CERN-Data Handling Division
DD/80/23
November 1980

CLUSTERS OF 16/32 BIT MINICOMPUTERS IN HIGH ENERGY PHYSICS EXPERIMENTS AT CERN

P.Bähler, A.Bogaerts, C.Eck, P.Ferran, D.Jacobs,
A.Lacourt, J.Ogilvie, H.Overas, J.O.Petersen

Presented at the Europhysics Conference on Computing in High Energy and Nuclear Physics, Bologna, Italy, 9-12 September 1980

DD-jt

CLUSTERS OF 16/32 BIT MINICOMPUTERS IN HIGH ENERGY PHYSICS EXPERIMENTS AT CERN

P.Bahler, A.Bogaerts, C.Eck, M.Ferran, D.Jacobs,
A.Lacourt, J.Ogilvie, H.Overas, J.O.Petersen

DD Division, CERN
CH 1211 Geneve 23

Abstract

A description is given of multi-processor on-line systems, based on NORD computers, for large physics experiments at CERN.

In complex experiments the data acquisition and control functions may be performed by different CPU's in a dual 16-bit NORD-100 computer system.

Some configurations include a 32-bit NORD-50 processor, allowing large analysis programs to be implemented and executed in the on-line system.

Several examples from currently running CERN physics experiments are given.

1. Introduction

In recent years on-line computer systems in large physics experiments have evolved to benefit from distributed processing: 16-bit minicomputers linked to front-end microprocessors are used for experiment control, monitoring and data acquisition; the addition of a 32-bit processor to the data acquisition computer provides the on-line system with data analysis power which was until now available only at larger computer centres.

The purpose of the present paper is to discuss some ways in which NORD computers can be used to implement such large data acquisition and processing systems.

Section 2 describes briefly the NORD-100 16-bit computer, the central component of all NORD computer configurations at CERN, with emphasis on the real-time aspects of the hardware and of the operating system SINTRAN III.

Section 3 deals with the CERN data acquisition package, which uses the real-time features of the NORD-100 to provide an efficient and modular data collection system adapted to a multi-processor environment.

A dual NORD-100 system, in which experiment control and data acquisition functions are performed by separate CPU's, is described in section 4. The example given is the European Hybrid Spectrometer at the CERN SPS. In this case the control computer monitors the complex detector system via front-end microprocessors. Detector information is collected and stored in a data base system, especially optimised for real-time use.

The subject of section 5 is on-line configurations including a NORD-50 32-bit processor. After a description of relevant hardware and software aspects of this computer, several examples will be given of its use in large experiments. Applications include complex on-line calibration, event filtering and complete event analysis in both on-line and off-line mode.

2. The NORD-100 computer system

The NORD-100 (or the older compatible model NORD-10S, also referred to as NORD-100 below) is a general purpose 16-bit computer, supporting large single- or multi-port memories with a maximum size of 32 Mbytes. It contains the paging hardware necessary for the implementation of a virtual memory system. The hardware basis for efficient interrupt handling and fast context switching is provided by a 16 level priority interrupt scheme, where each level has its own register file , including a program counter.

The operating system SINTRAN III is a multi-mode, multi-tasking, virtual memory system . The basic program concept is the segment, a contiguous virtual address area containing up to 64 pages, each 2 kbytes long. At program execution time, the segments, usually disk resident, are swapped between disk and memory on a page basis.

SINTRAN III supervises "standard" real-time programs, all running at the same hardware interrupt level and scheduled according to software priorities. In addition, SINTRAN III allows for the implementation of special, privileged tasks executing on higher interrupt levels, even above that of the operating system. In general, these so-called Direct Tasks have no access to the facilities of the operating system. The SINTRAN III support is limited to a few basic operations : triggering of a Direct Task from external interrupts and activation of "standard" real-time programs from Direct Tasks, and vice-versa. The data acquisition package described below is an example of a program system in which time critical components are implemented as Direct Tasks.

The term task is in the following used to mean either "standard" real-time program or Direct Task.

3.The Data Acquisition System

The CERN-developed data acquisition package (DAS) is a modular system, structured as a set of separate tasks which communicate via a common data area. There is one special task:the supervisor or DAS monitor.

DAS could be implemented to run under any multi-tasking operating system which provides a few, fairly common facilities:

- the possibility to activate and suspend the execution of a task from another task
- activation of a task by external interrupts
- communication using a common data area.

DAS runs on any NORD-100 configuration which supports the SINTRAN-III operating system. The use of Direct Tasks (sect.2) results in a fast and efficient data acquisition system, whose speed is essentially limited by the hardware.

Data is normally buffered in the NORD-100 memory. DAS supports data buffer sizes up to the NORD memory limits.

Multi-tasking makes DAS a very flexible system. Different functions such as CAMAC data acquisition , data recording, data reduction, and on-line (sample) analysis may be implemented as separate tasks; more tasks may be added as the need arises: up-grading of detector systems will usually require addition of new tasks rather than modification of the existing software which greatly enhances the stability of large systems.

Multi-processor configurations may easily be implemented: the function may be executed by another processor such as a microprocessor for CAMAC data acquisition or a NORD-50 for data reduction or on-line analysis. The corresponding task on the NORD-100 will then simply be a driver of the external processor.

The DAS monitor controls the flow of data through one or more memory buffers and supervises data access by tasks. Data, originating from data sources, is structured into records or events and put sequentially into a buffer by one or more tasks, called producers. Data in any buffer may be accessed by one or more tasks which are called consumers. Consumers may act as data sinks, e.g. magnetic tape recording and (sample) analysis, or they may put data into another buffer. Several consumers may access the same data in a buffer. Therefore, a consumer is not allowed to modify the buffer contents. If this is necessary, a task may copy the data and write the processed data into another buffer. Such a task acts as both a consumer and a producer.

Consumers do not have to read all data. Several modes are available:

- a task takes all data e.g. magnetic tape recording

- a task may take data on request only e.g. an event display

- a task takes a fixed fraction of all data

- a task takes a minimum fraction, or more, depending on the priority of the program

The latter modes are typically used for sample analysis.

Data sources and sinks are serviced by tasks, allowing a large range of applications. Currently available as standard tasks are general CAMAC input (which also may drive CAMAC data links); magnetic tape input (to run the system with simulated data or for event play-back); magnetic tape recording in a CERN standard format; NORD-50 data communication, as well as skeleton programs for event displays and sample analysis programs with histogramming.

The use of specialized hardware resulting in multi-processor configurations and links to other computers shifts a considerable amount of the traditional workload away from the minicomputer. This tends to reduce the work of DAS to supervising, co-ordinating and synchronizing the different tasks, regulating data-flow and controlling buffer access.

DAS systems that can be built may be complex. At present, each task may write data in up to five different buffers and may consume data from up to ten buffers. Each buffer can accomodate five producers and ten consumers. This is illustrated by fig. 1. However, these constraints are arbitrary and depend only on the size of certain tables. The total number of tasks and buffers is not limited. An example of a simple DAS system is given in fig. 2. It consists of a

CAMAC reading task (producer) and the consumer tasks: magnetic tape recording, an event display and sample analysis. An example of a more complex configuration with data preprocessing is given in fig. 3. It should be noted that a DAS system can easily support several of these buffer/task subsystems, thus implementing multi-user data acquisition systems.

4. A dual NORD-100 configuration

In large and complex experiments, a single NORD-100 computer system may be insufficient to perform all the required on-line functions; a possibility is then to assign equipment control and monitoring to a second NORD-100, linked to the data acquisition computer by the manufacturer-supplied NORDNET system.

An example of such a configuration is given by the European Hybrid Spectrometer (EHS) at the CERN SPS [1]. This complex experimental facility comprises a bubble chamber and a variety of electronic detectors used for measurement and identification of secondary particles. All detectors were developed by groups in external laboratories.

On the data acquisition computer is implemented a multi-user, multi-buffer DAS system of the type described above. The modularity of DAS allows independent, concurrent development of code for - and testing of - the various detectors.

The DAS computer uses several ESOP microprocessors for second level trigger calculations and data reduction. A full description is given in [2].

The complexity of the detectors, particularly of the bubble chamber, imposes a high load on the control and monitoring system. This is structured as a set of real-time tasks, all having access to a data base which holds control and monitoring information on all the detectors. The common data base permits task intercommunication and allows access by different tasks to the information for each piece of equipment. For example, set-points and measured values are available not only to the relevant control tasks, but also to the tasks which present or update status information.

The real-time requirement of fast data base access excludes the use of existing file-based data base systems. Instead, another structure has been developed, in which the standard file system is bypassed. The data base access routines are resident in memory and shared by all real-time tasks. The data base itself is stored as a set of up to 16 segments, each of 128 kbytes (see sect.2). Using segments takes advantage of the virtual memory features of the NORD-100, resulting in faster data access than can be obtained by using the standard file system. The paging mechanism led to a tree structure organisation of the data base which allows a data item to be retrieved with a maximum of three page (disk) accesses. A task carries out a data base transaction using two buffers resident in memory, one for specifying the access keys and the other for the data itself.

The control computer is relieved of repetitive tasks by front-end microprocessors linked to it via CAMAC. Some examples are: scanning of switch positions, high precision gas control for Cerenkov counters and ionisation measurement chambers, and HT supply control.

In discussing the support of such microprocessors, two phases of use must be distinguished: development and production. At EHS the first phase is catered for using the normal NORD text editor and file system for program source management, a cross-assembler and a system for downline loading of the microprocessor. At this stage the microprocessor operates with RAM and is connected to the NORD-100 by an RS-232C terminal line. A simple program in the microprocessor connects its terminal to this line in a transparent way, allowing the user to log on to the NORD-100 system in the usual manner. When downline loading is initiated, the microprocessor recognises the load block header and switches function to load its own memory. In practice this has only been used for the I8080, but the necessary software exists for the M6800 and the TMS9900.

For trouble-free use in production, the program of the microprocessor is in ROM with automatic initialisation on power up. Communication with the NORD-100 is then by exchange of messages. At EHS, two standards have been adopted. The first is used when the equipment builders require that the microprocessor be capable of stand-alone operation using a terminal instead of the NORD-100. In this case, the NORD-100 simply replaces the terminal via a CAMAC serial interface. It simulates the commands typed on the terminal and decodes the replies. The second standard is more efficient and uses 16-bit parallel transmission via a CAMAC I/O register. The data is sent in binary form with a simple message format of identifier, length, data and checksum.

5. NORD-100/NORD-50 configurations

The processing power of a NORD on-line system may be increased by adding a NORD-50 computer to the configuration. The NORD-50 is a 32-bit processor which is under complete control of a NORD-100 computer system. Actually the latter controls the operation of the NORD-50 CPU by the use of I/O instructions. The NORD-50 has no direct control over peripherals; I/O requests are passed to the NORD-100 which performs the data transfers (including DMA). The NORD-100 and NORD-50 have access to a multiport, multibank memory which holds programs for both processors and allows them to share data. Addressing of the multibank memory is arranged in hardware in such a way that the NORD-50 is able to access a 32-bit word in one memory cycle (interleaved memory); this is of importance, since the NORD-50 has no cache memory. The NORD-100/NORD-50 communication system - for control and data - is shown schematically in fig.4.

The performance of the NORD-50 CPU is compared to that of other computers in table 1. The first column gives the total execution time of a representative set of FORTRAN benchmark programs, while the second column shows the same information normalised to the performance of the IBM 370/168. In the table, 168/E stands for the IBM 370/168 emulator [3]. Also given are preliminary results for the NORD-500, the successor of the NORD-50.

Software available for the NORD-50 is limited to assembler, FORTRAN compiler, loader and monitor, all of which execute on the NORD-100. NORD-50 program images, prepared on the NORD-100, are loaded into shared memory and executed under supervision of the monitor which forms part of the SINTRAN III operating system. The NORD-50 is in practice operated in a single-program mode, which facilitates synchronisation and data exchange with programs running on the NORD-100. The synchronisation is readily implemented by algorithms using flags in a fixed common area of shared memory.

The NORD-50 structure has, in both hardware and support software, been kept simple, thus providing an analysis tool with a minimal system overhead.

5.1 Experiment NA4

The use of NORD-50 is illustrated by examples, the first being from the experiment NA4 at the CERN SPS, a large spectrometer used to study muon-induced reactions [4]. The NORD-100/NORD-50 on-line system is in this case linked to another front-end minicomputer which is responsible for data acquisition and thus relieves the NORD-100 of providing large multi-event buffers. The layout of physical memory is depicted in fig.5. The upper part is occupied by a data acquisition system as described in section 3 . Producer tasks read data either directly from the front-end computer via the link or - in replay mode - from magnetic tape. A consumer task extracts events from the DAS buffer and copies them into a fixed memory area accessible to the NORD-50. A large part of memory, about 300 kbytes, is occupied permanently by a NORD-50 program. This is essentially a copy of the full production off-line analysis program, which was developed on a CDC 7600 before being implemented on the NORD-50. The program reads data from a shared memory buffer (fig.5) and, when an event has been fully analysed, is suspended in a wait loop until new data are available. The coupling between the NORD-100 and the NORD-50 is loose, being limited to a one-way data transfer.

The system as described has been used for two main purposes. First for on-line tuning of the experiment; the information (tables, statistics, histograms etc.) based on completely reconstructed events available on-line, provides the physicist with a powerful tool to set experiment parameters. The second application is for off-line production analysis. About 25% of the total amount of data analysed so far has been processed by the NORD-50 with an overall throughput ratio of about 1/20 compared to the CDC 7600.

5.2 Experiment WA1

The next example is from the neutrino experiment WA1 at the CERN SPS, a large spectrometer set-up consisting of modules of magnetised iron interleaved with drift chambers and scintillator hodoscopes, the latter being equipped with a several thousand channel CAMAC ADC system [5]. Events registered during the very short beam spill (about 100 microsec.) are buffered at the CAMAC hardware level. Between beam bursts, cosmic ray events are collected for calibration purposes.

The event flow through the NORD-100/NORD-50 system is illustrated in fig.6. Cosmic ray events are read from CAMAC by the NORD-100 and copied to the NORD-50 which performs tracking and records information for hit scintillators. This information is used to build calibration tables which allow to follow the hodoscope performance as data collection progresses. Beam events are in a similar way made available to the NORD-50 which pre-processes the event by using the calibration tables computed on-line from the cosmic-ray events. The pre-processed event is sent back to the NORD-100 which, in parallel, has analysed the same event, and then written onto magnetic tape.

By performing the calibration and event pre-processing on-line , the number of magnetic tapes written is greatly reduced (by a factor of 3 to 15 depending on beam conditions) implying a considerable reduction of the off-line analysis work.

5.3 Experiment NA10

In the previous examples the NORD-50 processor was of secondary importance in the on-line system. In the case of experiment NA4, the NORD-50 program - used in on-line mode - can be considered as a powerful monitoring task executing in parallel with consumer tasks on the NORD 100. However, data taking with some quality control could still proceed without the NORD-50. Similarly, in the second example, the on-line system with no NORD-50 could still perform the basic data collection, although the overall performance would be degraded.

The last example is a NORD-100/NORD-50 configuration in which the NORD-50 has been given a more central role in the on-line system.

The experiment is NA10, a spectrometer for studying multimuon events produced in a pion beam at the CERN SPS [6]. Conceptually, the experiment is simple, consisting of a large system of proportional chambers and scintillator hodoscopes, placed in field-free space in front of and behind a bending magnet.

The basic idea in the on-line system is to let the NORD-100 perform I/O tasks (data acquisition, link handling, graphic displays etc.) and allocate the NORD-50 to CPU bound analysis programs. The layout of physical memory is shown in fig.7. The NORD-100 has access to 512 kbytes of memory which contains a data acquisition system and large buffers for CAMAC input, histograms and magnetic tape writing. The upper half of memory is occupied by a NORD-50 program, which is largely a copy of the FORTRAN off-line production program with time critical parts written in assembler language. The main task of the NORD-50 program is event reconstruction and monitoring. During the beam burst, data flows directly into memory. When the first events are available, the NORD-50 commences the event analysis - filtering and track segment reconstruction - and writes the results into the tape buffer area from where they are transferred to magnetic tape by the NORD-100. All events are thus channelled through the NORD-50, the processor speed of which ,therefore, is a critical parameter in the on-line system. If the NORD-50 is not able to process all the incoming data, it limits the data flow through the system.

6. Future developments

The previous sections illustrate the current trend in on-line processing for physics experiments at CERN towards multi-processor systems, in particular configurations including a 32-bit processor. For NORD computer based on-line systems this development will get further momentum with the introduction of the NORD-500 32-bit processor (see table 1). Soon, two early models of this computer will be integrated in CERN experiments, replacing existing NORD-50 processors and operating in a similar single task mode as described above. Later, the NORD-500 will be equipped with a memory management system, necessary for the support of an efficient multi-tasking system. This will allow a further shift of compute bound tasks - including some of those of the operating system - to the 32-bit computer, leaving the NORD-100 to perform I/O oriented tasks, comprising the data acquisition, control and monitoring functions of the on-line system.

REFERENCES

[1] Internal document CERN SPSC/75-15/P-42

[2] D.A.Jacobs, Uses of ESOP, a fast microprogrammed processor, in High Energy Physics experiments at CERN, presented at the Europhysics Conference on COMPUTING IN HIGH ENERGY AND NUCLEAR PHYSICS, Bologna, September 9-12 1980.

[3] P.F.Kunz, Nuclear Instruments and Methods 135 (1976) 435

[4] Internal document CERN SPSC/74-79/P-19

[5] M.Holder et al., Nuclear Instruments and Methods 148 (1978) 235.

[6] Internal document CERN/SPSC/77-110/P-91

[7] Private communication, Per Scharff-Hansen ,DD Division, CERN

FIGURE CAPTIONS

fig. 1 DAS task/buffer structure
a) DAS consumer/producer task
b) DAS buffer with task input/output

fig. 2 Single-buffer DAS configuration with consumer tasks in different sampling modes (see section 3)

fig. 3 Double-buffer DAS configuration with data preprocessing

fig. 4 Schematic diagram of NORD-100/NORD-50 communication system for control and data

fig. 5 Layout of (physical) memory in the NORD-100/NORD-50 on-line system for experiment NA4

fig. 6 Schematic flow of beam events and cosmic ray events through the NORD-100/NORD-50 on-line system of experiment WA1

fig. 7 Layout of (physical) memory in the NORD-100/NORD-50 on-line system for experiment NA10

TABLE CAPTIONS

Table 1 Performance of the NORD-50 and NORD-500 32-bit processors compared to that of other computers. 168/E stands for the IBM 370/168 emulator [3].

*) based on execution of large FORTRAN analysis programs [7]

**) preliminary result

COMPUTER	CPU TIME (secs)	RELATIVE TO IBM 370/168
CDC 7600	37.5	0.3
CDC 6500	579.9	5.1
IBM 370/168	113.4	1
168/E		~2.5 *)
NORD-50	871	8
NORD-500	323	3 **)
DEC PDP10	525	4.6
DEC VAX	444	3.9

table 1

a)

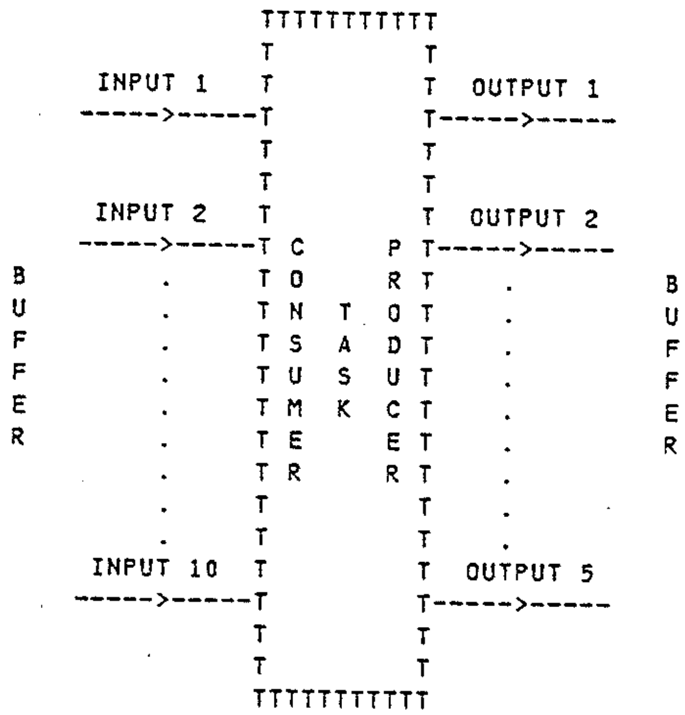
INPUT 1
INPUT 2
INPUT 10
OUTPUT 1
OUTPUT 2
OUTPUT 5
BUFFER
CONSUMER
TASK
PRODUCER
BUFFER

b)

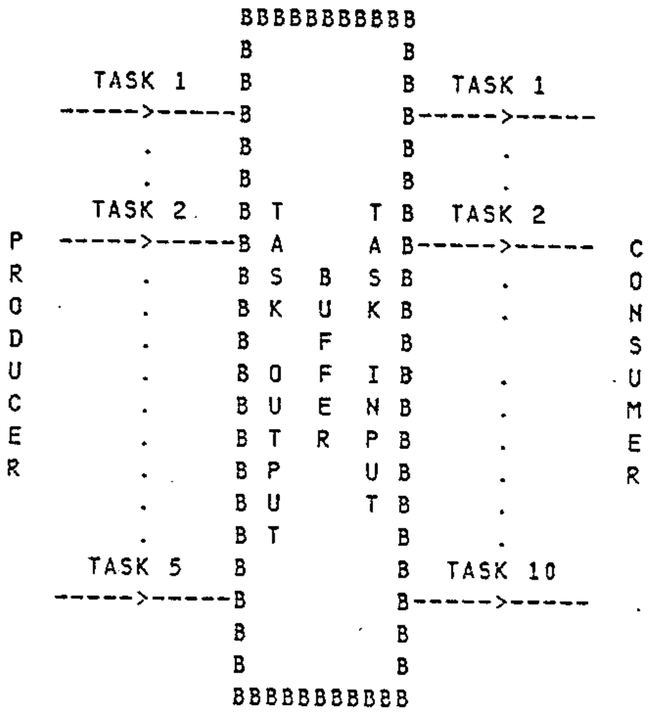
TASK 1
TASK 2
TASK 5
TASK 1
TASK 2
TASK 10
PRODUCER
TASK OUTPUT
BUFFER
TASK INPUT
CONSUMER

fig. 1

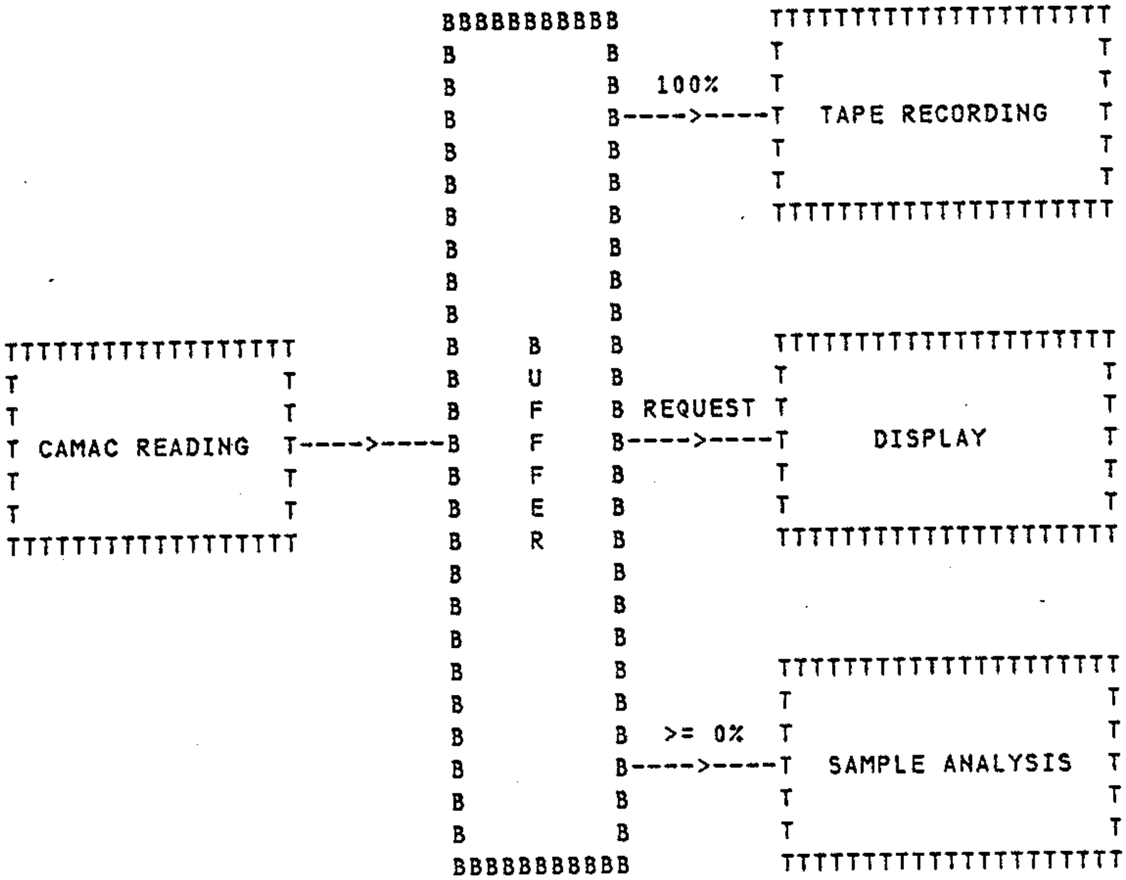
CAMAC READING
BUFFER
100%
TAPE RECORDING
REQUEST
DISPLAY
>= 0%
SAMPLE ANALYSIS

fig. 2

```
                BBBBBBB                        BBBBBBB
                B     B      TTTTTTTTTTTTT     B     B      TTTTTTTTTTTTT
                B     B      T           T     B     B      T           T
                B     B 100% T   DATA    T     B     B 100% T   TAPE    T
                B     B--->--T REDUCTION T-->--B     B--->--T RECORDING T
TTTTTTTTTTT     B     B      T           T     B     B      T           T
T         T     B  B  B      TTTTTTTTTTTTT     B  B  B      TTTTTTTTTTTTT
T  CAMAC  T     B  U  B                        B  U  B
T READING T-->--B  F  B                        B  F  B
T         T     B  F  B      TTTTTTTTTTTTT     B  F  B      TTTTTTTTTTTT
TTTTTTTTTTT     B  E  B      T           T     B  E  B      T          T
                B  R  B  REQ T   EVENT   T     B  R  B >=0% T  SAMPLE  T
                B     B--->--T  DISPLAY  T     B     B--->--T ANALYSIS T
                B     B      T           T     B     B      T          T
                B     B      TTTTTTTTTTTTT     B     B      TTTTTTTTTTTT
                B     B                        B     B
                BBBBBBB                        BBBBBBB
```

fig. 3

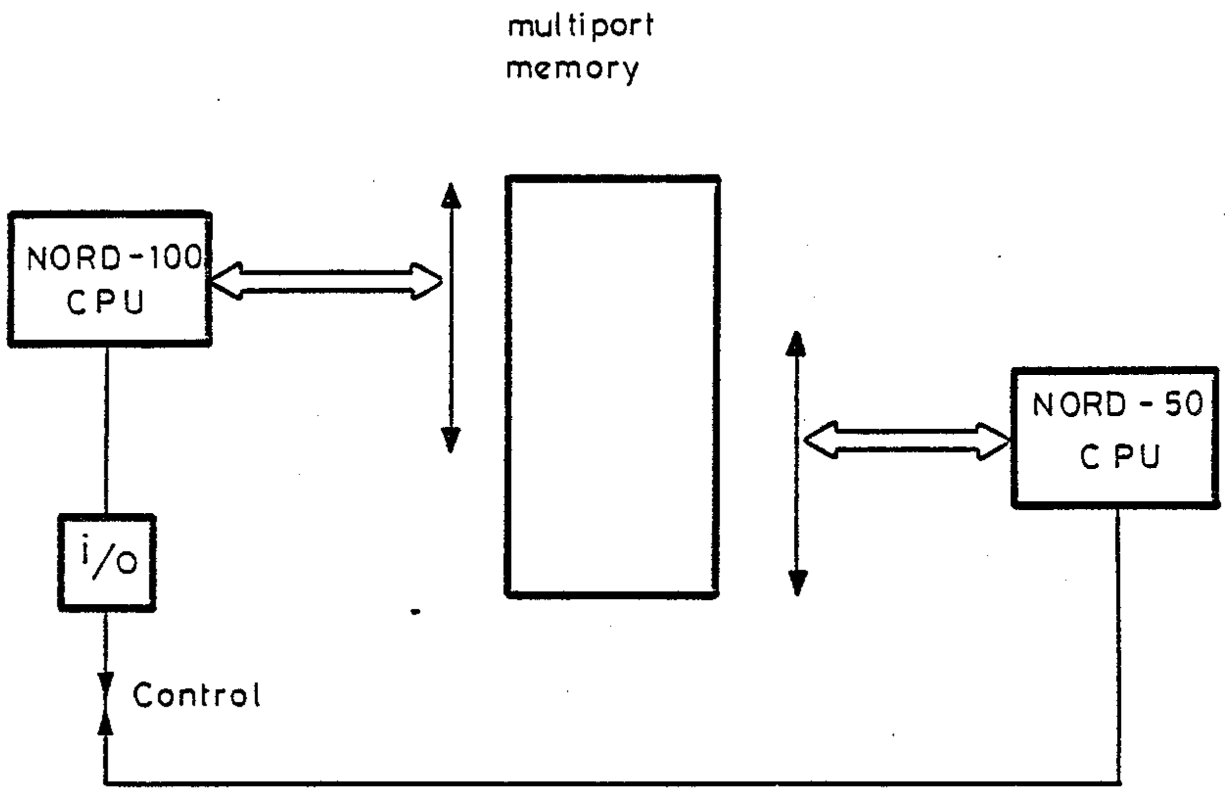

fig. 4

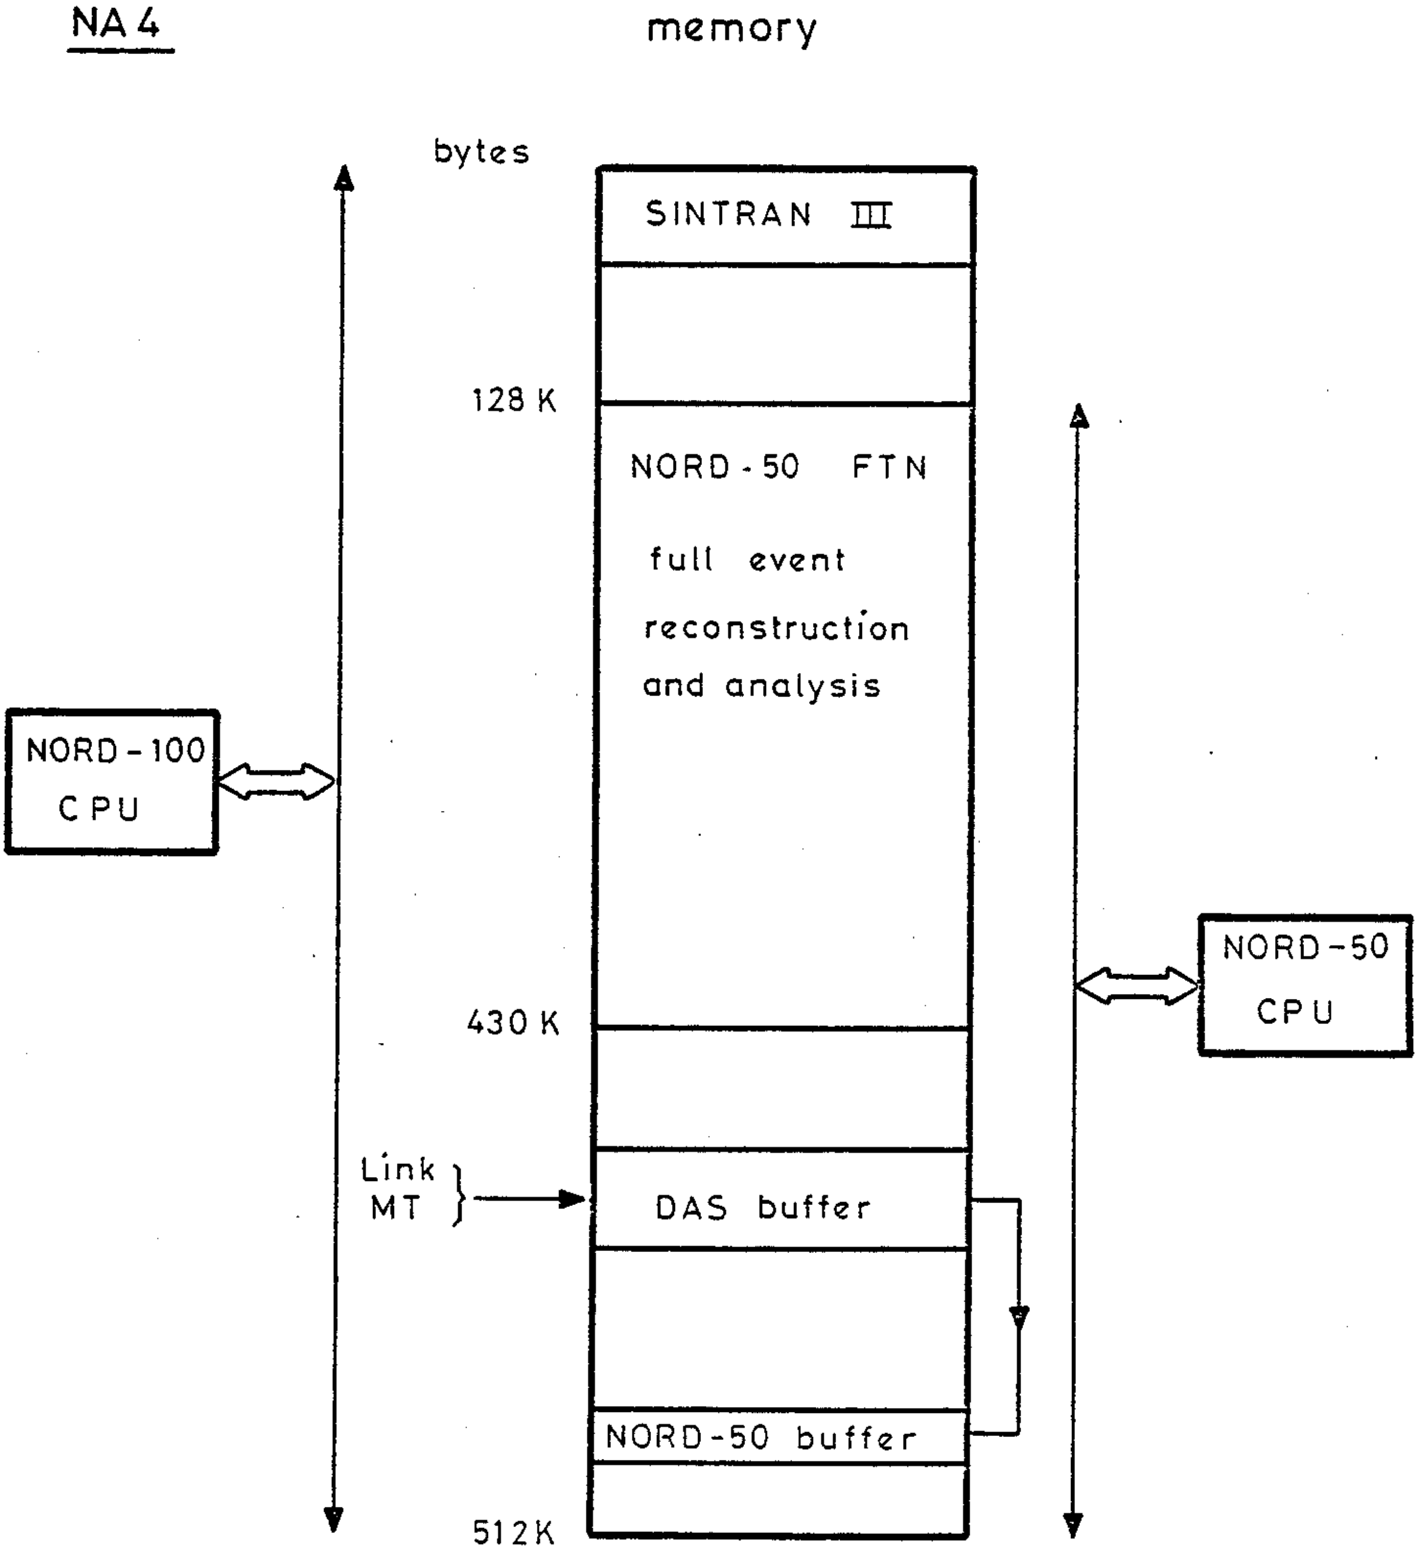
NA 4
memory
bytes
SINTRAN III
128 K
NORD - 50 FTN
full event
reconstruction
and analysis
NORD - 100
CPU
NORD - 50
CPU
430 K
Link
MT
DAS buffer
NORD - 50 buffer
512 K

fig. 5

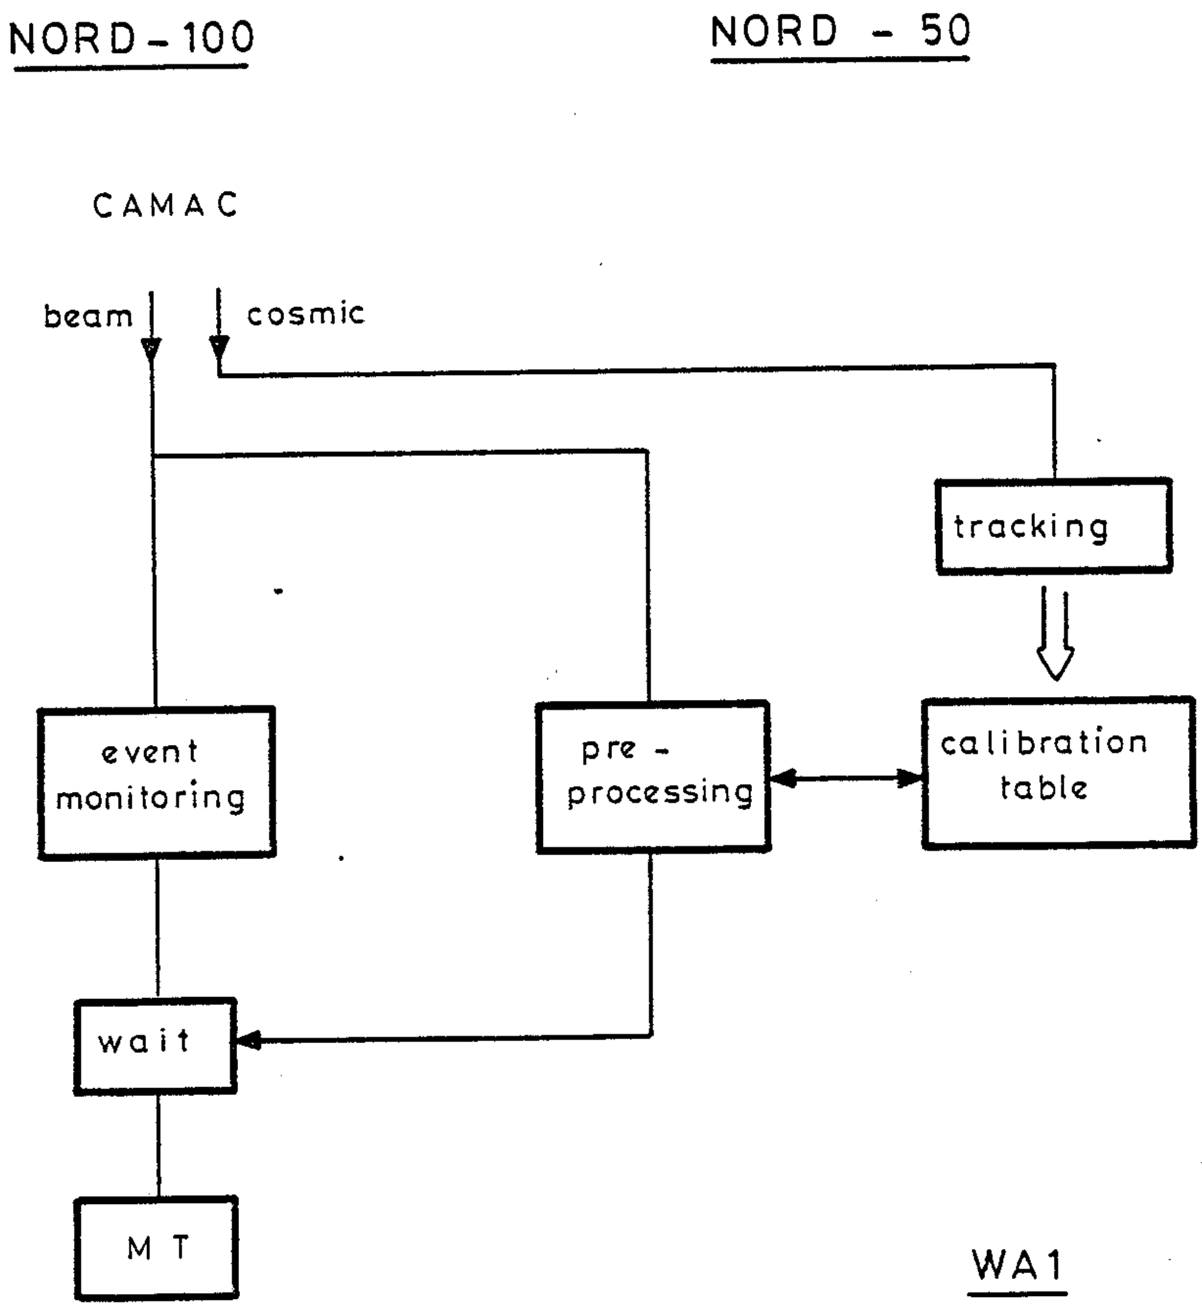
NORD-100
NORD - 50
CAMAC
beam
cosmic
tracking
calibration table
event monitoring
pre - processing
wait
M T
WA1

fig. 6

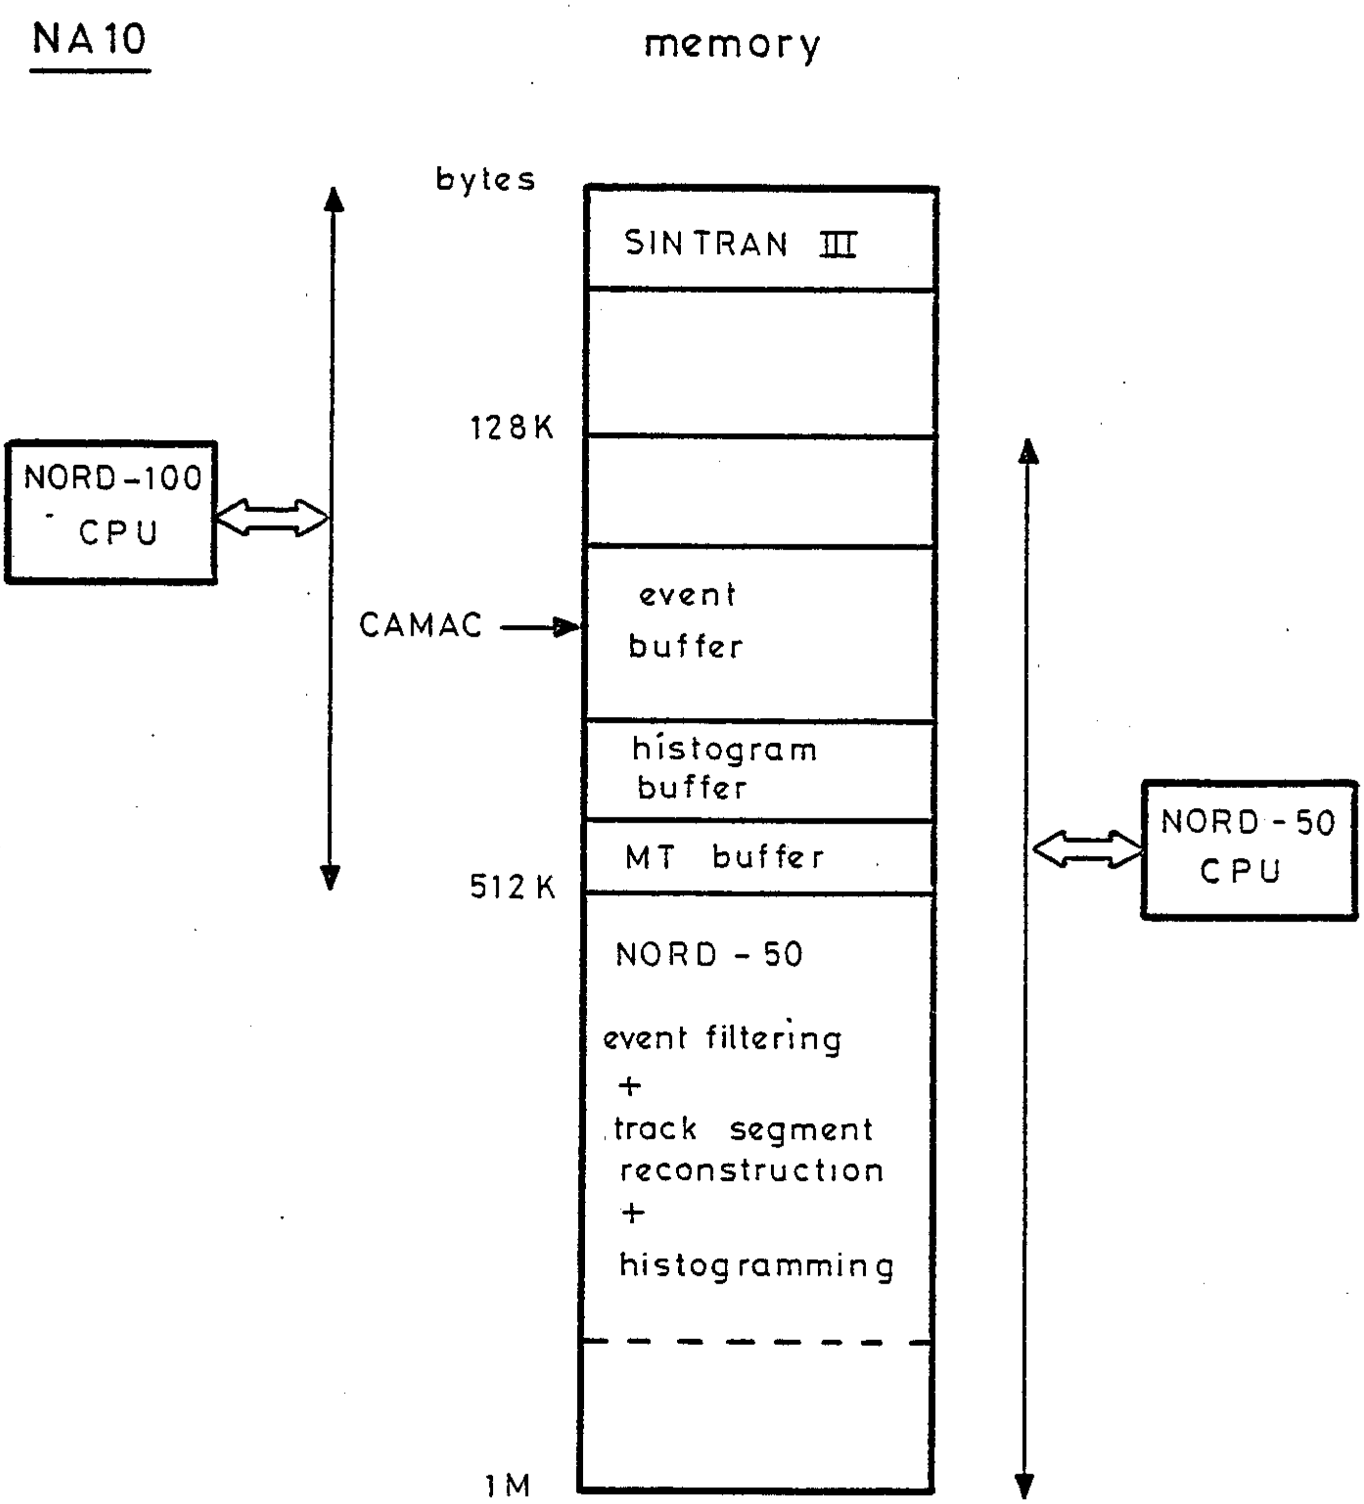
NA10
memory
bytes
SINTRAN III
128K
NORD-100 CPU
event buffer
CAMAC
histogram buffer
MT buffer
512K
NORD-50 CPU
NORD - 50
event filtering
+
track segment reconstruction
+
histogramming
1M

fig. 7