PRECISION TIMING CONTROL SYSTEM

Robert E. Meller
Cornell University, Ithaca, NY 14853 *

1 INTRODUCTION

Operation of the CESR storage ring with complex bunch patterns has made necessary more flexible timing controls. Multiple trains of bunches based on a fundamental spacing of 14 ns are used in both the injector and the storage ring. Hence, a generally programmable trigger pattern is needed for electron gun pulsing, beam detection, beam feedback gating, and RF phase control. To achieve this, an upgrade of the previous CESR timing system[1] using small-feature CMOS logic has been built. The new system has a time delay resolution of 10 ps and RMS jitter of 14 ps. The delay transfer function is linear to within 30 ps, and is monotonic over 32 bits of range. The use of programmable large-scale integrated circuits and the VME packaging system results in a system that is contained in 4 crates and occupies a fifth of the volume of its predecessor.

1.1 Design objectives

The previous CESR timing system had a resolution of 160 ps, which was soon found to be inadequate. It consisted of several hundred circuit cards, most of which were 8-bit timers. These were connected point-to-point with coaxial cables to obtain concatenated delays of sufficient length and precision. The precision functions were implemented with ECL logic, and hence there were a large number of power supplies which were also the main source of failure. The system had grown to occupy a volume of about 2 cubic meters, a quarter of which was power supplies.

The main objective of the timing upgrade was to improve the resolution and flexibility of the system. Resolution of 10 ps was chosen as a practical limit of the CMOS technology that is easily available. The system must also be able to generate arbitrary bit patterns at a 14 ns clock rate. Another objective was flexible configuration of the system from data files rather than by point-to-point connections, and testing of the system by software. Additional objectives were standard and compact packaging with a minimum number of power supplies.

2 SYSTEM DESIGN

2.1 Clock frequency

The fundamental period for CESR bunch loading is 14 ns, corresponding to a frequency of 71.4 MHz and 183 buckets per CESR revolution. This frequency is within reach of LSI logic devices, but it is very inconvenient for synchronous interconnection of a large system. A period of 42 ns was chosen for the timing system clock, corresponding to a frequency of 23.8 MHz. This frequency is the lowest common multiple of the revolution frequencies of CESR and the injection synchrotron. Multiplication of the clock to 71.4 MHz is only done where bit patterns are needed.

2.2 Logical organization

The standard timing channel is based on a 20 bit counter which operates synchronously at the 42 ns clock period. The counter is loaded from either of two 20 bit registers, which can control either pulse delay, width, or period. The counter has a start and stop input, and can be programmed to issue pulses singly or periodically when started, with or without a software enable. Both phases of the clock are used to give the output pulse 21 ns delay resolution. The range of the counter is about 44 ms, which exceeds the 16.7 ms injection cycle time. This unit provides most of the needed timing functions, and is implemented in a single CMOS logic array, with 8 channels on a single 6U VME card.

Eight trigger signals are bussed cratewide through the P2 backplane, and each standard timing channel has a pair of programmable gates which derive the start and stop inputs as coincidence functions of any of the 8 triggers. Each standard channel has the capability of driving a trigger signal on its own card, and each trigger signal is passed between the card and crate backplane by a bidirectional buffer. Hence, each trigger signal can be driven externally, or by a timing channel either cardwide or cratewide. This permits any portion of the system with complex cross-connections to be confined to a single card, without using cratewide connection resources.

The precision timing channel consists of a standard channel, followed by two optional modules for the precision functions: The precision delay performs a 12 bit phase shift of the 42 ns clock, which gives 10 ps delay resolution. The pattern generator multiplies the phase-shifted clock to 71.4 MHz and uses it to clock a bit pattern from a random-access memory. Two precision channels are provided on a single 6U VME card, along with a power converter which provides low-ripple +5 V for the precision components. The precision output drivers issue both the 71.4 MHz clock and the bit string so that the timing of complex bit patterns can be reconstructed by the receiver without the distortions caused by cable dispersion: When a pulse signal is sent down a cable, each pulse develops a tail which interferes with the timing of any closely following pulse. However, since the clock is periodic, the effect of dispersion is only a uniform phase shift. If the bit string is resynchronized to

* Work Supported by the National Science Foundation

0-7803-4376-X/98/$10.00 © 1998 IEEE

the clock at the destination, then the dispersion of the bit string signal is irrelevant as long as it satisfies a setup time relationship with the clock signal.

Each timing crate is self-contained, and supports between 32 and 128 timing channels. Hence, each crate is an independent timing system, and typically requires only a clock input and two trigger inputs to synchronize it with other timing crates. This allows the overall timing system to be substantially de-centralized.

2.3 Software control

The VME bus controller is a MVME162 microcomputer. This unit has a 68040 microprocessor with extensive resources of RAM and ROM. The microprocessor code is stored in battery-protected static RAM, and will run automatically when the crate is powered up. The entire configuration of the timing system is down-loaded from the CESR control system via a custom serial interface when the system is initialized. The configuration consists of mode controls and interconnections for the standard timing counters, formulae for the numbers to be loaded into the timing registers, bit patterns to be loaded into the pattern generator RAMs, and also test masks to be used to verify installation of the correct timing hardware. After the initial configuration, the serial interface transmits commands from the control system. These commands are usually adjustments to timing numbers or channel enables/disables, but can include reconfiguration of any parameters except the hardware tests.

The MPU code which performs real-time control of the timing channels is executed at 60 Hz, synchronous with the accelerator injection cycle. During the idle period between injections, the MVME162 program performs all necessary setup for the next injection cycle. These operations include enabling or disabling timing outputs, loading values into the timing registers, loading bit patterns, and possibly reconfiguring the modes and connections of the timing channels. During the injection cycle, the MVME162 services command requests.

2.4 Automatic testing

When the system is initialized, each timing channel is polled for a hardware identification number. This indicates the card type, and also identifies the optional components associated with it, such as the output driver and any precision timing modules. Each module socket has pins that are reserved for identification codes that are unique to each module type. Before a channel is initialized, its polled identification is compared against a number downloaded with the other initialization data. Hence, the possibility of unknowingly installing incorrect hardware is eliminated.

The output pulse drivers support two test functions: The pulse output is continuously checked against the logic input to detect short circuits or driver failures. The output is also applied to a coarse aperture (42 ns) sampling gate. The sampling trigger is generated centrally and routed to all channels via the P2 backplane. These are used to implement a logic analyzer which can be used to examine timing outputs from any control system terminal.

2.5 Mechanical design

The system is built on the VME 6U format, with both P1 and P2 backplanes extending over 21 card slots. Each crate has 16 slots reserved for timing cards and 4 slots for support functions. The output pulse drivers are on separate 3U VME cards which plug into the back of the P2 backplane connectors. This allows a variety of logic standards to be used with one type of timing card, and also simplifies repair and noise isolation.

Each crate has a self-contained power supply, and the entire system is operated on +5 VDC. The +/- 12 VDC supplies are used only by output pulse drivers which are required to issue non-TTL logic levels. The precision timing cards require forced air to operate, which is provided by a single-height fan pack mounted below the crate.

3 METHODOLOGICAL DETAILS

3.1 Phase shifter design

The essential component of the precision phase shifter is a linear phase detector. The phase of two clock signals is compared, and the phase of the slave clock is adjusted through a phase-locked loop to make the measured phase equal to the command phase. Conceptually, the phase detector consists of two latches: One latch is set by each of the clock signals, and when both are set, they are reset simultaneously. If the latches are implemented in CMOS logic, then the output levels are VDD and 0, and when the two outputs are subtracted in a differential amplifier, the result is VDD $\Delta\phi/(2\pi)$. The phase shift can be made independent of VDD by obtaining the phase command from CMOS data latches also: The 12 bit command is latched in a CMOS register and converted to analog by a R-2R ladder network. Hence the digital phase command N gives an analog result of VDD (N - 2048)/4096 in offset binary code. If this and the phase detector output are nulled in a bridge circuit, the resulting transfer function is $\Delta\phi = 2\pi$ (N - 2048)/4096, and is independent of voltage. This was implemented with a custom metal-film resistor network which combines the DAC ladder network with the necessary matched balancing resistors for the phase detector pulses.

The use of a precision bridge network means that the accuracy of the phase shifter only depends on the accuracy of the phase detector. Standard flipflops are generally not suitable for this use, however, since the internal timing paths are not constant. For example, a D-type flipflop may have different signal paths from reset to Q, depending on the level of the clock input. The design currently in use uses a high-speed logic array to anticipate the latch states and steer the correct signal edges through a CMOS multiplexer, thus emulating an ideal CMOS latch with constant

internal delays.

Assuming an ideal latch circuit, the limits on phase detection linearity are due to crosstalk and damping of transitions. Damping is significant because the logic transition waveforms subtract out of the phase detector function only if they damp out completely before the next transition. The damping time of the transitions depends mainly on the parasitic inductance of the integrated circuit package. With a 42 ns clock period, the damping time available in the worst case is 10 ns, and this is not sufficient even with surface-mount circuit packaging. The damping problem has been solved by only operating the phase detector on alternate cycles and doubling the gain.

Crosstalk between the two phase pulses is more problematical. The most significant coupling is through the power supply, and 6 decoupled VDD circuits are used in this design. However, there is a risk of capacitive coupling wherever there is a logic element shared between the two phase detector channels. This problem is compounded by general ignorance of the internal design of commercial integrated circuits. The author's opinion is that most of the nonlinearity displayed in figure 1 is due to this type of crosstalk. A custom integrated circuit implementation of an ideal linear phase detector is under design and nearing completion.

3.2 Pattern generator design

The pattern generator is straighforward except for the frequency multiplier. The multiplier uses a phase-locked loop design with a phase detector similar in principal to the one used in the phase shifter. This type of phase detector has the advantage that if a good differential amplifier is used to subtract the phase pulses, then the phase output is DC at zero phase, and hence free of subharmonics which would modulate the oscillator.

4 PERFORMANCE

The test results were obtained from two precision timing channels: One channel is used to phase-lock the low-noise crystal oscillator which serves as the CESR RF frequency source. The oscillator is 500 MHz, or 21 times the clock frequency. The second timing channel is used as a trigger for an HP54120A sampling oscilloscope. The data points are the mean and σ of 500 sample histograms taken of the RF waveform for various settings of the phase command to the oscillator. Hence, the phase mean is the transfer function of the first timing channel, but the noise is substantially from the second channel.

Figure 1 shows the RF phase as the phase shifter is advanced by whole wavelengths through a full clock cycle. An ideal linear transfer function would give identical phase for each measurement. Figure 2 shows the phase shifter wrapping from full count to zero, and hence that the gain is accurate to 12 bits and that the transfer function is monotonic over this transition.

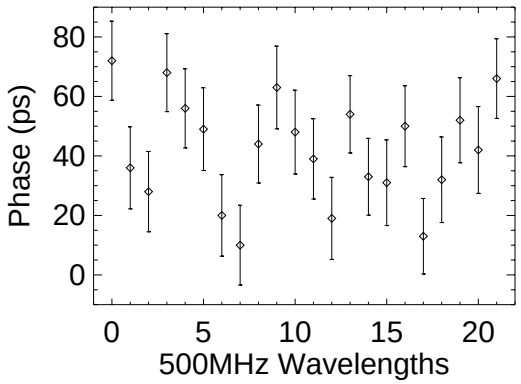

Figure 1: Phase of 500 MHz RF shifted by an integer number of wavelengths.

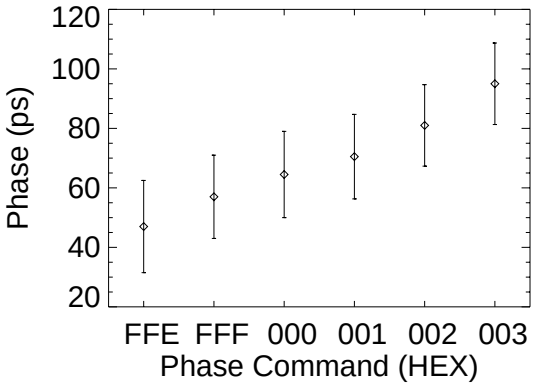

Figure 2: Phase shifter output across major carry transition.

5 ACKNOWLEDGEMENTS

Thanks to Meredith Williams for printed circuit design, and to Peter Quigley for integrated circuit design.

6 REFERENCES

[1] R. E. Meller, IEEE Trans. Nucl. Sci., Vol. NS-26, No. 3, p. 4152 (June 1979)