RCSL No: 31-D528

Edition: December 1978

Author: Peter Jørgensen

Title:

RC8000 DISC PROCESS

A/S REGNECENTRALEN
af 1979

Keywords:

RC8000, Disc, External Process.

Abstract:

This manual describes the conventions of an external process controlling a disc storage module (RC82xx) connected to an RC8000 Computer.
(24 printed pages)

Printed by A/S Regnecentralen af 1979, Copenhagen

42-i 1341

Disc Storage Module (RC82xx)

The Data Medium

The disc packs used on the disc storage modules consist of a number of coaxially arranged discs rotating together.

Data is stored in tracks that are concentric circles on the disc surface. A read/write head is allocated to each surface on the disc pack; these read/write heads are movable in radial direction by means of a common head activator. Hence the radial position of the heads defines a track on all surfaces, these tracks are in common called a cylinder. To select a certain track it is necessary to select a read/write-head corresponding to a certain disc surface.

The disc pack is further divided into 21 sectors i.e. pie-shaped sections. Hence the tracks are divided into arcs, each being the intersection of a track with a sector (referred to as a 'segment').

This three dimensional subdivision (cylinder, surface, sector) can be vizualized by the following figure (for the purpose of clarity, the surfaces have been rearranged):

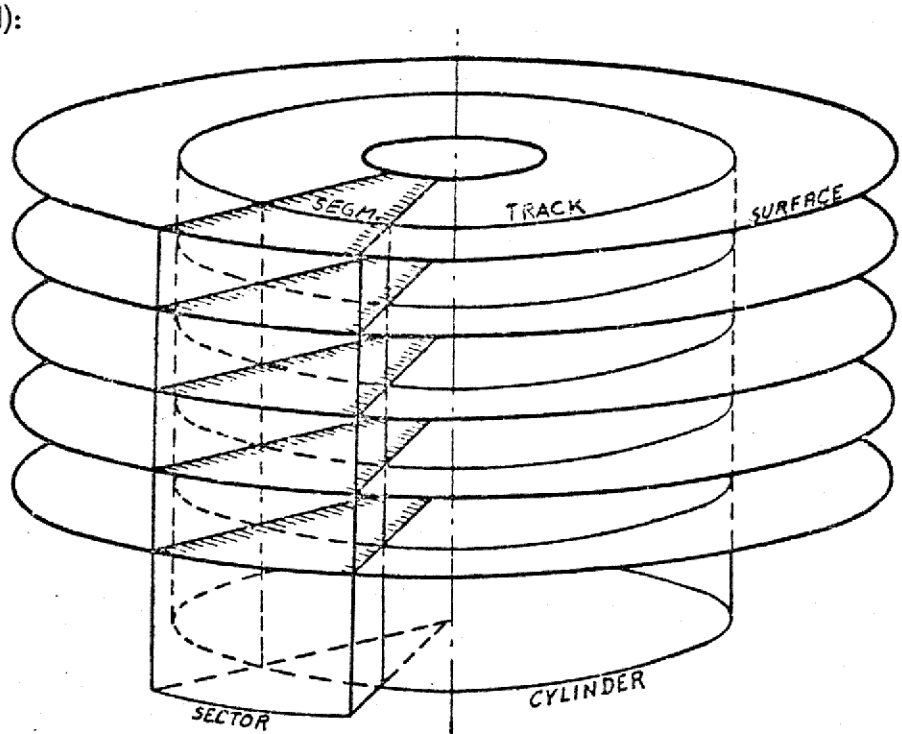

The actual surface allocation is shown in the following figure:

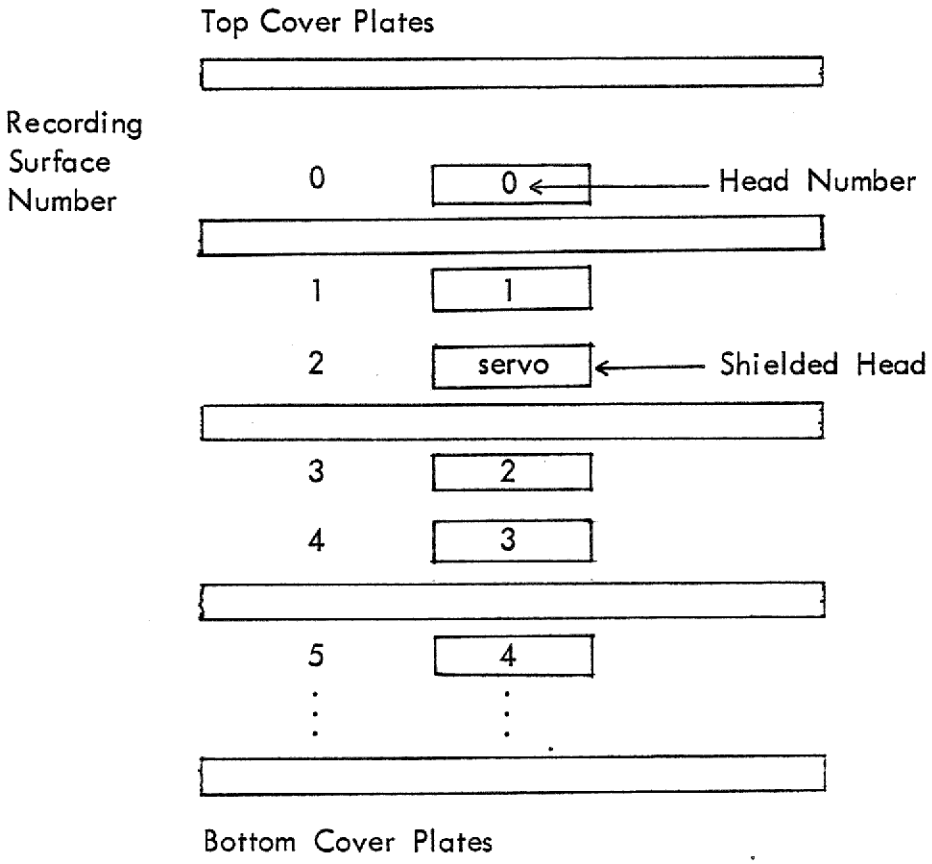

With the outer cylinder called cylinder 0 and sectors numbered 0-20 the following expression lays down the relation between a segment number and cylinder, surface given by head, and sector:

$$\begin{aligned}\text{segment number} = {} & \text{cylinder number} \times \text{number of heads} \times 21 \\ & + \text{head number} \times 21 \\ & + \text{sector number}\end{aligned}$$

Each segment is divided into two: an addressmark part, which is short in length (8 halfwords corresponding to 12 characters) and a longer datapart of 512 halfwords (768 characters). The address mark of a segment is automatically read and checked by hardware prior to any input or output operation to the segment.

During normal use, only the dataparts of segments are transferred to and from the computer.

An addressmark contains information about the position of the segment, and where to find next segment:

Character		0	1	2
Identification part		Cylindernumber		Headnumber
		3	4	5
		Sector	0	0
Next segment part		6	7	8
		Next cylindernumber		Next head
		9	10	11
		Next sector	0	0

To speed up operation, the sectors on odd cylinders (cylinder 1, 3, 5, ...) may be 'shifted' so as to avoid loosing a disc rotation when moving from one cylinder to the next. This can be illustrated as follows:

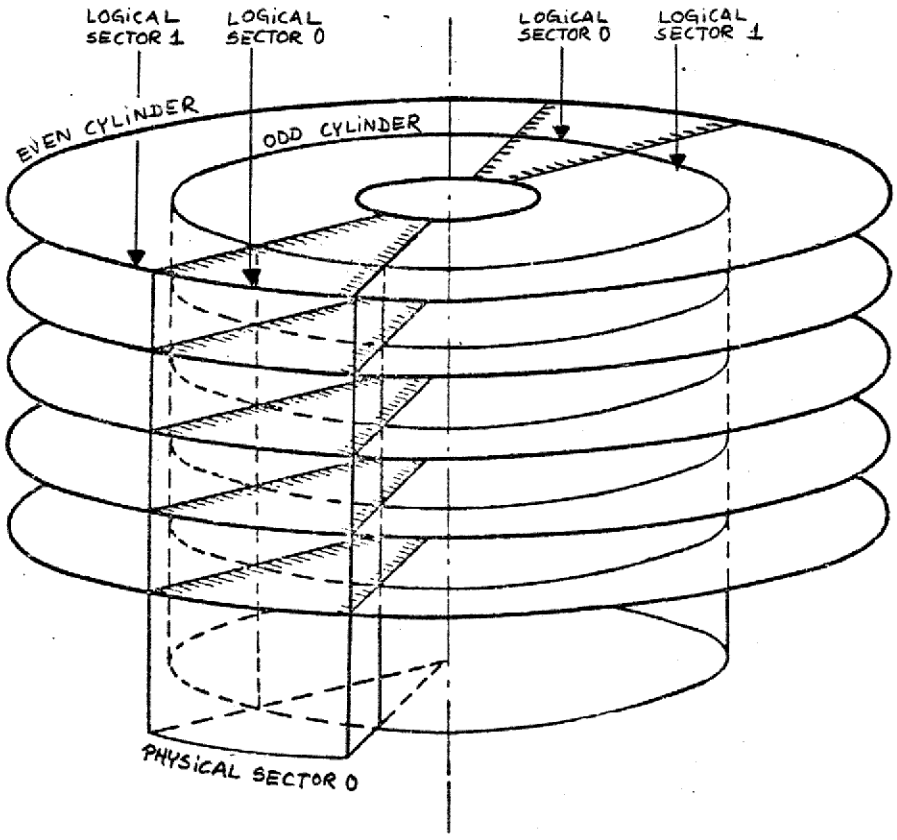

Use of odd cylinder shift will thus be reflected in the addressmarks:

- the next-sector field in the last addressmark on an even cylinder will be equal to the odd cylinder shift.

- the next-head field in an addressmark placed on physical sector 'odd cylinder shift - 1' on an odd cylinder will differ from the head-field in the identification part. The next-sector field will equal the odd cylinder shift except on the last surface on which it will be zero.

- the next-head field in an addressmark situated in the physically last sector on an odd cylinder will be equal to the head-field in the identification part and next-sector will be zero.

The formatting of a disc pack depends on parameters given to a disc formatting program.

For further information, consult the reference manual for the disc controller.

General Rules

A disc storage module may be accessed directly (i.e. not via area process) by sending a message to an external process controlling either a physical or a logical disc. The process kind is 62.

To each disc storage module corresponds always an external process (a physical disc driver) controlling the entire module (the physical disc). To each physical disc driver may exist none, one or more socalled logical disc drivers, a logical disc driver being an external process describing a part of the disc storage module called a logical disc (not to be mistaken for a logical backing storage). From the view of an internal process, each physical or logical disc driver represents an independent device identified by a unique device number.

Segment numbers supplied in messages sent to a physical disc driver specify absolute segments on the disc storage module, while segment numbers supplied in messages sent to a logical disc driver specify segments relatively to start of the logical disc. Odd cylindershift, if used, is automatically taken into account.

A message received by a logical disc driver will automatically be linked to the event queue of the physical disc driver controlling the disc storage module on which the logical disc in question is situated.

The division of one disc storage module into more logical discs is an installation dependent system option:

- if no division of a module is wanted, and if the module is not to be used as autoload device, then only a physical disc driver is required. If the module contains a logical backing store, then it also contains a catalog and a chain-table.

- if a module is to be used as potential autoload device and is to contain one

logical backing store, then a physical disc driver, a logical disc driver describing the autoload part, and a logical disc driver describing the backing store part are required. Only the latter part mentioned contains a catalog and a chaintable.

- if a module is to be divided into more independent logical backing stores, then each backing store contains a catalog and a chaintable, and each is represented by a logical disc driver.

Sense, position and input operations as well as extraction of statistics can be initiated by an internal process that is a user of the (physical or logical) device addressed. The device accepts this kind of messages simultaneously from more than one process provided no process has reserved it.

Initialize, clean track and output operations require that the device addressed has been reserved.

Operations specifying output of addressmarks, initialize and clean track operations can only be initiated on physical disc and will not be recognized by a logical disc driver.

A logical disc cannot be reserved if the corresponding physical disc is reserved by another internal process, and a physical disc cannot be reserved if one or more logical discs are reserved by other internal processes.

To use a logical disc, the corresponding physical disc driver need not have a name.

A disc driver representing a logical backing store is, during normal use, reserved by the anonymous process executing process functions.

Error Recovery

If a statusword indicates an intervention by the operator or a persistent lack of disc power, the names of the physical disc driver and associated logical drivers are removed (together with the present reservations) and all messages are answered with the result: receiver does not exist. Connections between area processes and involved devices will be cancelled.

Upon power restart and power failure restart, a loop, in which the disc is sensed regularly, is entered by those physical disc drivers which had had connection to their devices before power failed. If the disc does not become ready within a limited amount of time, the reactions will be as on operator intervention. During the loop execution, all messages will be left in the event queue, and an eventuel operation in progress at the time of power failure will be repeated when the disc becomes ready.

If an output or cleantrack operation is attempted on a write protected disc then the message will be answered with the result: message rejected.

If mode 2 (suppression of error recovery) is used, any other error will cause the message to be answered immediately.

If mode 2 is not used, the error recovery depends on operation and errorkind:

- Sense, output, position, and cleantrack operations are repeated up to three times before the error is considered irrecoverable. The same is performed for input operations provided the error is not a parity error.

- A parity error, which may be corrected by a hardware generated error correction code, occurring in the datapart of a segment input at an input operation causes the transferred data (but not the segment on the disc) to be corrected as if the parity error had not occurred.

- A parity error, which cannot be corrected by error correction code, will cause an input operation to be repeated a number of times utilizing the possibilities of 'strobing' in the segment a little earlier or later than nominally and of moving the read/write heads a little outside nominal position (offset).

  Three tries are made at each combination of strobe and offset (including nominal position), and the possibility of correcting the transferred data by error correction code is examined at each try.

  If mode 4 (limited error recovery) is used, only the 'best' combinations are tried and if the error is still unrecovered, the message is answered with a statusbit, disc error.

  In the normal mode all combinations of strobe and offset will be tried. The tries are divided into a number of portions; if the error is unrecovered when the tries of a portion have been performed then the message is moved to the end of the event queue, where it awaits the recovery process to resume until waiting messages have been processed.

A message answered with result: receiver malfunction, means that an error of the following kind has occurred: software timeout (disc operation not completed within one second), bus parity error, bus time out, bus communication error, interrupt error (something wrong during hardware status transfer), or no error indication (hardware status bit) at incomplete input/output operation.

The blocklength supplied in an answer to input and output messages corresponds to what was transferred before the error occurred.

Mode

The mode field of messages is used for specifiying specialactions, and consists of a sum of one or more of the following values:

1 specifies, in connection with input and output operations that the transfer concerns the addressmark parts of the segments.

2 specifies suppression of automatic error recovery. Will override mode 4.

4 specifies limited error recovery on parity error occurring at an input operation.

The normal mode (zero) implies:

- input and output operations concern the dataparts of segments.

- full error recovery.

Sense Operation

The device is sensed.

Input Operation

In the normal mode a number of consecutive segments of 256 words each are input to a storage area within the sending process. The first segment number wanted is given in the message.

The operation transfers the maximum number of segments for which there is room within the storage area, i.e.:

number of segments =
(last storage address + 2 - first storage address)//512

If the input block thus specified exceeds the upper limit of the (physical or logical) disc, input is performed only of that part of the block that is within the disc. In any case the actual number of halfwords transferred is given in the answer.

The number of halfwords transferred is a multiple of 512. The number of characters is defined as three times the number of words transferred.

If the first segment is outside the limits of the (physical or logical) disc, input is not initiated, but the answer contains a status bit, end medium, and the block length is zero.

In the addressmark mode the addressmark parts of a number of consecutive segments are input to a storage area while the dataparts of the segments are skipped. As each addressmark occupies 8 halfwords, the number of addressmarks input is given by:

number of addressmarks =
(last storage address + 2 - first storage address)//8

Apart from transferring a multiple of 8 halfwords (12 characters) the addressmark mode is equivalent to normal input mode.

Output Operation

Equivalent to the input operation.

In the addressmark mode the storage area must contain a number of addressmarks packed as previously described. The contents of the storage area is not checked. If the pack has not previously been formatted the clean track operation must be used prior to output of addressmarks (see also Initialize Operation).

Position Operation

The read/write heads are moved to the cylinder given by segment number in message. If the segment is outside the limits of the (physical or logical disc, no operation is initiated, but the answer contains a status bit, end medium.

Clean Track Operation

The whole track (cylinder, surface) on which the segment specified in message is situated will be cleaned.

If the segment is outside the limits of the (physical or logical) disc, no operation is initiated, but the answer contains a status bit, end medium.

The clean track facility is implemented for use the first time a disc pack is formatted.

Initialize Operation

This operation is used prior to formatting a new disc pack with addressmarks to prevent the automatical detection of odd cylindershift, number of heads and the like. No disc operation is initiated.

Disctype	0	graduated offset not possible on unit.
	1	graduated offset possible on unit.

Odd cylindershift 0 means no odd cylindershift.

Extract Statistics Operation

Dependent on the existence of a system option each physical disc driver maintains a table of statistical information reflecting the condition of the equipment.

The information may be input to a storage area within an internal process by its sending a message to a physical disc driver or one of its associated logical drivers. Having been moved, the table is reset.

If the storage area is too small, the answer contains a statusbit, block length error and the block length zero. Otherwise number of halfwords moved will be:

$$(\text{last storage address} + 2 - \text{first storage address})//2 \times 2$$

No disc operation is initiated.

The format of the table is described later in this section.

Status Bits

1	parity error or hard error
2	synchronization error (position error)
3	data overrun
4	blocklength error (extract statistics operation)
5	end medium
11	disc error (mode 4, input operation)

Messages and Answers

operation:	message:	answer:
sense	0<12 + mode	statusword
		0
		0
		i/o result
		current status
		event status
		detailed status
initialize	2<12	0
	number of heads	0
	odd cylindershift	0
	disctype	
input	3<12 + mode	statusword
	first storage address	number of halfwords
	last storage address	number of characters
	first segment number	i/o result
		current status
		event status
		detailed status
output	5<12 + mode	statusword
	first storage address	number of halfwords
	last storage address	number of characters
	first segment number	i/o result
		current status
		event status
		detailed status

operation:	message:	answer:
clean track	6<12 + mode	statusword
	irrellevant	0
	irrellevant	0
	segment number	i/o result
		current status
		event status
		detailed status
position	8<12 + mode	statusword
	irrellevant	0
	irrellevant	0
	segment number	i/o result
		current status
		event status
		detailed status
extract statistics	9<12	statusword
	first storage address	number of halfwords
	last storage address	number of characters
		undefined
		latest sensed current status
		latest sensed event status
		latest sensed detailed status

The information supplied in the answer from fourth word and on is normally only interesting to test programs.

The i/o result applies to the execution of the socalled channelprogram guiding the hardware:

i/o result	0	normal termination, the hardware has delivered an interrupt.
	1	bus reject, the device is not started.
	2	bus timeout, the device is not started.
	3	software time out.
	4	abnormal termination, the hardware has delivered an interrupt before the end of the channelprogram.
	5	wait program terminated.
	6	power restart.

Current status, event status, and detailed status are statuswords delivered by hardware. The statuswords are fully described in the reference manual for the disc storage controller.

Format of Statistical Information

halfword:	length:	contents:
0	1 word	number of times the device has been accessed.
2	-	number of times an error has occurred the first time an operation has been attempted (only repeatable errors counted).
4	-	number of times an error has been bypassed solely through use of error correction code.
6	-	number of times an error has been bypassed within three repetitions of the operation.
		number of errors bypassed through use of:
8	1 halfw.	negative offset
9	-	positive offset
10	-	(vacant)
11	-	late strobe
12	-	late strobe combined with negative offset
13	-	late strobe combined with positive offset
14	-	(vacant)
15	-	early strobe
16	-	early strobe combined with negative offset
17	-	early strobe combined with positive offset
18	-	(vacant)

halfword:	length:	contents:	
		number of errors bypassed through use of:	
19	1 halfw.	offset magnitude 1	
20	-	offset magnitude 2	
.	.	.	.
.	.	.	.
.	.	.	.
33	-	offset magnitude 15	
		number of occurrencies of:	
34	1 halfw.	event status bit 0	(intervention)
35	-	event status bit 1	(data error)
36	-	(vacant)	
37	-	event status bit 3	(data overrun)
38	-	event status bit 4	(hard error)
39	-	event status bit 5	(position error)
40	-	current status bit 0	(power low)
41	-	current status bit 1	(local)
42	-	current status bit 8	(write protect)
43	-	current status bit 9	(high density)
44	-	current status bit 10	(mode)
45	-	current status bit 5	(seek error)
46	-	(vacant)	
47	-	i/o result 6	(power restart)
48	-	i/o result 5	(wait prg. term.)
49	-	i/o result 4	(abnormal term.)
50	-	i/o result 3	(software timeout)
51	-	i/o result 2	(bustimeout)
52	-	i/o result 1	(bus reject)
53	-	i/o result 0	(normal term.)
54	-	event status bit 20	(bus comm. error)

halfword:	length:	contents:	
55	1 halfw.	event status bit 21	(interrupt error)
56	-	event status bit 22	(bustimeout)
57	-	event status bit 23	(bus parity error)
		number of occurencies of:	
58	-	detailed status bit 0	
59	-	detailed status bit 1	
.	.	.	.
.	.	.	.
.	.	.	.
81	-	detailed status bit 23	

table of potential error segments:

Each entry contains information about a segment, which could not be input or output in first attempt. When the table runs full, the registration of further detected error segments is ignored. To determine whether a segment has been included in the table due to media defect or a transient error a testprogram should be run. The table will typically consist of 10 entries each of 6 halfwords.

82	1 word	absolute segment no of 1st error segment.
84	1 halfw.	no of successful input operations (possibly using error re-recovery).
85	-	no of input operations requirering repetition.
86	-	no of successful output operations (possibly using error recovery).
87	-	no of output operations requirering repetition.
.	.	.
.	.	.
.	.	.

halfword:	length:	contents:
136	1 word	absolute segment no of 10th error segment.
138	1 halfw.	...
139	-	...
140	-	...
141	-	...

Please note that the table will not be updated when an error occurs in connection with a mode 2 (suppression of error recovery) operation.

RETURN LETTER

Title: RC8000 Disc Process RCSL No.: 31-D528

A/S Regnecentralen af 1979/RC Computer A/S maintains a continual effort to improve the quality and usefulness of its publications. To do this effectively we need user feedback, your critical evaluation of this manual.

Please comment on this manual's completeness, accuracy, organization, usability, and readability:

Do you find errors in this manual? If so, specify by page.

How can this manual be improved?

Other comments?

Name: _____ Title: _____

Company: _____

Address: _____

Date: _____

Thank you

42-i 1288

. . . . . Fold here . . . . .

. . . . . Do not tear - Fold here and staple . . . . .

Affix
postage
here

REGNECENTRALEN
af 1979

Information Department
Lautrupbjerg 1
DK-2750 Ballerup
Denmark