CMOS Programmable Delay Vernier

In the HP 9493 LSI test system, CMOS delay verniers replace the usual bipolar technology and are integrated with digital circuitry to produce a high-performance timing generator in a single monolithic CMOS VLSI formatter chip. This solution achieves bipolar-equivalent resolution, skew, and jitter performance with significantly lower power, cost, and circuit board space.

by Masaharu Goto, James O. Barnes, and Ronnie E. Owens

The HP 9493 is a mixed-signal LSI tester with a per-pin digital test resource architecture designed to offer the user test generation flexibility and ease of use. The timing vector generator is a key per-pin resource that generates and captures digital waveforms going to and coming from the device under test (DUT). Each DUT pin has an independent timing vector generator channel, allowing the user to select arbitrary timing, waveform format, and logic pattern without having to consider resource conflicts. Each timing vector generator consists of vector memory, formatter logic, and delay verniers, which are essentially high-precision programmable delay lines.

Traditionally, to meet the speed and timing performance requirements, high-precision delay verniers have been implemented with bipolar ECL technology.[1] However, these off-the-shelf delay verniers have the disadvantages of high cost and high power consumption. In addition, as separate packages, they increase board-level interconnections and package count. To address these problems in our system, the delay verniers were implemented with CMOS circuitry, allowing them to be integrated with the formatter logic into a single CMOS VLSI chip called ACCEL2 (Fig. 1). ACCEL2 was designed at the HP Integrated Circuits Business Division's Fort Collins Design Center and fabricated using HP's CMOS34 process.

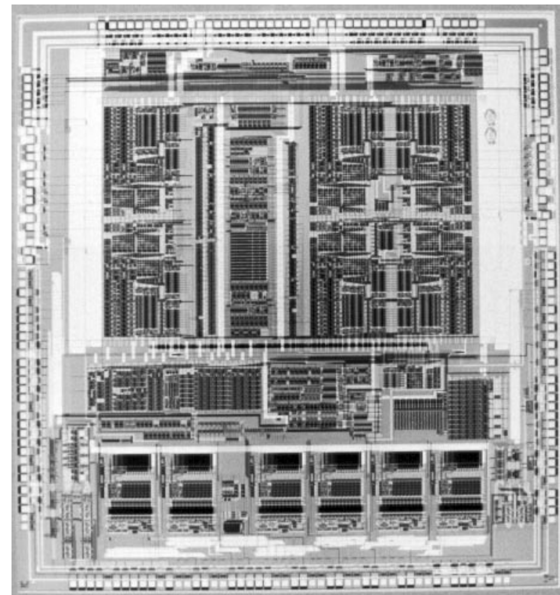

Fig. 1. ACCEL2 chip.

The benefits of this approach are evident, but there are many challenges in designing a CMOS timing system with the same level of performance as the bipolar counterpart. A stable, low-noise timing system is essential in a mixed-signal LSI tester. For example, testing state-of-the-art analog-to-digital converters (ADCs) and digital-to-analog converters (DACs) requires jitter performance of 20 ps rms or less. CMOS gate propagation delays are very sensitive to temperature and supply voltage, and CMOS implementations of standard circuits such as operational amplifiers typically have higher noise than their bipolar versions. Any use of CMOS in timing system applications must circumvent these limitations.

In previous attempts to use CMOS in this setting, feedback techniques such as phase-locked loops have been used to stabilize delays. However, such an approach often compromises jitter to an extent that is unacceptable in a mixed-signal environment. For ACCEL2, we developed a method of stabilizing delays while maintaining jitter at levels close to the theoretical minimum of the CMOS FETs. Our approach employs custom CMOS design in the delay verniers and on-chip dynamic power compensation to minimize the temperature delay sensitivity. In addition, the problems of supply and temperature variation were addressed at the system level. The combination of these techniques resulted in a timing vector generator module with performance equal to or better than a system with bipolar delay verniers, but with substantially lower cost and power dissipation. In a high-pin-count VLSI tester, these savings can make an important contribution to the performance and value of the product.

Test System Block Diagram

A block diagram of the HP 9493 mixed signal LSI test system is shown in Fig. 2 on page 43. The digital test subsystem consists of the sequencers, capture memory, timing vector generators, per-pin digital electronics, and per-pin dc electronics. The synchro pipe consists of special hardware to synchronize operation of the digital test subsystem and the analog test subsystems such as the arbitrary waveform generator and the waveform digitizer. The timing vector generator and the per-pin digital electronics are the per-pin digital test resources. Up to 256 digital channels can be installed.

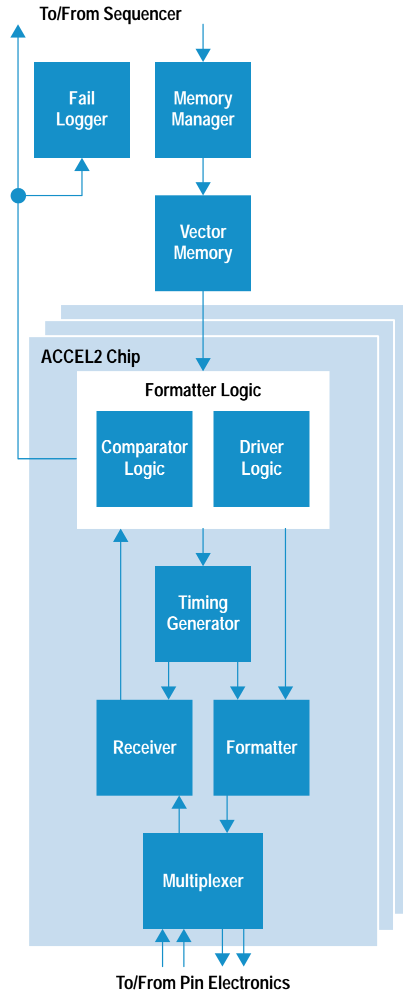

Fig. 2. Timing vector generator block diagram.

The timing vector generator block diagram is shown in Fig. 2. The memory manager, shared by eight timing vector generator channels, receives vector addresses from the sequencer and gives them to the vector memory. Test vectors stored in the vector memory are given to the formatter logic, which controls the formatter, receiver, and timing generators. A multiplexer is placed after the formatter and the receiver to support channel multiplexing and dual-test-head switching capability. The formatter logic, formatter, receiver, multiplexer, and timing generators are integrated in the ACCEL2 chip.

ACCEL2 Timing System

Fig. 3 is a simplified block diagram of the ACCEL2 timing generator. A coarse edge counter, implemented in standard CMOS digital logic, is used to generate a variable coarse delay ranging from 1 to 2^{14} periods of the system master clock. Because the counter terminal count output CE1 has significant jitter and skew, the edge is retimed by a flip-flop that we term the "last flip-flop." This block represents the final point along the path of the timing edge at which it has a fixed time relation to the system master clock. The last flip-flop is clocked by a clean version (MC) of the system master clock MCLK. After propagating the delay vernier, the delayed edge FE propagates to the drive or receive sections of the chip to control the timing of signals driven to the DUT or sampling of waveforms received from the DUT. There are six such timing generators on the ACCEL2 chip: four for the drive side and two for the receive side.

Fig. 3 illustrates a number of features that influenced the chip design. First, effective electrical isolation was needed between the standard digital part of the chip and the delay verniers and other timing-sensitive blocks such as the drive and receive circuitry. This required low-noise design techniques such as power supply isolation and control of rise and fall times in critical sections.

Second, a number of parasitic delays exist in any real implementation. Examples are the clock-to-output delay of the last flip-flop and the minimum delay through the time verniers. The sum of these delays makes up the intrinsic delay, which is specified in Table I. This delay must be stable.

Third, the design of the delay verniers is simplified if only one type of timing edge need be accurately delayed. In our case, it is the falling edge of the negative-going pulse from the coarse counter. As an illustration of this, the clock-to-output delay of the last flip-flop, which is a component of the intrinsic delay, is more easily controlled if only one transition is important. This is also true for the delay vernier itself.

Fourth, in this timing system, edges are separated by a variable time interval that depends on the setting of the coarse edge counter and the PCLK (period clock) frequency. This variation complicates the design of the delay verniers.

Fifth, since the coarse edge CE has a timing resolution of one master clock period, the delay range required of the delay vernier is related to this period. On ACCEL2, the required range is 16 ns.

Timing Specifications

The timing vector generator module must present highly accurate and precise waveforms to the device under test (DUT). This directly translates into stringent specifications for the on-chip timing system and delay verniers in particular. Table I gives the system requirements for the ACCEL2 timing system. These specifications are as stringent as those for the HP 9491A test system, which used external bipolar delay lines. Achievement of these specifications on a CMOS chip with a large amount of high-speed digital logic was the primary design challenge of the project.

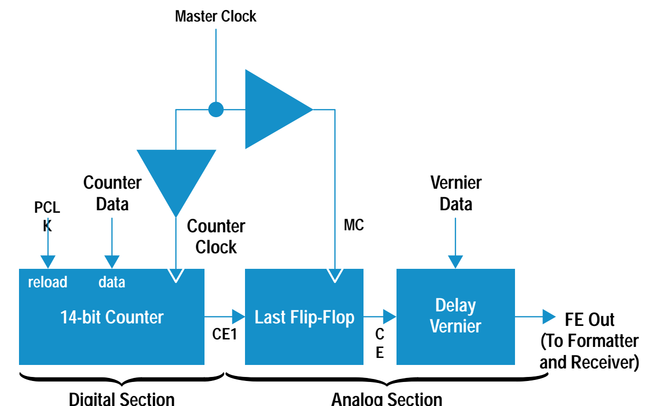

Fig. 3. Simplified block diagram of the ACCEL2 timing generator.

A few observations may be made regarding these specifications. The resolution shown is the user-programmable resolution, and linearity is given in terms of this step size. Because of the nature of the calibration scheme, as discussed later, the delay verniers have a minimum time step size equal to one-half the user-programmable resolution.

Table I
System Requirements for the Delay Vernier

Specification	Absolute	Relative	Typical CMOS
Span	16 ns		
Resolution	8-bit or 62.5 ps		
Linearity	< 1 LSB		
Skew:			
Temperature	±240 ps/5°C (ambient)	±0.16%/°C	±0.3%/°C
Power Supply	±10 ps/ 50 mV	±0.07%/ 100 mV	±0.8%/ 100 mV
Jitter	10 ps rms	0.03% rms	≈0.1% to 1%
Pulse Symmetry	not required		
Intrinsic Delay	< 8 ns		

In Table I and elsewhere in this article, skew is defined as the variation in delay caused by a change of any external influence, including environmental variations. Jitter is defined as delay variation for a series of edges propagating down the line with all external influences constant. The third column in the table shows the skew and jitter specifications as a percentage of the longest time-sensitive timing path on the ACCEL2 chip. This path occurs on the drive side. Since the magnitude of skew and jitter scales with the length of the path, this allows the specifications to be compared with typical performance seen on standard CMOS designs. This data is given in the final column; it is seen that the required skew temperature dependence and jitter performance for the ACCEL2 chip are significantly better than standard CMOS design practices would yield.

Intrinsic delay is simply the delay through the delay vernier with the delay setting at minimum. It is desirable that this delay be as short as possible, since a longer intrinsic delay can complicate the design or compromise the performance of other blocks in the system. However, achieving a shorter intrinsic delay requires greater on-chip power. This is true of other specifications in the table as well; for example, smaller step size requires more power.

Delay Vernier Architecture

The first step in the design of the delay vernier was to select a delay line architecture. A number of delay line implementations have been reported. These include ramped comparators,[2] charge-coupled devices,[3] ECL gate arrays,[1] and multiplexed and tapped delay lines. Resolution, range, jitter, and skew requirements eliminated most of these; the final choice came down to either a binary-weighted multiplexed delay line or a tapped delay line. These two designs are shown in Fig. 4.

Although the multiplexed delay line architecture requires less power and silicon area, it would require the line to be made up of elements of different delays (specifically a binary-weighted series of elements) with inferior skew and linearity performance compared to the tapped line. Because of device mismatches, a binary-weighted multiplexed delay line may have nonlinearities that cannot be corrected by a simple calibration. This is especially possible if there are gaps in the delay-versus-timing relationship because of mismatches. A tapped delay line scheme greatly reduces the effects of device mismatch and essentially avoids these problems entirely. There are other problems with the multiplexed delay line. Low intrinsic delay is difficult to achieve, since for an 8-bit line the minimum delay would need to pass through eight multiplexers. For these reasons, we chose the tapped delay line structure, Fig. 4b.

As shown in the figure, the tapped delay line consists of two types of delay elements: coarse and fine. The coarse elements are calibrated to an identical fixed delay of approximately 2 ns during an initial calibration procedure. An eight-input wired-OR circuit selects the edge at various points along the line, providing a timing adjustment to a resolution of one coarse element delay. The multiplexer is designed so that its propagation delay is nearly identical through all taps. This required putting an extra dummy load on the tap at the end of the line to equalize the capacitive load on all tap inputs. Despite these measures, however, an actual chip will have

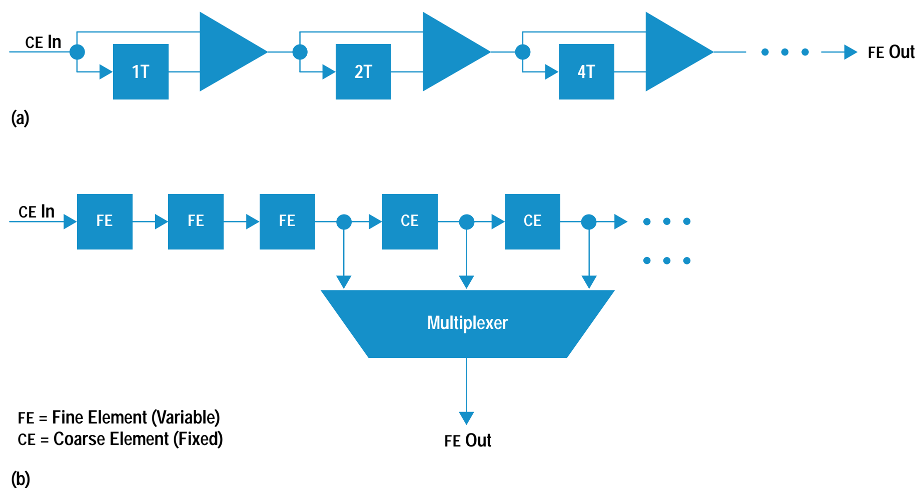

Fig. 4. (a) Multiplexed delay line. (b) Tapped delay line.

Theoretical Approach to CMOS Inverter Jitter

Metal-oxide-semiconductor or MOS technology is known to be inferior in noise performance to bipolar or junction FET (JFET) technology. For example, a CMOS operational amplifier has two to three orders of magnitude worse noise performance than a bipolar or JFET-input operational amplifier. How about jitter? Jitter can be defined as a timing uncertainty or noise. If the device is so noisy, won't the noise affect its jitter performance? We tried to clarify this question by a theoretical calculation.

Major noise sources in a CMOS device are flicker noise and Johnson noise. To the simple CMOS inverter circuit shown in Fig. 1, we applied a noise model and analyzed the thermal jitter.

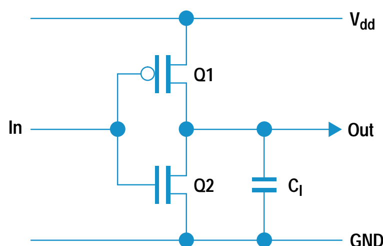

Fig. 1. Simple CMOS inverter with load capacitance C_l.

The jitter is evaluated at the mid-transition point of the output of the inverter (Fig. 2). We assume that the input has already reached the V_{dd} voltage. Q1 completely turns off and Q2 discharges the load capacitance C_l.

Fig. 3 shows the equivalent circuit of the flicker and Johnson noise model, where v_j represents the Johnson noise of the equivalent resistance of Q2 and v_f represents the input flicker noise.

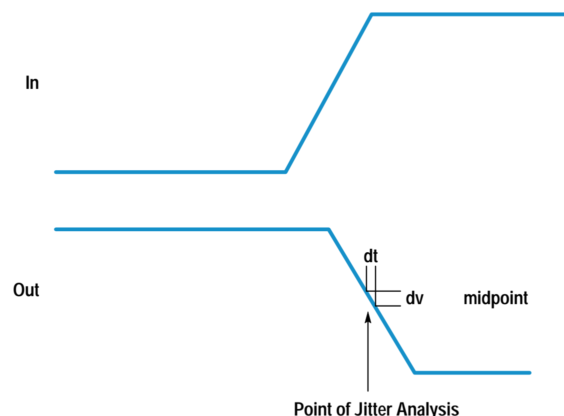

Fig. 2. Jitter evaluation point.

slight mismatches because of manufacturing variations. These can be calibrated out during coarse element calibration, as discussed later.

Finer resolution is provided by three fine delay elements. These have delay that can be varied over approximately a 1-ns range in 31-ps steps by turning on or off internal capacitors, as described later. On the ACCEL2 chip, the same element design is used for the coarse and fine elements. The element has a nominal delay of about 2 ns with the 5-bit digital delay control set at a default value. In the coarse elements, delay variability is used only for calibration.

Two design techniques were key to achieving the required linearity and skew performance: (1) use of essentially identical delay elements throughout the line and (2) use of "thermometer decoding" in control of the line. The benefits of these techniques will be discussed in more detail below. Thermometer decoding is defined as follows: as the input delay setting of the line is increased, delay elements are added to but never removed from the delay path. This guarantees monotonicity of the uncalibrated delay as a function of digital setting and improves uncalibrated linearity, thus making simplified calibration schemes possible. This characteristic is evident in the coarse delay part of the line, but is also employed in the internal structure of the fine elements.

Delay Vernier Element Design

The basic delay element is shown in Fig. 5. It consists of an input inverter, programmable banks of capacitors, and an output inverter. In its quiescent state, the element is presented with a high voltage at the input. An edge propagates through the element when the input voltage undergoes a high-to-low transition. The delay of the edge through the element is determined by the number of capacitors that are turned on and by the bias voltage applied to the gate. When the internal node voltage reaches the switching threshold of the output inverter, the output makes a high-to-low transition, which is applied to the next element in the delay line.

The bias voltage is generated and adjusted with a DAC, and is used to calibrate the delay such that process variations from chip to chip are nulled. The DAC also compensates for delay variations caused by temperature and supply voltage variations.

The capacitors are programmed using a 5-bit digital input. The higher-order capacitor banks are accessed in thermometer-decoded fashion while the lower-order banks are programmed in a straight binary fashion. The breakpoint in this decoding was chosen based on the expected variation in the capacitor elements. This approach minimizes the differential nonlinearity of the delay as a function of the digital capacitor setting. The capacitor array is sized such that the nominal delay through each element is 2 ns. The minimum time step is 31 ps.

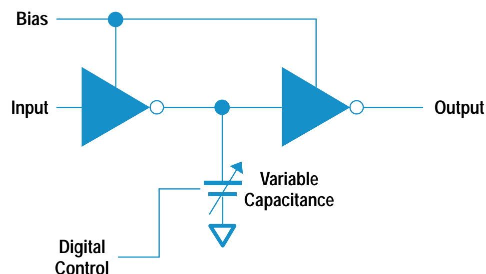

Fig. 5. Basic delay element.

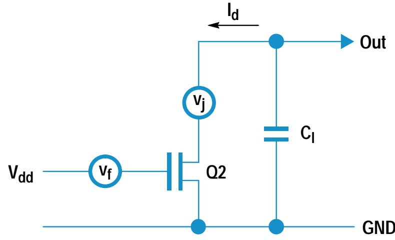

Fig. 3. Noise jitter model.

We calculated the voltage noise V_n and the slew rate dV/dt at the output and estimated the jitter T_{jit} as

$$T_{jit} = \frac{V_n}{dV/dt}.$$

Fig. 4 shows the result of calculations for various gate lengths L and gate widths W. T_{jit} does not exceed 0.3 ps rms even with the worst-case device. A series of 50 such inverters will only produce 2.1 ps rms jitter. While not perfectly accurate, this analysis provides valuable insight.

Another jitter source is coupling noise. This is considered to be a dominant source of jitter in standard CMOS design. In the ACCEL2 design, we avoided subdividing the test period or PCLK signals because subharmonics cause periodic jitter. When the ACCEL2 chip is generating a periodic waveform, all coupling sources are phase coherent with the timing edges. This means we will see the identical coupling waveform at every test cycle and it won't cause any jitter.

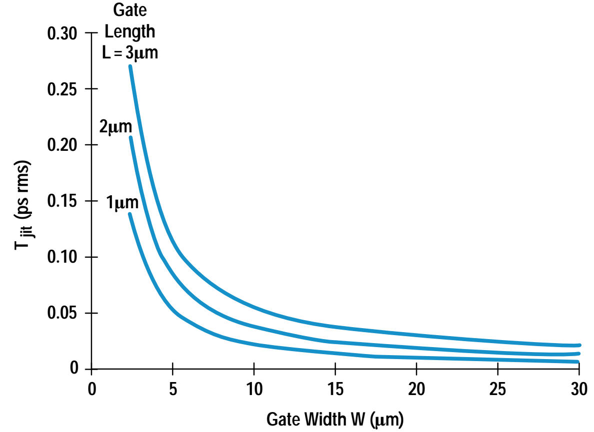

Fig. 4. Calculated CMOS inverter jitter.

Masaharu Goto
Design Engineer
Hachioji Semiconductor Test Division

Support Circuitry

As already discussed, a tapped delay line architecture was chosen for linearity considerations. As seen in Fig. 6a, the delay line contains twelve delay elements and supporting circuitry. Three delay elements are used for fine adjustment of the overall delay. During calibration, the digital settings required to generate a given delay are stored in a look-up table RAM. There is one calibration look-up table RAM per delay vernier. Three elements are used in the fine section to guarantee a span of at least 2 ns corresponding to the delay of one coarse element. The output of the third fine delay element is presented to seven coarse delay elements whose

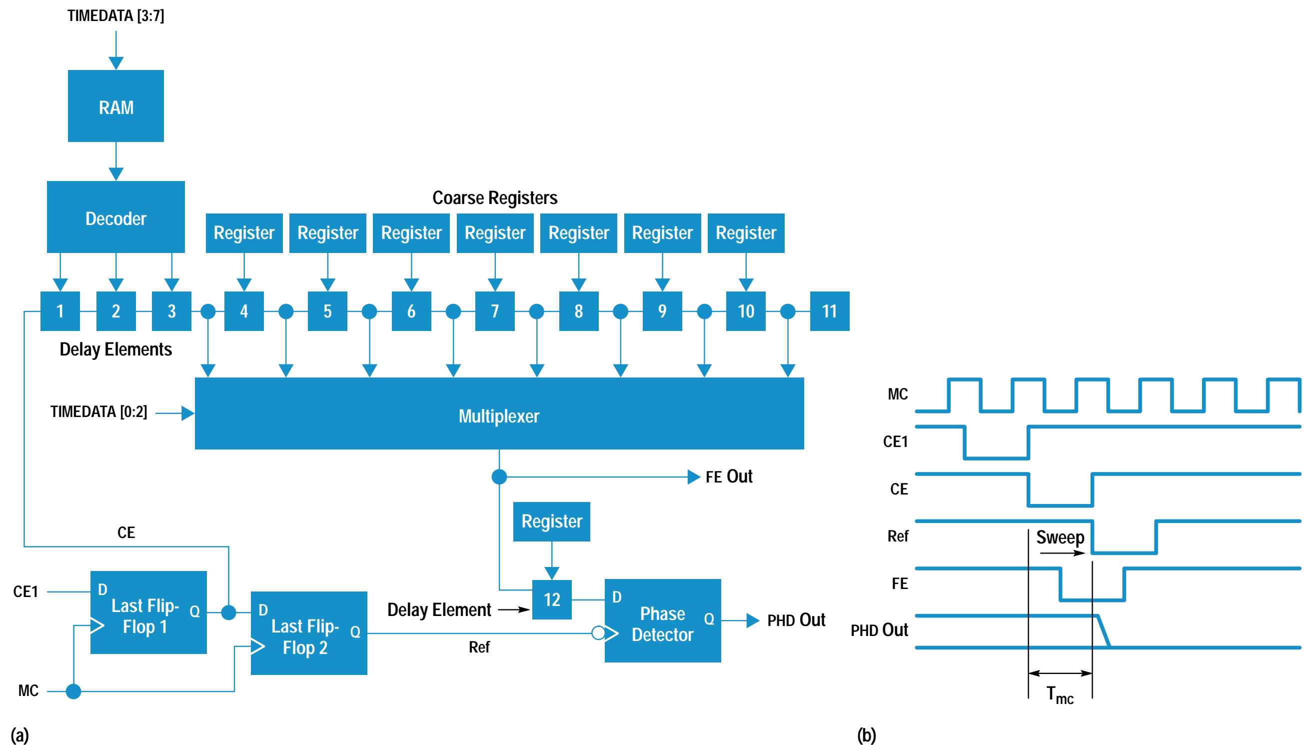

Fig. 6. (a) ACCEL2 delay vernier. (b) Calibration timing waveforms.

digital settings are adjusted during calibration to get a delay as close to 2 ns as possible. A dummy element at the end of the line is used to load the final element.

The delay line is driven by a flip-flop with stabilized propagation delay. This last flip-flop is driven by a so-called coarse edge signal generated by a counter in the digital section. The clock for the last flip-flop is a master clock that is buffered and distributed exclusively to the delay vernier section of the chip. This approach removes jitter on the coarse edge caused by noise generated in the digital section of the chip. To minimize its own contribution to jitter and skew, the last flip-flop is designed to have a short clock-to-output delay.

The multiplexer selects one of eight taps along the delay line to be directed to the output. The multiplexer elements are simple two-input NOR gates whose outputs are connected in a wired-OR arrangement and buffered by a final inverter. This circuit is designed to have minimum propagation delay without unduly loading the basic delay elements. Using this arrangement, the propagation delay varies from some minimum intrinsic delay, T_i, to T_i + 16 ns adjustable in 31-ps increments. Five bits of timing data are used to select one of the 32 fine delay look-up table entries and three bits of timing data are used to control the multiplexer setting. T_i is given by

$$T_i = 3T_{idl} + T_{imux}$$

where T_{idl} is the intrinsic delay of the basic element, and T_{imux} is the delay through the multiplexer. The multiplexer delay in the final design is 1.75 ns and the delay through the basic delay element with a digital setting of 0 is 1.75 ns. Thus the overall intrinsic delay of the line is 7 ns.

Used only during calibration are an additional delay element and a special flip-flop called the phase detector, which is designed for minimum setup time and is used to detect a match between the master clock frequency and the delay through the delay line. The operation of this additional calibration circuitry is described below.

Calibration Scheme

The HP 9493 system uses a high-quality frequency synthesizer as a master clock source. A master clock frequency from 4 kHz to 128 MHz can be programmed in 1-microhertz steps with extremely low jitter and high frequency stability. This translates to subpicosecond master clock period (T_{mc}) programmability. Therefore, we decided to use the master clock period T_{mc} as a timing reference for linearity calibration. A second last flip-flop delays the CE signal by one master clock period to create Ref. Because last flip-flop 1 and last flip-flop 2 are identical and are toggled by the same MC, the time interval between CE1 and Ref is equal to T_{mc}, which we can arbitrarily control from 7.8 ns to 250 μs (Fig. 6b). As described above, the intrinsic delay of the delay vernier is 7.00 ns. Because 7.8 ns is the shortest controllable time interval, we needed to put another delay element after FE. This element adds another 2 ns to the calibration path, resulting in a 9.00-ns minimum delay setting. This is within the range of control of the master clock period T_{mc}. The twelfth delay element is only used for linearity calibration. To prevent an increase in intrinsic delay and timing skew, FE does not go through this element.

The linearity calibration process for each individual delay vernier consists of three parts: bias DAC calibration, fine delay look-up table calibration, and coarse register calibration.

First, all delay elements are programmed to a default value. The bias DAC setting is then calibrated to adjust each element delay to approximately 2 ns. Process variation is roughly calibrated out by this step. Secondly, the fine delay look-up table RAM is calibrated for addresses 0 to 31. At RAM address 0, the master clock frequency is set to 111 MHz (T_{mc} = 9.0 ns) and the RAM data, which is applied to the fine delay elements, is incremented until the delay through the three fine elements equals this value. At the next RAM address, T_{mc} is incremented by 62.5 ps and the RAM data value is again incremented to cause the vernier delay to match the clock period. This is repeated 32 times to calibrate all RAM addresses with 62.5-ps resolution. Thirdly, each coarse delay element is calibrated, starting with the first coarse element. RAM address 0 is selected, the second multiplexer tap is selected, and T_{mc} is set to 11.0 ns, which is 2 ns more than the value used to calibrate RAM address 0. The coarse element register value is incremented until the delay matches this clock period. All the remaining coarse elements are calibrated in a similar manner with the master clock period increased by 2 ns for each successive element. The coarse element calibration compensates for slight variations in the multiplex delay through various taps in addition to calibrating the delay element itself.

The ACCEL2 chip contains a calibration sequencer block called the calibration logger, which supervises the per-pin parallel timing calibration. The calibration logger increments digital timing data for the three phases of calibration described above until the phase detector output of the particular time vernier changes state. At the setting that causes the phase detector to change, the calibration logger stops incrementing the timing data and logs the value of the digital input.

During each calibration operation, the calibration logger can average up to 256 pulses to prevent occasional noise from terminating the measurement prematurely. In addition to calibrating the delay verniers, the calibration logger performs other system calibration and deskewing operations. After all calibration logger operations are completed, the tester controller reads the logged value to get the measurement results. Since all timing vector generator channels in the system can operate in parallel, computer overhead is small and we can perform full linearity calibration of a 256-pin system within 30 seconds.

Dynamic Power Compensation

While the power dissipation of an ECL device remains approximately constant for all operating conditions, a CMOS device changes its power consumption drastically between the static and dynamic states. This is because the power required to charge and discharge internal capacitances when the nodes are toggling at a high frequency is much greater than the standby leakage power. This dynamic power variation is a problem when integrating precision analog circuits and a large amount of digital logic onto a single CMOS VLSI chip. Operation of the analog circuit is often sensitive to temperature, and junction temperature changes caused by

dynamic power variation can be a major source of inaccuracy. The ACCEL2 chip contains 20,000 gates of CMOS logic along with precision delay verniers. The dynamic power variation inherent in the logic cannot be neglected. For a given package thermal resistance, the junction temperature was estimated to vary by 4.5°C because of dynamic power variation. Even with the reduction of delay temperature sensitivity afforded by the custom time vernier design, this variation will unacceptably degrade the system timing accuracy. The dynamic power compensation circuit was developed to solve this problem.

Generally in CMOS logic design, power dissipation is almost proportional to clock frequency. Two on-chip clock networks dominate the ACCEL2 logic operation: the master clock MCLK and the period clock PCLK. From a previous design, we found that the dynamic power of the chip could be reasonably predicted from the frequency of these two clocks. The MCLK frequency is programmed by the tester controller and stays constant during critical operation, so the MCLK dependent power can be calculated by the tester controller. PCLK is an external data input latched by MCLK; it initiates a test period. Because research on customer needs indicated that the ability to change the length of the test period on the fly would be a useful feature, the PCLK frequency can change at any time during critical operations. Therefore, the PCLK dependent power cannot be estimated by the tester controller.

To compensate for dynamic power variation without any increase in total power dissipation, we designed a state machine that roughly monitors the amount of activity in the logic circuitry and dynamically turns an on-chip power compensation heater on and off. Fig. 7a is a conceptual schematic diagram of the dynamic power compensation circuit and Fig. 7b illustrates the principle of operation. When the test system is not in use and MCLK is not running, MCSTAT is reset (MCSTATN = 1). At this time, the X register value is fed to the heater. The X register is programmed by the tester controller as follows:

$$X = P_{pcmax} + P_{mcmax},$$

where P_{pcmax} is the PCLK dependent power at the maximum PCLK frequency f_{pcmax} (not shown in Fig. 7b) and P_{mcmax} is the MCLK dependent power at the maximum MCLK frequency f_{mcmax}. At this moment, the dynamic power dissipation of the ACCEL2 chip is zero. Therefore, the sum of the dynamic power and the heater power is simply $X = P_{pcmax} + P_{mcmax}$. The tester controller sets the Y and Z values such that:

$$Y = \left(P_{pcmax} + P_{mcmax}\right)\frac{f_{mcmax} - f_{mc}}{f_{mcmax}}$$

$$Z = P_{pcmax}\frac{f_{mc}}{f_{mcmax}}$$

When MCLK starts running at frequency f_{mc}, MCSTAT detects this condition and MCSTATN becomes 0. Because PCLK is not toggling, PCSTAT is reset (PCSTATN = 1) and Y+Z is fed to the heater. Dynamic power in this state is shown as P_{mcdyn} in Fig. 7b. The sum of the dynamic power and the heater power is equal to $P_{pcmax} + P_{mcmax}$.

When PCLK transitions, PCSTATN becomes 0 for N MCLK cycles, which is equal to the minimum PCLK interval, and then returns to 1. During this period of N MCLK cycles, the Z value is

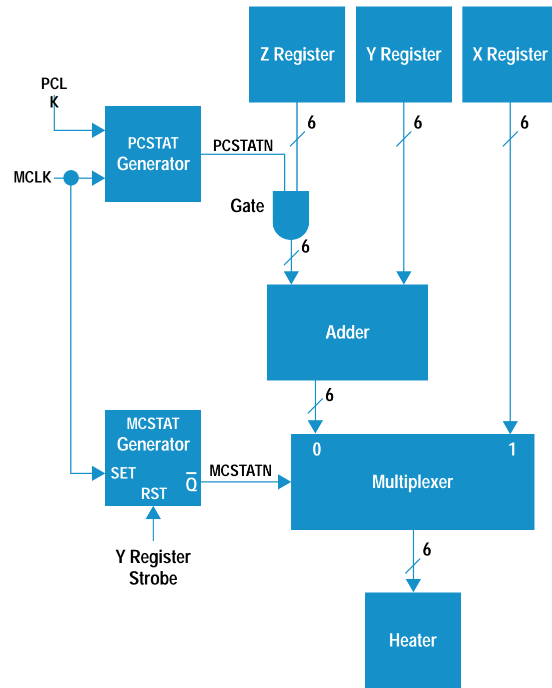

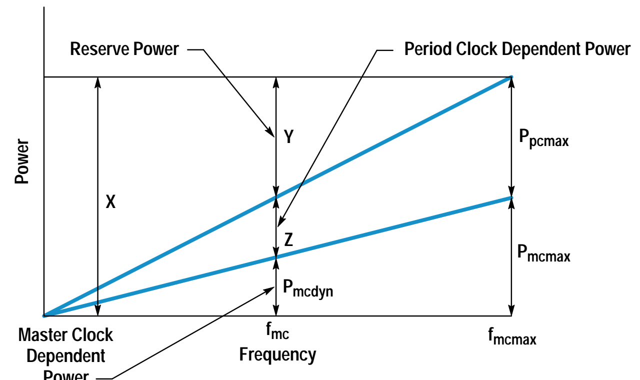

Fig. 7. (a) Dynamic power compensation circuit. (b) Power relationships.

gated off. This decreases the heater power by an amount equal to the dynamic power consumed by a single PCLK cycle. Therefore, the sum of the dynamic power and the heater power remains $P_{pcmax} + P_{mcmax}$ regardless of the PCLK frequency.

When MCLK is stopped by the tester controller, MCSTAT is reset immediately afterwards so no significant power glitch will occur.

In this way, the total power consumption of the chip is kept constant. This scheme greatly improves system timing accuracy.

Measured Delay Vernier Performance

Fig. 8 shows the raw (uncalibrated) integral nonlinearity of the three fine elements measured on the ACCEL2 chip. The

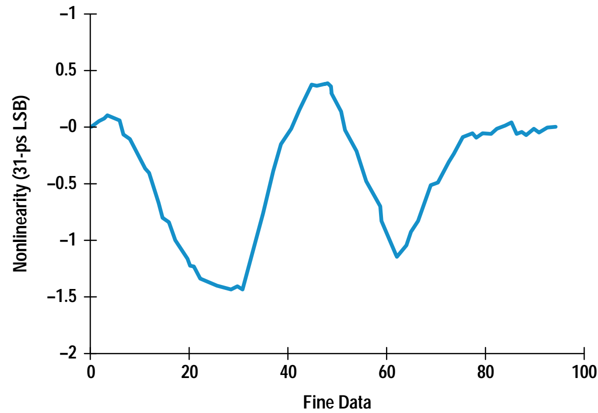

Fig. 8. Uncalibrated fine delay element linearity.

fine element nonlinearity is approximately +1.5 LSB of 31-ps raw resolution.

Fig. 9 shows calibrated delay vernier linearity. This curve covers the entire digital input range of 0 to 255, with a corresponding delay range of 0 to 16 ns. The nonlinearity is expressed in terms of the system LSB of 62 ps. As the curve shows, the linearity calibration guarantees less than ±1 LSB of integral nonlinearity in the system timing resolution over the entire delay range.

Jitter measurement was done using an HP 54121T digitizing oscilloscope with the delay line set at the maximum value of 16 ns. The measured jitter was 3.3 ps rms. Removing the intrinsic jitter of the HP 54121T (evaluated to be 1.2 ps rms) resulted in an estimated ACCEL2 jitter of less than 3.1 ps rms. This is about 0.01% of the path delay, which is one to two orders of magnitude better than typical digital designs and better than the specification by a factor of three (see Table I).

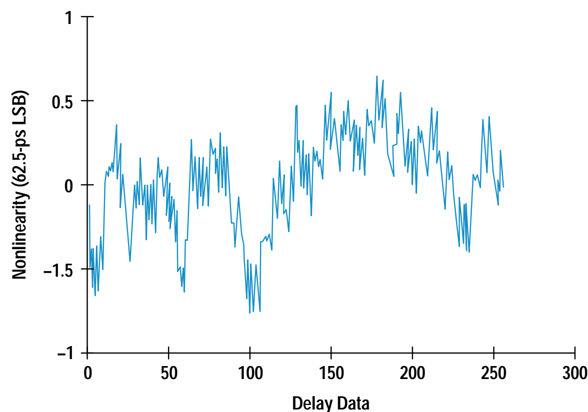

Fig. 9. Calibrated delay vernier linearity.

Table II shows the measured temperature coefficient of the propagation delay through the delay vernier circuitry and CMOS formatter block. As the table shows, the ACCEL2 delay vernier is about three times better than the custom CMOS formatter circuit. This temperature stability performance is equivalent to bipolar time verniers. In the ACCEL2 chip the overall temperature coefficient was measured as 30 ps/°C.

Table II
Temperature Coefficient of Propagation Delay

CMOS Formatter	0.15%/°C of path delay
Delay Vernier	0.058%/°C of path delay
Total Critical Path	30 ps/°C

Measurements were also made of the power supply dependence of the delay of a critical timing path . The propagation delay was extremely stable as the V_{dd} voltage changed.

Conclusion

Integration of delay verniers with formatter logic in the custom VLSI chip ACCEL2 was the key to achieving a low-cost, low-power LSI test system design. By moving the delay lines on-chip, the cost and power of the timing system were reduced by nearly an order of magnitude while at the same time providing ECL-equivalent timing performance.

Acknowledgments

The design was accomplished by a cooperative effort of the Integrated Circuits Business Division Fort Collins Design Center and the Hachioji Semiconductor Test Division. We would like to thank Albert Gutierriz and Chris Koerner for designing a major part of the delay vernier. Gary Pumphrey of the Colorado Springs Division contributed key ideas used in the design of the delay element. Keith Windmiller managed the project and provided guidance and direction. Barbara Duffner designed the high-performance CMOS formatter and kept track of many other details during the course of the project. Koh Murata and Kohei Hamada designed the surrounding system of the chip. We would also like to thank Larry Metz and Charles Moore for their invaluable consulting for the advanced CMOS analog design.

References

1. T.I. Otsuji and N. Naumi, "A 3-ns Range, 8-ps Resolution Timing Generator LSI Utilizing Si Bipolar Gate Array," *IEEE Journal of Solid-State Circuits,* Vol. 26, no. 5, May 1991, pp. 806-811.

2. T. Ormond, "Delay Lines Take on Timing Tasks," *EDN Magazine,* December 1991, pp. 108-112.

3. C.F.N. Cowan, J.W. Arthur, J. Mavor, and P.B. Denyer, "CCD Based Adaptive Filters: Realization and Analysis," *IEEE Transactions on Acoustics, Speech, and Signal Processing,* Vol. ASSP-29, no. 2, April 1981, pp. 220-229.