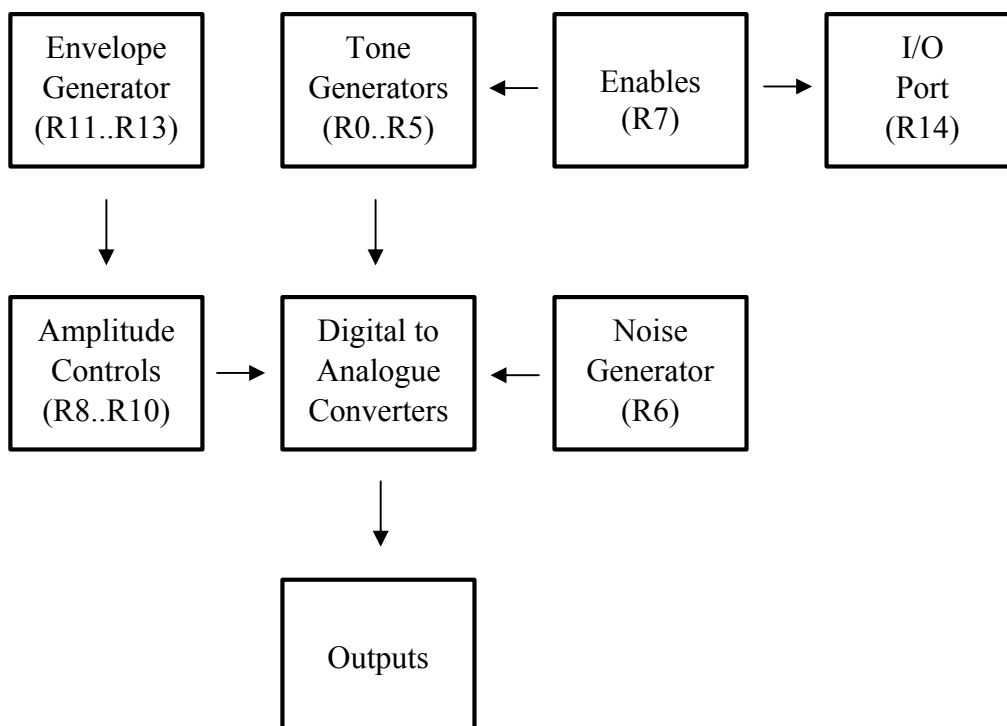
Appendix IX

The Programmable Sound Generator.

The programmable sound generator (PSG) is an AY-3-8912 chip. This is briefly described in section 7.1. The PSG has a number of registers which are described below. This information is provided for the interest of the user, particularly if hardware enveloping is to be used (in which case section (e) will be of special interest). However, the software enveloping provided by the Sound Manager can achieve all that the sound chip is capable of unless very short attacks or decays are required.

If the user is intending to drive the sound chip directly rather than by using the Sound Manager then the information presented is not complete and the user should consult the manufacturer's data sheet. The user is advised to call the routine MC SOUND REGISTER to write data to a sound chip register as this obeys the timing constraints on access to the sound chip.

The following diagram indicates the interactions between the various sections of the sound chip:



The sound chip data registers are as follows:

Register 0:	Channel A tone period fine tune.
Register 1:	Channel A tone period coarse tune.
Register 2:	Channel B tone period fine tune.
Register 3:	Channel B tone period coarse tune.
Register 4:	Channel C tone period fine tune.
Register 5:	Channel C tone period coarse tune.
Register 6:	Noise period.
Register 7:	Enables and I/O direction.
Register 8:	Channel A amplitude and envelope enable.
Register 9:	Channel B amplitude and envelope enable.
Register 10:	Channel C amplitude and envelope enable.
Register 11:	Envelope period fine tune.
Register 12:	Envelope period coarse tune.
Register 13:	Envelope shape.
Register 14:	Input from or output to port A.
Register 15:	Not used.

a. Tone Generators (Registers 0..5)

Each channel has two tone period registers associated with it. These set the period of the sound to be generated (in units of 8 microseconds) by that channel. The fine tune register stores the least significant 8 bits of the period; the coarse tune register stores the most significant 4 bits of the period. To include the tone in the output of a channel the appropriate bit in the enables register must be cleared.

b. Noise Generator (Register 6).

There is a single pseudo-random noise source. The output from this can be included in the output of any of the three channels (as specified by the enables register). The period of the noise generator is set by bits 0..4 of the noise period register. The period specifies the middle frequency of the noise produced in 8 microsecond units.

c. Enables (Register 7).

The enables register specifies whether tone or noise is to be included in the output from each channel. It also specifies whether the I/O port is to act in input or in output mode. The bits are allocated as follows:

Bit 0:	Channel A tone disable.
Bit 1:	Channel B tone disable.
Bit 2:	Channel C tone disable.
Bit 3:	Channel A noise disable.
Bit 4:	Channel B noise disable.
Bit 5:	Channel C noise disable.
Bit 6:	Port A output mode.
Bit 7:	Not used.

Note that port A is connected to the keyboard and joystick and so the port must always be in input mode. The user must ensure that bit 6 of the enables register is always set to zero.


d. Amplitude Controls (Registers 8.. 10).


Each channel has an amplitude control register associated with it. Bit 4 of this register specifies whether hardware enveloping is to be used for the channel. If the bit is set then the channel amplitude (volume) is under the control of the hardware envelope generator. If the bit is clear then the amplitude is set by bits 0..3 of the register - a value of 0 means no sound and a value of 15 means maximum volume.

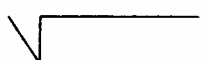
e. Envelope generator (Registers 11..13).

The sound chip has a single hardware envelope generator which can be used to control any combination of the three sound channels as specified by the channel's amplitude register (see (d) above). Bits 0 to 3 of register 13 control the shape of the envelope in a rather unobvious manner. The following table gives values required to generate each of the 8 hardware envelopes that are possible. Other values (0..7) duplicate envelopes 9 and 15.

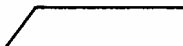


8:  Repeated jump up and ramp down.

9:  Jump up and ramp down once then hold at minimum volume (zero).

10:  Jump up then repeatedly ramp down and up again.

11:  Jump up and ramp down then jump up and hold at maximum volume (fifteen).

12:  Repeatedly ramp up and drop down.

- 13:  Ramp up then hold at maximum volume (fifteen)
- 14:  Repeatedly ramp up and down again.
- 15:  Ramp up and drop down once then hold at minimum volume (zero).

The length of each of the ramps, upwards or downwards, is set by the envelope period. The envelope period is a full 16 bit value whose less significant byte is stored in register 11 and whose more significant byte is stored in register 12. The period is given in 128 microsecond units and is the time between steps in the ramp. Since the ramp has 16 steps (corresponding to the 16 volume settings) the total time taken for the ramp is the envelope period times 1024 microseconds (i.e. the envelope period approximately sets the total time for the ramp in milliseconds).

f. I/O Port (Register 14).

The mode of operation of the PSG port is set by a bit in the enables register (see section (c) above). However, since port A is dedicated to reading the keyboard and joysticks it should always be operated in input mode. The port may be read by reading the contents of register 14. However, scanning the keyboard is a complex action and is best left to the Key Manager which provides ample facilities for access to the keys.

References to port B in the manufacturer's data sheet should be ignored as the AY-3-8912 is a version of the chip that does not have port B.

Appendix X

Kernel Block Layouts.


The user provides a number of blocks to the Kernel for various puposes. The layouts of these blocks are described below, mainly for the interest of the user. There are very few occasions when the user is allowed to write to one of these blocks. Routines are provided to perform most actions that the user could wish to perform (see KL INIT EVENT, KL ADD TICKER, KL NEW FRAME FLY, KL NEW FAST TICKER and KL DISARM EVENT). These routines set values into the block from registers. The user should not write to the blocks, except as noted below.

All the following blocks must lie in the central 32K of RAM (otherwise the Kernel will be unable to access them).

On the CPC6128 the user is advised to keep these blocks in RAM block 2 if any bank switching is being performed (see section 2.5).

a. Event Blocks.

See section 12 for a general discussion of events and event blocks. An event block is laid out as follows:



Chain is a system pointer which must never be written to by the user. It is used to store events on the various event queues.

Class records the type of the event. It should not be written to by the user.

- Bit 0: 1 = Near address, 0 = Far address.
- Bits 1..4: Synchronous event priority.
- Bit5: Must be zero.
- Bit 6: 1 = Express event, 0 = Normal event.
- Bit 7: 1 = Asynchronous event, 0 = Synchronous event.

Note that many system queues are kept in priority order and so the block must be requeued if the priority is changed, it is not sufficient merely to change the priority in the event block.

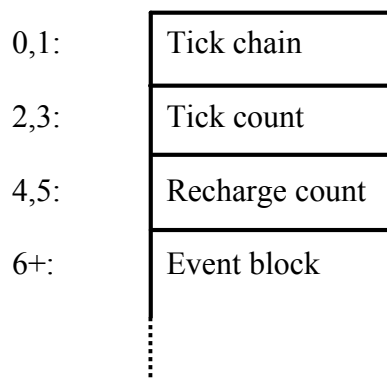
Count is the event count - a record of how many kicks are waiting to be processed or whether the event is disabled. See section 12.2 for a full discussion of the use of the event count.

Routine address and ROM make up the far address of the event routine. If the near address bit in the event class is true then the event routine is at a near address - the ROM select byte (byte 6) is ignored and the event routine is called directly. If the near address bit is false then the event routine is at far address - bytes 4,5 and 6 make up the far address to call to run the event routine. The user may write to the routine address and ROM fields (and to the near address bit in the class byte as well) provided that the operation is performed indivisibly (i.e. interrupts should be disabled).

The user fields are optional. They may be used to provide a data area specific to the event block so that a single event routine may be shared between a number of different event blocks (the event routine is passed the address of the user fields).

b. Ticker Queue Blocks.

See section 11 for a general discussion of ticker interrupts and the ticker queue. A ticker queue block is laid out as follows:



Tick chain is a system pointer which must never be written to by the user. It is used to store the block on the ticker queue.


Tick count is a count of the number of ticks before the next kick occurs. A tick count of zero means that the tick block is dormant and will not generate any kicks. (Ideally a dormant block should be removed from the ticker queue to avoid wasting time). The user may write to this field if required providing this is done indivisibly.

Recharge count is the value that the tick count is set to after each kick. If the recharge count is zero then the ticker block will become dormant after generating one kick. The user may write to this field if required providing this is done indivisibly.

Event block is a standard event block as described in section (a) above.

c. Frame Flyback Queue Blocks.

See section 11 for a general discussion of frame flyback interrupts and the frame flyback queue. A frame flyback queue block is laid out as follows:




Frame chain is a system pointer which must never be written to by the user. It is used to store the block on the frame flyback queue.

Event block is a standard event block as described in section (a) above.

d. Fast Ticker Queue Blocks.

See section 11 for a general discussion of fast ticker interrupts and the fast ticker queue. A fast ticker queue block is laid out as follows:



Fast chain is a system pointer which must never be written to by the user. It is used to store the block on the fast ticker queue.

Event block is a standard event block as described in section (a) above.

Appendix XI

The Alternate Register Set.

The Z80 microprocessor has two sets of registers - the normal set (AF, BC, DE and HL) and the alternate set (AF', BC', DE' and HL'). Unless the techniques outlined in this appendix are implemented the user is prohibited from using the alternate register set. This is because the alternate register set is used by the firmware (the Kernel in particular) for storing certain system values and flags. Providing that the user never enters the firmware then the alternate register set may be used without restriction. Of course this would mean that the user would be unable to use any facilities provided by the firmware. Furthermore, the user would also have to disable interrupts as interrupts cause firmware routines to be executed.

In the sections below a number of different methods are described that allow the user to overcome these restrictions. The method chosen will depend on what use is to be made of the alternate register set.

a. The firmware's use of the alternate register set.

The Kernel stores a couple of system variables in the alternate register set. This allows the Kernel to access these variables easily and thus speeds up a number of operations (particularly entry to and exit from firmware routines). Only BC' and the alternate carry flag (carry') are used to store values, however, routines do make use of the other alternate registers and so firmware routines may corrupt them.

B' is used to store the I/O address of the gate array (# 7F). C' is used to store the value required to set the current ROM state and screen mode:

Bits 0..1:	Set the screen mode.
Bit 2:	Disables the lower ROM.
Bit 3:	Disables the upper ROM.
Bits 4. .7:	System value to select gate array function.

By changing the ROM state bits and performing an OUT (C),C instruction the user can enable or disable ROMs. (N.B. The Z80 OUT (C),r and IN (C),r instructions use B as the top 8 bits of the I/O address. The hardware uses these top bits for decoding the I/O address, it ignores the bottom 8 bits!) OUT (C),C may be used to change the ROM state during the interrupt path when the normal Kernel entries (e.g. KL U ROM ENABLE) may not be called because they enable interrupts.

Carry' is normally false. When carry' is true this indicates that the firmware is in the interrupt path. The firmware occasionally uses this flag to allow it to take a different action when it is in the interrupt path to the action it takes when it is not in the interrupt path (usually avoiding enabling interrupts).

b. Simple use of the alternate register set.

The technique described in this section allows use of the alternate register set providing that no firmware routines are called and that interrupts are disabled.

After disabling interrupts registers A', DE' and HL' may be used as required. If registers BC' or F' (in particular carry') are used then their original contents must be restored before interrupts are re-enabled. The user may alter bits in C' (as described in (a) above) and need not restore the original value provided that an OUT (C),C is performed to keep the hardware abreast of the current state. The machine will not function correctly if the hardware and the value in C' are out of step when interrupts are enabled.

This technique requires interrupts to be disabled for the duration of the operation being performed. This is acceptable if the operation is short but not if the operation is lengthy. Disabling interrupts for a lengthy period will stop many firmware functions such as timers (and hence ink flashing, sound generation and keyboard scanning). If the operation to be performed is lengthy then it might be better to consider the use of one of the techniques described in sections (c) or (d) instead.

Example.

The user might want to provide a routine that performs an LD A,(BC) from RAM (similar to the RAM LAM pseudo-instruction provided by the firmware).

The code for this routine could be written as follows:

```
A_FROM_BC:
    PUSH BC
    DI                ; * * About to use alternate registers
    EXX
    POP HL           ;Transfer BC to HL'
    LD A,C
    SET 2,A          ;Set the disable lower ROM bit
    SET 3,A          ;Set the disable upper ROM bit
    OUT (C),A        ;Tell the hardware
    LD A,(HL)        ;Read the value from RAM
    OUT (C),C        ;Restore the old ROM state

    EXX
    EI                ; * * End of use of alternate registers
    RET
```

N.B. This routine must be RAM resident or disabling the ROMs will have an unfortunate effect!

c. Use of the alternate register set with interrupts enabled.

The technique described in this section allows the alternate register set to be used and interrupts to be enabled. It does not allow firmware routines to be called.

The simplistic use of the alternate register set by disabling interrupts as described above is unsatisfactory if this results in interrupts being disabled for an extended period of time. By patching INTERRUPT ENTRY in the low Kernel jumpblock interrupts can be trapped and appropriate action to restore the firmware registers can be taken. The actions that must be performed are as follows:

Before starting to use the alternate register set the firmware's BC' is saved and INTERRUPT ENTRY is patched so that the user's interrupt routine is used.

When the user has finished with the alternate register set the firmware's BC' and carry' are restored and INTERRUPT ENTRY is patched back to the firmware's interrupt routine.

When an interrupt occurs the user's alternate registers are saved, the firmware's BC' and carry' are restored and INTERRUPT ENTRY is patched back to the firmware's interrupt routine. The latter is done in case a second interrupt occurs whilst processing the events kicked from the interrupt path of the first interrupt (remember that the event processing is performed with interrupts enabled).

After interrupt processing has finished the firmware's BC' is saved, the user's alternate registers are restored and INTERRUPT ENTRY is patched back to the user's interrupt routine again.

Note that when INTERRUPT ENTRY is patched it is vital to ensure that the lower ROM is disabled and remains disabled. It is impossible to patch the ROM version of INTERRUPT ENTRY! If an interrupt occurred whilst the lower ROM was enabled then the firmware would jump straight into its interrupt routine without restoring its alternate registers first.

Example.

The following routines implement the scheme described above:

```
;  
; The following storage locations are used  
;  
FIRM_BC:          DEFS 2 ;Two bytes to store the firmware's BC'  
FIRM_INT:         DEFS 2 ;Two bytes to store the address of the  
                  ;firmware's interrupt routine  
USER_AF:          DEFS 2 ;Two bytes to store the user's AF'  
USER_BC:          DEFS 2 ;Two bytes to store the user's BC'  
USER_DE:          DEFS 2 ;Two bytes to store the user's DE'  
USER_HL:          DEFS 2 ;Two bytes to store the user's HL'
```

```

;
; This routine sets up the environment so that the
; user may make use of the alternate register set.
; N.B. Interrupts are enabled.
;
;
USER_ALTERNATE:
    DI                                ;An interrupt would be disastrous
    EX  AF,AF'
    EXX                                ;Swap to the alternate register set
;
    LD  (FIRM_BC), BC                 ;Save the firmware's BC'
;
    LD  HL,(INTERUPT_ENTRY+1)
    LD  (FIRM_INT), HL                ;Save the firmware's interrupt routine
;
    LD  HL, USER_INTERRUPT            ;Use the replacement interrupt routine
    LD  (INTERRUPT_ENTRY+1), HL
;
    LD  HL, (USER_AF)
    PUSH HL
    POP  AF                            ;Restore user's AF'
    LD  BC, (USER_BC)                 ;Restore user's BC'
    LD  DE, (USER_DE)                 ;Restore user's DE'
    LD  HL, (USER_HL)                 ;Restore user's HL'
;
    EXX                                ;Swap back to the standard register set
    EX  AF, AF'
    EI                                ;We have finished with the alternate regs
    RET
;
; This routine restores the environment for the firmware to use the alternate
; register set. N.B. Interrupts are disabled and not re-enabled.
;
;
FIRM_ALTERNATE:
    DI                                ;An interrupt would be disastrous
    EX  AF,AF'
    EXX                                ;Swap to the alternate register set
;
    LD  (USER_HL),HL                  ;Save user's HL'
    LD  (USER_DE), DE                 ;Save user's DE'
    LD  (USER_BC), BC                 ;Save user's BC'
    PUSH AF
    POP HL
    LD  (USER_AF), HL                 ;Save user's AF'
;

```

```

LD    HL, (FIRM_INT)      ;Restore the firmware's interrupt routine
LD    (INTERRUPT_ENTRY+1), HL

LD    BC, (FIRM_BC)      ;Restore the firmware's BC'
OR    A, A                ;Set the firmware's carry' to be false

EXX                      ;Swap back to the standard register set
EX    AF,AF'
RET                                ;N.B. May be about to enter the interrupt
                                ;path so no EI.
;
;
;
; This routine replaces the firmware's interrupt routine
; when the user is using the alternate register set
;
;
;
USER_INTERRUPT:
CALL FIRM_ALTERNATE      ;Switch the environment to the firmware
CALL INTERRUPT_ENTRY    ;Run the normal interrupt routine
JP USER_ALTERNATE      ;Switch the environment back to the user

```

To start using the alternate register set the user obeys the instruction:

```
CALL USER_ALTERNATE
```

To finish using the alternate register set the user obeys the instructions:

```
CALL FIRM_ALTERNATE
EI
```

d. Calling firmware routines whilst using the alternate register set.

The technique described in this section extends the technique described in section (c) to allow the user to call firmware routines whilst using the alternate register set.

To call a firmware routine requires exactly the same action as is required for the interrupt routine:

Before calling a firmware routine the user's alternate registers are saved, the firmware's BC' and carry' are restored and INTERRUPT ENTRY is patched back to the firmware's interrupt routine. The latter is done in case an interrupt occurs whilst executing the firmware routine.

After running the firmware routine the firmware's BC' is saved, the user's alternate registers are restored and INTERRUPT ENTRY is patched back to the user's interrupt routine again.

As indicated in section (c) it is vital to ensure that the lower ROM remains disabled while the alternate register set is in use since INTERRUPT ENTRY in the ROM is not patchable and jumps straight to the firmware's interrupt routine.

Using the routines defined in section (c) a firmware routine may be called by using the following sequence:

```

CALL FIRM_ALTERNATE    ;Switch the environment to the firmware
EI                     ;FIRM ALTERNATE disables interrupts
CALL firmware          ;Run the firmware routine
CALL USER..ALTERNATE  ;Switch the environment back to the user

```

The above code is rather long if a lot of firmware calls are to be made (10 bytes per call). The following routine takes the address of a firmware routine to call to as an inline parameter (and only uses 5 bytes per call).

```

;
; This routine saves the user's alternate registers, calls a
; firmware routine whose address is passed inline and then restores the user's
; alternate register set afterwards.
;
;
FIRM_ROUTINE:
CALL FIRM_ALTERNATE    ;Switch to the firmware environment
EXX                    ;N.B. Interrupts are disabled
POP HL                 ;Recover address of routine to call, uses
LD  E,(HL)             ;firmware's DE' and HL' which may be
INC  HL                 ;corrupted
LD  D,(HL)             ;Get routine to call into DE'
INC  HL
PUSH HL                ;Put the real return address back
;
LD  HL,USER_ALTERNATE ;Restore the user environment when
;the firmware returns by putting a
PUSH HL                ;fake return address on the stack
;
PUSH DE                ;Save the routine to call
EXX
EI
RET                    ;Jump to the routine to call

```

To call a firmware routine using the above routine the following sequence should be used:

```

CALL FIRM_ROUTINE
firmware                ;Address of routine to call
..                      ;FIRM_ROUTINE returns here

```


Appendix XII

The Hardware and Hardware Variants

A. Processor.

The processor is a Z80A running at a clock frequency of 4.00 MHz ($\pm 0.1\%$). There is logic that stretches $\overline{\text{MREQ}}$ and $\overline{\text{IORQ}}$ using the $\overline{\text{CPU WAIT}}$ facility so that the processor can only make one access to memory each microsecond.

The processor NMI pin is pulled up and made available on the expansion bus. However, a non-maskable interrupt may cause the firmware to violate various timing constraints and so its use is not recommended.

The processor interrupt pin is driven by a flip-flop in the video gate array. This flip-flop is set during every vertical flyback and every 52 scan lines thereafter until the next vertical flyback. The interrupt is arranged to occur approximately 2 scans (125 microseconds) into the 8 scan (500 microsecond) vertical flyback signal. The interrupt latch is cleared by the processor acknowledging the interrupt, or explicitly, using a software command. The top bit of the divide by 52 scan counter is also cleared when the processor acknowledges an interrupt occurring after this counter has overflowed. This allows the interrupt system to be expanded.

B. Memory.

ROM

A single 32K byte ROM is present on the processor board, but is mapped onto two blocks of 16K in processor address space. The lower half of the ROM occupies addresses #0000 to #3FFF and the upper half occupies addresses #C000 to #FFFF. These two halves of the ROM can be separately enabled and disabled by two control latches in the video gate array. On power-up or other system reset both halves of the ROM are enabled.

An expansion port signal can be used to disable this internal ROM and allow external ROM(s) to be accessed instead. These are selected by output instructions and replace the upper half of the on-board 32K byte ROM when selected.

RAM

On the CPC464 and CPC664 64k bytes of dynamic RAM are fitted to the processor board at addresses #0000 to #FFFF. On the CPC6128 128k bytes are fitted and this can be bank switched to be accessed at addresses #0000 to #FFFF (see section 2.5). The lowest 16K and the top 16K are overlaid when ROM is enabled. Whether the ROM is enabled or not affects where data is read from, it has no effect on write operations which will be correctly performed 'through' the enabled ROM to the underlying RAM.

VDU SCREEN MEMORY

The display uses 16k of the processor RAM memory as screen refresh memory. On the 464 and 664 the 16k used can be switched between the blocks starting at #0000, #4000, #8000 and at #C000 by the top two bits (bits 12 and 13) programmed into the HD6845S start address register. On the 6128 the top two bits select RAM banks 0, 1, 2 or 3 to be used (see section 6.4 for further details).

The arrangement of data in the VDU screen memory is dependent on the VDU mode currently selected. In all modes the memory can be considered as consisting of 8K 16 bit words. Each word contains either 4, 8 or 16 pixels (P0..Pn) of 1, 2 or 4 bits (B0..Bm) depending on the mode as follows:

A0	Bit	Mode 0	Mode 1	Mode 2
0	D7	P0 B0	P0 B0	P0 B0
0	D6	P1 B0	P1 B0	P1 B0
0	D5	P0 B2	P2 B0	P2 B0
0	D4	P1 B2	P3 B0	P3 B0
0	D3	P0 B1	P0 B1	P4 B0
0	D2	P1 B1	P1 B1	P5 B0
0	D1	P0 B3	P2 B1	P6 B0
0	D0	P1 B3	P3 B1	P7 B0
1	D7	P2 B0	P4 B0	P8 B0
1	D6	P3 B0	P5 B0	P9 B0
1	D5	P2 B2	P6 B0	P10B0
1	D4	P3 B2	P7 B0	P11B0
1	D3	P2 B1	P4 B1	P12B0
1	D2	P3 B1	P5 B1	P13B0
1	D1	P2 B3	P6 B1	P14B0
1	D0	P3 B3	P7 B1	P15B0

Data for lines 0,8,16,24.. on the display are packed into the first 2K byte block of the memory, lines 1,9,17,25.. are packed into the corresponding places of the next 2K byte block of memory, with lines 7,15,23,31.. occupying the top 2Kbyte block of the 16k memory area.

The bottom 10 bits of the HD6845SP start address register define where within these 2K blocks the screen starts. The offset from the start of the 2K byte block is always even and is calculated as twice the register contents modulo 2K bytes. When data has to be displayed from beyond the end of a 2K byte block wrap around occurs to the beginning of the same 2Kbyte block. See section 6.4 for a fuller description.

C. AY3-8912 Programmable Sound Generator.

The PSG is accessed using ports A and C of the μ PD8255 device. Note that when writing or loading address to the AY-3-8912 the maximum duration of the write or load address command with BDIR high is 10 microseconds. The clock input to the sound generator is exactly 1.00 MHz. The BC2 signal is tied permanently high. On power-up the I/O port should be programmed to input mode.

The user is advised to use the firmware routine MC SOUND REGISTER to write to the PSG.

D. HD6845S CRT Controller (HD6845S CRTC).

The character clock to the CRTC occurs for every two bytes fetched from memory, i.e. every 1.0 microseconds. The first byte of a pair has an even address, the second has an odd address. In normal operation the internal registers are set up as follows:

Register	Function	PAL	SECAM	NTSC
0	Horizontal Total	63	63	63
1	Horizontal Displayed	40	40	40
2	Horizontal Sync. Posn.	46	46	46
3	Vsync., Hsync. widths	#8E	#8E	#8E
4	Vertical Total	38	38	31
5	Vertical Total Adjust	0	0	6
6	Vertical Displayed	25	25	25
7	Vertical Sync. Posn.	30	30	27
8	Interlace and Skew	0	0	0
9	Max. Raster Address	7	7	7
10	Cursor Start Raster	X	X	X
11	Cursor End Raster	X	X	X
12	Start Address (H)	X	X	X
13	StartAddress(L)	X	X	X
14	Cursor(H)	X	X	X
15	Cursor(L)	X	X	X

In the above table the numbers for PAL and SECAM standards are identical.

For information on the NTSC Register 5 see Appendix XIII.

Note that X indicates that software may vary these numbers during device operation. The firmware only makes use of the start address register which is used to set the screen base and offset.

E. Video Gate Array.

The software must access this device in order to control the enabling and disabling of ROMs, the mode of operation of the VDU and also to load colour information for 'inks' into the palette memory. One I/O channel is used for all commands, the top two bits of data specifying the command type as follows:

Bit 7	Bit 6	Use
0	0	Load palette pointer register.
0	1	Load palette memory.
1	0	Load mode and ROM enable register.
1	1	Bank Switching Register on CPC6128.

MODE AND ROM ENABLE REGISTER

This write-only register controls the VDU mode and ROM enabling as follows:

Bit 7:	1
Bit 6:	0
Bit 5:	** Reserved ** (send 0)
Bit 4:	Clear raster 52 divider.
Bit 3:	Upper half ROM disable.
Bit 2:	Lower half ROM disable.
Bit 1:	VDU Mode control MC1.
Bit 0:	VDU Mode control MC0.

Writing a 1 to bit 4 clears the top bit of the divide by 52 counter used for generating periodic interrupts.

Modes are defined by the mode control pins as follows:

MC1	MC0	Mode
0	0	Mode 0, 160 x 200 pixels in 16 colours
0	1	Mode 1, 320 x 200 pixels in 4 colours.
1	0	Mode 2, 640 x 200 pixels in 2 colours.
1	1	**Do not use**

The gate array hardware synchronises mode changing to the next horizontal flyback in order to aid software that requires different parts of the screen to be handled in different modes.

On power-up and other system resets, the mode and ROM enable register is set to zero, enabling both halves of the ROM.

PALETTE POINTER REGISTER

This write-only register controls the loading of the VDU colour palette as follows:

Bit 7:	0
Bit 6:	0
Bit 5:	** Reserved ** (send 0)
Bit 4:	Palette pointer bit PR4.
Bit 3:	Palette pointer bit PR3.
Bit 2:	Palette pointer bit PR2.
Bit 1 :	Palette pointer bit PR1.
Bit 0:	Palette pointer bit PR0.

Bits PR0 to PR3 select which ink is to have its colour loaded, providing bit PR4 is low. If bit PR4 is high then bits PR0 to PR3 are ignored and the border ink colour is loaded.

PALETTE MEMORY

This write-only memory controls the VDU colour palette as follows:

Bit 7:	0
Bit 6:	1
Bit 5:	**Reserved ** (send 0)
Bit 4:	Colour data bit CD4.
Bit 3:	Colour data bit CD3.
Bit 2:	Colour data bit CD2.
Bit 1:	Colour data bit CD1.
Bit 0:	Colour data bit CD0.

The ink entry pointed at by the palette pointer register is loaded with the colour sent on this channel. The number of colours that need to be loaded ranges from 2 colours in mode 2 to 16 colours in mode 0. In addition to loading the colours an extra colour data byte must be sent to this channel to define the border colour. On power-up and other system resets the contents of the palette are undefined, but the border colour is set to BLACK, to avoid unsightly effects on power-up.

The 32 colour codes are decoded to drive the RGB signals, producing 27 different colours. The hardware colours are listed in Appendix V.

BANK SWITCHING REGISTER (CPC6128 only)

This write-only register controls the layout of the bank switchable RAM as follows:

Bit 7: 1
Bit 6: 1
Bit 5: ****Reserved**** (send 0)
Bit 4: 0
Bit 3: 0
Bit 2: x
Bit 1: x
Bit 0: x

The xxx appearing on Bits 0, 1, and 2 is the code for the selected bank layout, detailed as follows:

Address	Code	0	1	2	3	4	5	6	7
#C000:		3	7	7	7	3	3	3	3
#8000:		2	2	6	2	2	2	2	2
#4000:		1	1	5	3	4	5	6	7
#0000:		0	0	4	0	0	0	0	0

Selecting Code 0 (the number sent to the Bank Switch Register would be #C0) will switch to the normal default bank layout, i.e. Banks 0..3.

The other 64k can be accessed by selecting codes 4..7. These switch one of the other 16k blocks into area #4000. See section 2.5 for more details.

F. μ PD8255 Parallel Peripheral Interface.

The PPI as well as the 8 port pins on the PSG are used to interface to the keyboard and to control and sense miscellaneous signals on the processor board. Port A must be programmed either to input or to output in mode 0 since this port is used for reading and writing to the PSG. Port B must be programmed to input in mode 0. Port C must be programmed to output in mode 0 on both halves.

Circuitry is provided around the PPI to reset it during system reset. For details of the operation of the μ PD8255 see the NEC product specification.

CHANNEL A (Input or Output)

Bit 7:	Data/Address DA7 connected to AY-3-8912.
Bit 6:	Data/Address DA6 connected to AY-3-8912.
Bit 5:	Data/Address DA5 connected to AY-3-8912.
Bit 4:	Data/Address DA4 connected to AY-3-8912.
Bit 3:	Data/Address DA3 connected to AY-3-8912.
Bit 2:	Data/Address DA2 connected to AY-3-8912.
Bit 1:	Data/Address DA1 connected to AY-3-8912.
Bit 0:	Data/Address DA0 connected to AY-3-8912.

CHANNEL B (Input Only)

Bit 7:	Datacorder cassette read data.
Bit 6:	Centronics busy signal.
Bit 5:	Not expansion port active signal.
Bit 4:	Not option link LK4.
Bit 3:	Not option link LK3.
Bit 2:	Not option link LK2.
Bit 1:	Not option link LK1.
Bit 0:	Frame flyback pulse.

The option links, LK1...LK4 are factory set. LK4 is fitted for 60 Hz TV. standards and omitted for 50 Hz standards.

CHANNEL C (Output Only)

Bit 7:	AY-3-89i2 BDIR signal.
Bit 6:	AY-3-89i2 BC1 signal.
Bit 5:	Datacorder cassette write data
Bit 4:	Datacorder cassette motor on.
Bit 3:	Keyboard row select KR3.
Bit 2:	Keyboard row select KR2.
Bit 1 :	Keyboard row select KR1.
Bit 0:	Keyboard row select KR0.

G. Centronics Port Latch.

This latch is loaded with data by output commands to the correct I/O channel. It cannot be read. Note that the timing requirements on Centronics interfaces generally specify that the data must be present on the seven data lines at least 1 microsecond before the strobe is made active and must remain valid for at least 1 microsecond after the strobe returns inactive. The duration of the strobe must be between 1 and 500 microseconds. The busy signal can be inspected as soon as the strobe is inactive in order to determine when more data can be sent.

- Bit 7: Centronics strobe signal (1 = active).
- Bit 6: Data 7 to Centronics port.
- Bit 5: Data 6 to Centronics port.
- Bit 4: Data 5 to Centronics port.
- Bit 3: Data 4 to Centronics port.
- Bit 2: Data 3 to Centronics port.
- Bit 1: Data 2 to Centronics port.
- Bit 0: Data 1 to Centronics port.

On power-up and other system resets the outputs of this latch are all cleared.

H. Keyboard and Joysticks.

The keyboard and joystick switches are sensed by selecting one of ten rows using the four control bits on channel C of the PPI and reading the data from the PSG parallel port using port A of the PPI.

The keyboard and joystick switches are arranged in a 10 by 8 matrix. One of ten rows is selected using the code on KR0...KR3 and the eight bits of data are then read in parallel on a parallel port as described above. A switch is active (closed) if the corresponding data bit is a logic 0.

The key number associated with each key (see Appendix I) is constructed as follows:

Bit:	7	6	5	4	3	2	1	0
	0	Row number				Bit number		

Thus the key that is associated with bit 5 in row 4 has key number $37(4*8+5)$.

I. Disc Interface

Floppy Disc Controller

The controller uses an NEC type μ PD765A Floppy Disc Controller IC to connect to the disc drives. Only two disc drives are supported, since the US1 line from the μ PD765A is ignored. This results in the two disc drives being accessed as drives 0 (zero) and 1 (one) and again as drives 2 and 3. The controller supports both single and double sided and single and double density mini-floppy disc drives. Note that the clock frequency supplied to the μ PD765A CLK pin is 4.00 MHz rather than the 8.00 MHz used with larger disc drives.

The full facilities described in the NEC data sheet for the μ PD765A are available with the exception of interrupts and DMA which are not supported.

Expansion ROM

The disc ROM is normally number #07, but may be set to #00 by cutting option trace LK1 in the case of the DDI-1, LK7 in the case of the 664, LK201 in the case of the 6128.

The EXP signal (pin 48), of the 50 way expansion connector, is grounded when the disc ROM number is #07 in order that this address can be avoided by other expansion peripherals. A 200 nanosecond 27128 type EPROM or ROM is normally used, and may be fitted in a DIL socket in some machines.

Option traces LK2 and LK3 are manufacturing options for write pre-compensation. They should not need alteration by the user.

Motor Control

Writing to channel #FA7E starts and stops the disc drive motors. Writing #00 will stop the motors, #01 will start the motors. On power-up and other system resets the motors are not stopped.

Electrical Levels

All electrical levels on the controller are TTL compatible. Signals originating in the drives are terminated by 6800 resistors to +5v at the controller and received with gates with input hysteresis. The maximum permissible cable length is .075 metre.

Using Other Disc Drives

It is possible to use other disc drives with the CPC range of computers (or the DDI-1), in particular 5¼" drives. Some hardware knowledge will be required. The following gives some advice and information which should assist in using a different drive.

If Drive A: the 5V power should be supplied to pins 2,4,6 and 14 of the 5¼" drive after ensuring that any existing connections to the drive circuitry have been removed.

If Drive B: no terminating resistor should be installed.

The drive MUST have a 'READY' signal on pin 34.

The drive will require its own power supply.

The extra cabling should be as short as possible and should consist of a cable-mounting male connector (to mate with the female socket connector on the cable from the interface), all 34 conductors and, normally, a 34 way double-sided card edge connector (to mate with the 5¼" drive).

The step rate, motor on and off timeout may have to be changed, see, in the case of CP/M 2.2, the SETUP utility.

A drive such as the Shugart 201 is not suitable since it does not have a ready signal, but one such as the Chinon F 051-MD is suitable.

Pin Arrangement

All odd numbered pins are pulled to ground (GND). All signals are active low.

PIN No.	SIGNAL
2	+5V
4	+5V
6	+5V
8	INDEX
10	Drive select 0
12	Drive select 1
14	+5V
16	Motor On
18	Direction Select
20	Step
22	Write Data
24	Write Gate
26	Track 0
28	Write Protect
30	Read Data
32	Side 1 Select
34	Ready

J. Serial Interface

The BIOS supports a two channel asynchronous serial interface. The interface is an optional extra for the CPC 464/664/6128 computers. This section describes the recommended hardware configuration for the serial interface such that, if fitted, it can be driven by the BIOS.

The interface consists of a Zilog Z80A SIO/0 or Z80A DART together with an Intel 8253 programmable interval timer. The 8253 is used as a Baud rate generator as follows:

Timer 0 generates the transmit clock for channel A of the 510.

Timer 1 generates the receive clock for channel B of the 510.

Timer 2 generates both transmit and receive clocks for channel B of the 510.

It is assumed that the CLK inputs to all three channels of this device are driven by a 2.0(0.1%)MHz clock signal derived from the 4.0 MHz CPU clock. The GATE inputs to all three channels are tied permanently high.

NOTE: This is an optional extra for all machines.

K. I/O ports

The following is a detailed list of the Z80 I/O used by the CPC range of computers.

PORT	OUTPUT	INPUT
#7Fxx	Video Gate Array	**Do not use**
#BCxx	HD6845 CRTC address	**Do not use**
#BDxx	HD6845 CRTC data	**Do not use**
#BExx	**Do not use**	Reserved for CRTC status
#BFxx	**Do not use**	HD6845 CRTC data
#DFxx	Expansion ROM select	**Not used**
#EFxx	Centronics latch	**Do not use**
#F4xx	μPD8255 port A data	μPD8255 port A data
#F5xx	μPD8255 port B data	μPD8255 port B data
#F6xx	μPD8255 port C data	μPD8255 port C data
#F7xx	μPD8255 control	**Undefined**
#F8xx	Expansion bus	Expansion bus
#F9xx	Expansion bus	Expansion bus
#FAxx	Expansion bus	Expansion bus
#FBxx	Expansion bus	Expansion bus
#FFxx	**Not used**	**Not used**

Disc Interface

#FA7E	Motor Control	**not used**
#FB7E	**not used**	μPD765A Status Register
#FB7F	μPD765A Data Register	μPD765A Data Register

Serial Interface

#FADC	510 Channel A Data	510 Channel A Data
#FADD	510 Channel A Control	510 Channel A Control
#FADE	510 Channel B Data	510 Channel B Data
#FADF	510 Channel B Control	510 Channel B Control
#FBDC	8253 Load Counter 0	8253 Read Counter 0
#FBDD	8253 Load Counter 1	8253 Read Counter 1
#FBDE	8253 Load Counter 2	8253 Read Counter 2
#FBDF	8253 Write Mode Word	**not used**

Expansion bus I/O channels in the address range #F800 to #FBFF are reserved as follows:

ADDRESS (AO..A7)	USE
#00...#7B	** Do not use **
#7C...#7F	Reserved for Disc Interface
#80...#BB	** Do not use **
#BC...#BF	Reserved for future use.
#C0...#DB	** Do not use **
#DC...#DF	Reserved for communications interfaces.
#E0...#FE	Available for user peripherals.
#FF	Reset peripherals.

L. Hardware Variants

There have been slight differences with the circuitry concerning the ULA in some of the CPC range of computers, but as long as legal values are used when accessing them, there should be no noticeable difference.

The firmware has been altered for a few foreign computers. The major changes have been to the character set (to provide a few new national characters) and to the characters available on the keyboard (to allow new characters to be entered), also some messages and version numbers have been changed. These foreign changes are detailed below.

Spanish System

Character Set

#A1 matrix is set to Ñ
 #AB matrix is set to ñ
 #A3 matrix is set to Pt

Keyboard

* now generates Ñ
 : now generates ñ
 + now generates :
 { now generates *
 } now generates +
 £ now generates Pt

(with control now generates {
) with control now generates }
 ! with control now generates ¡
 ? with control now generates ¿

Start Up Message

The V1 or V3 message at start-up has been changed to S1 or S3.

Identification

The ROM modification byte (at location #C003 in the on-board ROM) is now 1, not 0.

Printer

Certain Codes have been changed for the printer translations. CPC6128 printer codes are unchanged

CPC464 translations changed are : #A1 -> #5C
#AB -> #7C
#AE -> #5D
#AF -> #5B

Danish System

Character Set

#30 matrix is now	0	(i.e. slash removed from zero)
#5B matrix is now	Æ	
#5C matrix is now	Ø	(i.e. letter O with slash through it)
#5D matrix is now	?	(i.e. letter A with ° accent)
#7B matrix is now	æ	(i.e. lower case of #5B above)
#7C matrix is now	ø	(i.e. lower case of #5C above)
#7D matrix is now	å	(i.e. lower case of #5D above)

Keyboard

	now generates ?
@	now generates å
{	now generates □
[now generates @
*	now generates Æ
:	now generates æ
+	now generates Ø
;	now generates ø
}	now generates *
]	now generates :
□	now generates +
\	now generates ;

Control versions have moved with keys/chars (as expected).

Start Up Messages

The V1 or V3 message at start-up has been changed to d1 or d3.

Identification

The ROM modification byte (at location #C003 in the on-board ROM) is now 2 and not 0.


French System

Character Set

#40	matrix is now	à	(i.e. a grave)
#5C	matrix is now	ç	(i.e. c cedilla)
#5E	matrix is now	ˆ	(i.e. circumflex)
#7B	matrix is now	è	(i.e. e grave)
#7C	matrix is now	ù	(i.e. u grave)
#7D	matrix is now	é	(i.e. e acute)
#A2	matrix is now	°	(i.e. degrees sign)

Keyboard

So many characters have been moved around that it is easiest to show the changes with the following diagram. The control versions have moved with the keys/chars (as expected).



Start Up Message

The V1 or V3 message at start-up has been changed to f1 or f3.

Identification

The ROM modification byte (at location #C003 in the on-board ROM) is now 3, not 0.

Appendix XIII - Hints, Tips, and Workarounds

Following are a number of 'WORKAROUNDS' for the 464 machine, that is, routines which allow the 464 to act as specified in the '464 FIRMWARE SPECIFICATION'.

Soft End Of File

Reading characters from the disc using CAS IN CHAR when it is redirected to the AMSDOS routine can run into problems caused by the routine returning an error when it reads the end of file character #1A. This can be avoided by patching the jumpblock so that the end of file error is detected and ignored. The following program does this.

```
        SAVE_ENTRY:      DEFS 3          ;Space to save jumpblock entry
        ;
INSTALL:
        LD   A, (CAS_IN_CHAR +0)
        LD   HL, (CAS_IN_CHAR +1)
        LD   (SAVE_ENTRY + 0), A
        LD   (SAVE_ENTRY + 1), HL      ;Save original contents
        ;
INTERCEPT:
        LD   A,#C3
        LD   HL, NEW_CAS_IN_CHAR
        JR   PATCH                    ;A/HL = jump to new routine
        ;
RESTORE
        LD   A, (SAVE_ENTRY +0)
        LD   HL, (SAVE_ENTRY + 1)     ;A/HL = jump to original routine
        ;
PATCH
        LD   (CAS_IN_CHAR + 0), A
        LD   (CAS_IN_CHAR + 1), HL
RET
        ;
NEW_CAS_IN_CHAR
        PUSH HL
```



```

;
CALL RESTORE                ;Put original jump back because
                            ;AMSDOS requires it to be executed
                            ;in its original position!
CALL CAS_IN_CHAR           ;Read the character
PUSH AF
CALL INTERCEPT           ;Continue intercepting jumpblock
POP AF
;
POP HL
RET C                       ;Quit if OK
RET Z                       ;Quit if ESC
CP #1A
SCF
CCF
RET NZ                      ;Quit if a real error
OR A

SCF
RET                          ;Pretend OK if soft EOF

```

Before reading from the file the user should call INSTALL and from then on CAS IN CHAR will return character #1A just like any other character. Note that INSTALL must only be called once, otherwise the original contents of the jumpblock entry will be lost! The patch will be lost if the external commands TAPE, TAPE.IN, DISC or DISC.IN are executed.

MODE switching on V1.0 firmware

Some programs, such as SORCERY+, run with different parts of the screen in different modes. It is possible to do this using the firmware but it is necessary to intercept the SCR MODE CLEAR indirection so that the screen is not cleared each time the mode is changed. On V1.1 firmware all that is necessary is to patch a RET instruction into the first byte of the SCR MODE CLEAR indirection. On V1.0 firmware this would result in all the inks being set to the background colour! This can be overcome by using the following routine to install the users own ink refresh routine.

```

LD HL,SCR_MODE_CLEAR
LD (HL),#C9

LD HL,EVENT_BLOCK          ; Points to 9 byte block
LD DE,INK_ROUTINE         ; points to Interrupt routine
LD B,#81
CALL KL_NEW_FRAME_FLY     ; Add Frame Flyback Interrupt
RET

```

```
INK_ROUTINE: LD DE,INK_VECTOR      ; Points to INK settings
             JP MC_SET_INKS
```

```
EVENT_BLOCK: DEFS 9
```

```
INK_VECTOR:  DEFB 4,4,10,19,12,11,20,21,13,6,30,31,7,18,25,26,5
```

Detecting the DDI-1 on a 464

If you are writing a program which needs to detect whether a DDI-1 interface is connected to the computer or not, then the way to accomplish this is to issue a DISC command, if this is not found then the DDI-1 is not connected.

The following program shows how to do this under machine code.

```
FIND_DISC:  LD HL,DISC_COMMAND
            CALL KL_FIND_COMMAND
            SBC A,A
            RET
```

```
DISC_COMMAND:  DEFM 'DIS' ,'C' +#80
```

Calling FIND DISC will return A = #FF if a DDI-1 is connected or A = 0 if a DDI-1 is not connected).

Express Asynchronous Events

A problem was discovered with Express Asynchronous Events, in that, the COUNT byte should always be reset (to 0) upon termination of the Event, otherwise the event will not be kicked again. This can be done simply by adding the following program to the event routine:

```
LD HL,EVENT_BLOCK+2
LD (HL),#00
RET
```

Printing characters above 127 and suppressing the double line feed

The following patches are to enable characters greater than 127 to be printed, and that the line feed character automatically sent after a carriage return can be suppressed and so stop the infamous double line feed problem on printers.

First the main patch code which is used in both programs should be set up as follows:

```
INITIALISE:  LD      A,(MC_WAIT_PRINT)
             LD      HL,(MC_WAIT_PRINT+1)
             LD      (NEW_PRINT),A
             LD      (NEW_PRINT+1),HL ; patch the jumpblock to end of new code

             LD      A,#C3
             LD      HL,PRINT._PATCH
             LD      (MC_WAIT_PRINT),A
             LD      (MC_WAIT_PRINT+1),HL ; redirect the jumpblock to the new code

             RET
```

Then the PRINT PATCH code for printing characters above ASCII 127 is as follows:

```
PRINT_PATCH: CP      128
              JR      C,NEW_PRINT ; character below 128 then print it
              AND     #7F ; mask off bit 7
              LD      (CHAR),A ; store away the character
              PUSH    HL ; HL is preserved on exit
              LD      HL,ESCAPE ; point HL at escape sequence
              LD      B,5
PRINT_LOOP:  PUSH    BC
              LD      A,HL ; get character from sequence
              CALL    NEW_PRINT ; send character to printer
              INC     HL ; bump HL to next character
              POP     BC
              DJNZ    PRINT_LOOP ; do this 5 times
              POP     HL ; restoreHL
              RET     ; return to calling program
```

```

NEW_PRINT:    DEFS 3                ; Storage foF MC_WAIT_PRINTER
                                           ; jumpblock

ESCAPE       DEFM    27,">"        ; Sets the printer to alternate character set

CHAR         DEFB    0              ; storage for original character being sent
                                          
                                           DEFM    27,"="        ; Sets the printer into normal character set

```

NOTE the escape sequences (ESC= and ESC>) are used with the AMSTRAD DMP2000 printer and may have to be changed for other printers, (e.g. a number of printers use ESC4 and ESC5 instead). Check with your printer user manual.

The PRINT PATCH code to suppress the line feed after a carriage return appears below:

```

PRINT_PATCH:  CP          10
              JR          NZ,NOT_LF    ; jump if character not a line feed
              LD          A,(CHAR)     ; restore last character printed
              CP          13           ; test if last char was a 13(CR)
              LD          A,0
              LD          (CHAR),A     ; zeroise the last char. printed store
              SCF              ; set carry flag, char. printed OK
              RET          Z           ; Return to calling program if last char.
                                           ; printed was a 13 (CR)
              LD          A,10         ; If last char. was not a 13 (CR) then print
                                           ; 10 (Line Feed)
NOT_LF:       LD          (CHAR),A     ; store character
NEW_PRINT:    DEFS 3                ; Execute printing of character
                                           ; and return to calling program
CHAR:        DEFB 0                ; storage for last characterprinted

```

So to use either the LF suppressor or the printing characters above 127 routines, the code INITIALISE and the re'evant PRINT PATCH should be tied together and then assembled at a convenient location in RAM. Once initialized they will work until a call to MC RESET PRINTER which re-initializes the original indirection jumpblock at MC WAIT PRINTER.

American CPC6 128: Frame Flybacks and Interrupts

This technical note discusses the relationship between frame flybacks and interrupts on the NTSC version of the CPC464/664/6128. Currently the only production model affected is the American version of the CPC6128 - all other markets use PAL/SECAM models and this note does not apply to them.

There was an error in the original hardware specification for the CPC464/664/6128 in that the value to be loaded into the Vertical Total Adjust register in the HD6845 (register 5) was incorrectly given as 6 for the NTSC version whereas it should have been 4. Thus the various ROMs produced for the CPC464/664/6128, when used with an NTSC system, load an incorrect value into the 6845 whenever a full reset occurs; for example when the machine is first powered on, or the RESET_ENTRY firmware call is made.

PAL/SECAM systems work correctly and, fortunately, the only effect of the incorrect value on NTSC systems is to cause the interrupt associated with frame flyback to occur at exactly the same time as the frame flyback pulse becomes true. With the correct NTSC value the interrupt will occur 125 microseconds after frame flyback becomes true -corresponding to PAL/SECAM systems and the description given in Section 1.

What this means in practice is that on an American CPC6 128 a program which tests the PPI Frame Flyback signal (for example by calling MC_WAIT_FLYBACK) will not see frame flyback become true before the interrupt occurs, but must rely upon frame flyback still being true when the processing associated with the interrupt is complete. If the interrupt processing takes too long, the program will appear to 'Lock Up' because it never sees frame flyback set true.

In order to minimise the possibility of 'Lock Ups' occurring Amstrad have ensured that all American CPC6 128 machines are fitted with the type of 6845 which does not have a programmable frame flyback pulse length. Therefore the frame flyback will last for a fixed 1000 microseconds rather than the programmed time of 500 microseconds. The 500 microseconds period is in fact quite sufficient for the system routines invoked by the frame flyback interrupt; 1000 microseconds will allow a number of user routines to also occur at that time without any difficulty.

Clearly though, software for American CPC6128s which has much to do at frame flyback time, or which wishes to avoid flickering effects on the top few lines of the screen, must arrange to place the correct NTSC value into register 5 of the 6845 using code equivalent to the following:

```
SET_NTSC          LD      B,#F5
                  IN      A,(C)          ; Read PPI port B
                  AND     #10           ; Inspect LK4
                  RET     NZ            ; Return if not YTSC No action required
                  DI      ; Need exclusive access to CRTC
                  LD      BC,#BC05
                  OUT     (C),C         ; Set CRTC address to register 5
                  LD      BC,#BD04
                  OUT     (C),C         ; Set Vertical Total Adjust to 4
                  EI      ; End of exclusive access
                  RET
```

Using Interrupts with Z80 Peripherals

Z80 support chips such as PIO, SIO, DART, DMA, and CTC have an elaborate interrupt priority system involving the connection of the IEO output of one chip to the IEI input of the next in a daisy-chain.

When a chip wishes to interrupt it inspects its IEI input. If this is '1' then no higher priority device is interrupting and the chip may pull on the interrupt request signal. It will also set IEO to '0' so that lower priority devices are aware of its request. If IEI into a chip goes to '0' then the chip will not interrupt until the higher priority devices have been serviced.

When the CPU is actually interrupted an interrupt acknowledge bus cycle occurs and the highest priority interrupting device (the one with IEI = 1 and a reason to interrupt) assumes that it is being serviced and disables its interrupt. This means that interrupt service routines have the option to issue an EI instruction to allow immediate response to higher priority interrupts.

When interrupt servicing is complete a RETI instruction must be issued. This causes a support chip with an interrupt under service to redetermine its interrupt status and the state of IEI and to set IEO accordingly.

Section 11.2 clearly describes the scheme used for external interrupting devices. This relies upon the external device continuing to interrupt during the service routine so that it can be distinguished from the internal ticker interrupt which is automatically disabled as soon as the interrupt acknowledge bus cycle occurs. From the above description it will be clear that Z80 support chips do not meet this condition and thus their interrupts cannot be used. Hardware designers should also note that interrupts should be disarmed by an OUT to the RESET PERIPHERALS channel (#F8FF).

The code given below does two things, it arranges for a RETI instruction to be issued after every interrupt to ensure that all chips which assumed that they were being serviced will reassert their interrupt request. This is important where several Z80 support chips are involved because there is no provision for IEO-IEI connections between add-on devices. Secondly, a RETI is issued immediately before deciding whether an interrupt is internal or external which will mean that Z80 support chips will renew their interrupt request and the firmware will correctly determine that the interrupt is external.

Set up an external device service routine by intercepting the indirection at #003B. Remember to make this interception code relocatable.

Add the following Z80 support chip code only if a mark 1 ROM versions 0, 1 or 2 is fitted (This means all existing CPC464/664/6128 machines - see KL PROBE ROM.)

#0038 (ROM or RAM) originally contains:

```
JP      ADDRESS_X
```

ADDRESS_X + 5 will be in RAM, not under a ROM, and originally contains:

```
LD      A,C
SCF
EI
EX      AF,AF'
DI
```

replace the five bytes at ADDRESS_X + 5 by:

```
CALL    NEW_CODE
RETI
```

replace NEW CODE et seq (which must not be under a ROM) by:

```
LD      A,C
SCF
EX      AF,AF'
CALL    LABEL_1
DI
JP      ADDRESS_X +10
```

LABEL_1:

```
EI
RETI
```

Note that there is no suitable indirection or jumpblock into which the new code can be added, so that it is necessarily somewhat more contorted than the usual sort of code which one adds. Note also that this code is only for use with existing ROMs such as are fitted to the CPC464/664/6 128. Any future compatible machines will not support it - so it is most important that the program to install the code checks the ROM version number before proceeding. Steps will be taken to ensure newer ROM versions will not need alteration in this way.

Note that the old interrupt code and the indirection at #003B will be replaced when KL CHOKE OFF, MC BOOT PROGRAM or MC START PROGRAM are run (viz when a new foreground or background program is executed). Fortunately these are also the routines which issue RESET PERIPHERALS request.

Appendix XIV

Printer Translation Table

A facility is provided whereby special characters which may appear on the screen and which are supported by the AMSTRAD DMP-1, will be printed even though the character codes for the screen and printer may be different. The majority of these symbols will only be available when the printer is switched to one of its foreign language modes.

For example, if a circumflex is to be printed then the character code for the screen would be &A0 but on the printer it is &5E, if an &A0 was sent to the printer it would be translated into a &5E thus printing a circumflex

The following table shows all the default translation codes:

character	Tran.	US/UK	France	Germany	Spain				
#A0 Circumflex	:	#5E Circumflex	:	Circumflex	:	Circumflex	:	Cirumflex	:
#A1 Acute	:	#5C Backslash	/	C cedilla	ç	UC Oumlaut	Ö	UC Ntilde	Ñ
#A2 Umlaut	:	#7B Open Curly	{	E acute	É	LC Aumlaut	Ä	Umlaut	:
#A3 Pound	£	#23 Hash/pound	#/£	Hash	#	Hash	#	Hash	#
#A6 Section	§	#40 At	@	A grave	À	Section	§	At	@
#AB Plus/minus	±	#7C Vertical bar		U grave	Û	LC Oumlaut	ö	LC Ntilde	ñ
#AC Division	÷	#7D Close curly	}	E grave	È	LC Uumlaut	ü	Close curly	}
#AD Not	¬	#7E Tilde	~	Umlaut	¨	Scharfes S	ß	Tilde	~
#AE Inverted	¡	#5D Close square]	Section	§	UC Uumlaut	Ü	Inverted	¡
#AF Inverted	¿	#5B Open square	[degree	°	UC Aumlaut	Ä	Inverted	¿

The default translation table only translates the additional characters in the character set (#A0..#AF). It does not translate any of the standard ASCII characters or the graphics characters.