

Precise Delay Generation Using Coupled Oscillators

John G. Maneatis and Mark A. Horowitz

Abstract—A new delay generator based on a series of coupled ring oscillators has been developed; it produces precise delays with subgate delay resolution for chip testing applications. It achieves a delay resolution equal to a buffer delay divided by the number of rings. The coupling employed forces the outputs of a linear array of ring oscillators oscillating at the same frequency to be uniformly offset in phase by a precise fraction of a buffer delay. The buffer stage used in the ring oscillators is based on a source-coupled pair and achieves high supply noise rejection while operating at low supply voltages. Experimental results from a 2- μm *N*-well CMOS implementation of the delay generator demonstrate that it can achieve an output delay resolution of 101 ps while operating at 141 MHz with a peak error of 58 ps.

I. INTRODUCTION

PRECISE delay generation is a necessary function in state-of-the-art single-chip testers [1]. When testing digital integrated circuits, it is necessary to supply digital waveforms as input, which requires accurate delays referenced to some clock signal. The delay resolution needed in order to accurately measure parameters such as setup and hold times is often finer than that of an intrinsic gate delay of the device under test. Presently, this fine delay control is obtained using higher speed integrated circuit technology for the tester than for the device under test. A more cost-effective approach would be to limit the IC technology used for the tester to one no more advanced than that used for the device under test. However, generating precise delays with significantly finer resolution than an intrinsic gate delay has been difficult to achieve in this manner. This paper describes an array oscillator comprised of a series of coupled ring oscillators that can achieve a delay resolution equal to a buffer delay divided by the number of rings [2]. Using a 2- μm *N*-well CMOS technology, a delay resolution of 101 ps is achieved with a peak error of 58 ps at a frequency of 141 MHz.

Because an array oscillator is based on a series of ring oscillators, this paper will begin with a description of precise delay generation using ring oscillators. The concept of a ring oscillator will then be extended to an array oscillator in Section III. This section will also include a description of the general issues related to the operation and implementation of an array oscillator. The generation of precise delays requires low-noise buffer stages to prevent an effective loss of precision due to jitter in the output signals. Section IV will describe the buffer circuit design used in an implementation of the array oscillator for high supply noise immunity while being able to operate at

Manuscript received May 19, 1993; revised July 24, 1993. This work was supported by the Advanced Research Projects Agency under Contract N00039-91-C-0138.

The authors are with the Center for Integrated Systems, Stanford University, Stanford, CA 94305.

IEEE Log Number 9212118.

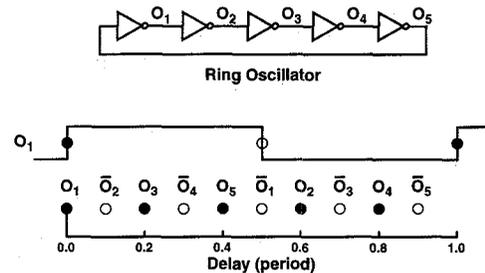


Fig. 1. Phase relationship among ring oscillator outputs for a ring with five buffers.

low supply voltages. Another issue critical to the overall precision of the array oscillator is the method in which the various outputs are read from the array core. Section V will present the output channel circuits and related implementation issues. The paper will also present experimental results demonstrating the ability of an array oscillator to produce precise delays with a resolution equal to one seventh of a buffer delay.

II. RING OSCILLATOR DELAY GENERATORS

Precise delays can be generated with ring oscillators by taking advantage of the symmetry in a ring. Since all buffer stages are identical, the relationship between a buffer delay and the period is set by the number of stages. Once phase locked to an established clock period, the delay between the ring outputs will be precisely known. Different delays can be generated by accessing different ring outputs with multiplexers. Fig. 1 illustrates the phase relationship among individual buffer outputs. Rising transitions are indicated by dots, and falling transitions are indicated by circles. If a ring contains five differential buffers as shown, utilizing both inverted and non-inverted outputs, ten different output phases are available, which uniformly span the output period. The limitation of using ring oscillators as delay generators is that the delay resolution is limited to a buffer delay. The only way to add more output phases is to add more buffers, which in turn decreases the maximum oscillation frequency. Thus, the delay resolution remains unchanged. Ideally, it would be desirable to be able to add more buffers to a ring-like structure without changing the oscillation frequency and thereby increase the delay resolution to a fraction of a buffer delay.

III. ARRAY OSCILLATOR

An array oscillator is a structure based on a series of coupled ring oscillators. By coupling several rings together it is possible to break the dependence of the oscillation frequency on the number of buffers. With the oscillation frequency

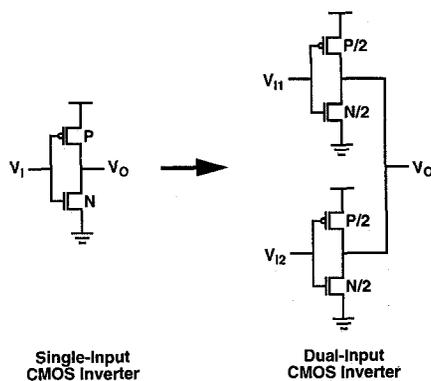


Fig. 2. Example of a dual-input inverting buffer. A CMOS inverter can be converted into a single-ended dual-input buffer by shunting the outputs of two half-sized CMOS inverters.

unaffected, the delay resolution can be increased simply by adding more rings. The basic idea is to force several rings oscillating at the same frequency to be uniformly offset in phase. The coupling between the rings that generates this uniform spacing is the key to the design of the array oscillator, as it causes corresponding outputs from each ring to divide a buffer delay into several equal delay intervals.

A. Dual-Input Inverting Buffer

To couple rings together, array oscillators utilize a new kind of inverting buffer. This buffer is similar to a single-input inverting buffer except that it has two inputs of the same polarity, one referred to as the ring input and one referred to as the coupling input. An example of such a dual-input buffer is shown in Fig. 2. It is constructed from a static CMOS inverter by shunting the outputs of two half-sized static CMOS inverters. Both the ring and coupling input transition times determine when the output transition will occur. In an array oscillator, the delay between the ring and coupling input transitions is always small, so the transitions overlap to some extent. Although neither transition in isolation may be able to cause a complete transition at the output, the overlapping transitions allow both the ring and coupling inputs to affect the time of the output transition. The coupling input transition can advance or retard the output transition relative to the ring input transition. Early coupling inputs reduce the buffer delay, while late coupling inputs increase the buffer delay.

In general, a dual-input buffer can be made from any single-ended or differential inverting buffer by splitting the input devices in half, with the inputs for each half of these devices forming the two new inputs. There are no requirements on the linearity or strength of the coupling input for the dual-input buffers used in an array oscillator. The only requirement is for the coupling input to have a monotonic effect on the buffer delay. Section IV will describe in detail the dual-input differential buffer used in the implementation of the array oscillator.

B. Array Structure

An array oscillator is structured as a two-dimensional array of dual-input inverting buffers as shown in Fig. 3. Rings extend

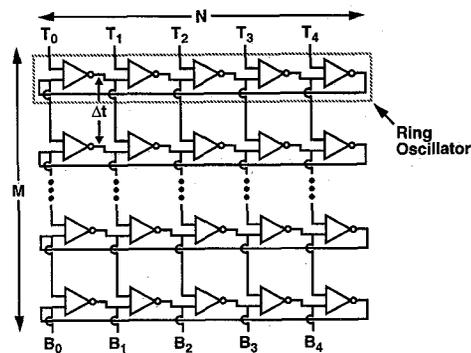


Fig. 3. Array oscillator structure.

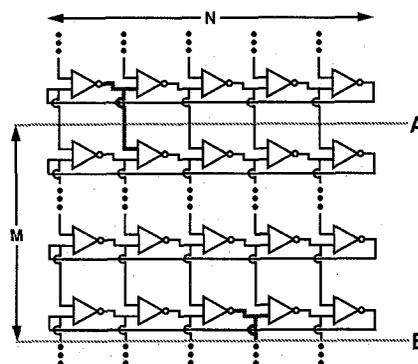


Fig. 4. Infinite array oscillator.

horizontally and are coupled together vertically through the coupling inputs. The top array nodes are connected to the bottom array nodes in a unique manner to form a closed structure. The coupling inputs and closing connections are critical to the operation of the array oscillator. The coupling inputs force the rings to oscillate at the same frequency while maintaining a precise phase relationship to one another. The closing connections force the delay spanned by the rings to equal some multiple of the buffer delay.

C. Array Operation

The operation of the array oscillator can be most easily explained with a simplified structure—an infinite series of coupled rings, as shown in Fig. 4. Suppose all rings are oscillating in phase so that the phase difference between the ring input and coupling input of each buffer is zero. The delays of all buffers will then be the same, so each ring will oscillate at the same frequency. Thus, the phase difference between the ring and coupling inputs of all buffers will remain zero and not change with time, leading to a consistent state for the array. Although this state is consistent, it is not very interesting since the outputs from each ring are exactly aligned in phase to the corresponding outputs from all other rings, leading to a delay resolution no better than that of a simple ring oscillator.

To improve the delay resolution of this structure, suppose instead that there is a fixed phase difference between the ring and coupling inputs of all buffers. The delays of all buffers will still be the same, since they experience an identical phase difference between their ring and coupling inputs. With equal buffer delays, the oscillation frequency of each ring will then

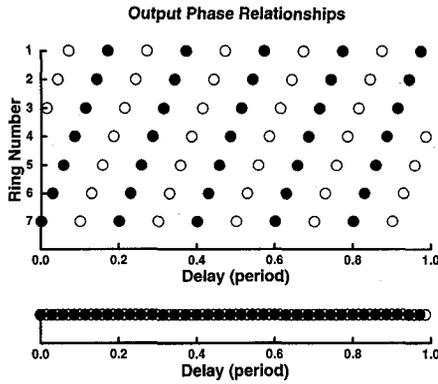


Fig. 5. Phase relationship among array oscillator outputs for an array with seven rings, each with five buffers.

also be the same. Thus, the phase difference between the ring and coupling inputs of all buffers will remain fixed with time, again leading to a consistent state. This state is more interesting because the outputs of adjacent rings will be skewed by a fixed delay so that the outputs at a particular ring position will uniformly span delays in time.

After some number of rings M , the phase of the ring outputs could identically match those of a previous ring, but not at the same ring positions. Suppose that the phases of the ring outputs along line A in Fig. 4 identically match those along line B, but shifted to the right by two buffers. The two highlighted nodes will then be at the same phase. This situation would be the same as if the ring inputs along line A connected directly to the ring outputs along line B, but shifted to the right by two buffers, thus forming a closed structure. Closing the array with a non-zero buffer shift forces a phase difference between the top and bottom nodes of the array. Because of the symmetry in the array, a phase difference forced at the boundary of the array will in turn force a small uniform phase shift between adjacent rings.

Suppose that the array in Fig. 3 is closed as described above, where top array nodes T_i connect to bottom array nodes B_{i+2} . In the simplest case the phase difference across all corresponding ring nodes will uniformly span, from the top to the bottom of the array, -2 buffer delays in phase. The phase difference between corresponding nodes in adjacent rings is -2 buffer delays divided by the number of rings. The plot in Fig. 5 illustrates phase relationship among the individual buffer outputs for such a closed array containing seven rings, each with five differential buffers. With all buffers considered, utilizing both inverted and noninverted outputs, 70 different output phases are available that uniformly span the output period with a resolution of one seventh of a buffer delay as shown.

D. Modes Of Oscillation

The previous section discussed the operation of the array oscillator with the implicit assumption that the phase shift between rings is small, just small enough so that the delay spanned by corresponding ring nodes will equal the two buffer delays established by the closing connections. If the phase shift between rings is larger, so that the total delay is equal

to one period plus two buffer delays, the boundary conditions will still be satisfied, and once again the array will be in a consistent state. This section will generalize the possible phase relationships among the buffer outputs in the array and discuss their impact on the array performance.

The phase distribution discussed in the previous section is for an array closed with the top array inputs connected to the bottom array outputs shifted to the right by two buffers. However, an array can be closed with the bottom array outputs shifted by any number of buffers. Suppose the array in Fig. 3 is closed by connecting the top array inputs to the bottom array outputs shifted to the left by k buffers, so that nodes T_i connect to nodes B_{i-k} , where k is the number of buffer delays establishing the array boundary conditions. Since the buffer stages are inverting, k can be odd only if the closing connections are wire inverted by crossing differential signals to cancel out the odd number of buffer inversions. In the closed array, the delay spanned by all corresponding ring nodes along each column is bounded by the k buffer delays established by the closing connections. Because signals in the array are periodic, this spanned delay can include integer multiples of the oscillation period and still satisfy the boundary conditions. If the array contains M rings each with N buffers oscillating with period T and buffer delay $D = \frac{T}{2N}$, then

$$M\Delta t = kD + xT \quad (1)$$

where Δt is the delay between corresponding ring nodes in adjacent rings as indicated in Fig. 3, kD is the phase shift forced by the boundary conditions, and x is an integer representing the number of extra periods spanned by corresponding ring nodes. Equivalently, solving for Δt ,

$$\Delta t = \frac{xT + k\frac{T}{2N}}{M} \quad (2)$$

so that

$$\frac{\Delta t}{T} = \frac{C}{2NM} \quad (3)$$

where

$$C = k + x2N \quad (4)$$

C is defined to be the array coupling factor or, equivalently, the mode of oscillation, and is equal to the number of buffer delays spanned by all corresponding ring nodes along each column. Thus, for each value of x , the array will oscillate in a different mode defined by a different coupling factor C and will exhibit a different period fraction $\frac{\Delta t}{T}$ as the delay between corresponding ring nodes in adjacent rings. The magnitude of $\frac{\Delta t}{T}$ is not the resultant delay resolution since adjacent output phases do not necessarily come from corresponding nodes in adjacent rings. The ordering of consecutive phases in the array depends on the mode of oscillation.

The nature of the boundary conditions suggests that an array oscillator will support an infinite number of modes with coupling factors C periodically spaced by $2N$ in both the positive and negative directions. In actuality, the number of modes that the array will support is limited. Oscillations in the array will not occur in modes with coupling factors far

from zero because of the large absolute delay between the ring and coupling input transitions. SPICE simulations show that in practice the number of stable modes is typically $\frac{M}{N}$, with equally sized ring and coupling inputs.

The oscillation frequency of the array changes with coupling factor C due to changes in the buffer delay of all buffers. Because the array will not oscillate when the delay between ring and coupling input transitions becomes large, a simple linear model can be used to approximate how the buffer delay changes with the phase difference between the ring and coupling inputs. With equally sized ring and coupling inputs, the buffer delay, from ring input to output, will be equal to the delay of a buffer with simultaneous ring and coupling inputs, less one half of the time the coupling input transition occurs before the ring input transition. Thus

$$D(C) = D(0) - \frac{1}{2}\Delta t(C) = D(0) - \frac{1}{2}D(C)\frac{M}{N} \quad (5)$$

so that

$$D(C) = \frac{D(0)}{1 + \frac{1}{2}\frac{C}{M}} \quad (6)$$

where $D(C)$ is the buffer delay, from ring input to output, as a function of the coupling factor C . The oscillation period $T(C)$ is equal to $2ND(C)$.

The resolution achieved by an array can be worse than a buffer delay divided by the number of rings since for some values of the number of rings M and coupling factor C , the array outputs will not oscillate at unique phases. For differential buffers, if C and M share common factors, each column will contain nodes offset by an integer number of buffer delays. Since the rings constrain the nodes in each column to have a single buffer delay offset from the nodes in adjacent columns, the phase of corresponding ring nodes along a single column of the array will be identical to those in other columns. The number of unique phases for an array composed of differential buffers is then

$$\frac{2MN}{GCD(C, M)} \quad (7)$$

where both inverted and noninverted outputs are utilized. Since single-ended buffers do not have complementary outputs, the number of unique phases for an array composed of single-ended buffers is

$$\frac{MN}{GCD(\frac{C}{2}, M)} \quad (8)$$

where C is even.

Because the oscillation frequency and phase ordering of array outputs change with each mode, it is necessary to be able to initialize the array oscillator into a known mode. In order to selectively reset the array into a particular mode with coupling factor C , the phase relationship among the nodes of the array must be initialized so that it is closer to the phase relationship for this particular mode than for any other mode. The two closest neighboring modes to C are $C - 2N$ and

$C + 2N$. Thus, the fractional delay between corresponding ring nodes in adjacent rings $\frac{\Delta t}{T}$ must satisfy the inequality

$$\frac{C - N}{2NM} < \frac{\Delta t}{T} < \frac{C + N}{2NM} \quad (9)$$

for the array oscillator to enter the mode with coupling factor C after the reset operation. The reset operation is most easily accomplished by switching off the bias voltages in one or more of the rings so that the array of coupled rings no longer forms a closed loop. Switching the buffer outputs is undesirable because these switches would need to be added to all buffer outputs in order to maintain the symmetry of the array, which would add excessive loading to the buffer output nodes. The desired boundary conditions can then be forced on the array by adjusting the bias voltages in the first ring. Modes with C close to zero are readily achieved with buffer stage designs based on differential pairs or current switching in general. With these buffer stages, all of the rings in an open array tend to oscillate in phase without any adjustment to the bias voltages in the first ring. The inactive coupling inputs on the first ring will not reduce its oscillation frequency from a ring with no coupling inputs, since the delay of the buffers does not depend on the size of the input devices switching the currents. Such being the case, the other rings will oscillate with their ring and coupling inputs at the same phase.

E. Array Core Accuracy

An important consideration in the selection of the array dimensions is the manner in which the delay precision of the array core changes with the array size. Static errors in the delays of the output phases from the array core occur as a result of random device and capacitance mismatches in the buffers causing random variations in the buffer delays. For a ring oscillator, the RMS error in the delays of the output phases can be derived from the RMS error in the buffer delays by considering a ring of buffers, each with independent random delays. The result of such a derivation indicates that

$$\Delta P = \frac{\sqrt{3}}{6} \sqrt{N - \frac{1}{N}} \Delta D \quad (10)$$

where ΔP is the RMS error in the delays of the output phases and ΔD is the RMS error in the buffer delays. For an array oscillator, an exact expression for the RMS error in the delays of the output phases cannot be easily derived. Random-process statistical simulations show that the RMS error in the delays of the output phases of an array oscillator approximately scales with the square root of both array dimensions, similarly to a ring oscillator, and is minimized when equally sized ring and coupling inputs are used. Thus, the absolute accuracy of the array core is ideally similar to that of a simple ring oscillator. However, the actual accuracy of the complete array oscillator may be limited by the output channel, as will be discussed in Section V.

F. Layout Issues

Like a simple ring oscillator, the array oscillator's operation depends on all of the buffer delays being identical. In an array

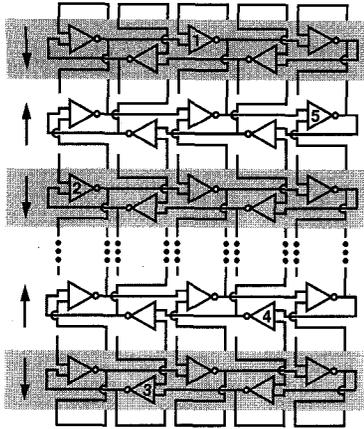


Fig. 6. Floor plan of the array core, illustrating the interleaved buffers in both horizontal and vertical directions and the single buffer shift in every ring. The numbered buffers indicate consecutive buffers along a single logical column.

oscillator, however, the required matching is more stringent since the subbuffer delay resolution requires extremely high precision.

The array oscillator is most naturally laid out as a two-dimensional array of buffer cells with rings extending in one dimension and arrayed in the other. In order for all of the buffer delays to be the same, the interconnect capacitance at each buffer output node must be carefully balanced. This requirement implies that both the buffers in each ring and the rings in the array should be interleaved so that adjacently connected buffers are separated by a single buffer and adjacently connected rings are separated by a single ring, as illustrated in Fig. 6.

The interconnect capacitance at the array closing connections performing the shift by k buffers and at all

other output nodes in the array must also be balanced. This problem can be solved through shifting by a single buffer in every ring of the array so that no shift is necessary in the interconnect at the boundary of the array. The shift is accomplished by connecting the coupling inputs of the buffers in one ring to the ring inputs of the buffers shifted one buffer forward in the previous ring, or equivalently to the ring outputs at the same buffer positions as illustrated in Fig. 6. In order to achieve a net shift of k buffers through M rings after possibly wrapping around the N buffer rings an arbitrary number of times, the number of rings M must be constrained so that

$$M = yN - k \quad (11)$$

for some positive integer y . With the single buffer shift, all connecting wires travel only in the horizontal and vertical directions between adjacent interleaved buffers in all rows and columns of the array also illustrated in Fig. 6.

G. Summary

In summary, in contrast to a simple ring oscillator, the array oscillator has achieved a delay resolution equal to a buffer delay divided by the number of rings and a number of period divisions equal to two times the total number of buffers in the array independent of the desired oscillation frequency. The oscillation frequency is determined primarily by the number

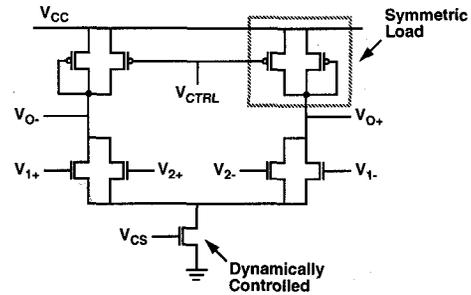


Fig. 7. Schematic of the dual-input differential buffer stage, containing symmetric loads and a dynamically biased current source.

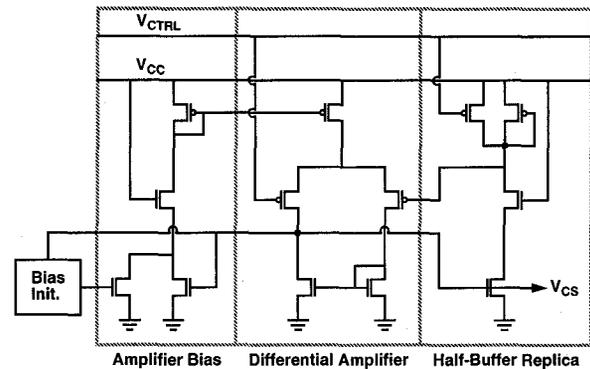


Fig. 8. Schematic of the self-biased replica-feedback current source bias circuit.

of buffers per ring and is largely independent of the number of rings in the array. More output phases can be added, and the delay resolution can be increased simply by adding rings to the array. In addition, the precision of the array oscillator is ideally similar to that of a simple ring oscillator.

The coupled-ring structure of an array oscillator addresses only some of the issues that must be resolved to make a precise delay generator that has a delay resolution equal to a fraction of a buffer delay. The oscillator must be able to operate over a large frequency range and provide high supply noise immunity. These issues must be addressed by the buffer design.

IV. BUFFER DESIGN

An array oscillator can be realized with any single-ended or differential inverting buffer. In order to provide precision delays at high resolution, the buffer outputs must have low phase jitter, since phase jitter can reduce the effective precision and resolution of any delay generator. Low phase jitter, however, is difficult to achieve in the noisy environment of a digital integrated circuit, as it requires high supply noise immunity. In addition, state-of-the-art digital technologies typically have limited supply voltages due to thin gate oxides. The differential buffer stage described in this section is designed to have high supply noise immunity while being able to operate at low supply voltages. The key components of the buffer stage design that achieve these objectives are the symmetric load elements and the self-biased replica-feedback current source bias circuit shown in Fig. 7 and Fig. 8, respectively.

Supply noise sensitivity has both static and dynamic components. Static supply sensitivity is dominated by the output

resistance of the current sources used. Achieving high static supply rejection is typically incompatible with low-voltage circuit design since it usually requires cascoding to achieve high output impedances. The current source bias circuit described in this section enables the buffer stages to achieve high static supply rejection without cascoding through the use of self-biasing and replica-feedback techniques. The dynamic supply sensitivity is dominated by the load structure of the buffer stages and the coupling capacitance to the buffer outputs. The symmetric load elements used in the buffer stages provide for high dynamic supply rejection through a first-order cancellation of noise coupling.

A. Differential Buffer Stage

The buffer stage used in the array oscillator is based on an NMOS source-coupled pair with symmetric load elements and a dynamically-biased simple NMOS current source, as shown in Fig. 7. The coupling input is formed from an additional source-coupled pair sharing the same loads and current source. The bias voltage of the simple NMOS current source is continuously adjusted in order to provide a bias current that is independent of supply and substrate voltages. With the output swings referenced to the top supply, the current source effectively isolates the buffer from the negative supply so that the buffer delay remains constant with supply voltage. The load elements are composed of a diode-connected PMOS device in shunt with an equally sized biased PMOS device. They are called symmetric loads because their I - V characteristics are symmetric about the center of the voltage swing. The control voltage, V_{CTRL} , is the bias voltage for the PMOS device. It is used to generate the bias voltage for the NMOS current source and provides control over the delay of the buffer stage.

Fig. 9 contains simulated symmetric load I - V characteristics at low and mid-range bias voltages. With the top supply as the upper swing limit, the lower swing limit is symmetrically opposite at the bias level of the PMOS device, V_{CTRL} . The dashed lines show the effective resistance of the load and illustrate the symmetry of their I - V characteristics. The buffer delay changes with the control voltage since the effective resistance of the load also changes with the control voltage. The buffer bias current is adjusted so that the output swings vary with the control voltage rather than being fixed in order to maintain the symmetric I - V characteristics of the loads.

Linear resistor loads are most desirable for achieving high dynamic supply noise rejection. Because they provide differential-mode resistance that is independent of common-mode voltage carrying the supply noise, the delay of the buffers is not affected by this common-mode noise. Unfortunately, adjustable resistor loads made with real MOS devices cannot maintain linearity while generating a broad frequency range. Symmetric loads, though nonlinear, can also be used for achieving high dynamic supply noise rejection. Nonlinear load resistances normally convert common-mode noise into differential-mode noise, which affects the buffer delays. With symmetric loads, however, the first-order noise coupling terms cancel out, leaving only the higher order

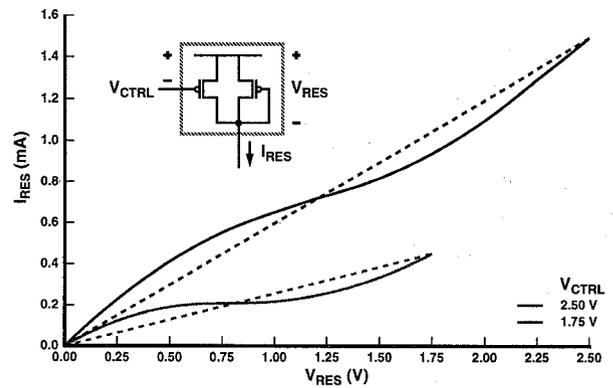


Fig. 9. Simulated symmetric load I - V characteristics at low and mid-range bias voltages. The dashed lines show the effective resistance of the loads and reveal the symmetry of the I - V characteristics.

terms and, therefore, substantially reducing the jitter caused by common-mode noise present on the supplies. SPICE simulations of an array with worst-case coupling of buffer output interconnection capacitance, and with the control voltage fixed relative to the top supply, show that an instantaneous 500-mV supply voltage step results in a total phase error of only 0.5% of an oscillation period with static effects factored out.

The MOS realization of symmetric loads has additional advantages. The quiescent biasing point of a buffer with symmetric loads, the point of symmetry at the center of the output voltage swing, is the point where the buffer's gain is largest. As a result, the oscillation frequency range will typically be very broad. Furthermore, because the load resistance of symmetric loads decreases toward the ends of the voltage swing, the transient swing limits will always be well defined near the dc swing limits, resulting in reduced noise sensitivity. Also, because the load elements contain two equally sized transistors, their sizes will be similar to the sizes of the differential pair and current source devices used in the buffer.

B. Current Source Bias Circuit

The current source bias circuit, shown in Fig. 8, achieves two functions. First, it sets the current through a simple NMOS current source in the buffers in order to provide the correct symmetric load swing limits. Second, it dynamically adjusts the NMOS current source bias so that this current is held constant and highly independent of supply voltage. It uses a replica of half the buffer stage and a single-stage differential amplifier.

The amplifier adjusts the current output of the NMOS current source so that the voltage at the output of the replicated load element is equal to the control voltage, a condition required for correct symmetric load swing limits. The net result is that the output current of the NMOS current source is established by the load element and is independent of the supply voltage. As the supply voltage changes, the drain voltage of the NMOS current source transistor changes. However, the gate bias is adjusted by the amplifier to keep the output current constant, counteracting the effect of the finite output

impedance. Compensation for the bias circuit is accomplished at the output with the loading of the simple NMOS current source gates. The output load should be limited to about ten buffer stages containing devices of the same size as the corresponding devices in the bias circuit in order to prevent the output recovery time of the bias circuit from limiting the dynamic supply noise rejection of the buffers.

In order for the supply voltage requirements of the current source bias circuit to match the low supply voltage requirements of the buffers, an amplifier based on a self-biased PMOS source-coupled pair is used. The amplifier bias is generated from the same NMOS current source bias through a stage mirroring the half-buffer replica so that amplifier supply voltage requirements are similar to those of the buffers and the amplifier bias current is highly independent of supply voltage. This replica bias stage is necessary because otherwise the input offset of the amplifier will vary with supply voltage, causing the output current of the NMOS current source to also change with supply voltage. The operation of the amplifier depends on its output, so an initialization circuit is needed to bias the amplifier at power-up. This initialization circuit prevents the NMOS current source bias from completely turning off the current sources.

With no required swing reference voltage, the only bias voltage that is required is the control voltage itself. Although no device cascoding is used, the resultant static supply noise rejection is equivalent to that achievable by a buffer stage and a bias circuit with cascoding, without requiring extra supply voltage. The total supply voltage requirement of the buffer stage and bias circuit is slightly less than a series NMOS and PMOS diode voltage drop.

V. OUTPUT CHANNEL

In order for a system to make use of the precise delays generated in the array core, an array oscillator must provide a means for accessing the internal array nodes. This function is accomplished by the output channel. Since a delayed signal implies the existence of a reference signal, the array oscillator must provide at least two output ports, one to be used as a reference and the other to address the desired delay. In order to maintain the high precision provided by the array core, each port of the output channel must have an equal delay from any internal array node. Such a constraint requires careful matching and balancing of the interconnect capacitance in and among the output ports.

A. Channel Design

The output channel is organized similarly to that of a multiported memory containing word lines and the counterpart of bit lines. A simplified block diagram of a single port of the output channel is shown in Fig. 10. The word lines select a row of buffer cells containing a single buffer from each ring. The column multiplexer for each output port in each buffer cell is isolated from the array buffer through an additional buffer to prevent changes in the loading at the array buffer output. The outputs from the selected buffers are multiplexed on to the bit lines. After an additional buffer to increase the signal

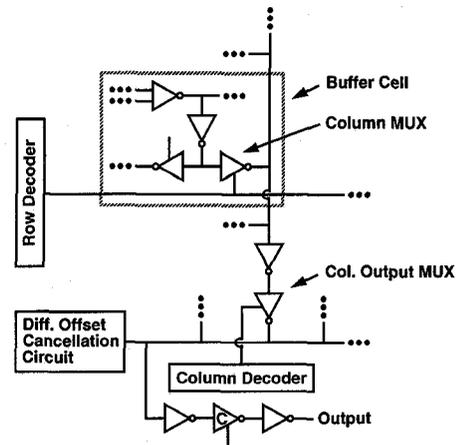


Fig. 10. Simplified block diagram of a single port of the output channel. The column and column output multiplexers are distributed structures and are represented as buffers with output enables. All wires in the signal path represent differential signals.

swings, the bit lines from the selected column are multiplexed on to a single pair of output wires. These output wires are followed by a buffer, a conditionally inverting multiplexer that can swap the differential signals, and a final output buffer. To produce a single delay only one of two output ports needs to be completely addressable. As such, only one set of column bit lines are necessary, but two column output multiplexers are still required to match the output port delays.

All buffers and multiplexers used in the output channel are based on single-differential-input versions of the buffer stage discussed in Section IV. The schematic of a multiplexer is shown in Fig. 11. For each input, there is a differential pair with a biasing current source and a pair of NMOS devices acting as switches connected to the differential pair outputs. The multiplexers are distributed structures with the differential pair, current source, and switches for each input located near the origin of the input. The switch devices selectively switch one of the differential pairs onto a pair of shared output wires with a single pair of shared load elements. The conditionally inverting multiplexer is composed of a single differential pair with a biasing current source that is connected or cross connected by two pairs of NMOS switch devices to a pair of load elements.

B. Bandwidth Limitations

As the number of buffers per ring is reduced, the required bandwidth from the output channel will increase due to the increase in oscillation frequency. In addition, as the number of rings is increased, the bandwidth of the column output multiplexer will decrease as a result of the increased output loading. Care must be taken to prevent an output channel bandwidth limitation from distorting the delays of the output phases and reducing the overall precision of the array oscillator.

Such bandwidth limitations can cause the voltage swings at the column and column output multiplexers to be less than their static limits. With incomplete swings, random device mismatches in the multiplexers result in address-dependent differential-mode offsets. As the bandwidth decreases, these

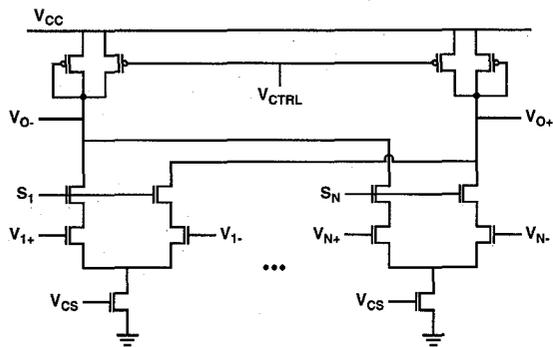


Fig. 11. Schematic of the multiplexers used in output channel.

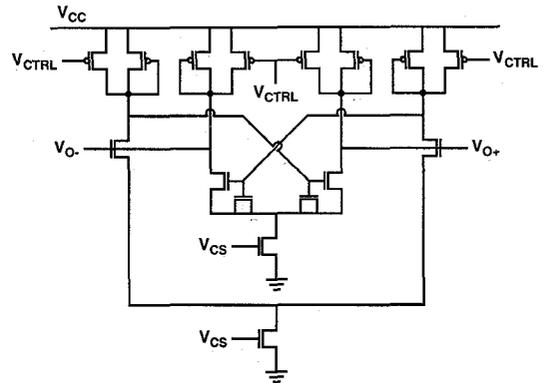


Fig. 13. Schematic of the differential offset cancellation circuit.

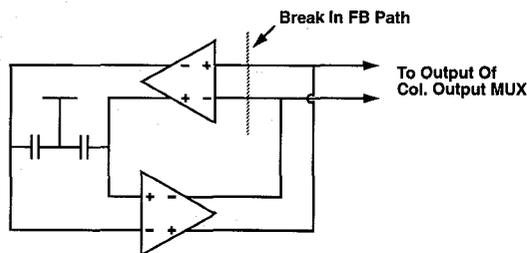


Fig. 12. Block diagram of the differential offset cancellation circuit.

differential-mode offsets will become a larger fraction of the decreasing swing. In addition, differential-mode offsets are amplified by the dc gain of the multiplexers, unlike the output signals, which experience little if any amplification, suggesting that differential offsets originating early in the output channel will result in larger differential offsets later in the output channel. When signals with differential-mode offsets are amplified to the static swing limits by low-fanout buffers, the differential-mode offsets are converted into differential duty cycle variations. The result is that bandwidth limitations in the output channel allow random device mismatches to cause address-dependent duty cycle variations among the output phases. Since the delays of the output phases are referenced at the output transitions, these duty cycle variations cause the delays to have an address-dependent error.

C. Differential Offset Cancellation Circuit

As long as the output channel is linear, no information about the delays of the output phases will be lost due to address-dependent differential-mode offsets. The output channel will be very linear as long as the signal swings are small. These facts suggest that a circuit can be added to the output of the column output multiplexer to cancel out the random differential-mode offsets and prevent them from turning into duty cycle variations. Such a circuit would allow the delays of the output phases to be referenced at the output transitions without an address-dependent error.

The block diagram for such a differential offset cancellation circuit is shown in Fig. 12, while the schematic is shown in Fig. 13. It contains two differential buffers in a feedback loop with two NMOS capacitors to remove the ac signal components and allow feedback only for the dc signal components. The operation of the circuit is most easily analyzed with the feedback path broken at the input to the first buffer as indicated

in Fig. 12, so that only the output of the second buffer connects to the output of the column output multiplexer. If the second buffer is of the same size as the column output multiplexer, the offset at the output of the second buffer will be equal to one half the original differential-mode offset at the output of the column output multiplexer. When referred to the input of the first buffer, this offset will be equal to one half the original differential-mode offset divided by the product of the gain of the two buffers. With the feedback path closed, the negative feedback will drive the output of the column output multiplexer to this input-referred offset. With the second buffer the same size as the column output multiplexer, its gain will be one half that of the first buffer. Thus, the differential-mode offset at the output of the column output multiplexer is reduced by the square of the buffer gain with this circuit.

VI. EXPERIMENTAL RESULTS

An array oscillator with seven rings of five buffers per ring has been fabricated in a 2- μm *N*-well technology with nodes T_i connected to nodes B_{i+2} , as previously described. This configuration gives rise to two possible modes of oscillation with different frequencies, where the phase difference across all corresponding ring nodes is -2 or -12 buffer delays. The array is selectively reset in a particular mode by switching the ring bias lines. A micrograph of the fabricated array oscillator with a superimposed floor plan is shown in Fig. 14.

The differential offset cancellation circuit was not designed in time to be a part of this 5×7 array oscillator implementation. However, on a comparison chip containing two identical arrays of a larger size with one utilizing the differential offset cancellation circuit, the peak error in output delays was reduced by more than a factor of two. Further improvement was limited by other sources of delay error in the output channel and array core.

A. Output Accuracy

The measured output accuracy of the 5×7 array oscillator is summarized in Fig. 15. The plot shows the error in delays for the 70 output phases as a function of the percentage of the period. The results indicate that with a period of 7 ns and an LSB of 101 ps, the peak error is slightly greater than a half

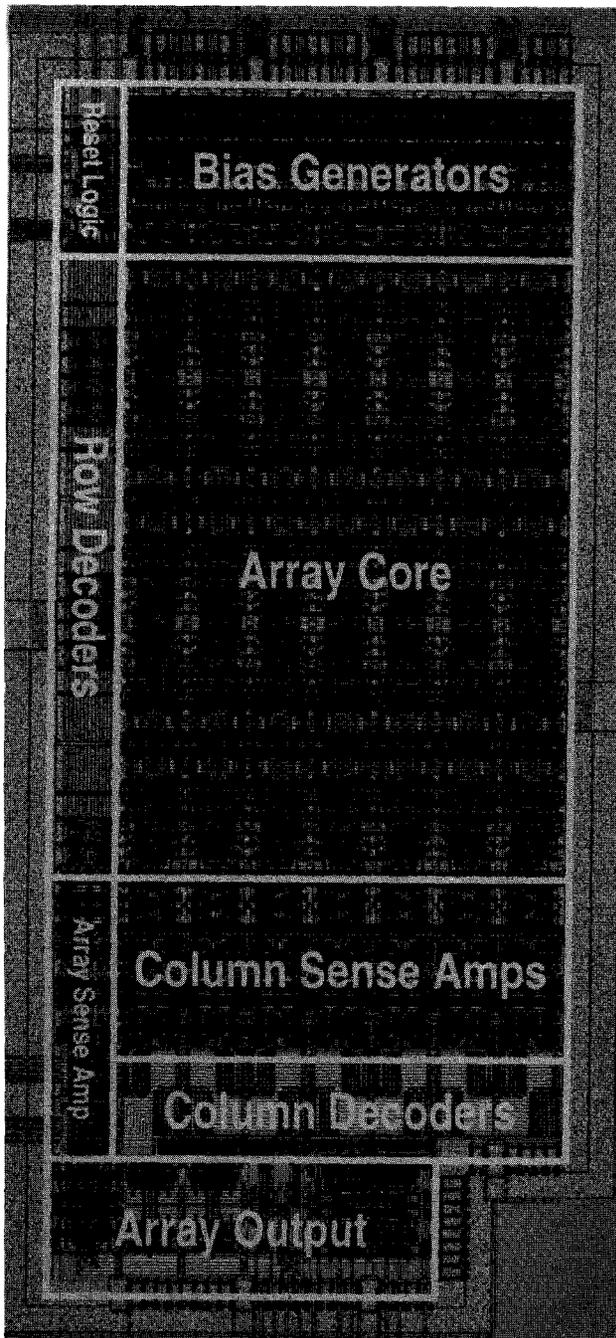
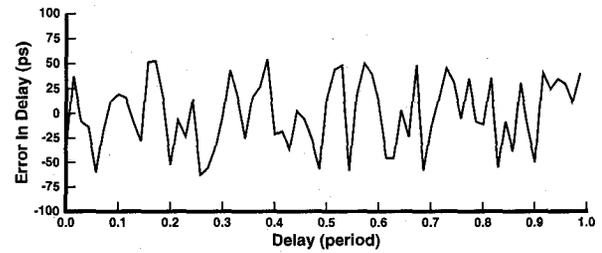


Fig. 14. Die micrograph of the 5×7 array oscillator.

LSB. In terms of differential nonlinearity, the peak is about one LSB.

B. Overall Results

The measured frequency as a function of static supply voltage for several control voltages referenced to top supply is shown in Fig. 16. These curves are very flat, indicating very low frequency sensitivity to supply voltage, yet the minimum supply voltages are relatively small. The exact frequency sensitivity is a little distorted by changing die temperatures, especially at higher control voltages where thermal effects dominate the results due to increased power dissipation. In



Vctrl 2.75 V LSB 101 ps
 Freq 141 MHz Peak INL 58 ps
 Period 7.07 ns Peak DNL 99 ps

Fig. 15. Measured output accuracy of the 5×7 array oscillator. The plot shows the error in delays for the 70 output phases as a function of the percentage of the period.

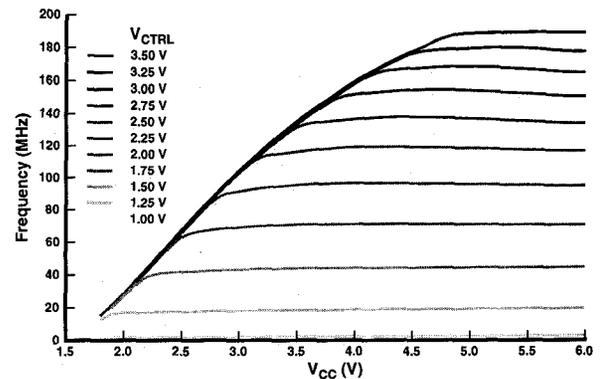


Fig. 16. Measured frequency as a function of static supply voltage for the 5×7 array oscillator. At higher control voltages, thermal effects dominate the results due to increased power dissipation.

TABLE I
 PERFORMANCE SUMMARY OF THE 5×7 ARRAY
 OSCILLATOR AS A VOLTAGE-CONTROLLED OSCILLATOR

Frequency Range, Sensitivity	5–190 MHz, 80 MHz
Static Supply Sensitivity	0.25%/V @ 70 MHz
Minimum Supply Voltage	2.5 V @ 70 MHz
Power Dissipation	62 mA @ 70 MHz
Die Area	3.27 mm ²
Technology	2- μ m N-well CMOS

addition, the frequency varies linearly with control voltage except at high control voltages where thermal effects once again dominate.

Table I summarizes the overall specifications of the array oscillator as a voltage-controlled oscillator. The high supply current and large die area result from maximizing the oscillation frequency obtainable in the 2- μ m technology. The measured frequency sensitivity to static supply voltage is less than 0.25% over most of the operating range for each mode, less than the 0.7% previously reported [3].

VII. CONCLUSION

This paper has described a new delay generator based on a series of coupled ring oscillators for producing delays with resolution equal to a buffer delay divided by the number of rings. By coupling several ring oscillators together, the delay

generator breaks the dependence of the oscillation frequency on the number of buffers in a ring oscillator. The delay generator achieves a precision intrinsically as high as that for a ring oscillator. However, the precision actually realizable can be limited if care is not taken to address such external issues as jitter performance and delay errors due to bandwidth limitations in the output channel. The delay generator also utilizes a differential buffer stage with high supply noise rejection while operating at low supply voltages. The dynamic supply noise sensitivity of the differential buffer stage is very small, making its static supply noise sensitivity the dominant factor in the phase-locked jitter performance of the delay generator. Experimental results from a $2\text{-}\mu\text{m}$ N -well CMOS implementation of the delay generator indicate that it can achieve an output delay resolution of 101 ps while operating at 141 MHz with a peak error of 58 ps. These results confirm that the delay generator can be used in single-chip testers as a cost-effective solution for producing precise delays with high resolution to test chips designed with higher speed integrated circuit technologies.

ACKNOWLEDGMENT

The authors thank T. Chanak and D. Ramsey for their assistance during the final stages of the layout effort for the fabricated chip.

REFERENCES

- [1] J. Gasbarro and M. Horowitz, "A single-chip, functional tester for VLSI circuits," in *ISSCC 1990 Dig. Tech. Papers*, pp. 84–85, Feb. 1990.
- [2] J. Maneatis, and M. Horowitz, "Precise delay generation using coupled oscillators," in *ISSCC 1993 Dig. Tech. Papers*, pp. 118–119, Feb. 1993.

- [3] I. Young *et al.*, "A PLL clock generator with 5 to 110 MHz lock range for microprocessors," in *ISSCC 1992 Dig. Tech. Papers*, pp. 50–51, Feb. 1992.



John G. Maneatis was born in San Francisco, CA, on November 7, 1965. He received the B.S. degree in electrical engineering and computer science from the University of California, Berkeley, in 1988 and the M.S. degree in electrical engineering from Stanford University, Stanford, CA, in 1989. He is currently a Ph.D. candidate in electrical engineering at Stanford University.

He worked at Hewlett-Packard Laboratories, Palo Alto, CA, during the summer of 1989 on high-speed analog-to-digital conversion and monolithic clock recovery, and at Digital Equipment Corporation Western Research Laboratory, Palo Alto, CA, during the summer of 1990 on CAD tool development and ECL circuit design. His research interests include high-performance circuit design for phase-locked loops, microprocessors, and data conversion. He is a Registered Professional Electrical Engineer in the State of California.

Mr. Maneatis is a member of Tau Beta Pi, Eta Kappa Nu, and Phi Beta Kappa.



Mark A. Horowitz received the B.S. and M.S. degrees in electrical engineering from Massachusetts Institute of Technology, Cambridge, MA, in 1978 and the Ph.D. degree in the same field from Stanford University, Stanford, CA, in 1984.

He is currently an Associate Professor of Electrical Engineering at Stanford University, where his research interests are in digital integrated circuit design. He has led a number of processor design projects at Stanford, including MIPS-X, one of the first processors to include an on-chip instruction cache, and TORCH, a statically scheduled superscalar processor. In 1990 he took leave from Stanford to help start Rambus, Inc., a company designing high-bandwidth memory interface technology. His current research includes work in both high-speed and low-power circuits, memory design, processor architecture, and IC CAD tools.