Fermi National Accelerator Laboratory

FERMILAB-Conf-80/40-EXP
2100.000

DUMAND DATA ACQUISITION WITH TRIGGERING*

A. E. Brenner and D. Theriot
Fermi National Accelerator Laboratory, Batavia, Illinois 60510

and

R. H. March
Physics Department, University of Wisconsin, Madison, Wisconsin 53706

March 31, 1980

---

*Submitted to the Workshop on DUMAND Signal Processing, University of Hawaii, Honolulu, February 11-16, 1980.

Operated by Universities Research Association Inc. under contract with the United States Department of Energy

DUMAND Data Acquisition With Triggering

A. E. Brenner and D. Theriot, Fermi National Accelerator Laboratory, and R. H. March, University of Wisconsin

Introduction

We present here a data acquisition scheme for the "standard" DUMAND array[1] which includes a simple triggering scheme as a fundamental part of the system. Although there are a number of not yet fully understood parameters including the expected signal rate and also the various backgrounds, we proceed here with the assumption that thresholds can be set in such a manner as to give rise to a triggered signal which is not so dominated by randoms that it gives a substantial decrease in the data acquisition rate over that which would be required by a non-triggered system. The assumption is also made here that the triggering logic is relatively simple and does not need major computational capabilities for a trigger logic decision. With these assumptions, it is possible to generate the trigger at the array and restrict the data transfer to shore. However, with a not unreasonable delay of 200 micro-seconds, it is even possible to transmit the information for the trigger to shore and perform all that logic on the shore. The critical point here is to send the minimum amount of information necessary to construct the trigger such that one need not send all the possible information in all detectors of the array continuously to shore.

In defining this scheme, certain natural assumptions have been made. It is assumed that each detector unit has its own photo tube, base, A to D converter, and also a small amount of digital storage to serve as a short-term buffer for the results of the A to D conversions. It also has at least a minimum amount of intelligence so as to debug itself, to set thresholds and to otherwise make configuration changes within the detector under control of a command string sent from shore. It is also assumed that there must be at least a single simplex signal cable which sends signals from the detector through the necessary logic and repeater electronics to the shore. This cable from the detector modules to the first logic/repeater unit is nominally capable of better than 1 MHz bandwidth transfer, including all of the driving electronics, and is very likely to be a fiber optics cable. There is one cable from each detector, which terminates at the sea bottom at a proper connection point. In addition, there is a second cable which daisy chains from one detector to the next and which brings in power to each detector module. Over this cable, it is possible to superimpose relatively slow control signals on top of the power level such that one can communicate in a "broadcast mode" to all the detectors at a modest rate, <10KHz.

Further, a desirable goal is to have as little electronics as possible between each individual detector and the shore processor. Each detector already must have enough internal sophistication that any additional modest enhancements in the detector module will not substantively change either the cost or the reliability of these modules. Furthermore, single detectors as they fail can be "turned off", and a reasonable number can be so lost without detrimentally affecting the array. Of course, any failures on shore can be handled by standard techniques and also do not represent any serious burden. On the other hand, any electronics situated between the detectors and the shore will, by its very nature, be responsible for at least a string of detectors and in some cases for much more. Consequently, a failure of any intermediate electronics would have a major deleterious effect on the performance of the array. Of course, it is not possible to completely eliminate such electronics and, therefore, that risk will exist. However, if it is kept as simple as possible, it will be more readily possible to duplicate or otherwise make such units redundant.

In the description of a system that embodies these ideas which follows, a number of almost arbitrary choices have been made for the sake of definiteness. A more careful study should be made as one proceeds down this route as to where the optimum breakpoints occur for the most reliable and most cost-effective system organization.

Topology

Figure 1 is a schematic in the topological connections of the proposed array. From each detector along any given string, a single fiber optics cable is strung from the detector along the string to the bottom of the array and, without a junction, it is strung from there along the base of its plane to the leg of the Y to which that plane communicates. Thus, at that Y junction, the plane or end station, there are 18 modules times 21 strings of fiber cables all entering a junction box containing a "plane processor". The function of this processor is to collate and otherwise organize all of the signals coming from all the detectors in that plane, reshape them, serialize them and retransmit them. The retransmitted signal from the plane processors is the logical "or" of all the signals feeding it. Its output is a single optics fiber which is strung from the plane processor along the Y leg to the center point of the array where the "central processor" is positioned.

The central processor itself drives three cables to shore for the final collection point of all data. Of course, in the case of the central processor and probably also of the plane processors, the system will be so organized with enough redundancy that a failure of any active component will not jeopardize the experiment. Not shown in the diagram is the power cable which at the central processor is made to branch into three separate legs and which at the plane processor is broken one more time into a number of segments. Each of these segments, daisy chains a number of detectors. The precise organization is not critical for this discussion, but it might be one power cable coaxial daisy chain along one string through all 18 detectors on that string.

Signals

It is assumed that the real event rate might be of the order of 30 per second. It is also expected, although it must be tested more precisely with experiment, that the largest single random background rate is due to the beta decay of K40. At the one photo-electron level, the counting rate in each detector might be 200K per second. It is assumed that some threshold can be set greater than one photo-electron so that the single counting rate in each detector is of the order of 1000 counts per second. This will define the lowest level signal that can ever be recorded by any detector in the system. A second threshold is set, probably at 6 or more photo-electrons, which define the triggering level for muons. Finally, a high discrimination level is set at possibly twice that number to define a trigger for hadronic showers. In all cases, the length of the discriminated pulse is measured and a cutoff is imposed such that all signals must be short enough so that bioluminescent activity with relatively long-time constants ($\sim$msecs) from biological specimens can be distinguished and, unless otherwise programmed in, may be excluded from contributing to the trigger.

All detectors are calibrated in situ such that their pulses reach the central processor isochronously. To accomplish this during startup, each detector unit is given a delay time by which it must delay its information so that its signal reaches the central processor at the same time that any other signal in the array reaches that processor for a signal that leaves each detector at the same absolute time. This is analagous to the timing cables that are usually used in high energy physics experiments. Here, however, it is done digitally with a digital countdown in each detector unit which effectively delays the signal.

There are a number of signals described in greater detail in the next sections which are necessary. These are:

Trigger signal
Stop signal
Data signal
Clock signal
Control signal

The trigger, data and clock signals all use the high-speed optical fiber link, whereas, the stop and control signals use the much lower bandwidth, daisy chained cable also used to distribute the power. Each of these will be discussed in turn here.

Trigger Signal

Each detector processor strobes onto its output signal fiber optics cable, a properly delayed, shaped pulse of about one micro-second duration for each PMT signal which exceeds the trigger threshold. A much larger signal (ten times) is placed on that cable whenever a hadronic level event occurs in a single PMT. These signals leave each detector delayed by the detector processor by an amount determined during setup which assures that all signals from all detectors arrive at the central processor at equal times. The leading edge of these signals defines the time of the occurence. The one micro-second pulse is chosen so as to make the driving electronics in the detector as low power and as reliable as possible. Pulse pairs that come within that one micro-second interval from the single detector are necessarily lost.

The signals proceed down their cable from the detectors to the plane processor responsible for each plane. At the plane processor, all 378 signals which enter it are reshaped to 100nsec width pulses and ORed together on to a single output line. Since the transit time between adjacent detectors (40 meters) exceeds 100nsec, those pulses caused by a single traversing particle or shower will give rise to detector pulses which are separated in time by that amount. Thus, the output from the plane processor with 100 nsec pulses will have a string of pulses which if they came from adjacent detectors would appear as separate pulses on the output cable from the plane processor.

Since there are only 60 plane processors, it is feasible both from a reliability and power point of view to have much faster and redundant electronics in each of these. Thus, the requirement of an ECL power-hungry circuit with enough redundancy such that no single failure causes them to become inoperative to reshape the longer pulses into 100 nsec pulses is acceptable. Thus, the philosophy is not to let a plane processor fail, although single detectors may fail and be turned off as necessary with only minor consequences for the total array.

The outputs of all plane processors now reshaped and regenerated progress down each leg of the Y in bundles that grow to 20 along each leg to the central processor. At the central processor, all cables coming in are ORed together and a relatively modest circuit counts all pulses in a sliding window of fixed length. The time width of this window is typically ~ 5 micro-seconds, that being the transit time of a particle through the whole array, but is program-adjustable to suit. The circuit searches within that window for a detector hit count which exceeds a predetermined, but program-adjustable, minimum number. A trigger pulse is generated whenever the hit count within the window exceeds the threshold number. The central processor, which also generates the central clock pulses, strobes the time which is defined by the first pulse in the triggering window. It should be noted that if there is a second level trigger for hadronic showers which might require two adjacent detectors to have fired, a second windowing action requiring two of the larger detector pulses within a 200 nanosecond window (that being the time to include transient between any two adjacent detectors) can also be constructed in parallel.

The determination of this trigger signal generates the stop signal. The time of the event as determined by the central processor is stored and defines the event.

If the trigger rate is too high using the entire array to make the sum of hit counts, a crude topological trigger[2] can be made by ORing all three adjacent planes of detectors to make the sum for the detector hit count. This would require that all detectors that make the sum be in the same local area, i.e., a shower; it would also reduce the enormous accidental rate arising from 30,000 modules.

Stop Signal

Once this trigger has been generated, the central processor now distributes by broadcast over the power communications link at about a 10K baud rate the time of the event. This is the stop signal. The resolution for the time of the event need be good to only about one microsecond. Consequently, about 10 or 12 bits of time information need be sent back to the detector processors. Because of the relatively low speed of this communications link, it takes about a millisecond for the signal to propagate to all detectors.

Data Signal

Once the detector processors receive the stop signal time message, each one searches its internal memory for any stored data which occurred in the 10 microsecond window around that time. Any detector which finds such data then transmits, for each hit in that time window, all the necessary information associated with that hit down its high speed output fiber optics cable. The information transmitted includes the time of the event, the detector address, the pulse height of the signal and finally any additional flags that might otherwise clarify the nature of the event. In all, a total of about 40 bits is expected. To distinguish this data stream from a trigger data stream, a prefix of five sequential pulses is used to initiate the transfer of each such word. This pulse train also gives the strobing frequency for the data that follows.

Each individual detector performs this operation asynchronously. The data transmitted is all data that exceeded the minimum threshold (not the trigger threshold) within the selected time window that would have been digitized during the previous several milliseconds. When this data is received at the plane processors, it is buffered and multiplexed on the single output line which connects each plane processor with the central processor. The central processor in turn multiplexs and reshapes these signals from the array of plane processors and transmits the signal stream to shore. Although the details are not critical here, very likely three separate cables (each redundant) carry the information from each leg separately to shore. In fact, the total bandwidth is estimated to sufficiently small that all information can pass down a single fiber optics transmission line with all signals being multiplexed by the central processor onto that line if any major cost savings could be realized by such an approach. It appears sensible at this point, however, to keep them separate and eliminate one multiplexing circuit in the central processor.

If it does become necessary to reduce the data being transmitted to store, a topological trigger processor[3] can be used to reduce the accidental triggers once the data is stored in the central processor.

It should also be noted here that it would be possible to send the whole stream of trigger signals from the central processor to the shore, perform the trigger logic on shore and retransmit back via a full duplex line to the central processor the results of all trigger logic. This additional transmission to shore requires of the order of 200 microseconds which hardly changes the delay time for the relatively slow transmission of the event time which occurs between the central processor and each of the detectors. In this approach, some relatively simple logic is removed from the central processor and performed on the shore for what appears to be a 20% additional dead time transmission time in returning the trigger signal back to each detector. This variant is only of minor interest, but is indicative of a wide range of options possible as to where the various logical operations are performed. A more detailed engineering analysis should be made to make decisions on this and a number of other related matters.

Clock Signal

Maintaining the timing in an absolute manner in all modules is a very difficult task. There are two levels of requirements. First is the requirement to identify those signals in each detector which have appeared in a trigger window. These times need not be known to better than one microsecond. The second requirement is to know the time of a hit in absolute timing system with a knowledge of the time to better than 10 nanoseconds. If one can record the time to that level of precision, then one is able to improve the fits of muon tracks as they traverse the full detector. As that timing information degrades and moves closer to ∿ 50 nanoseconds, then the timing information adds very little much more than the fundamental spatial information which indicates that a track appeared at a certain detector in the trigger time window. To achieve 10 nanosecond precision timing in an absolute sense without going to a full duplex, high-speed communications link between each detector module and the central processor may be very difficult to obtain.

To accomplish the timing, every element in every processor in the system will have a quartz controlled clock with a frequency of about 5 MHz and a phase detectable resolution of about 10 nanoseconds with the stability that comes from the constant temperature environment of about one part in 10^9. In each detector a scaler keeps track of local time for up to about a total period of 10 seconds; a total of 30 bits for the scaler. On every overflow of that scaler, i.e., about once ever 10 seconds, the detector module sends a pulse stream of its standard pulses (equivalent to its trigger pulses) a string of 10 such pulses to indicate that this is clock information. The plane processor recognizes this string as a timing synchronizing signal, and by using the leading edge of the first of the 10 pulses will be able to ascertain the time sent by that detector to within a fraction of a microsecond, but certainly greater than 100 nanoseconds.

If, as is planned, all detectors in a plane are isochronous at the plane processor, then the plane processor would expect to receive all timing pulses from all its detectors at the same instant in time once they had been so adjusted. Since all clocks will not be in perfect synchronism, but rather will slowly move out of phase one with respect to the other, it is the plane processor's responsibility to recognize which detectors have moved out of phase relative to its brethren. Whenever that happens, the plane processor sends back to the erring detector over the control lines a control signal to increment or decrement that detector's clock by some number of pulses. Since the resolution of the system is only good to something greater than 100 nanoseconds, these increments or decrements will typically be in units of 100's to 1000's of nanoseconds. All such changes sent to each detector is also sent to the central processor and eventually to shore to keep track of the adjustments.

A similar procedure occurs for all the plane processors on each leg; thus, whenever each plane processor's clock overflows, presumably in step with the clocks in all its detectors, it sends a similar clock pulse stream to the central processor which performs a similar adjustment on all the plane processors in the system. Of course, every time the central processor adjusts a plane processor during the next cycle, many of the detectors controlled by the adjusted plane processor will probably require additional adjustment. The system, however, should not be unstable. Since the land processor will have all information as to adjustments made to each processor in the system, it will be able to keep track of the absolute time over long periods for each individual detector. It will also be able to calculate with great precision the actual average frequency of each clock at each detector module and in every plane processor. This system is certainly good enough for trigger window timing, but is inadequate for the 10 nanosecond timing to refine the spatial fitting procedures.

It should be noted that the communications for adjustments between the central processor and the plane processors comes over a return high-speed, fiber optics cable between the central processor and the plane processors, whereas, the adjustments between the plane processors and its detectors is over the slower communications control system superimposed upon the power distribution cable.

The whole clock system, of course, must be initiated and restarted occasionally. This is a complex procedure which also must take into account the mechanism of making all detectors isochronous. These procedures which may be done periodically or possibly very infrequently will require a setup mode which makes use of program procedures and the use of the control lines discussed in the next section.

The most natural crystal oscillators run at about 5 MHz. With phase lock, closed loop procedures, it is possible to clock down to the 10 nanosecond level. Typically, under constant temperature conditions, the stability of such oscillators can be expected to be good to one part in 10^{10} . Therefore, if one can get the clock synchronized at the right level and reset them every few seconds, it may be possible to achieve the 10 nanosecond local timing. On the other hand, some more work is required to devise a scheme that can really maintain synchronization down to the 10 nanosecond level. At the 1 microsecond level, there are no complications.

Control Signals

The control signals which are received by each of the detector processors are superimposed on the power cable system. The signal rate is approximately 10 KHz and it is used to communicate control, initialization, and test information to each of the detector modules. These control signals are generated by the plane processors, sometimes on their own initiative and at other times, because of signals transferred to them, from the central processor. Much of the activity is initiated from the land processor which communicates on a high-speed fiber optic cable to the central processor. The central processor communicates control signals to each of the plane processors over a fiber optics cable. Thus, the connections between the central processor and each of the plane processors is in both directions a fast fiber optics cable. The transmission rate between the land processor and the central processor exceeds 100MHz and between the central processor and each of the plane processors is better than 10MHz.

The land processor uses these control lines for a number of important functions. The initialization of the whole system is under the control of the land processor, and it must load every processor in the system with appropriate programs. Some of the programs will be locally fixed by ROM, but as much of the program as possible in each processor will be RAM-driven, and the program to be stored in those memories must be fixed at initialization time. Thus, at initializaion time the nature of the triggers or the levels of discriminator settings, or the logic for triggers is all set by the land processor. When necessary, the land processor can load or otherwise cause to be executed diagnostic procedures in any selected detector or processor. If a diagnostic session is unable to repair a failed module, it can signal that detector module to turn itself off. Also, if any special calibration procedures are required for any detector modules, either on an _ad_ _hoc_ basis or periodically, the land processor will control the initialization of such procedures. Finally, the synchronization of the clocks is also a function of the control signals, although the actual control once initialized is vested in the central processor.

Controls signals make use of standard communications protocols which are self-clocking, and each separate type of message has a well-defined header word. For clock synchronization and other short-term controls, such as turn on and turn off, only a few 10's of bits must be sent to address the module and convey the function. For loading a processor, whereby several thousand bytes possibly might be loaded, clearly the data stream to be transmitted is much larger. Certain broadcast commands may be defined such that a plane at a time or possibly even the whole array can receive the same information simultaneously. Thus, even though it may take some minutes to load a single module, by loading all modules in parallel, it is possible to keep the time for such procedures to a reasonable level.

One Important Option

The average distance between a detector and its plane processor down a string and along a plane line is approximately 1 kilometer. At ∿ $.10 per foot, a reasonable estimate for the cost of fiber cables by the middle of the decade the cost projects to about $100 for each cable connection. Thus, with about 30,000 detectors that comes to a total cost ∿ $3M--a reasonable, but not happy cost. Furthermore, the cost of each transmitter/receiver pair might approach the same level. If one chose to use full duplex between the detectors and the plane processors in the future, which would simplify some of the protocols, the cost of the transmission system would increase by possibly 50% of its cost. Although this cost is not prohibitive, it is not trivial. On the other hand, such a system would simplify some of the other communications protocols and both improve and simplify the clock synchronization problems. Furthermore, it would allow the possibility of working at higher bandwidth. However, that might require the use of ECL circuitry, which would add considerably to the cost and power consumption. At the same time, it is very likely that the reliability will probably go down. Thus, there are some important engineering considerations to be carefully studied to consider all costs and tradeoffs before these decisions can be made.

Conclusions

The proposal made here starts with the premise that the detectors are active elements of reasonable cost and must be sufficiently intelligent to self-calibrate, self-threshold, to discriminate and to locally store signals and to transmit both trigger and data signals. Even with the most drastic simplifications, a detector which must contain the phototube, its circuitry and discriminators cannot be made arbitrarily inexpensive or simple. Thus, the philosophical point of view taken here is that one should rather add to the already non-trivial detector module as much circuitry as possible to minimize the need for sophisticated units elsewhere in the array in the ocean. Furthermore, the more complex logic that one can move out of other elements and put into the detectors, the less likely one has failure complications each one cannot get around. Thus, although one must be sensitive to the costs of the detector elements, it is clear that small increments of processing power or memory to utilize the greater flexibility and independence installed in the module will not substantively change its cost. On the other hand, if other elements, such as the plane processor or the central processor, could be made simpler since they cannot be turned off without turning off the whole plane or the whole array, one gets a potentially more reliable system. Of course, for these processors one could afford to make them highly redundant. The philosophy here is that the single or a small collection of detectors may fail or be turned off with no ill consequences whereas any high level processor which must be turned off starts becoming a burden which might jeopardize the experiment.

In the abstract case where there are no active elements any place except in the detector modules or in the land processor, one has the ultimately most reliable system. Since this is realistically impossible to achieve, one must make appropriate compromises. In the description given in this paper, the functions each processor performs can be certainly moved from one to another and a more detailed engineering analysis will be required to make the necessary cost and reliability tradeoffs.

REFERENCES

1. Arthur Roberts and George Wilkins, The 1978 Dumand "Standard" Array, Proceedings of the 1978 Dumand Summer Workshop, Vol.3, p.9 (1979).

2. Carl Akerlof, "Why DUMAND Array Triggers Cannot be Simple"; these Proceedings.

3. Carl Akerlof, "A Topological Trigger Processor for the DUMAND Optical Array"; these Proceedings.

FIG. 1. Topology of the "Standard" DUMAND Array

Shown here is a schematic depiction of the topology of the proposed triggered data acquisition array. The detector modules are each separately connected to a plane processor with a single optical fiber cable and also a daisey-chained (not shown) power and control cable. Each of the 60 plane processors are connected to the central processor with one (or possibly two) fiber optics cables. The central processor is connected to shore with three fiber optics cables.

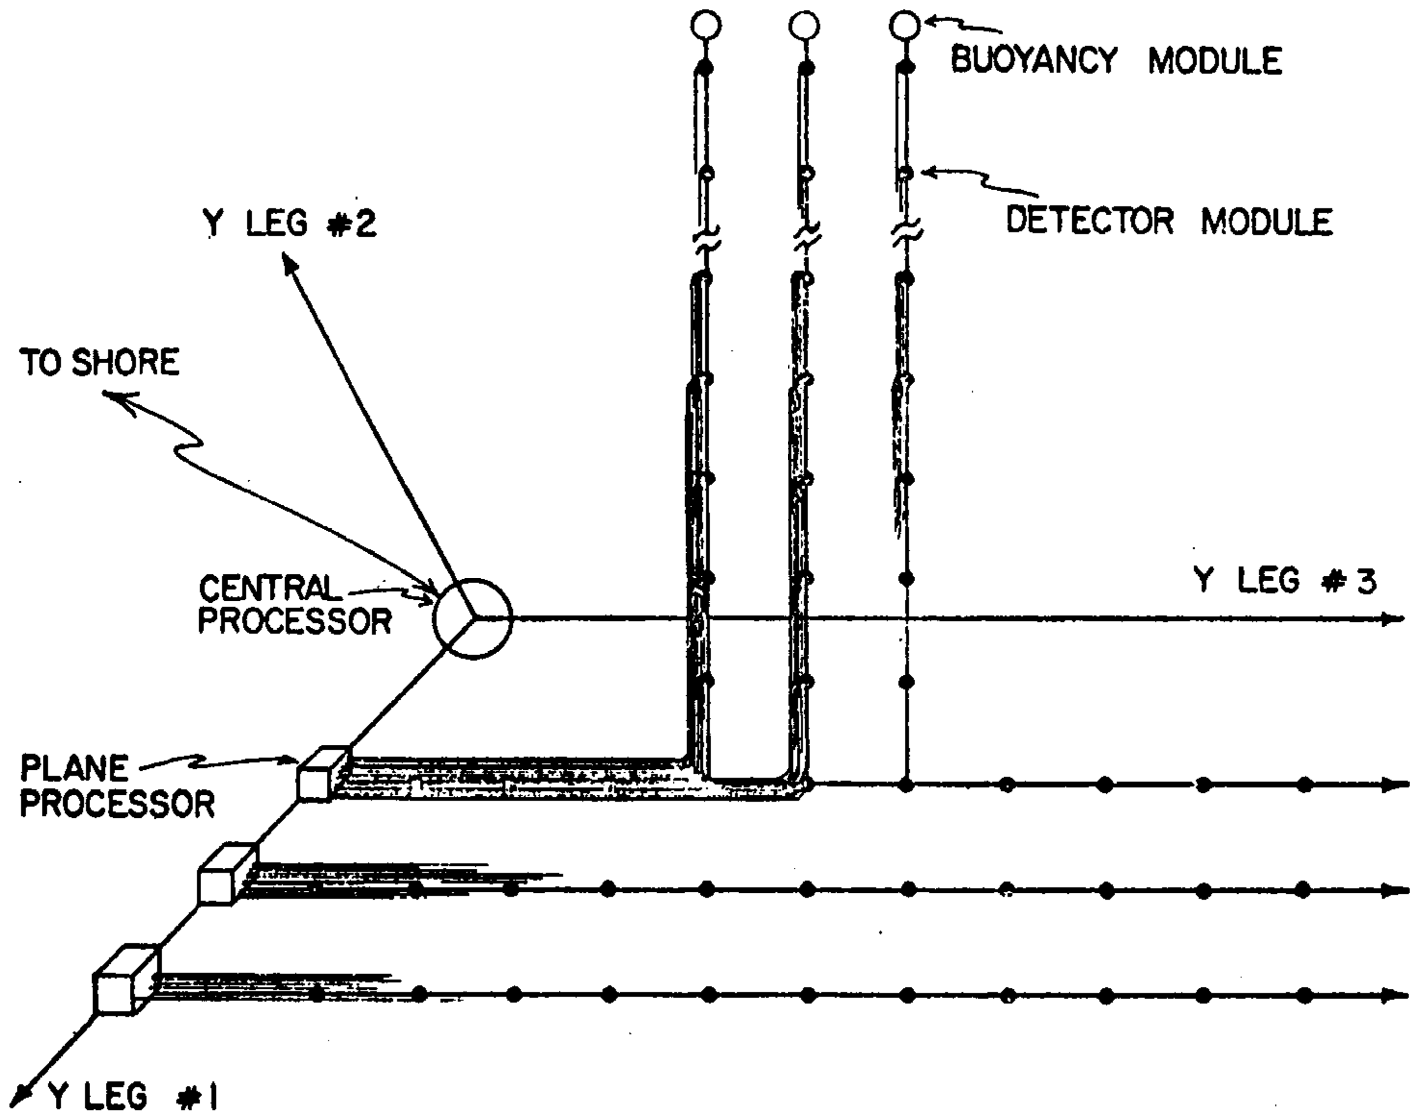
BUOYANCY MODULE
DETECTOR MODULE
Y LEG #2
TO SHORE
CENTRAL PROCESSOR
Y LEG # 3
PLANE PROCESSOR
Y LEG #1