Fermi National Accelerator Laboratory

The Fermilab DØ Master Clock System

C. Rotolo, M. Fachin, S. Chappa, M. Rauch, C. Needles and A. Dyer

Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510

November 1991

MASTER

DEC 1 6 1991

Disclaimer

THE FERMILAB D0 MASTER CLOCK SYSTEM[1]

C. Rotolo, M. Fachin, S. Chappa, M. Rauch, C. Needles, A. Dyer

Fermi National Accelerator Laboratory
P.O. Box 500
Batavia, IL 60510

Abstract

The Clock System provides bunch crossing related timing signals to various detector subsystems. Accelerator synchronization and monitoring as well as timing signal generation and distribution are discussed. The system is built using three module types implemented in Eurostandard hardware with a VME communications interface. The first two types of modules are used to facilitate synchronization with the accelerator and to generate 23 timing signals that are programmable with one RF bucket (18.8 ns) resolution and 1 ns accuracy. Fifty-four of the third module type are used to distribute the timing signals and two synchronous 53 MHz and 106 MHz clocks to various detector subsystems.

SYSTEM OVERVIEW

The D0 Clock System is an outgrowth of the clock system at CDF [1]. Although the D0 Clock has retained some of the basic circuit concepts of the CDF Clock, its architecture is significantly different and allows much more flexibility in its operation and in the distribution of timing signals. A block diagram of the system is shown in Figure (1). The system is housed in six 9U by 280 mm Eurostandard crates plus three NIM bins located in seven racks distributed throughout the D0 counting houses and platform. The system is built using three module types called the Phase Coherent Clock (PCC), Sequencer (SEQ), and Selector Fanout Module (SFM) which respectively perform the functions of synchronization with the accelerator, timing signal generation, and timing signal distribution. The 53 MHz Tevatron RF, which is, of course, coherent to the rotating bunches in the accelerator, oscillates 1113 times in the time it takes a bunch to make one revolution around the ring ($\approx$ 21 μs). A processed version of the Tevatron RF and a once-a-revolution synchronization pulse from the accelerator called TEV SYNC are input to the Phase Coherent Clock module. The module produces a 53 MHz clock called PCLK and a 106 MHz clock called MCLK both of which are coherent to TEV SYNC, and hence with the bunches in the machine.

In addition to PCLK and MCLK, a once-a-revolution SYNC pulse is output from the Phase Coherent Clock to the

Sequencer module which uses PCLK and SYNC to establish 1113 states per revolution corresponding to RF buckets. The Sequencer derives 23 timing signals from memory data whose addresses correspond to the 1113 states. With the memory being accessed sequentially at the PCLK rate, programmable timing signals are produced that have one RF bucket (18.8 ns) resolution relative to any bunch. This mechanism has no limitation as to bunch spacing or location so long as the pattern is reproduced once per revolution. A reference pulse called BC_REF is produced for each anticipated bunch and compared to an input pulse (BC_INP) derived from a beam pickup. The coincidence of these two signals is used to establish the absolute reference for the Clock and to monitor Clock performance. The 23 timing signals are referred to as Timing Lines (TL<0:22>). Those timing signals produced by the Master Sequencer (MAS SEQ) are output to a custom backplane along with PCLK and MCLK for distribution to Selector Fanout Modules.
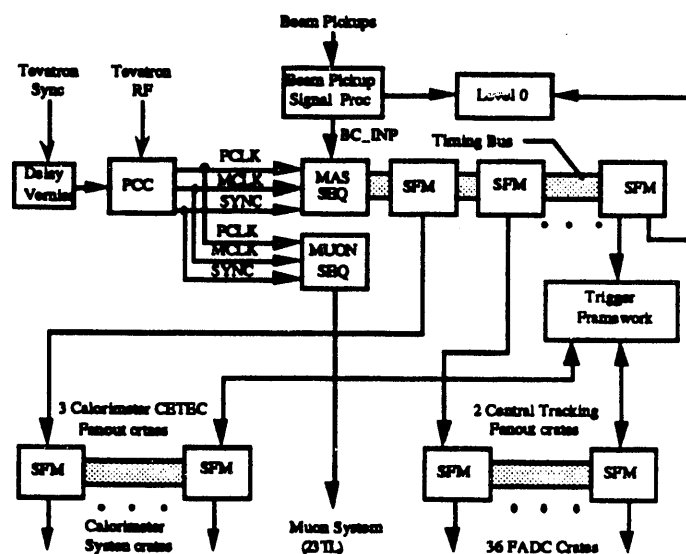
Figure (1) - D0 Clock System Block Diagram

Most of the timing signals are distributed to detector subsystems with a two tiered approach. Selector Fanout Modules, located in the crate with the Sequencer called the Clock crate, select timing signals and clocks from the backplane and output them to Selector Fanout Modules in other crates called fanout crates. One or more Selector Fanout Modules in fanout crates can accept signals from the Clock crate and other sources such as the D0 Trigger Framework [2]. These modules are then made to drive the input signals to

selected timing lines on the backplane of the fanout crate. Additional Selector Fanouts in the fanout crate can then each select and output four of these timing line signals with a maximum fanout of three for output to various detector subsystems. Beside the four selected timing signals, the Selector Fanout can output either MCLK or PCLK to the detector subsystem and accept a Busy signal from the subsystem and route it back to the Trigger Framework.

The second Sequencer shown in Figure (1) is identified as the MUON SEQ and generates an additional set of 23 timing signals that are used exclusively by the Muon system. The Muon Sequencer runs synchronously to the Master Sequencer because they both receive the same clock and synchronization signals from the PCC. Although the Muon Sequencer resides in the Clock crate, its timing line outputs to the backplane are disabled and its outputs are taken from its front panel. However, another Sequencer generating additional timing signals could reside in another crate with Selector Fanout Modules and operate in parallel with the Clock crate.

PHASE COHERENT CLOCK MODULE (PCC)

The PCC is composed of two commercial PC boards manufactured by Berkeley Nucleonics Corporation (BNC) piggy backed on a single PC board which contains interface, control, and error monitor logic circuits. Figure (2) is a block diagram of the PCC. The BNC C-1000 Trigger Coherent Oscillator, shown as the Phase Lock Board in Figure (2), is capable of producing a stable 106 MHz output frequency, f_0, coherent with a random trigger input. However, by supplying a variable reference frequency to the BNC C-1010 Freq Lock Board, the coherent output frequency of the C-1000 will track

the reference over a very small range (10 ppm). These two boards are well suited to the PCC's application since within limits they allow independent control over the phase and frequency of the output signal. The PCC receives the Tevatron RF transported to D0 over the CATV link which cannot be relied upon for its long term phase stability. However, its frequency is that of the real Tevatron RF which changes with energy level of the machine and serves as the reference input to the C-1010. The phase of the 106 MHz output is adjusted once each revolution by the Tevatron Sync pulse (TEV SYNC) which becomes the trigger input to the C-1000 Phase Lock Board.

The output frequency of the C-1000 (106 MHz) becomes MCLK that is used by the Central Tracking detector systems. PCLK is the primary clock frequency of the Clock system and is derived by dividing a version of MCLK by two such that it is equal to the 53 MHz Tevatron RF. SYNC OUT is produced once a revolution and is derived from a 1113 countdown of PCLK. Although the edges of SYNC OUT are derived from PCLK, synchronism with TEV SYNC is insured by presetting the countdown on each revolution in a transparent manner. The primary output signals from the PCC are MCLK, PCLK, and SYNC OUT and each has a fanout of three for the purpose of driving multiple Sequencers.

The TEV SYNC input to the PCC is derived from a Camac 279 module which receives the coded Tevatron Beam Sync Clock (TVBS) from the accelerator which is synchronous to the Tevatron RF. The 279 module produces a programmable once-a-revolution pulse, coherent to the bunches in the machine, that has a resolution of seven RF buckets. A programmable delay in the PCC called SYNC DLY allows for the adjustment of the phase of PCLK relative
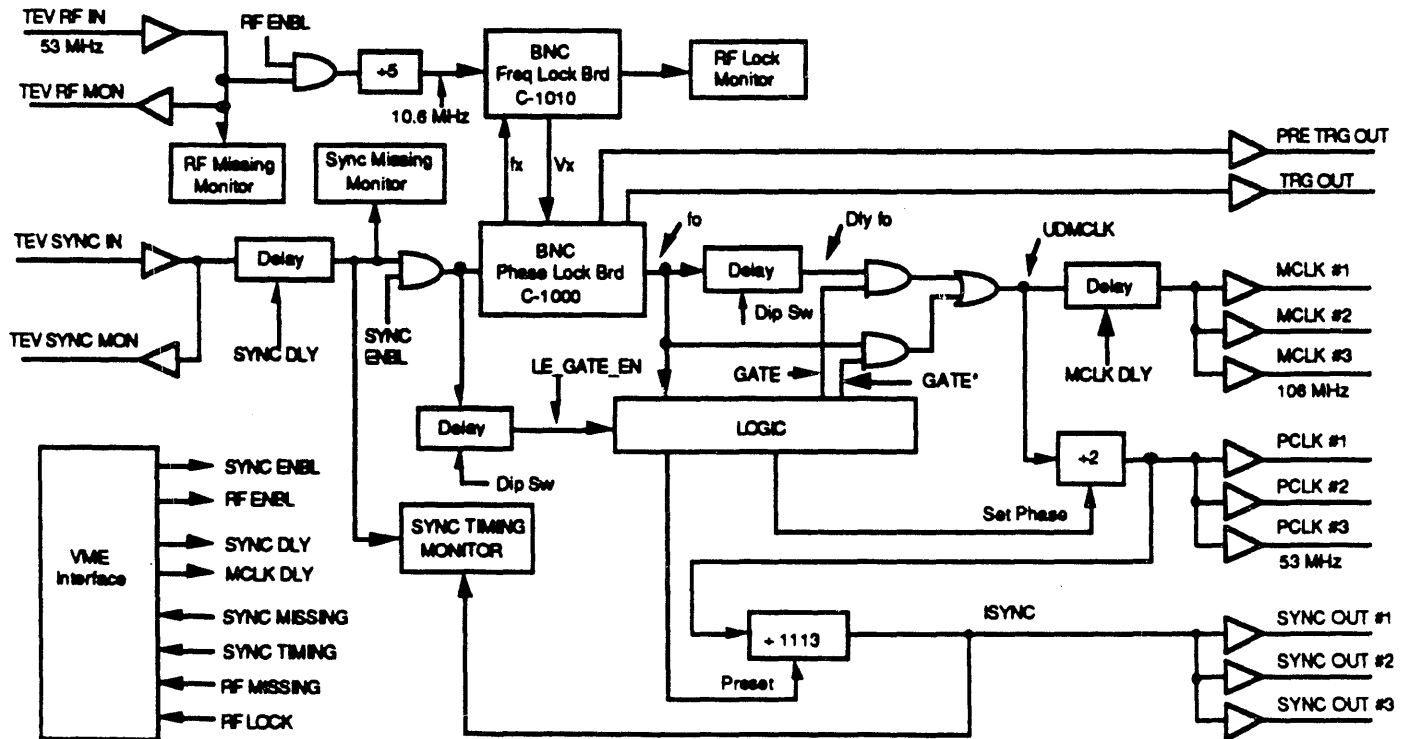

Figure (2) - Phase Coherent Clock Block Diagram

to the bunches in the machine with a resolution of 1 ns and a range of 16 ns. Adjustment of SYNC DLY will change the relative phase between all the timing signals produced by the Clock and the bunch crossings. A Sync Missing error is produced whenever TEV SYNC is missing or slow, and a Sync Timing error is produced whenever TEV SYNC and SYNC OUT are not coherent (± 3 ns). When TEV SYNC is enabled to trigger the C-1000 Phase Lock Board, the output, f_O, drops out for 30 ns. After the drop out, f_O restarts at the same frequency that existed prior to the trigger, but with a coherent phase relationship to the trigger. A delayed version of f_O is inserted into the f_O pulse stream to fill in the gap caused by the drop out. As a result, a version of MCLK called UDMCLK (undelayed MCLK) is created and appears uninterrupted despite the trigger. When TEV SYNC is disabled, f_O runs continuosly and remains coherent with the most recent trigger. The phase of MCLK relative to PCLK is adjustable with a resolution of 1 ns and a range of 16 ns using a programmable delay called MCLK DLY.

The frequency locking range of the C-1010 in terms of the Tevatron RF is 1.13 KHz and covers Tevatron energies from 150 GeV to 1 TeV. A rate detector monitors the TEV RF input and produces an RF Missing error if more than one or two pulses are missing, and the C-1010 produces an RF Locking error if it is unable to lock onto its reference input. With TEV RF disabled and the reference input to the C-1010 removed, the C-1000 oscillates at a fixed free running frequency $f_O = 106.20842$ MHz with an aging specification of ±0.2 ppm/week and temperature stability of ±1 ppm from 0 to 50° C. The PCLK free running frequency is 53.10421 MHz and corresponds to a Tevatron energy of approximately 210 GeV.

MODE	SYNC ENBL	RF ENBL
Free-Run	0	0
Freq-Lock	0	1
Sync-Lock	1	0
Normal	1	1

Table (1) - PCC Modes of Operation

The PCC has four modes of operation called Normal, Sync-Lock, Freq-Lock, and Free-Run and are controlled by setting two control bits SYNC ENBL and RF ENBL as shown in Table (1). In Free-Run mode, both TEV SYNC and TEV RF are disabled and the PCC free runs independent of the accelerator with all its error monitors disabled. In Sync-Lock mode, only TEV SYNC is enabled and the frequency of PCLK in between TEV SYNC pulses is fixed at the free running frequency. Since for energy levels other than 210 GeV the time between TEV SYNC pulses changes, a fraction of a pulse more or less than 1113 is produced between "perfectly timed" TEV SYNC pulses. This fraction of a pulse is constant for a given energy level and has a maximum value of 500 ps which is made up within a single cycle on each revolution near the time of TEV SYNC. The end result is that the phase of PCLK changes relative to the Tevatron RF throughout each revolution with an abrupt discontinuity near the time of TEV SYNC as shown in Figure (3). In Normal mode, both TEV

SYNC and TEV RF are enabled and the PCLK frequency between TEV SYNC pulses tries to track TEV RF. Hence, one would expect the deviation in phase over one revolution to be less than in Sync-Lock mode. In practice, the required phase adjustment in Normal mode also has a maximum value of 500 ps at 1000 GeV and near 1 ns in the opposite direction at 150 GeV. The reasons for this are not understood, but can be attributed to the internal response behavior of the C-1000.


Figure (3) - PCLK Phase Discontinuity at TEV SYNC

In Freq-Lock mode, TEV SYNC is disabled and TEV RF is enabled. PCLK is phase locked to the TEV RF and runs without the necessary phase adjustment on each revolution because it tracks the TEV RF at all Tevatron energies. In order to operate in this mode, however, TEV SYNC would periodically have to be enabled for a short time to establish a consistent phase relationship between PCLK and the bunches in the machine. This would be accomplished by placing the PCC in Normal mode for a second or so approximately every ten minutes. Because TEV SYNC does not trigger the C-1000 Phase Lock Board in this mode, the relative phase of PCLK does not exhibit any discontinuity at TEV SYNC as it does in Sync-Lock and Normal modes. Although TEV SYNC is disabled from triggering the C-1000 in this mode, functions dependent on TEV SYNC such as Sync Timing and Sync Missing error monitors remain operable. If a Sync Timing error occurs, synchronization must be re-established in Normal mode.

Beside the constant phase discontinuity in PCLK, an additional random discontinuity in the phase of PCLK occurs at the same time due to jitter in the TEV SYNC signal from the accelerator of ≈400 ps. Hence, in Normal and Sync-Lock modes, the jitter in PCLK relative to a fixed bunch crossing from revolution to revolution will be ≈400 ps rms as shown in Figure (3). In Freq-Lock mode, TEV SYNC is inhibited from triggering the Phase Lock Board and hence, there is no random discontinuity on every revolution. The phase of PCLK relative to the Tevatron RF would remain constant during the period that the PCC was in Freq-Lock mode, and change to a new constant value during subsequent Freq-Lock periods.

Although Freq-Lock mode minimizes jitter and eliminates the PCLK phase discontinuity on each revolution, the current precision requirements of the experiment do not merit the additional operating complexity of using Freq-Lock mode. The effects of discontinuity in the phase of PCLK exhibited in Normal and Sync-Lock modes are minimized by forcing the discontinuity to take place between bunch crossings when data
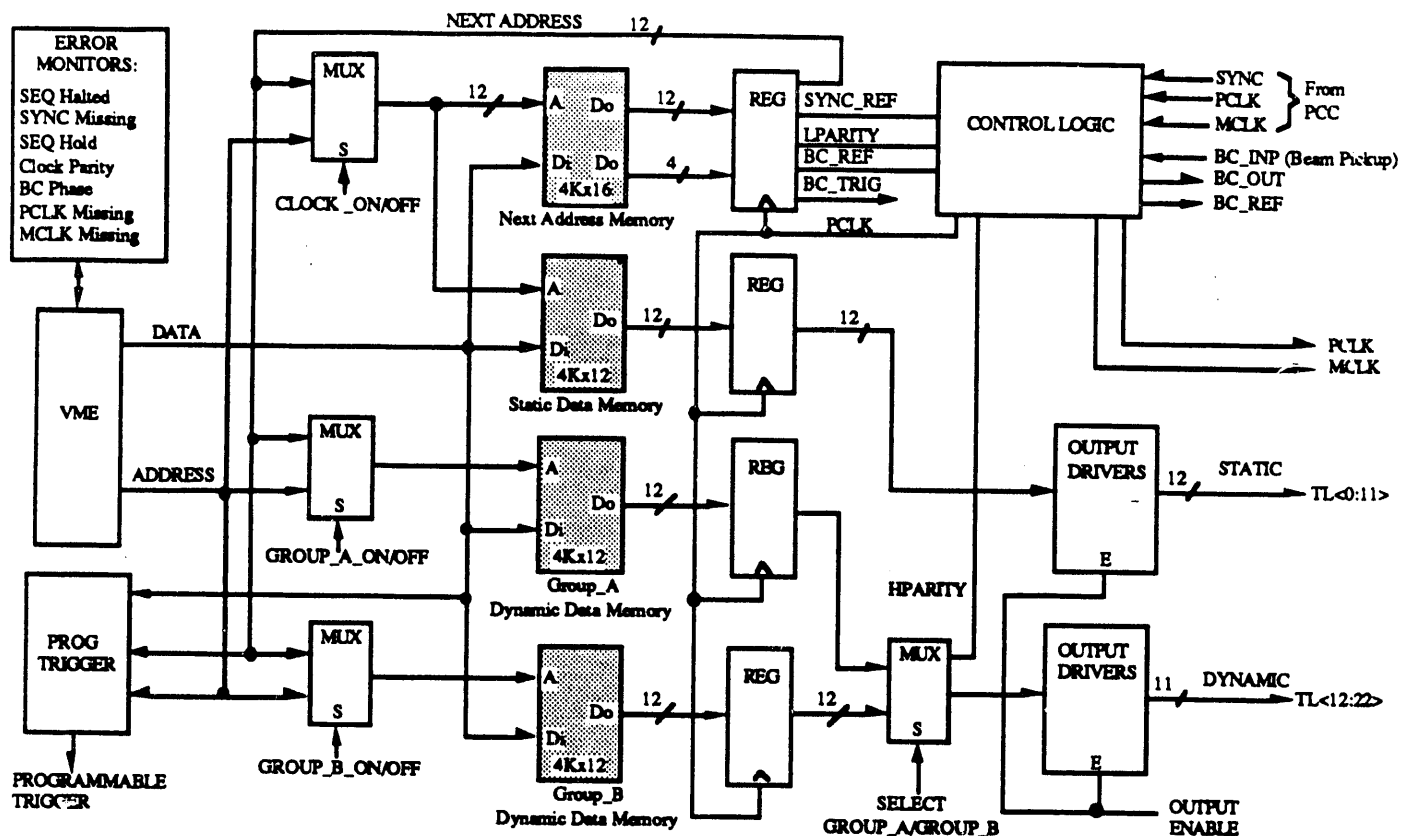
Figure (4) - Sequencer Block Diagram

acquisition is quiet. The jitter due to TEV SYNC in Normal and Sync-Lock modes from revolution to revolution is non-systematic and small enough so as not to be a problem. The measured jitter in PCLK at any fixed bucket in all modes of operation using a near perfect TEV SYNC and TEV RF as input is <100 ps rms.

SEQUENCER MODULE (SEQ)

Both the Master Sequencer and Muon Sequencer receive PCLK, MCLK, and SYNC from the PCC and are operated in the same manner, each producing 23 distinct timing signals with a resolution of one RF bucket or PCLK cycle. The difference between them is that the Muon Sequencer's timing signals are output to its front panel only, and the Master Sequencer outputs its timing signals to the custom J3 backplane and its front panel. Figure (4) is a block diagram of the Sequencer and shows four memory blocks called Next Address Memory, Static Data Memory, Group_A Dynamic Data Memory, and Group_B Dynamic Data Memory. Each of the memories has a registered output that is updated on each PCLK cycle. With 12 bits of the Next Address Memory pointing to the address of the next sequencer state, the memories are cycled at the PCLK rate and their output data become the timing signals. The memories in the Sequencer are 4K deep and hence capable of a repeat cycle of 4,096 states. However, a software limit of 1113 states is imposed to correspond to Tevatron RF buckets. Although MCLK is input to Sequencer, it is not used in any manner to generate timing

signals. It is, however, monitored and output to the front panel and the backplane as are the timing signals and PCLK.

The 23 timing signals called timing lines are split into two groups. The first group of 12, TL<0:11>, are referred to as Static and are derived from the Static Data Memory. The second group of 11, TL<12:22> are referred to as Dynamic and are derived from either the Group A or Group B Dynamic Data Memories. In order to download the Next Address Memory or the Static Data Memory, the Clock ON/OFF MUX must select the VME address which turns the Sequencer OFF, Selecting the NEXT ADDRESS turns the Sequencer ON. The Static Data Memory contains data corresponding to the Static timing lines which are referred to as statically programmable because the Sequencer must be turned off in order to change them. The Group A and Group B Dynamic Data Memories are functionally identical, and outputs from one or the other are selected for output to the Dynamic timing lines. The Mux ahead of each memory separately determines whether that memory is ON and is sequencing by selecting the NEXT ADDRESS, or is OFF by selecting VME. While one of the memories is ON outputting data to the timing lines, the other can be turned OFF and downloaded with new data for one or more timing lines. Then by turning ON and selecting the newly downloaded memory, timing signals can be changed on the fly without turning the Clock off. Hence, these timing lines are called dynamically programmable.

One of the control bits of the Next Address Memory called BC_REF is programmed to "1" at each location corresponding to the bucket number in which the bunch resides. The pulse

that is produced is compared to the BC_INP pulse which is derived from a beam pickup located in the accelerator near D0. The coincidence of these two pulses is used to establish the absolute reference for the Clock. Their continued coincidence is monitored internal to the Sequencer to within ±3 ns to insure synchronization with the accelerator. A BC Phase error is generated when anti-coincidence occurs on two successive revolutions to monitor phase drift while being immune to TEV SYNC jitter.

Another control bit of the Next Address Memory called SYNC_REF is programmed to "1" at a single location and is compared to the SYNC input pulse once each revolution to verify synchronism. If the SYNC input is found not to occur within the same bucket as SYNC_REF, the Sequencer is placed in a hold condition at the SYNC_REF bucket until the next SYNC input occurs. In operation, the Sequencer goes in and out of the hold condition in the same bucket once each revolution in a manner transparent to the user. This mechanism forces SYNC_REF to automatically become aligned with the SYNC input. If synchronization is lost and the Sequencer is forced into the hold condition beyond one PCLK cycle, a Sequencer Hold error is generated. However, recovery from a temporary loss of synchronization is automatic and occurs within one revolution. The position of SYNC_REF can be programmed relative to other signals. Hence, it can be used to establish the timing of the outputs relative to the SYNC input and the bunch crossing to within one bucket. The Sequencer can also be made to Free-Run, in which case the Sequencer continually recycles independent of

the coincidence of SYNC_REF and the SYNC input.

Two parity bits corresponding to the Static and Dynamic timing lines are written to the memories. A Parity error is generated if incorrect parity is detected at any bucket. Three additional error monitors detect whether any of the PCLK, MCLK, or SYNC signals are missing or slow. These errors, along with those previously discussed, are latched upon their occurrence and hence may not indicate the current status of the Sequencer. However, the error status of the Sequencer and the PCC will be monitored and cleared by the D0 Alarms and Limits system on a regular basis.

Diagnostic features implemented on the Sequencer include: front panel outputs of key signals, onboard LED display of timing line states, a Step Mode where VME cycles are substituted for PCLK, and two front panel trigger pulses. The first, called BC_TRIG, is programmed as part of the Next Address Memory independent of the timing lines. It is intended to be used as a scope pre-trigger to monitor Bunch Crossing synchronization but can be used for other purposes. The second front panel pulse, called the Programmable Trigger, is a once-a-revolution pulse that can be programmed through VME to occur at any bucket without disrupting the operation of the Clock.

SELECTOR FANOUT MODULE (SFM)

The Selector Fanout Module (SFM) has the function of distributing Sequencer timing signals to various detector systems. Figure (5) is a block diagram of the SFM and shows
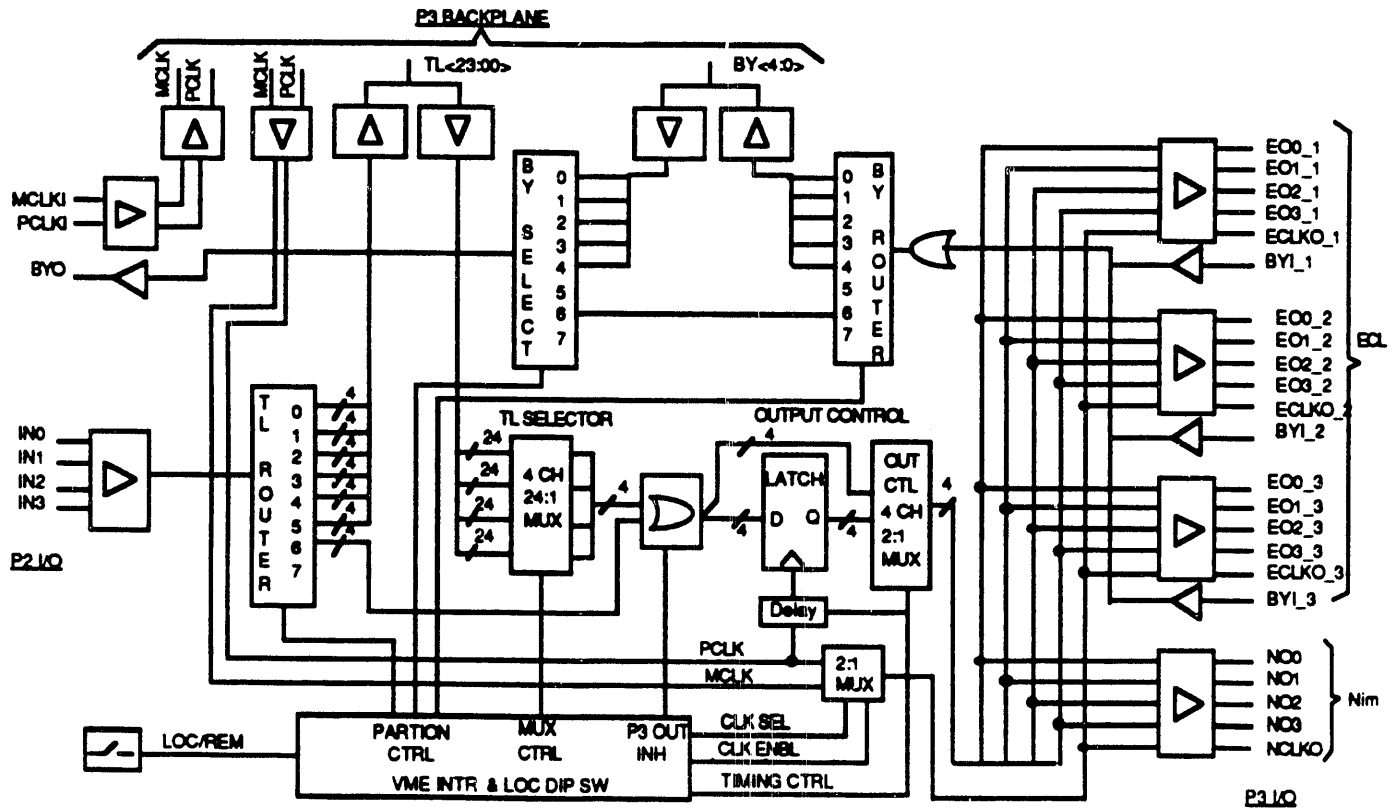

Figure (5) - Selector Fanout Module Block Diagram

its P2 I/O on the left, P3 backplane connections on the top, and P3 I/O on the right. SFMs are used for a variety of functions in the system which use different sections of the module. In the Clock crate, SFMs only receive signals from the P3 backplane and drive them at their P3 I/O. In Fanout crates, SFMs drive and receive signals at their P2 and P3 I/O, in addition to driving and receiving signals at the P3 backplane. A group of SFMs along with the custom J3 backplane can be thought of functioning as a crosspoint switch. Up to four signals can be input at the P2 connector of each SFM and be made to drive specific bussed backplane timing lines in groups of four via the TL Router. Up to six such modules, each routing signals to specific timing lines on the backplane, would put as many as 24 timing signals on the bus. The TL Selector of each module is then capable of selecting up to four signals from any of the timing lines on the backplane. These signals plus either of the two clock signals, MCLK or PCLK, are duplicated and output at the P3 I/O. Thus, signals input at P2 on various boards can be mixed and matched and output at the P3 I/O of other boards. Additionally, up to three Busy signals can be input at P3, summed and routed via the BY Router to one of five bussed BY<0:4> lines on the backplane. The BY Selector from the same or other SFMs can select any of the bussed BY lines and output the selected Busy signal at P2.

A front panel Local/Remote switch determines the source of data for module control registers. In Remote, all registers are readable and writeable from VME, whereas in Local the contents of these same registers are determined by onboard dip switches. The Local option permits SFMs, having constant configurations, to be operated in areas where VME is not available.

SFMs located in the Clock crate will normally use only the TL Selector and fanout portions of the circuit. The Sequencer resides in slot 7 of the Clock crate and drives timing lines, TL<0:22>, along the bussed J3 backplane to SFMs which can occupy any slots 8 thru 20. The TL Selector of a given module selects signals from the backplane for distribution through its P3 I/O to fanout crates or directly to detector systems. Actual outputs are taken from rear mounted auxiliary boards having suitable connectors for the application. Both PCLK and MCLK are driven by the Sequencer onto the J3 backplane where they are actively fanned out and input directly to each module. Although SFMs in the Clock crate normally only use the TL Selector, nothing prohibits the use of other sections of the module so long as the applications are compatible.

SFMs in fanout crates can reside in any slots 7 thru 20. However, only the SFM in slot 7 is able to accept PCLK and MCLK and distribute them to other modules in the crate. Timing signals from the Clock crate as well as signals from other sources such as the Trigger Framework can be input at P2 of any module in the fanout crate. They are input to the SFM from rear mounted auxiliary boards having varied connector configurations depending on the type of cable used. Signals input at P2 are treated as a group and can be inhibited, routed to the backplane, or routed directly to the P3 output

control logic. In the latter case, they do not consume any timing lines on the backplane, but other SFMs in the crate can not gain access to these signals. Four TL Selectors in each module select timing signals from the backplane to be output to P3 with a Differential ECL fanout of up to three and a Nim fanout of one. MCLK and PCLK are input to each module either of which can be selected and fanned out along with the timing signals. Each SFM has the ability to retime the selected timing signals prior to being output by latching them with a delayed version of PCLK having a resolution of 1 ns and a range of 16 ns. This allows for the adjustment of the phase of the timing line signals relative to the CLK (MCLK or PCLK) output.

SOFTWARE

A PC based software package for controlling all three module types is available for use in the laboratory using a PC to VME interface directly, or in the experimental hall using a PC to token ring interface and the D0 control system. The program is called the D0 Clock Control and Test Program (CT) [3] [4] and requires the use of an additional program called the D0 Clock Timing Language Translator [5]. The translator accepts an ASCI source file using specific syntax to describe the desired behavior of the timing lines and converts it to a file that can be downloaded to the Sequencer with the CT program. Interactive control over all three module types including individual SFM by device name is implemented. Two additional PC based programs provide for auto testing of the Sequencer and the SFM.

VAX-based control software for the Phase Coherent Clock and the Sequencer is implemented using a client/server model allowing access from multiple clients without interference [6]. Users can be privileged or non-privileged and can reserve specific timing lines that no other user can alter. The program automatically executes the necessary VME commands to the Sequencer to modify timing lines. Modification of the static timing lines requires that the Clock be turned off and hence can only be executed by privileged users.

REFERENCES

[1] C. Rotolo, M. Bowden, R. Kwarciany, and S Chappa, "The Collider Detector Master Clock System", CDF Note #514, October 1986.

[2] M. Abolins, et al.,"A High Luminosity Trigger Design for the Tevatron Collider Experiment in D0", IEEE Transactions on Nuclear Science, Vol 36, No. 1, February 1989, pp 384-389.

[3] B. Graves,"Users Guide D0 Clock Control and Test Program v1.03", Fermilab D0 Note #922, January 31, 1990.

[4] R. Angstadt,"D0 Clock Control and Test Program v3.02", D0 Note #1259, April 30, 1991.

[5] B. Graves,"Users Guide D0 Clock Timing Language Translator v2.3", Fermilab D0 Note #919, January 25, 1990.

[6] Laura A. Paterno and J. Frederick Bartlett,"User Manual For The D0 Clock Interface v1.0," Fermilab D0 Note #1029, October 25, 1990.

END

DATE FILMED
01/24/92