© 1979 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.
IEEE Transactions on Nuclear Science, Vol. NS-26, No. 3, June 1979

HIGH RESOLUTION DIGITAL TIME BASE*

R. E. Meller†

Summary

A time base has been developed to supply general timing signals for the Cornell Electron Storage Ring, including beam bunch transfer triggers and beam detector triggers. The system consists of 8-bit counters and digital phase shifters of 8-bit resolution, under control of a data bus tied to the CESR control computer interface. The counters use conventional Schottky TTL techniques at a clock rate of 23.8 MHz, and the phase shifters are based on an ECL phase-locked loop, giving a resolution of 150 ps/LSB, jitter of 50 ps, an offset T.C. of 10 ps/^{o}C, and a gain T.C. of 100 ppm/^{o}C.

System Organization

A modular organization is required so that delays and frequency division functions can be cascaded at will and so that timing signals can be fanned out or combined in coincidence or multiplex functions. Each timing function occupies a single 4.5" by 6.5" printed circuit card. Timing inputs and outputs are brought to Lemo jacks at the front of the cards so that timing signal connections can be programmed by plugging coaxial cables directly between the timing functions. Counting numbers are loaded directly from an 8-bit data bus wired in the backplane, driven from the CESR control computer. The inputs for enabling, multiplexing, and mode controls are wired to a taper block interconnect area, where an input can be wired to a permanent high or low level or to a binary latch loaded from the data bus. In this way the system structure can be modified by computer control.

After a timing signal has been generated by counting operations or by synchronization of external signals, a phase shifter can optionally be inserted before the signal is sent to the line driver for external use.

Clock Distribution:

The linac injector for CESR produces bunches of electrons or positrons at 23.8 MHz rate, corresponding to a period of 42 ns. This is chosen as the frequency of the master timing clock. The fundamental frequency for the beam chopper, linac, synchrotron, and CESR is 11.9 MHz derived from a high stability crystal oscillator, so a frequency multiplication of 2 is required. This is done with a phase-locked loop instead of a harmonic multiplier in order to get good rejection of the 11.9 MHz component in the output without fancy tuned circuits. The clock is distributed as a square wave individually to each timing function by a tree network using terminated 50-ohm coax. 74S140 TTL level driver I.C.'s are used, except that 10210 ECL level driver I.C.'s are used to distribute to the phase shifters.

Timing Signals:

Timing signals are represented by pulses of one clock period duration. Interconnection between timing functions is done with terminated 50-ohm coax driven by 74S140 line drivers. Inverted logic levels are used to take advantage of the 74S64 TTL AND-OR-INVERT gate as an enabled-input coincidence gate without an additional inversion.

In order to prevent ambiguity as to which cycle of the clock a timing signal refers to, it is necessary to guarantee that a timing signal initiated at one positive edge of the clock in one timing function be received by the next timing function on exactly the next clock edge. To this end, the edges of the timing signal are not used, and only the level of the timing signal at the next clock edge is used. In this way, the exact arrival time of the timing signal is not important as long as it is between the correct two clock edges. This means that the total transmission delay between two timing functions must be less than the 42 ns clock period, and the clock skew between functions must be less than the minimum transmission delay, about 5 ns.

The transmission delay is minimized by synchronizing the timing signal to the clock before transmitting, and then synchronizing to the clock again upon receiving. This minimizes the hardware delay, and leaves the maximum window for cable delays, which may be as large as 15 ns with safety. Also, the window is standard and independent of delays or setup times internal to the timing function. The minimum insertion delay for a timing function is 2 clock periods. 74S74 dual D-type flip-flops are used as the synchronizing elements by applying the signal to be synchroonized to the D input and the master clock to the clock input.

Timing Functions

Synchronizer:

This device receives external triggers, places them in synchronism with the master clock, and fans them out as timing signals. It has two timing inputs which can be programmed as a coincidence gate or as a multiplex. Also, it has status latches which can be set by timing signals and read and cleared by the data bus. It is used for introducing asynchronous signals into the system, single pulse gating, multiplexing, and signal fanout.

Counter:

This device uses an 8-bit binary counter to implement the functions of divide by N, delay by N, divide by M with phase N, and N-wide gate, where N is a number loaded by the data bus, and M is a number wired on the board with jumpers, usually 60 or 61, the synchrotron and CESR revolution periods respectively. The counter, consisting of two cascaded 74S163 counter I.C.'s, is loaded with the complement of N. It is wired to produce an output whenever it passes into the zero state. When started by a timing input, the counter counts upward through full count to zero. producing an output delayed by N. If the counter is allowed to continue, it will reach M and be reset to zero repeatedly, producing an output with period M, where N is characterized as a phase. If the counter is programmed to delay by N and restart automatically, divide by N results. An N-wide gate is generated by allowing an output continuously from when the counter is started until it stops at zero. In the divide by M mode, the counter can be used as a marker on a specific beam bunch in an accelerator, where M is the accelerator period and N selects the bunch. The default counting rate is the master clock, but an input is provided for a submultiple counting rate.

Phase Shifter:

A fine timing adjustment is necessary after the selection of a specific clock cycle by a timing signal. This adjustment must have a resolution and stability of at least 1 ns. Since this is the last element of the timing chain that references the clock, special care is taken to distribute the clock to the phase shifters with a minimum of intermediate stages. When the timing signal is synchronized to the clock at the phase shifter, the resulting signal is as accurate as the clock at that point in the system, so that the clock to the earlier stages, i.e. counter and synchronizer functions, can be distributed less carefully.

The technique of using a phase-locked loop as a phase shifter was chosen over a box of switched delay

*Work supported by the National Science Foundation.
†Laboratory of Nuclear Studies, Cornell University, Ithaca, NY 14853.

0018-9499/79/0600-4152$00.75 © 1979 IEEE

cables because of the volume and cost of a cable box. The interesting problems arising from the PLL scheme are stability, capture of the loop at initial power-up, and synchronization to a local clock of arbitrary phase.

In a basic PLL, the master clock and the local clock are sent into a phase detector. The local clock is a voltage tuned oscillator, and it is tuned by the output of the phase detector. When the loop is in lock, the local clock must be at the same frequency as the master clock. If a high-gain amplifier is introduced between the phase detector and the voltage tune input, then the local clock must be in phase with the master clock, i.e. there must be zero output from the phase detector. In this application, a linear phase detector is used[1], and a set-point signal is summed with the phase detector output before amplification. Therefore, the local clock must adjust itself in phase until the phase detector output is equal and opposite to the set-point signal, resulting in a zero sum. Since the phase detector is linear, the local clock phase is a linear function of the set-point signal. The input time signal, which was originally synchronized to the master clock, can now be resynchronized to the local clock, resulting in an adjustable delay of the timing signal.

The local clock oscillator must be chosen carefully, since the jitter in an oscillator output is directly related to the width of the oscillator frequency spectrum. An L-C tuned oscillator is indicated over a multivibrator type oscillator because of its intrinsically narrower spectrum. The 1648 I.C. oscillator used has the additional advantage of an internal AGC network, which allows operation of the oscillator in a linear mode, thereby reducing modulation sidebands in the output. The output is a buffered ECL level square wave. The external tank circuit consists of a 0.4 uH inductor of type 4C4 ferrite, and an MV1401 varactor tuning diode. A typical spectrum bandwidth is 10 kHz[2]. Under voltage control, the approximate tuning gain of the oscillator is 2.6 MHz/V. The oscillator was found to have a T.C. of -7 kHz/^{o}C, or -2.7 mV/^{o}C referenced to the tuning input. This was compensated approximately by obtaining the varactor reference bias from a zener diode of equal and opposite T.C.

The linear phase detector is in principle a flip-flop which is set by the rising edge of the master clock and cleared by the rising edge of the local clock. This gives a pulse train of duty factor proportional to the phase difference. This has a range of about 300^{o}, and asymmetric imperfections in the logic level transitions will cause an error in the pulse width. Therefore, the phase detector is symmetrized by using two flip-flops, one of which is set by the master clock and the other of which is set by the local clock. When both are set, a gate detects this state and clears both simultaneously. Then the two output pulses are subtracted in a differential amplifier. One of the pulses is a dummy which only serves to subtract out the effects of the flip-flop rise and falltimes. Since this configuration has a signed output, the range is twice as great as the one flip-flop scheme, or 600^{o}. ECL logic is essential for this application because of the need for well defined levels. The TTL high level is poorly defined, and may continued to rise by several volts after a high-going transition. The 12040 ECL phase detector I.C. implements the desired function, except that the flip-flops are cleared so quickly that the dummy pulse never achieves a full logic level transiton. This spoils the cancellation symmetry, and also causes a non-linearity in the phase detection characteristic at zero phase. Instead, similar function was made with a 10231 dual D-type flip-flop and a 10102 NOR gate. The delay of the gate is sufficient to allow the flip-flop to change state completely before clearing. The resulting phase detector has a gain of 130 mV/Rad, or .8V/Cycle.

ECL levels have a T.C. of approximately 1 mV/^{o}C, which corresponds to a phase gain T.C. of 1200 ppm/^{o}C. This can be mostly cancelled by introducing the set-point signal as an ECL level signal also, so that the two temperature coefficients cancel. This is done with an 8-bit binary rate multiplier which is loaded from the data bus. The BRM produces a train of N pulses out of a possible 256. When filtered, this pulse train gives a voltage in 3 mV increments, corresponding to 150 ps phase increments. The BRM is implemented in TTL for convenience, then translated to ECL level and synchronized to the master clock before summing, to take advantage of the master clock period as an accurate reference for the BRM pulse length. It was found that the outputs of ECL gates of the same type on the same chip have levels that are matched to within 2%, where gates on different chips differed by 8% or more. In addition, gates on the same chip are thermally linked. Both of the phase detector pulses and the BRM pulses are sent through gates on the same 10102 I.C. solely for the purpose of matching the levels before filtering and summing. The residual gain T.C. is 100 ppm/^{o}C, and the residual offset T.C. is 10 ps/^{o}C. The temperature coefficients were measured over a 40^{o}C difference with a Tektronix 7904 real-time oscilloscope.

Since the BRM contains frequencies other than the clock frequency, great care must be taken to prevent pollution of the local clock by cross-talk from the BRM. In fact, most of the measured 50 ps jitter is due to such cross-talk. The power supplies in common to the BRM and PLL have additional decoupling, and the PLL is located on the far side of the circuit board from the BRM and shielded with a continuous ground plane.

The sidebands of the local oscillator must be removed in order to have a stable output. This is done by the loop amplifier in its normal process of comparing and correcting the phase of the local oscillator. The loop amplifier must have a high gain over the bandwidth where the oscillator has significant sidebands, which is roughly 100 kHz. This is done with an operational amplifier based integrator. The total phase lag of the feedback loop is now 180^{o}, due to the 90^{o} lag of the integrator and the 90^{o} lag of the phase detector. This results in an unstable feedback loop, so the total lag of the loop amplifier must be corrected to less than 45^{o} at the unity gain frequency of the loop to insure a good stability margin. This is done by summing the integrator output with the output of a flat-response one-transistor AC amplifier. The flat-response amplifier dominates the integrator at frequencies above 100 kHz. At frequencies below this the integrator provides the very high gain necessary for sideband suppression. The unity-gain frequency is 1.5 MHz. At that frequency the most important parasitic phase lag is involved in driving the .005 uF bypass capacitor at the voltage tune port of the local oscillator. A LH0002 buffer amplifier with 6 ohm output impedance is used to keep the phase lag within 20^{o}.

When the circuit is powered up, the local oscillator is initially operating at a frequency at an extreme of its tuning range. In this state, the phase detector produces the beat frequency between the local and master clocks. To insure capture by the loop, the beat frequency must be within the loop bandwidth. Capture was found to be reliable if the maximum excursion of the local clock was limited to one-third of the unity-gain frequency, i.e. 500 kHz. However, if the set-point signal were absent, capture would be unconditional regardless of the loop bandwidth, because the type of phase detector used produces a DC

output component when the frequencies are different.

The main problem in synchronizing a timing signal to the local clock is that if the phase of the local clock is increased continuously, at some point the output will move to the previous clock edge. This is due to the fact that a clock that has been shifted by 360^{o} cannot be distinguished from the unshifted clock. The phase detector is signed, however, so it is possible to distinguish a clock shifted by $+180^{o}$ from one shifted by -180^{o}. A continuous output range is obtained by synchronizing first to the master clock, then to the trailing edge of the local clock flip-flop in the phase detector, and then to the local clock. The phase detector flip-flop output has the property that when the local clock is leading in phase, the trailing edge comes after the master clock, setting up an output on the very next edge of the local clock. When the local clock is lagging in phase, the trailing edge comes after the local clock, setting up an output on the local clock edge after that, or the second edge after the timing signal is received. This gives the phase shifter a range of 180^{o} to 540^{o}, covering a span of 360^{o}.

Line Drivers:

These devices provide a full 5 volt level in positive logic configuration for transmission to remote locations. Two types of line drivers are provided: A type that produces 10 us duration pulses limited to the 60 Hz accelerator repetition rate for triggering pulsed magnets, and a type that produces the 42 ns timing signals unmodified, at any repetition rate for triggering beam detectors and sample-and-hold gates.

Typical Application

Transfer of a beam bunch from the synchrotron to CESR requires that the accelerators be in a certain relative phase at time of transfer. In other words, both the bunch to be transfered and the bucket to be filled must be lined up with the transfer line. A divide by 61 counter can be used to mark the bucket in CESR to be filled. The synchrotron injection trigger is introduced by way of a synchronizer. This timing signal is used to simultaneously gate a bunch of charge into the synchrotron and start a divide by 60 counter to record the position of the charge bunch in the synchrotron. The phase relations for beam transfer are correct whenever there is a coincidence between the two counters. Therefore, the synchrotron extraction trigger is used to enable a synchronizer which waits for the coincidence and then starts a set of delay by N counters which provide adjustments on the triggering times of the different beam transfer magnets. Line drivers are then used to send the triggers to the transfer line.

The system must also be used to coalesce positron bunches,[3] where the synchrotron is used as a phasing loop for e^{+} bunches already in CESR. A delay by N counter is used to count outputs from the divide by 60 counter to initiate transfer back to CESR after N revolutions in the synchrotron. Multiplexing is used to switch between the two modes.

References:

1. Motorola Semiconductor Products, Inc., MECL Integrated Circuits Data Book, 1973, 6-39.
2. Motorola Semiconductor Products, Inc., MECL Integrated Circuits Data Book, 1973, 6-3.
3. M. Tigner, Improved Method for Filling Electron Storage Rings from a Synchrotron, Particle Accelerator, 1975.