file - Multics
DSS 181/190

DATE: JUNE 26, 1972

TO:
C. T. CLINGEN
F. J. CORBATO'
R. C. DALEY
J. W. GINTELL
R. F. MABEE
N. I. MORRIS
J. H. SALTZER
R. B. SNYDER
D. R. VINOGRAD
S. H. WEBBER

FROM: L. SCHEFFLER

SUBJECT: GENERAL DESCRIPTION OF DSS-181 AND DSS-190 SUBSYSTEMS

Attached please find a draft namely, "General Description of DSS-181 and DSS-190 Subsystems".

Will you please read this document and direct any questions, comments, or gripes to me at Honeywell, X229, or send mail to Scheffler, Multics.

L. Scheffler

L. Scheffler

att.

1. Introduction

This document provides a general description and introduction to the equipment, features, and terminology of the on-line disk subsystems to be used with the follow-on Multics installation at MIT. It is essentially a summary of various Honeywell EPS-1 documents, and more detailed information may be obtained from the documents referenced on the last page.

2. Basic Disk Subsystem Components

Two new Disk Storage Subsystems (DSS) will be used with the Multics follow-on: DSS-181 and DSS-190. Both subsystems include several distinct components:

PSIA - Peripheral Subsystem Interface Adapter - This is the IOM side of the cable that goes to the disk subsystem controller.

LA - Link Adapter - This is the controller side of the cable from the PSIA at the IOM.

MPC - Microprogrammed Peripheral Controller - Microprogramming in the MPC converts instructions from the IOM into instructions for individual disk units, processes status and interrupts and maintains statistics.

CA - Controller Adapter - This is the MPC side of the cable that goes to each disk unit.

DLI - Device Level Interface - This is the disk unit side of the cable from a CA.

DSU - Disk Storage Unit - This is the actual disk drive and associated electronics that execute instructions from the MPC and transfer data.

Dual Port Option- This allows two cables, from separate CAs (and possibly separate MPCs), to hook up to a single disk unit so that there are two data paths to/from the disk unit.

The minimum configuration for either DSS-181 or DSS-190 is shown in figure 1.

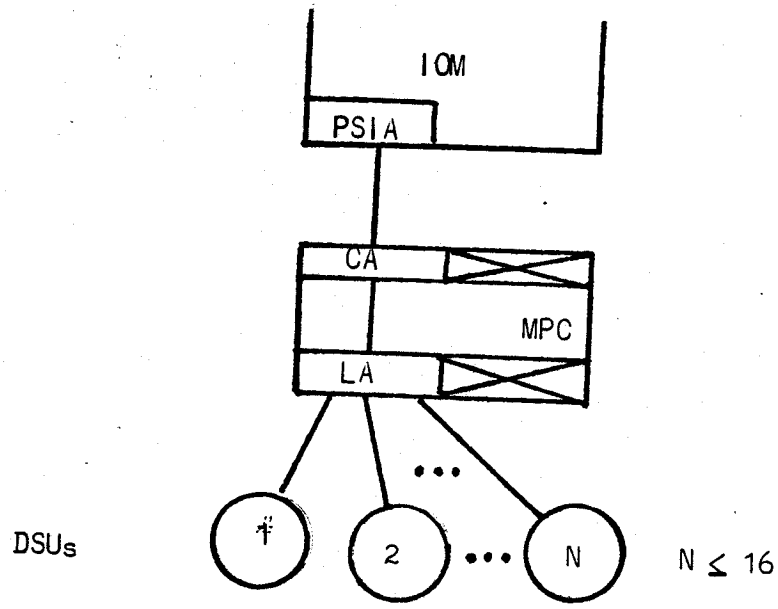

Figure 1 - Minimum possible DSS-181 or DSS-190 configuration, known as 1X1XN.

An MPC may contain up to two LA's. Each LA may support one PSIA. (A special LA called a switched channel can support two PSIAs, but executes channel programs over only one PSIA at a time. This option is intended for disk subsystems serving two separate functional entities, such as a CPU and a mass storage processor.)

An MPC may have up to two CAs. MPC microprogramming may allow either LA to communicate with either CA, resulting in a limited crossbar configuration.

Each CA has 16 ports to which disk units may be attached. Thus, a single controller can have up to 32 disk units (16 units on each CA), each on-line with only one path through are CA.

Alternatively, with the Dual Port Option in each disk unit, each unit may be connected to one port on each of two different CAs, resulting in a full crossbar configuration, with a maximum of 16 disk units are MPC.

Figure 2 Two channel full crossbar 1x2*N disk subsystem.

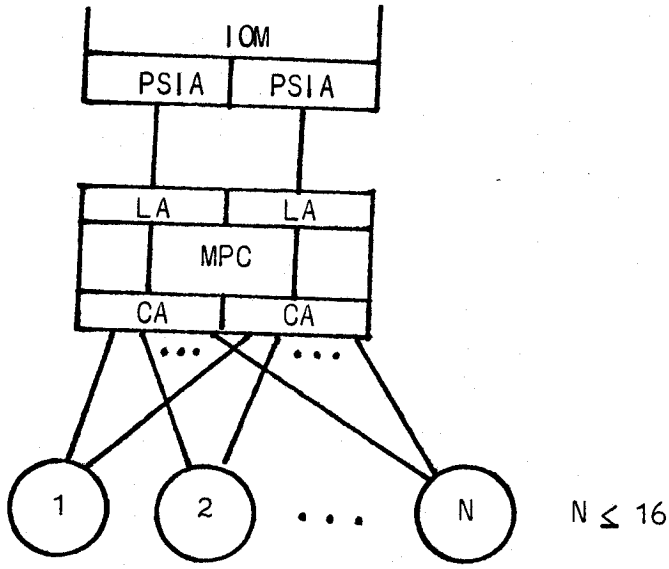

A second MPC may be used if additional throughput capacity or accessibility is necessary. It is clear that there are a very large number of possible configurations, presumably to satisfy a wide range of size, speed, reliability, and cost requirements. The major information about various configurations can be expressed in a compact shorthand notation. As an example, figure 3 shows a single controller with two PSIA/LA channels, internal limited crossbar microprogramming in the MPC, and 32 disk units without Dual Port Option.

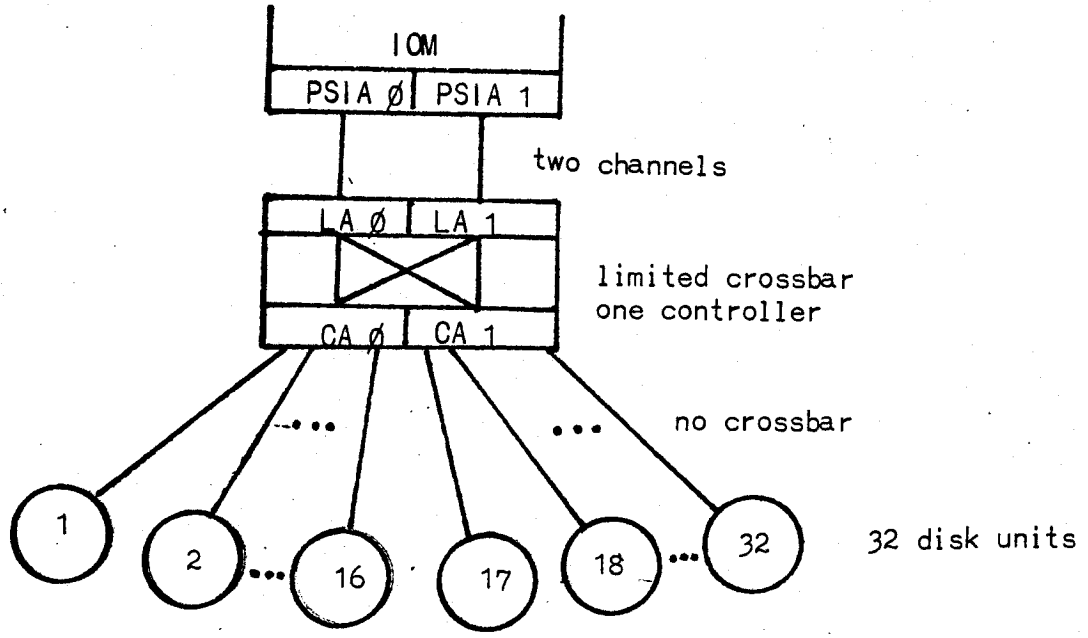

Figure 3

In our notation, this configuration is expressed by:

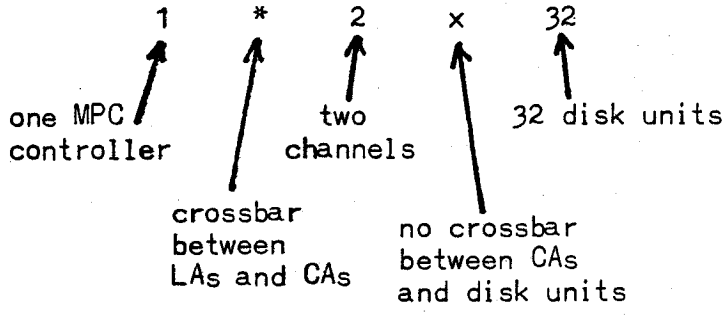

The notation may also include factoring. The same subsystem, but without the internal crossbaring in the MPC (i.e., LA1 can access only CA1), is expressed as

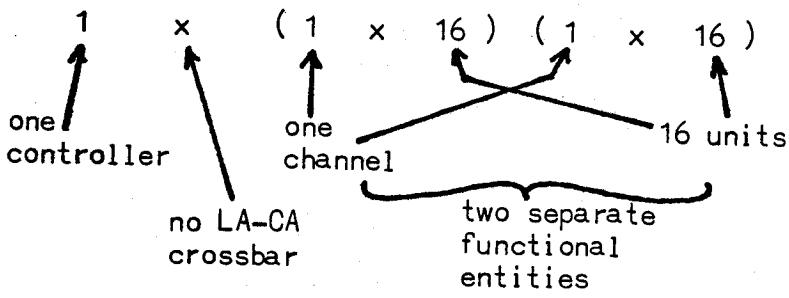

3. Disk Units

The DSU-181 disk packs are the industry standard IBM-2314 type packs. They have 202 cylinders and 20 recording surfaces per cylinder. Depending on formatting, each pack can hold from 4000 to 5252 Multics 1024-word pages.

DSU-181 spindle speed is 2400-rpm, or 25 ms rotation time. Read/write head seek time is somewhat improved over IBM 2314, with a mean seek time of 34 milliseconds. This results in a mean page fetch/write time (not including queueing delays) of about 66 ms, as compared to about 95 ms for IBM 2314 units.

The DSU-190 packs are a bit larger than the 2314 packs, and are compatible with IBM-3330 packs in size. There are 409 cylinders, and 19 recording surfaces/cylinder. (The 20th surface is pre-recorded with patterns used in rotational position sensing.) Depending on formatting, a single DSU-190 pack can hold between 14700 and 19200 Multics 1024-word pages.

Rotational speed on the DSU-190 is expected to be 3600 rpm, or 16.7 ms per revolution. The target for mean seek time (taken as the sum of the times of all possible seek operations divided by the number of these operations) is 30 ms. This results in a mean page fetch/read time of about 46 ms (not including queueing delays).

In addition, the DSU-190 units also incorporate a feature known as RPS, or Rotational Position Sensing, by which the MPC controller may investigate each unit to find out its current angular position (relative to the index mark that defines the start of each track). As described later, the MPC can make use of this information, when several data transfer requests are pending, to transfer from the unit whose head is closest to its data.

4. MPC Controller

The MPC controller, a new standard Honeywell product, will be used to control all IOM-based peripheral devices for Multics (i.e., disks and tapes). The MPC is a computer in its own right, with microprogrammed instructions from its own memories providing DCW interpretation, status formatting, interrupt reporting, data transfer, error handling, statistics keeping, and many other features.

The MPC has two memory units: a high speed Control Store, part of which is writeable, and a lower speed main memory. It has several registors whose purposes are defined primarily by its microprogramming for a particular application.

The MPC for the DSS-181 and DSS-190 subsystems will include several features:
Dual Simultaneous Transfers - This means that data transfers to/from two different units may proceed simultaneously. The speed capabilities of this kind of operation are sufficient to accomodate dual simultaneous transfers and from the very high speed DSU-190 units. (This is of course limited to configuration with 2 PSIA/LA channels.)

Command Stacking - Command stacking means that the MPC will, in most interesting cases, accept and hold a second instruction for a given unit even though the first instruction is still executing.

Device Reservation - Many commands cause a disk unit to become "reserved" to a particular logical channel (data path) for subsequent operations. This is of use in multi-channel configurations to prevent interference between channels attempting to access a particular disk. With the Dual Port Option, this reservation extends to the disk unit itself so that the unit can be reserved to one of its two DLI lines.

Command Execution - The MPC accepts "macro" commands like seek, read, write, and translates them into sequences of lower-level instructions to be sent to the disk units.

Interrupt Reporting - Disk units notify the MPC of significant events (such as seek complete, end-of-cylinder) via a DEN (Device Event Notification) signal. The MPC selectively reports events to the IOM/SYSTEM according to the IOM conventions. Special Interrupts (IOM level 7 interrupts) are reported, when a channel program is not active, for events such as seek completion, power down, etc.

Status Formatting - Associated with all interrupts are status codes to be referred to the system for software processing. The MPC translates signals from devices into status codes along established conventions. Status is also stored for Special Interrupts, so that system software can determine the exact cause of the Special Interrupt.

Error Retry - When the MPC detects an error in the execution of most DCWs, before reporting the error, it will attempt to retry the DCW three times by moving the IOM list pointer back to that DCW. On seek retries, a restore of the read/write head assembly to cylinder 0 is first executed.

Detailed Status Reporting/Statistics Gathering - The 181/190 MPC command repertoire includes several commands to read and reset status registers. The MPC Control Register includes a bit pair for each device that reflects seek complete/incomplete and device busy status. The Control Register for each device includes several resettable counters for numbers of seek, read, write operators, errors, etc. The Status Register contains detailed status information about the controller and disk units.

Automatic Track Switching - When a data transfer encompasses more than one track, the MPC automatically chooses the next head for the next track in the cylinder without losing a disk revolution or interrupting data transfer.

Multiple Sector Sizes - Standard MPC microprogramming allows formatting any single track in one of three formats: 64 word sectors, 320 word sectors, and one sector per track. The choice is made by a two-bit code in the data for the format track DCW.

Track formats may be mixed within a pack. When automatic end-of-track operations attempt to continue a data transfer from a second track whose format is different, an error is reflected to the IOM.

Stranger Packs - Packs formatted in other than standard formats can still be read. An entire track may be read, including all auxiliary header, count, key, and data fields. It is left to system software to interpret the data from the track.

The DSS-190 controller in addition supports several throughout - enhancing and reliability features:

Rotational Position Sensing (RPS) - The MPC interrogates disk units with data transfer requests pending to determine angular position, and therefore, time until data to be read/written comes under the head.

Block Multiplexing (BM) - The controller allows each physical channel (PSI A/LA) several logical channels. A logical channel is an area in the MPC where several DCWs for a particular device may be stacked. The logical channel - device relationship is established when a channel program issues a seek instruction for the device. The system issues several connects, one receiving a different logical channel.

Each connect sets up a complete DCW list (seek, seek data, read/write, and data transfer DCW's) for one device. The controller continues the instructions it has stacked for the several devices. When several data

transfers are ready (i.e., they are only waiting for the disks to rotate the data under the heads), the MPC checks its RPS data and then executes the data transfer for the channel program whose data is closest to the heads.

Error Detection And Correction (EDAC)

In the interests of preserving the integrity of data stored on disk, a 56-bit cyclic error detection and correction code is added to each field of data stored. Generation of the code and error checking is performed by feedback shift registers in the CAs. The code used has a high probibility of detecting burst errors.

When a field is read and the EDAC code bits read do not match those computed, several retrys are first attempted. If successive trys produce the same error an attempt is made to locate and correct the error. If successive trys result in different errors, a maximum of four retrys are made, with an error correction attempted on the last retry.

End of Cylinder Processing - When a data transfer continues beyond the last sector in a cylinder, the MPC suspends data transfer at the end of the cylinder, advances the read/write heads to the next cylinder automatically, and continues the data transfer with the first sector on the second cylinder.

Alternate Track Processing - When tracks are formatted, a pair of track-indicator (TI) bits in the track heads indicate whether the track is good or bad, and if good, whether it is an alternate track for some bad track. When the MPC seeks to a track that is formatted as a bad track, it looks up the alternate track address in the header (if an alternate has been assigned) and automatically re-seeks to the alternate track.

5. Documents

The following are the presently available references for further information on the DSS-181 and DSS-190 subsystems.

MPC Controller	EPS1 43A177875 Microprogrammed Peripheral Controller
	EPS1 43A232230 DSC 181/DSC190 Controller
MPC Microprogramming	MPC-1 Microprogrammed Peripheral Controller Microprogram Reference Mannual
DSS-181, DSS-190	EPS1 43A239851 DSS181 and DSS190 removable Media Disk Storage Subsystems

PSIA	EPS1 43A177874 Peripheral Subsystem Interface
LA	EPS1 43A177879 MPC - PSI Link Adapter