Hewlett-Packard

Computer Peripherals Bristol

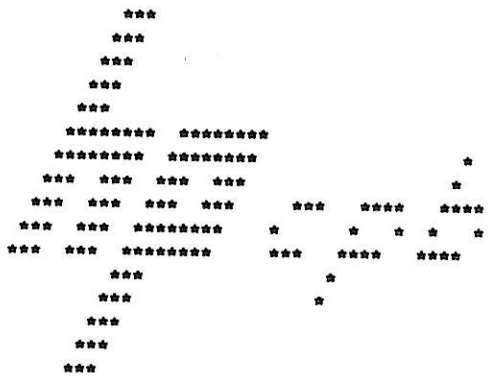

DC Control Proposal for Sequential Access devices with special regard to DDS DAT drives.

Document : X3T9.2/90-119 Rev 1

Date : 29th October 1990

Author : Steve Krupa, HP CPB

Hewlett Packard, Computer Peripherals Bristol,
Filton Rd,
Stoke Gifford,
Bristol, BS12 6QZ

Tel : + 272 799910
Fax : + 272 236091

INTRODUCTION

This is the first revision of the document X3T9.2/90-119.

The initial document was dated 16th July 1990.

A Proposal for a SCSI MODE SENSE/SELECT page to control and report on the operation of a DC DDS drive, with specific regard to Data Compression components. It is intended for, but not necessarily limited to, support of DDS-DC drives which make use of lossless compression algorithms which are based on substitution; e.g. those of the Lempel-Ziv family.

There are a number of issues addressed by this document with regard to data compression.

a) How does a host use the available functionality in SCSI-2 to control data compression ?

b) What additional features could a device which supports data compression offer to the host and how should these be made available to the host ?

c) If a device which does not support data compression encounters a piece of media containing compressed data, what action should it take ?

d) What form does the host-device interface take in order to allow the host to perform software decompression if required ? This is an issue for both current functionality in SCSI-2 and for any additional feature set.

Document Structure :

Part 1 covers a simple control scheme for a limited subset of DC control requirements which are, nevertheless, sufficient for most purposes.

Part 2 describes an additional, richer scheme for controlling and monitoring a wider set of DC parameters. A device may optionally implement this scheme in addition to that of Part 1.

Part 3 is an Application Note which describes the host/device interaction at boundaries between compressed and uncompressed data using either scheme.

PART 1
****

DATA COMPRESSION CONTROL using DEVICE CONFIGURATION MODE PAGE

This section describes the way in which Data Compression is controlled using the SDCA (Select Data Compression Algorithm) field in the Device Configuration Mode Page for DDS drives. It also addresses the issue of software decompression and non-DC drives.

1.1 DC support in X3T9.2 SCSI-2 Rev 10

The only support currently available in the SCSI set is a one-byte field, the SDCA field (byte 14), in the Device Configuration Page (page 10h). This page is specific to Sequential Access devices. The use of this field is in many ways vendor-specific - the QIC manufacturers have reached a separate agreement on its meaning.

1.1.2 X3T9.2 definition

Byte value	Description
00h	Disable Compression
01h	Select targets default compression algorithm
02h - 7Fh	Select compression algorithm #
80H - FFh	Vendor Specific

1.1.3 QIC definition

7	6	5	4	3	2	1	0
DC on	On Drive	QIC-Approved		DC	Algorithm		

Byte Value	Description
00h	Disable Compression
01h - 7Fh	Invalid
10xxxxxxb	DC on - host performs algorithm # xxxxxxb
11xxxxxxb	DC on - drive performs algorithm # xxxxxxb

As can be seen from the QIC definition, QIC chooses to ignore the X3T9.2 definition of byte values 01h-7Fh.

1.2 Controlling Data Compression

The format of the SDCA field for DDS drives is :

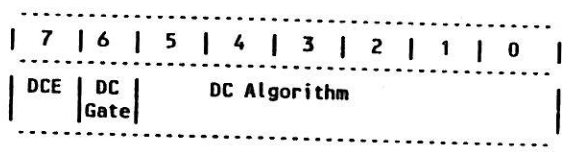

DCE - Data Compression Enable :

On a MODE SELECT command this field allows the host to enable or disable data compression.

If the host sets this field to 1 then if the device supports data compression, it will compress any subsequent data sent to it by the host before writing it to the media. The algorithm used to compress the data will be the one specified in the DC Algorithm field (providing that the DC Gate field is 1). If the device does not support data compression then it will reject a value of 1 in this field with a CHECK CONDITION status and a sense key of ILLEGAL REQUEST.

If the host sets this field to 0 then the device will disable data compression, and any subsequent data sent to the device by the host will be written to the media uncompressed. If the device does not support data compression then this field must be 0.

On a MODE SENSE command, this field allows the host to determine whether compression is enabled or disabled. A value of 1 indicates that compression is enabled, and a value of 0 indicates that compression is disabled. For devices which do not support data compression, this field will always be 0.

DC Gate - Data Compression Gate :

On a MODE SELECT command this field allows the host to control the device's ability to decompress data which it is reading from the media.

If the host sets this field to 1 then the DC Algorithm field refers to the algorithm the device will use to compress data sent to it by the host before writing it to the media. If the host sets this field to 0 then the DC Algorithm field refers to the type of data item (uncompressed data item - ie record, or compressed data item) which the device will send to the host in response to a READ command. This is to allow the host to perform software decompression if necessary. Note that some devices may not support the ability to return compressed data items to the host for software decompression. See Part 3 of this document for a discussion of software decompression issues with regard to host/device interaction.

On a MODE SENSE command this field allows the host to determine whether or not the device supports data compression. A value of 1 will be returned by devices which do support data compression, whilst a value of 0 will be returned by devices which don't support data compression.

DC Algorithm :

On a MODE SELECT command the meaning of this field is determined by the value of the DC Gate field.

If the DC Gate field is 1 then the DC Algorithm field determines the algorithm the device will use to compress data sent to it by the host. A value of 0 will direct the device to deselect all compression algorithms and is only valid if the DCE field is 0. A CHECK CONDITION status with ILLEGAL REQUEST sense key will otherwise be generated. A value of 1 will direct the device to select it's default compression algorithm whilst a value of N, where N > 1, will direct the device to select compression algorithm N. If the device does not support the indicated algorithm then it will issue a CHECK CONDITION status with ILLEGAL REQUEST sense key. Note that for devices which don't support data compression, values of 0 or 1 are legal for this field.

If the DC Gate field is 0 then the DC Algorithm field determines the type of data item which the device will send to the host in response to a READ command. A value of 0 directs the device to return uncompressed records to the host and to generate a CHECK CONDITION if it encounters any other data item type. A value of 1 directs the device to return uncompressed records to the host and if possible to decompress any compressed data items which it encounters and return the data to the host. If any compressed data items which the device cannot decompress are encountered, a CHECK CONDITION will be generated. A value of N, where N > 1, will direct the device to return compressed data items of type N to the host undecompressed but to generate a CHECK CONDITION if it encounters any other data item type (including uncompressed records) on the media. See Part 3 of this document for a discussion of software decompression issues with regard to host/device interaction.

On a MODE SENSE command, the meaning of this field is determined by the value of the DC Gate field.

For a device which supports data compression, the DC Gate field is always 1 on a MODE SENSE command to indicate to the host that it supports data compression. For this reason, the DC Algorithm field contains the algorithm identifier which the device will use to compress data sent to it by the host.

For a device which doesn't support data compression, the DC Gate field is always 0 on a MODE SENSE command to indicate to the host that it does not support data compression. For this reason, the DC Algorithm field contains the data item type which the device will return to the host in response to a READ command.

Because the SDCA field is not large enough for all the necessary information, it is not possible for a device which supports data

compression to indicate to the host via a MODE SENSE command, the data item type which the device will return to the host in response to a READ command. However as the host must constantly be aware of whether the device is sending it compressed or uncompressed data in response to a READ command (so that it can perform software decompression as and when required) it is not a major problem.

Due to the limited support for DC given in the Device Configuration page, DDS-DC drives can optionally implement and support the Data Compression Characteristics page defined in Part 2 of this document, as well as the SDCA field of the Device Configuration page. Note that hosts also have this choice.

1.2.1 Compression Control Example

Consider :

A) A device which supports data compression connected to a host which doesn't support software decompression.

The only data compression control the host has to worry about is enabling and disabling compression and selecting the compression algorithm. If the host issues a MODE SENSE command to the device to read the Device Configuration Mode Page, the SDCA field will, for the sake of argument, contain 01nnnnnn. This indicates to the host that

a) Data Compression is disabled (DCE = 0)
b) The device supports data compression (DC Gate = 1)
c) The device has currently selected algorithm N as its compression algorithm (DC Algorithm = nnnnnn)

The host can send the following values of SDCA in the MODE SELECT Device Configuration Page without causing the device to issue a CHECK CONDITION :

a) 01nnnnnn - no change
b) 01000000 - deselect compression algorithm
c) 01000001 - select default compression algorithm
d) 11nnnnnn - enable compression with algorithm N
e) 11000001 - enable compression with default algorithm

B) A device which supports data compression connected to a host which supports software decompression.

The only data compression control the host has to worry about is enabling and disabling compression and selecting the compression algorithm. For software decompression the host must concern itself with the algorithm number of the data item type the device will return in response to a READ command. If the host issues a MODE SENSE command to the device to read the Device Configuration Mode Page, the SDCA field will, for the sake of argument, contain 01nnnnnn. This indicates to the host that

a) Data Compression is disabled (DCE = 0)
b) The device supports data compression (DC Gate = 1)
c) The device has currently selected algorithm N as its compression algorithm (DC Algorithm = nnnnnn)

The host can send the following values of SDCA in the MODE SELECT Device Configuration Page without causing the device to issue a CHECK CONDITION :

a) 01nnnnnn - no change
b) 01000000 - deselect compression algorithm
c) 01000001 - select default compression algorithm
d) 00000000 - return only uncompressed records in response to a READ command. Any other data item type to generate a CHECK CONDITION
e) 00000001 - return uncompressed records and if possible decompress any compressed data item encountered on the media, in response to a READ command. Any other data item types to generate a CHECK CONDITION
f) 00nnnnnn - return only undecompressed data items of type N. Any other data item type to generate a CHECK CONDITION
g) 00mmmmmm - return only undecompressed data items of type M. Any other data item type to generate a CHECK CONDITION
h) 11nnnnnn - enable compression with algorithm N
i) 11000001 - enable compression with default algorithm
j) 10000000 - return only uncompressed records in response to a READ command. Any other data item type to generate a CHECK CONDITION.
Enable compression using currently selected compression algorithm, N.
k) 10000001 - return uncompressed records and if possible decompress any compressed data item encountered on the media, in response to a READ command. Any other data item types to generate a CHECK CONDITION
Enable compression using currently selected compression algorithm, N.
l) 10nnnnnn - return only undecompressed data items of type N. Any other data item type to generate a CHECK CONDITION
Enable compression using currently selected compression algorithm, N.
m) 10mmmmmm - return only undecompressed data items of type M. Any other data item type to generate a CHECK CONDITION.
Enable compression using currently selected compression algorithm, N.

C) A device which doesn't support data compression connected to a host which supports software decompression.

For software decompression the host must concern itself with the algorithm number of the data item type the device will return in response to a READ command. If the host issues a MODE SENSE command to the device to read the Device Configuration Mode Page, the SDCA

field will, for the sake of argument, contain 00000000. This indicates to the host that

a) Data Compression is disabled (DCE = 0)
b) The device doesn't support data compression (DC Gate = 0)
c) The device will only return uncompressed records to the host in response to a READ command. (DC Algorithm = 0)

The host can send the following values of SDCA in the MODE SELECT Device Configuration Page without causing the device to issue a CHECK CONDITION :

a) 00000000 - no change
b) 00000001 - return uncompressed records and if possible decompress any compressed data item encountered on the media, in response to a READ command. Any other data item types to generate a CHECK CONDITION. Note that in effect this is identical to a).
c) 00mmmmmm - return only undecompressed data items of type M. Any other data item type to generate a CHECK CONDITION
d) 01000000 - deselect compression algorithm
e) 01000001 - select default compression algorithm

23

PART 2
****

DATA COMPRESSION CHARACTERISTICS MODE PAGE
****

This section describes a new Mode Page to be used for control of data compression in DDS drives. It adds new support for DC in the SCSI standard which is not possible using the current features.

2.1 Possible Information Requirements for DC page

Below is a list of possible DC information requirements of a host. Some of this information will be specified in the DC Mode Page. Some MAY be more suitable for a DC Log Page, whilst some info may be omitted altogether for being of limited or no use to the host. All items which are of POTENTIAL interest to a host are included here. This list is finally distilled into a Mode Page.

2.1.2 For compatibility with byte 14 Device Configuration page

It is important that any new Mode Page which adds new features on top of existing features must provide compatibility with those existing features. Therefore the support given by the SDCA field in the Device Configuration Mode Page must be catered for in the new Mode Page.

2.1.2.1 DC algorithm # (RW)

This field defines the Data Compression algorithm to use when writing data to the tape. Each target will normally only support 1 algorithm, so this will generally not change. A value of 01h indicates default algorithm selection, and a value of N indicates specific algorithm selection.

2.1.2.2 DC write on/off control (RW)

This field enables/disables compression. A value of 0 turns compression off, whilst a value of 1 turns it on. Reading this field indicates the current status of compression (ie on or off). Note that this field has no effect on decompression.

2.1.2.3 Boundary read algorithm # (RW)

A write to this field allows the drive to return entities of the specified type to the host without issuing any warnings about entity type. A read of this field allows the host to determine which

entity types will be returned undecompressed without host intervention. This allows the host to perform software decompression. This is explained in more detail in the section on DC page definition and in Part 3 of this document.

2.1.2.4 Decompress on drive/on host (RW)

This field defines whether undecompressed entities will be returned to the host or whether the entity should be decompressed by the drive. A value of 1 indicates that the drive will decompress the data, whilst a 0 indicates that the host will decompress it.

2.1.3 Auto-reset control

Some implementations of compression algorithms include optimisation features which attempt to maximise compression ratio and minimise overhead. Support for these features is given by the following fields.

2.1.3.1 Adaptive algorithm on/off (RW)

This field defines whether or not the device will dynamically monitor the compression ratio during writing in order to attempt to optimise it. Writing a 01h to this field will enable compression ratio monitoring. A 00h will disable this feature.

2.1.3.2 Enable Dynamic adaptive algorithm (RW)

This field defines whether or not the device will dynamically monitor the rate of change of compression ratio and attempt to use this data for optimisation purposes. Writing a 01h to this field will enable this feature, 00h will disable it.

2.1.3.3 Adaptive algorithm sample period (RW)

This field defines when the compression ratio will be sampled in order to attempt to optimise it. The section on DC page definition describes this in more detail.

2.1.3.4 Adaptive algorithm threshold ratio (RW)

This field defines the threshold level for the compression ratio below which the device will take optimising action to improve it. The section on DC page definition describes this in more detail.

2.1.4 Byte/Record/Entity/Object relationships

2.1.4.1 Max # of records in entity (RW)

This field defines the maximum number of records to be bundled together in one entity.

2.1.4.2 Min # of records in entity (RW)

This field defines the minimum number of records to be bundled together in one entity.

2.1.4.3 Max # of compressed bytes in entity (RW)

This field defines the maximum number of bytes of compressed data which can share an entity.

2.1.4.4 Min # of compressed bytes in entity (RW)

This field defines the minimum number of bytes of compressed data which can share an entity - this applies to multi-record entities only. Single record entities must contain the whole record so this field has no effect in this case.

2.1.4.5 Max # of uncompressed bytes in entity (RW)

This field defines the maximum number of uncompressed bytes which after compression can share an entity.

2.1.4.6 Min # of uncompressed bytes in entity (RW)

This field defines the minimum number of uncompressed bytes which after compression can share an entity.

2.1.4.7 Max # of records in object (RW)

This field determines the maximum number of records which will share the same dictionary. This field is used to force the creation of a new dictionary after the specified number of records. If the adaptive algoritm is enabled, then these dictionary creations will be in addition to any automatically mandated by the adaptive algorithm.

2.1.4.8 Min # of records in object (RW)

This field defines the minimum number of records to be bundled together in one object.

2.1.4.9 Max # of compressed bytes in object (RW)

This field defines the maximum number of bytes of compressed data which can share an object.

2.1.4.10 Max # of uncompressed bytes in object (RW)

This field defines the maximum number of uncompressed bytes which after compression can share an object.

2.1.4.11 Min # of compressed bytes in object (RW)

This field defines the minimum number of bytes of compressed data which can share an object.

2.1.4.12 Min # of uncompressed bytes in object (RW)

This field defines the minimum number of uncompressed bytes which after compression can share an object.

2.1.4.13 Force new entity (W)

This field is write-only. On Read it will always be returned as 0. If the host needs to force the drive to close the current entity and start a new one, then it sets this field to 1. This allows the host maximum flexibility when it comes to handling entities. At any time, the host may set this field and so force a new entity. (See also Force Dictionary Reset below).

Note : If multiple maximum (or minimum) values are specified for the fields above then the first boundary reached will be the limiting condition.

2.1.5 Compression ratio

2.1.5.1 Cumulative compression ratio (RW)

This field returns the cumulative compression ratio achieved by the drive since the last power-on, SCSI device reset or write to this field, whichever occurred most recently. A write to this field will clear it to an undefined state (ie 0). Compression Ratio data will then be accumulated afresh.

2.1.5.2 Number of entities written (R)

This field contains a count of the total number of entities written to tape since power-on/SCSI device reset. This field is read-only. It is ignored on writes. A power-on or SCSI device reset will clear this field.

2.1.5.3 Number of entities decompressed (R)

This field contains a count of the total number of entities read from tape and decompressed by the device since power-on/SCSI device reset. This field is read-only. It is ignored on writes. This field is cleared by a power-on or SCSI device reset.

26

2.1.5.4 Number of entities sent for host decompression (R)

This field contains a count of the total number of entities read from tape and sent to the host for decompression since power-on/SCSI device reset. This field is read-only. It is ignored on writes. This field is cleared by a power-on or SCSI device reset.

2.1.6 Dictionary and Buffer/FIFO

In order to smooth performance, many implementations of compression algorithms have either buffers or FIFOs on their input. Information about the buffers or FIFOs may help the host to set the appropriate maximum burst length in order to tune performance. Some implementations require the use of a data dictionary which is used to compress and decompress the data. Information about the dictionary may help the host in setting other optimisation parameters for specific types of data.

2.1.6.1 SRAM size (RAM type) (R)

This read-only field defines the SRAM size (in 1K blocks) used by the DC engine.

2.1.6.2 Buffer/FIFO size (R)

This read-only field describes the buffer/FIFO size (in 1K blocks) used by the DC engine.

2.1.6.3 DC Engine status (R)

This read-only field describes the status of the internal buffering within the DC engine. A value of 01h means that the the DC engines internal buffers have been flushed, whilst a value of 00h means that there is data held within the DC engine.

2.1.6.4 Dictionary size (R)

This read-only field describes the size of dictionary (in 1Kbyte blocks) used by the DC engine.

2.1.6.5 Input Buffer/FIFO size

This read-only field defines the size of the input buffer/FIFO used by the compression engine.

2.1.6.6 Dictionary status (R)

This flags field indicates the dictionary/engine status.

Dictionary Full
Dictionary Half-full
Dictionary Frozen

2.1.6.7 Force Dictionary Reset (W)

This field is write-only. On Read it will always be returned as 0. If the host needs to force the drive to reset the dictionary to its empty state, then it sets this field to 1. This allows the host a degree of control over dictionary use. Along with the dictionary status information and the compression ratio this field allows the host to have additional influence over the drives performance. Using this field along with the Force New Entity field allows the host to create access points at the beginning of entities.

2.1.7 Miscellaneous

2.1.7.1 Expansion minimisation on/off (RW)

Expansion minimisation is similar to compression ratio monitoring in that it checks to see if the compression ratio has fallen below the threshold of 1:1, and turns off DC if it has. Setting this field to 1 will allow the drive to disable DC if the compression ratio falls below 1, and to re-enable it if it rises above 1. Clearing this field disables this function.

2.1.7.2 Parity checking/generation on/off (RW)

This field defines the status of parity checking/generation on the input and output ports which are used to pass data through the DC engine. Port A is defined as the SCSI-side port and port B is defined as the Buffer-side port.

2.2 Proposal for SCSI DC Characteristics Mode Page

2.2.1 DC Page Definition

The Data Compression Characteristics Mode Page is defined as :

Bit / Byte	7	6	5	4	3	2	1	0
0	Reserved (0)			Page Code (0Fh)				
1	Page Length (16h)							
2	DCE	Res(0)		DC Algorithm				
3	Res(0)	Res(0)		Boundary Read Algorithm				
4	Reserved (0)					EME	AAE	DME
5	Adaptive Algorithm Threshold							
6	Adaptive Algorithm Sample Period							
7	Maximum Records in Entity							
8	Reserved (0)						DSE	Res(0)
9	(MSB) Cumulative Compression Ratio							
10	(LSB)							
11	Dictionary Size							
12	Buffer/FIFO Size							
13	Reserved (0)							
14	Reserved (0)							
15	Reserved (0)							
16	Reserved (0)							
17	Reserved (0)							
18	Reserved (0)							
19	Reserved (0)							
20	Reserved (0)							
21	Reserved (0)							

2.2.1.1 DCE : Data Compression Enable

On a MODE SELECT command this field allows the host to enable or disable data compression.

If the host sets this field to 1 then the device will compress any subsequent data sent to it by the host before writing it to the media. The algorithm used to compress the data will be the one specified in the DC Algorithm field.

If the host sets this field to 0 then the device will disable data compression, and any subsequent data sent to the device by the host will be written to the media uncompressed.

On a MODE SENSE command, this field allows the host to determine whether compression is enabled or disabled. A value of 1 indicates that compression is enabled, and a value of 0 indicates that compression is disabled.

2.2.1.2 DC Algorithm :

The DC Algorithm field determines the algorithm the device will use to compress data sent to it by the host. A value of 0 will direct the device to deselect all compression algorithms and is only valid if the DCE field is 0. A CHECK CONDITION status with ILLEGAL REQUEST sense key will otherwise be generated. A value of 1 will direct the device to select it's default compression algorithm whilst a value of N, where N > 1, will direct the device to select compression algorithm N. If the device does not support the indicated algorithm then it will issue a CHECK CONDITION status with ILLEGAL REQUEST sense key.

2.2.1.3 Boundary Read Algorithm :

The Boundary Read Algorithm field determines the type of data item which the device will send to the host in response to a READ command. A value of 0 directs the device to return uncompressed records to the host and to generate a CHECK CONDITION if it encounters any other data item type. A value of 1 directs the device to return uncompressed records to the host and if possible to decompress any compressed data items which it encounters and return the data to the host. If any compressed data items which the device cannot decompress are encountered, a CHECK CONDITION will be generated. A value of N, where N > 1, will direct the device to return compressed data items of type N to the host undecompressed but to generate a CHECK CONDITION if it encounters any other data item type (including uncompressed records) on the media. See Part 3 of this document for a discussion of software

decompression issues with regard to host/device interaction.

2.2.1.4 EME : Expansion Minimisation Enable

This field is used to ensure that the drive limits the actual expansion of data which it is trying to compress. If the data written by the host is random, then algorithms designed to compress the data may actually expand it. In order to minimise this, some drives are able to monitor the compression ratio independent of the adaptive algorithm and be able to disable DC if expansion is detected. Such a drive continues to monitor CR, and if it rises above 1 : 1 (ie real compression occurs) then DC will be re-enabled. On MODE SELECT, setting this field to 1 will enable expansion minimisation, whilst setting it to 0 will disable expansion minimisation. Drives which do not support this feature will return CHECK CONDITION with ILLEGAL REQUEST sense key when this field is set to 1. On MODE SENSE, this field allows the host to determine whether expansion minimisation is enabled or disabled.

2.2.1.5 AAE : Adaptive Algorithm Enable

This field is used to enable/disable the use of a device-specific Adaptive Algorithm for monitoring compression ratio. The drive can use this algorithm to optimise CR for the data currently being written. On MODE SELECT, setting this field to 1 will enable the Adaptive Algorithm using the parameters specified in the Adaptive Algorithm Threshold and Adaptive Algorithm Sample Period fields described below. If the drive does not support this feature then it will return CHECK CONDITION with ILLEGAL REQUEST sense key. On MODE SENSE, this field allows the host to determine whether the Adaptive Algorithm is enabled or disabled.

2.2.1.6 DME : Dynamic Modification Enable

This field is used to allow/disallow the drive to dynamically modify the parameters used by the Adaptive Algorithm in order to make it more responsive to the nature of the data being written. This introduces a higher level of optimisation control than is possible using a preset Adaptive Algorithm. On MODE SELECT, setting this field to 1 will enable Dynamic modification. If this field is set to 1, then the AAE field must also be set to 1. If the drive does not support this feature then it will return CHECK CONDITION with ILLEGAL REQUEST sense key. On MODE SENSE, this field allows the host to determine whether Dynamic Modification is enabled or disabled. Note that even though the device may modify the Adaptive Algorithm Threshold and Adaptive Algorithm Sample Period fields, a MODE PARAMETERS CHANGED CHECK CONDITION will not be generated.

2.2.1.7 Adaptive Algorithm Threshold :

This field sets the threshold level of CR, below which the Adaptive Algorithm will react accordingly in order to maximise the

compression ratio. This field is defined as :

CR threshold = 256/(256 - Adaptive Algorithm Threshold)

This gives values of

Adaptive Algorithm Threshold	CR Threshold
0 -> 128	1 -> 2
129 -> 170	2 -> 3
171 -> 192	3 -> 4
193 -> 224	4 -> 8
225 -> 240	8 -> 16
241 -> 255	16 -> 256

For this field to be valid the AAE field must be set to 1. On MODE SELECT, the host shall set the field as required to one of the above values. If the drive does not support this feature then it will return CHECK CONDITION with ILLEGAL REQUEST sense key. On MODE SENSE, this field allows the host to determine what the Adaptive Algorithm Threshold is. If the DME field is set to 1 then this field will be dynamically modified by the drive.

2.2.1.8 Adaptive Algorithm Sample Period :

This field sets the sample period, in terms of uncompressed bytes, at which the Adaptive Algorithm takes a snapshot of the CR, in order to determine whether or not to take action to maximise the compression ratio. This field is in 128 byte increments. For this field to be valid the AAE field must be set to 1. On MODE SELECT, the host shall set the field as required to one of the above values. If the drive does not support this feature then it will return CHECK CONDITION with ILLEGAL REQUEST sense key. On MODE SENSE, this field allows the host to determine what the Adaptive Algorithm Sample Period is. If the DME field is set to 1 then this field will be dynamically modified by the drive.

2.2.1.9 Maximum Records in Entity :

This field allows the host a degree of control over how many records are compressed into each entity. If the host is writing small records then this field can be used to tune performance to record size. This field is defined as :

Max Recs in Entity	Actual max recs
0	No Maximum
n	2*n

This gives a maximum value of 510 records per entity. On MODE SELECT, the host shall set the field as required to one of the above

values. If the drive does not support this feature then it will return CHECK CONDITION with ILLEGAL REQUEST sense key. On MODE SENSE, this field allows the host to determine what the maximum number of records allowed in each entity is.

2.2.1.10 DSE : Data Sink Enable

This field allows the host to determine what the actual compression ratio of an input data stream is without actually writing data onto the media. The data will be sinked through the DC engine giving the host access to all the information it would have if the data was actually written to the media. On MODE SELECT, setting this field to 1 enables Data Sink mode, whilst setting it to 0 disables it. On MODE SENSE, the state of Data Sink mode will be reported to the host.

2.2.1.11 Cumulative Compression Ratio :

This field allows the host to monitor the cumulative CR since device reset. This field is in binary-coded decimal and is 100x the compression ratio quoted to 2 decimal places. For example, if this field contains 0374, then the CR is 3.74:1. On MODE SELECT, setting this field to 0 will restart the Cumulative Compression Ratio calculation. Any other value will result in a CHECK CONDITION status with ILLEGAL REQUEST sense key. ON MODE SENSE, the cumulative compression ratio will be returned to the host. If the drive does not support this feature, then this field will be set to 0. Note that even though this field is updated by the device and therefore changes, a MODE PARAMETERS CHANGED CHECK CONDITION will not be generated.

2.2.1.12 Dictionary Size :

This field allows the host to determine the size of dictionary used by the compression engine. This field defines the number of 1Kbyte blocks allocated to dictionary size/compression RAM. So a value of 16 represents 16Kbytes. On MODE SELECT, this field has no meaning and is ignored. On MODE SENSE, this field takes one of the values as specified above.

2.2.1.13 Buffer/FIFO Size :

This field allows the host to determine the size of buffer/FIFO used by the compression engine. This field defines the number of 1Kbyte blocks allocated to input buffer/FIFO. On MODE SELECT, this field has no meaning and is ignored. On MODE SENSE, this field takes one of the values as specified above.

PART 3
****

APPLICATION NOTE
****

3.1 SCSI Protocol Issues

On a typical DDS-DC tape, two different data item types may be encountered. A data item is either :

a) An uncompressed record
b) An entity compressed using algorithm L

Entities are written by DDS drives which support on-board data compression. An entity is made up of a number of same-sized records compressed using the device's compression algorithm and prefixed by an entity header which is an uncompressed descriptor containing information about the data within the entity.

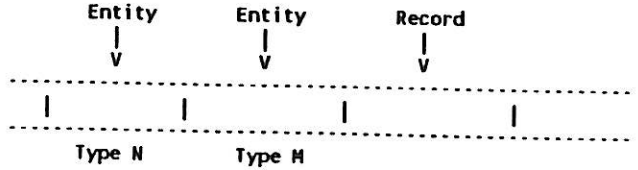

A device which supports compression algorithm N can decompress entities of type N when it encounters them on tape, and return the decompressed data to the host transparently. It can also return uncompressed records to the host as it finds them on the tape.

If such a device encounters an entity of type M (written by a different device using a different compression algorithm) then it doesn't generally know how to decompress it. Similarly if a device which doesn't support data compression at all encounters an entity on tape then it can not decompress it.

Some hosts may support software decompression, where they themselves are capable of decompressing entities. This requires that the device be able to return an entity UNDECOMPRESSED to the host. The host must also be aware that the data it is receiving is COMPRESSED data and not a normal uncompressed record.

Consider the following descriptions of the use of both the SDCA field in the Device Configuration Page, and the Data Compression Characteristics Page by both DDS and DDS-DC drives.

3.2 Host/Drive Interaction for Software Decompression

3.2.1 DC drive - initial operating mode example

After drive reset, a drive which supports data compression may, for example, power up with compression enabled and algorithm N selected. The SDCA field of the Device Configuration Page contains 11nnnnnn. The Data Compression Characteristics Page contains the following values (only those relevent to software decompression shown)

DCE = 1 :

Compression enabled.

DC Algorithm = N :

Compression algorithm N selected.

Boundary Read Algorithm = 1 :

Uncompressed records returned to host and entities of type N decompressed and the data returned to the host.

3.2.2 Non-DC drive - initial operating mode example.

A drive which doesn't support data compression may, for example, power up with the SDCA field of the Device Configuration Page containing 00000000. The Data Compression Characteristics Page contains the following values (only those relevent to software decompression shown) :

DCE = 0 :

Compression disabled.

DC Algorithm = 0 :

Compression algorithm deselected.

Boundary Read Algorithm = 1 :

Uncompressed records returned to host.

3.2.3 Example of Software Decompression control

The host issues a number of READ commands to the drive, all of which are successful. On the next READ command, however, the drive detects an entity of type M on the media, where M is an unsupported compression algorithm. The drive cannot decompress the entity. It is necessary for the drive to inform the host that it has encountered a

data item on the tape which it cannot decompress. It does this by issuing a CHECK CONDITION to the host and setting up the sense data as follows :

Valid = 1

The Information field contains residual information from the failed READ command.

Sense Key = NO SENSE

There is no specific error condition associated with the CHECK CONDITION.

Information = READ residue

The READ command failed with a residue as given in this field.

Command-Specific Information = Number of records in data item

The number of records in the entity is obtained from the entity header which the drive can read.

Additional Sense Code = CHANGE OF DATA ITEM TYPE ENCOUNTERED (60h)

This ASC indicates the reason for the CHECK CONDITION.

Additional Sense Code Qualifier = 0pmmmmmm

Where p=0 because the drive can't decompress the entity and mmmmmm identifies the algorithm used to compress the entity.

The drive is now positioned on the EOT side of the entity. The host must decide whether or not it wishes to read the entity and decompress it itself. If it doesn't then it can continue issuing READ commands which will succeed as long as the drive continues to encounter uncompressed records or entities of type N (where N is a compression algorithm which is supported by the drive) on the media. As soon as another entity of type other than N is encountered, the drive will issue a CHECK CONDITION and initialise the sense data as shown above.

3.2.3.1 Moving to Software Decompression mode

If the host wishes to read the entity, it must tell the drive that it should send back the entity UNDECOMPRESSED to the host. First, the host must reposition the drive to the beginning of the entity by issuing a SPACE record command with the count field set to the two's complement of the number of records in the entity (as given in the Command-Specific Information field of the sense data). The drive will space backwards over the entity.

Next the host issues a MODE SENSE command to read either the Device Configuration Page (if it is going to use the SDCA field) or the Data Compression Characteristics Page. If it is to use the Device Configuration Page then it sets the DC Gate field to 0, and the DC Algorithm field to mmmmm. If it is to use the Data Compression Characteristics Page then it sets the Boundary Read Algorithm field to mmmmm.

The host does not know how long the entity is in terms of bytes, so in order to be sure that it reads all the data, it must select variable record mode by setting the Block Length field in the Mode Parameter Block Descriptor to 0.

The host then issues a MODE SELECT command to write the page back to the drive. This will do two things :

a) Put the drive in variable mode
b) Change the mode of the drive so that it will only send data items of type M back to the host. If it encounters ANY OTHER data item type during a READ command, then it will issue a CHECK CONDITION.

From now on every READ command that the host issues to the device will result in an entity of type M being read, until a different data item type is encountered on the tape. Note that in order to guarantee that the host is able to retrieve the whole entity, it must issue a variable mode READ command with the SILI bit set and the Transfer Length set to OFFFFFFh.

If the drive encounters a different data item type on the media then it will issue a CHECK CONDITION status to the host. This is necessary as the host is only expecting undecompressed entities of type M. The drive initialises the sense data to contain the following information :

Valid = 1

The Information field contains residual information from the failed READ command.

Sense Key = NO SENSE

There is no specific error condition associated with the CHECK CONDITION.

Information = READ residue

The READ command failed with a residue as given in this field.

Command-Specific Information = Number of records in data item

If the data item is an uncompressed record then this field will be 1. If the data item is an entity then the number of records in the entity is obtained from the entity header which the drive can read.

Additional Sense Code = CHANGE OF DATA ITEM TYPE ENCOUNTERED (60h)

This ASC indicates the reason for the CHECK CONDITION.

Additional Sense Code Qualifier = 0pqqqqqq

Where

p=0 if the drive can't decompress the new entity encountered

and

p=1 if the drive has encountered an uncompressed record or an entity which it can decompress

and where

qqqqqq is 0 if an uncompressed record is encountered or non-0 if a compressed entity is encountered.

The drive is now positioned on the EOT side of the data item.

If the Additional Sense Code Qualifier has bit 6 (p) = 0 then the host must make the same decision it made earlier about whether it wishes to read the undecompressed entity.

3.2.3.2 Moving to In-drive Decompression mode

If the Additional Sense Code Qualifier has bit 6 (p) = 1 then the host can return the drive to its default mode of operation because the drive either can decompress the entity (if the data item is an entity) or has encountered an uncompressed record in which case no decompression is required anyway.

To do this, it must reposition the drive to the beginning of the data item encountered by issuing a SPACE record command with the count field set to the two's complement of the number of records in the data item (as given in the Command-Specific Information field of the sense data). The drive will space backwards over the data item.

Next the host issues a MODE SENSE command to read either the Device Configuration Page (if it is going to use the SDCA field) or the Data Compression Characteristics Page. If it is to use the Device Configuration Page then it sets the DC Gate field to 0, and the DC Algorithm field to 1. If it is to use the Data Compression Characteristics Page then it sets the Boundary Read Algorithm field to 1.

The host then issues a MODE SELECT command to write the page back to the drive. This will return the drive to its default mode so that it will send uncompressed records back to the host and will decompress entities of type N and send the data back to the host in response to a READ command.

3.3 Summary

The above example illustrates the requirements from both a host and a drive if software decompression is to be supported within a system. The interaction depends upon the drive being aware of, and being able to report, compression/non-compression data boundaries, upon the host having the capability to perform software decompression and upon drives ability to return compressed data items to the host undecompressed. The rationale behind the use of the MODE SENSE/SELECT control protocol is that it enables the use of a standard READ command to get the compressed data from the drive to the host.

GLOSSARY

Access Point :

A point in the compressed data stream at which a DC engine can start decompression without having any prior knowledge of the nature of the compressed data stream.

Adaptive Algorithm :

An algorithm which is self-modifying in response to a defined set of external stimulii. In the case of data compression the external stimulus is the nature of the input data stream.

Compression Ratio (CR) :

The ratio of input byte count to output byte count for a DC engine which is performing data compression. This term specifies the amount of redundancy removed from the input data stream.

Cumulative Compression Ratio (CCR) :

For a number of instances of input data to a DC engine, the cumulative compression ratio is the running average of the individual compression ratios for each input stream.

Dictionary :

A dynamically built data structure which describes repeating patterns in a input stream of data with the aim of allowing the substitution of a referencing codeword in place of the pattern in the output data stream in order to remove redundancy.

Dictionary Reset :

An object boundary - the point at which one object finishes and another begins.

Entity :

A data set including 0 or more object boundaries, 0 or more record boundaries and beginning with a format-specific header. This is the smallest logical compressed data item accesible to a DC device.

Expansion Minimisation :

The minimisation of the amount of time at which the compression ratio is less than 1 : 1.

Port A :

SCSI-side i/o port of DC engine.

Port B :

Device-side i/o port of DC engine.

Object :

A data set sharing a compression/decompression dictionary and delimited either logically within a media format, or physically by a specific data pattern within the data stream.

Record :

A data set delimited by EOR markers either represented logically within a media format, or physically by a specific data pattern within the data stream.