

THE CEBAF FAST SHUTDOWN SYSTEM\*

J. Perry and E. Woodworth  
 Continuous Electron Beam Accelerator Facility  
 12000 Jefferson Avenue, Newport News, VA 23606

Abstract

Because of the high power in the CEBAF beam, equipment must be protected in the event of beam loss. The policy that has been adopted is to require a positive permissive signal from each of several inputs in order to operate the gun that starts the beam. If the permissive is removed, the gun shuts off within 20  $\mu$ s.

The inputs that are now monitored include (1) radiation monitors that detect beam loss directly, (2) vacuum monitors (which also observe the status of various in-line valves), and (3) general input from the rf system, which combines detection of klystron failure, arcs, and rf window high temperature. The system is expandable, so other fault detectors can be added if experience shows their necessity.

I. Introduction

The CEBAF beam carries great destructive power, and the accelerator must be protected from it. The CEBAF fast shutdown system (FSD) must provide critical services to the accelerator control system:

- Primarily, the system must shut a stray beam off before it can puncture the vacuum wall ( $\approx 50 \mu$ s).
- Secondly, the system must provide:
  - a watchdog timer to ascertain that the control system is in constant communication with every part of the accelerator
  - readback registers for system integrity checks at system startup
  - first fault trace in case of a shutdown
  - current fault traces after a shutdown trace.

The system is explicitly designed for highly reliable shutdown, easy expandability, and fast fault tracing.

II. Architecture

The high power (800 KW) and small diameter (200  $\mu$ m) of the CEBAF beam imply that a stray beam will burn through the vacuum wall in an estimated 50  $\mu$ s – 100  $\mu$ s. Since there are already up to 21  $\mu$ s of beam stored in the accelerator, the time limit for shutdown is actually about 30  $\mu$ s. We chose to design to a 24  $\mu$ s limit, apportioning 10  $\mu$ s to alarm detection functions, 10  $\mu$ s to logic functions, and 4  $\mu$ s to signal propagation time.

Preliminary estimates of required inputs to the FSD system were of the order of 400–600 alarm sources, among which were

- beam loss monitors
- vacuum loss monitors
- fast vacuum valves
- arc detectors of various kinds.

Since the needs were not well defined when it was necessary to start the design of the fast shutdown system, a substantial overcapacity is available to be sure of avoiding an overloaded system at the end of construction.

To ensure both adequate capacity and adequate speed, the FSD incorporates a tree structured system (Figure 1), which allows both a hierarchical organization and liberal expandability. Since a tree structure allows exponential increases in input capacity with more levels, we were able to trade off input capacity for shutdown speed and gain a logical organization. Figure 1 shows the structure for initial accelerator testing. As more segments of the accelerator come on line, their FSD branches will be attached at the indicated points, and their inputs unmasked.

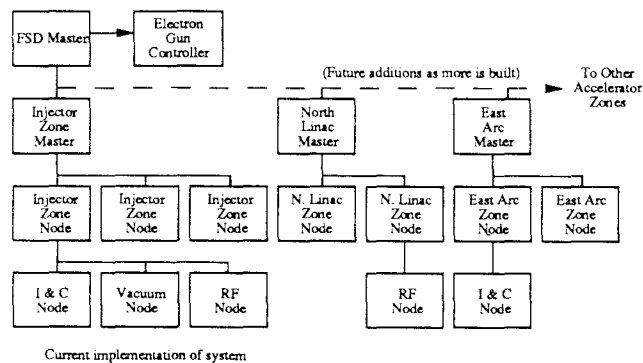


Figure 1. Fast Shutdown Tree Structure.

The requirement for high reliability demanded a permission signal that would fail, if it must, in a safe manner. We took a two-pronged approach to the fail-safe problem:

- First, we require the presence of a permission signal to run, rather than the absence of an inhibit signal.
- Second, we use a 5 MHz trapezoidal waveform as the permission signal to eliminate unsafe “stuck at ...” failures.

Figure 2 shows the permission signal and its characteristics. A normal fault is defined as a low state; it will be detected by the permission decoder within three cycles (600 ns). Only FSD system failures will cause the permission signal to remain high, and the permission decoder will detect this and remove permission within about 1.8  $\mu$ s.

\* This work was supported by the U.S. Department of Energy under contract DE-AC05-84ER40150.




Figure 2. FSD Permission Signal - One Level of Tree.

Shutdowns must be checked out, and the FSD system provides for two types of fault traces:

- a first fault trace, if there has been a shutdown;
- a current fault trace, if there has been no shutdown since the last reset.

Figure 3 illustrates a fault trace. Upon shutdown, each FSD node module in the shutdown chain holds the state of its inputs at the instant of shutdown, allowing determination of first loss of permission within about 2  $\mu$ s. The control system may at any time query any module about its permission status, and the module will provide the state of every permission input either at the time of shutdown, or at the last time of reset if there was no intervening shutdown.

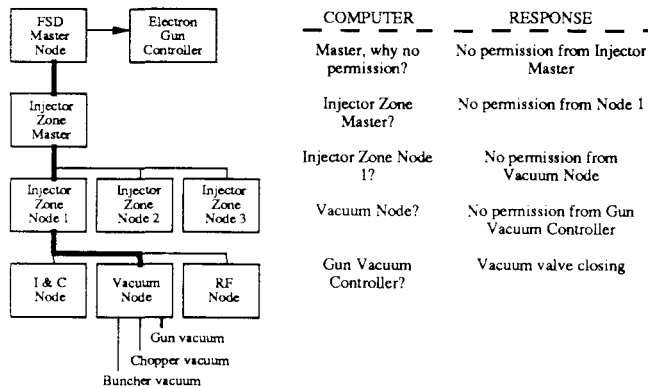


Figure 3. Fast Shutdown Fault Trace.

### III. Node Module

The FSD node module (Figure 4) is a single-width CAMAC module which incorporates the functions listed above. There are 16 input fault sources in four groups:

- seven fiber optic permission inputs for tree expansion
- seven optically isolated permission inputs for fault detection sources
- watchdog timer for communications integrity monitoring
- bussed P1 line connection for intra-crate signals.

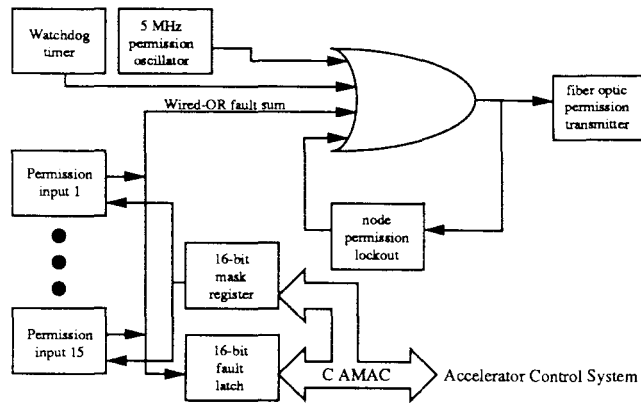


Figure 4. FSD Node Module Block Diagram.

We were able to design a highly reliable permission signal decoder (Figure 5) that would remove permission to the next level of the tree within 600 ns of a fault, and that permitted us to define a convenient number of seven fault inputs and seven structural (tree) inputs. This kept the cost of a module reasonable and allowed a small number of levels to keep the logic propagation time down. The present organization of five levels allows  $7^5 + 7^4 + 7^3 + 7^2 + 7 = 19,607$  fault inputs at a worst case normal delay of  $5 \times .6 = 3 \mu$ s, plus propagation time.




Figure 5. FSD Permission Signal Decoder - One of 15 Per Node.

The watchdog timer in the FSD node module supplies a system integrity need for the control system. The highly distributed control system at CEBAF incorporates a number of relatively independent control functions. It is possible that certain functions could continue to operate in the event of a failure in communications with the control system, resulting in segments of the system working at cross purposes. The watchdog timer offers a switch-programmable check on control scan interval from 13 ms to 426 s in 16 binary steps (.013, .026, ..., 215, 430). If the control system takes more than the programmed time limit to send a CAMAC command to the module, the watchdog timer times out and removes permission.

The last of the 16 permission inputs is the P1 line, which is a dual purpose input/output line for fault sensors that are either alone in a CAMAC crate with the FSD node module, or for which the order of faulting is unimportant. The node module drives the P1 line with a slow transition time 5 MHz permission signal, and a fault detector must short the P1 line to ground to remove permission from the system. If a source on another input removes permission from the node module, the node module also removes the signal from the P1 line, so that other users may monitor local fault conditions.

There is a fault latch and a mask register for each permission source. The fault latches hold the state of each permission source at the instant a fault is detected, and the mask registers keep unused inputs from interfering with the system. The fault latches and the mask registers correspond bit for bit with each of the 16 permission sources.

The fault latches have two functions:

- They latch if a fault is received by a previously clear module.
- They latch all current faults if a reset command is received after a fault.

The fault latches may be read at any time via CAMAC commands, and writing to their address resets the latches and updates them to the current input fault status (the specific datum written has no significance or effect). New faults cannot be latched until a previous fault has been reset; therefore resetting latches any new faults. If there are no existing faults, the node is armed to detect future faults.

Masks may be set at any time from the control system. A masked input will neither remove permission from the system, nor set its fault latch. We have taken the approach that only rigorously enforced administrative procedures are effective in preventing undisciplined maintenance violations of system integrity, and therefore only password protection in the control system software in combination with enlightened operating procedures will be used to prevent unauthorized masking of critical inputs.

#### IV. Reliability

Since a vacuum accident in the superconducting linacs would cause severe damage, reliability must be very high; therefore the design emphasized reliability both from the point of view of part by part integrity and redundancy of critical paths. The critical paths were carefully analyzed and optimized; however, noncritical paths were designed according to ordinary industrial design standards to reduce cost of the system.

Within a node module, the only critical components, those that can cause a failure to remove permission, are R1, D1, and Q1 in Figure 5. These have been made redundant, and have been specified at a sufficiently high reliability level that we expect no critical failures during the life of the accelerator. Furthermore, we are incorporating into the control system software regular automated integrity checks to identify any failures that might occur in spite of our efforts.

Finally, certain extremely sensitive inputs are being considered for full system-level redundancy to avoid the possibility of poisoning the expensive superconducting cavities in case of a beam loss in the vicinity of an undetected failed beam loss monitor or vacuum fault.