PRECISE DELAY GENERATION USING COUPLED OSCILLATORS

A DISSERTATION

SUBMITTED TO THE DEPARTMENT OF ELECTRICAL ENGINEERING

AND THE COMMITTEE ON GRADUATE STUDIES

OF STANFORD UNIVERSITY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

John George Maneatis

June 1994

Abstract

This thesis describes a new class of delay generation structures which can produce precise delays with sub-gate delay resolution. These structures are based on coupled ring oscillators which oscillate at the same frequency. One such structure, called an array oscillator, consists of a linear array of ring oscillators. A unique coupling arrangement forces the outputs of the ring oscillators to be uniformly offset in phase by a precise fraction of a buffer delay. This arrangement enables the array oscillator to achieve a delay resolution equal to a buffer delay divided by the number of rings. Another structure, called a delay line oscillator, consists of a series of delay stages, each based on a single coupled ring oscillator. These delay stages uniformly span the delay interval to which they are phase locked. Each delay stage is capable of generating a phase shift that varies over a positive and negative range. These characteristics allow the structure to precisely subdivide delays into arbitrarily small intervals.

The buffer stages used in the ring oscillators must have high supply noise rejection to avoid losing precision to output jitter. This thesis presents several types of buffer stage designs for achieving high supply noise rejection and low supply voltage operation. These include a differential buffer stage design based on a source coupled pair using load elements with symmetric I-V characteristics and a single-ended buffer stage design based on a diode clamped common source device. The thesis also discusses techniques for achieving low jitter phase-locked loop performance which is important to achieving high precision.

Based on the concepts developed in this thesis, an experimental differential array oscillator delay generator was designed and fabricated in a 1.2-µm N-well CMOS technology. The delay generator achieved a delay resolution of 43ps while operating at 331MHz with peak delay error of 47ps.

Acknowledgments

I would like to acknowledge the contributions of a number of people who were very helpful and supportive in my research.

I extend my sincere appreciation to my principal advisor, Professor Mark Horowitz, whose wisdom and expertise enabled the successful completion of this research. I am particularly grateful to him for providing me considerable freedom in defining and directing my research. I am also indebted to my associate advisor, Professor Bruce Wooley, and Professor Gordon Kino for reading this thesis and to Professor Simon Wong for serving on my oral examination committee.

I would like to extend my gratitude to Tom Chanak, Don Ramsey, and Drew Wingard for their continued support and helpful advice.

I wish to thank MOSIS for their support in prototype fabrication and Rambus, Inc. for allowing me access to their glass cutting laser. This research was supported by Advanced Research Projects Agency Contract No. N00039-91-C-0138.

Finally, I would like to express my greatest appreciation to my family for their constant support and encouragement.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

Ring oscillators and buffer delay lines are used in a variety of integrated circuit applications because of their ability to generate delays with very high precision at high operating frequencies. These applications include clock generation and synchronization [1], time digitization [2], and clock recovery [3].

Ring oscillators can be used to generate precise delays because of their inherent high delay linearity. If all the constituent buffers of a ring oscillator have the same delay, they will divide the oscillation period into precise delay increments. Thus, they achieve this high linearity through matched delays that do not rely on the specific characteristics of the identical elements to generate the delays [4]. Such a reliance leads to poor linearity for large delay ranges. However, to achieve high precision, the delay generator must not only have high linearity, but it also must produce delay increments that are referenced to some known delay.

In order to set the delay of the buffer stages to a known value, a phase-locked loop can be used to reference the oscillation frequency of the ring oscillator to the frequency of an established clock signal. The closed loop feedback forces the output of the ring oscillator to be equal to the input clock signal so that the delay of each buffer stage is then equal to the clock period divided by two times the number of buffers in the ring oscillator. Figure 1-1 shows a ring oscillator delay generator with a phase-locked loop (PLL) [5, 6, 7, 8]. The basic feedback loop consists of a phase comparison that generates an error signal based on the phase difference between the clock signal and the ring oscillator output. This error signal is integrated and filtered and then fed back to the ring oscillator as the control voltage. The control voltage adjusts the delay of the buffer stages in the ring oscillator and, thereby, controls the operating frequency of the ring oscillator. While a non-zero phase difference exists, the operating frequency of the ring oscillator will be adjusted by the PLL until it matches the frequency of the clock signal. Once the two frequencies are equal, the phase difference will go to zero and the control voltage will reach a constant value. At this point in time, the ring oscillator output will be phase locked to the clock signal. With the delay increments a precise fraction of a clock period, the generated delays will be of very high precision.

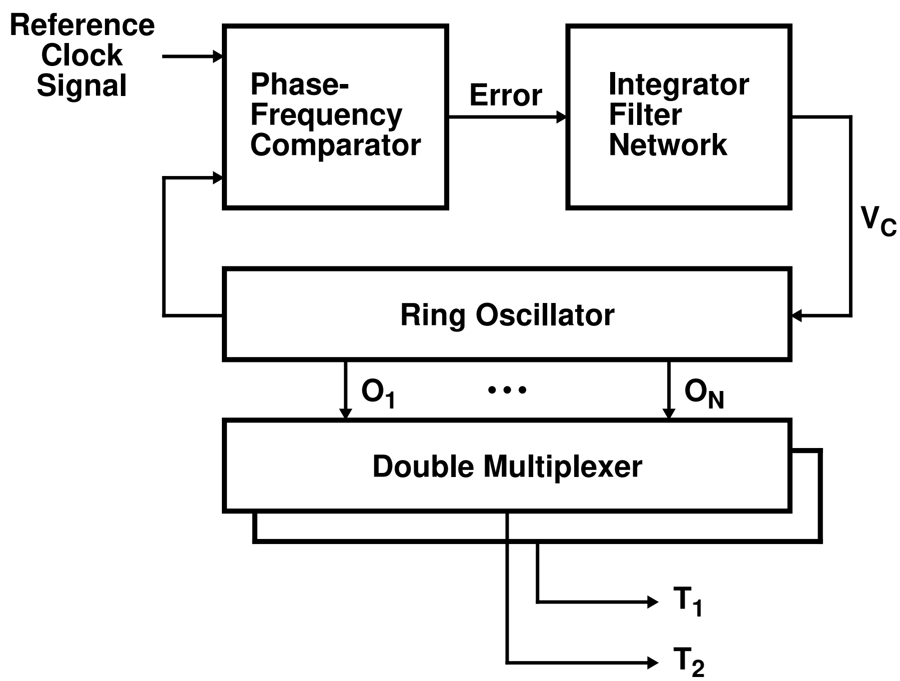

Figure 1-1: Ring oscillator delay generator with PLL.

With the ring oscillator phase locked, different fractions of the clock period can be obtained by selecting different ring outputs with multiplexers. Figure 1-2 illustrates the phase relationship among individual buffer outputs. Rising transitions are indicated by dots and falling transitions are indicated by circles. If a ring oscillator contains five differential buffers with both their inverted and noninverted outputs utilized as shown, ten different output phases are available which uniformly span the output period.

Although ring oscillators can provide high linearity and high precision, their delay resolution is limited to a buffer delay. Their delay resolution cannot be improved by adding more output phases. The only way to add more output phases is to add more buffers which decreases the maximum oscillation frequency. As a result, the delay resolution remains unchanged. In a 1.2-µm N-well CMOS technology [9], the delay resolution is limited to about 300ps.

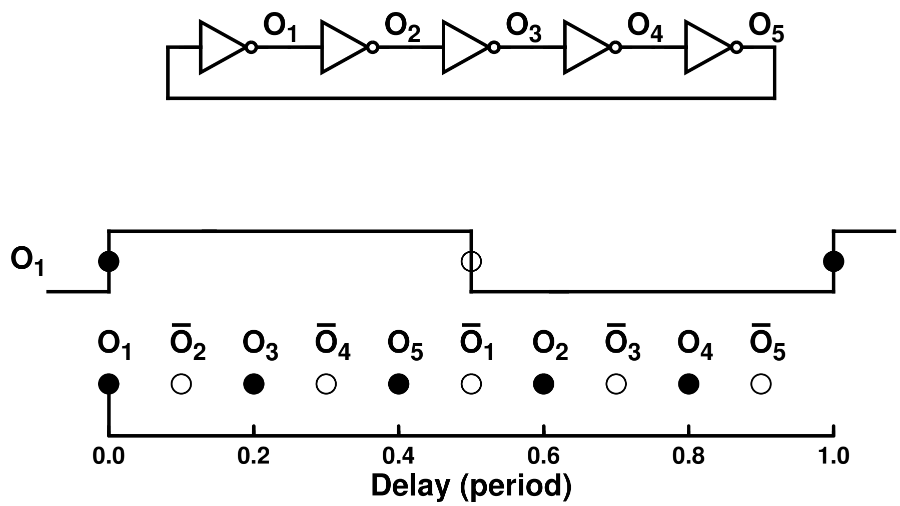

Figure 1-2: Phase relationship among ring oscillator outputs for a ring with five buffers.

While ring oscillators and buffer delay lines can be used effectively in a wide range of integrated circuit applications, they cannot be used successfully in some applications because of their inability to achieve delays with high resolution. Their delay resolution is limited by the minimum delay of the constituent buffer stages. Such applications, including digitally controlled clock recovery [10] and high resolution time digitizers [11], sometimes employ interpolation techniques to extend the resolution achievable by ring oscillators. However, these techniques lack linearity and require careful calibration to achieve high precision.

Another important application that cannot use ring oscillators is precision waveform timing generation for state-of-the-art single-chip testers [12]. Creating the timing reference for such a waveform generator was the original problem that inspired this work. When testing digital integrated circuits, it is necessary to supply digital waveforms as input which require accurate delays referenced to some clock signal, as illustrated in Figure 1-3. The delay resolution that is 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 can be obtained with ring oscillators only by using a 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, ring oscillators cannot be used in such a cost-effective approach because of their limited delay resolution.

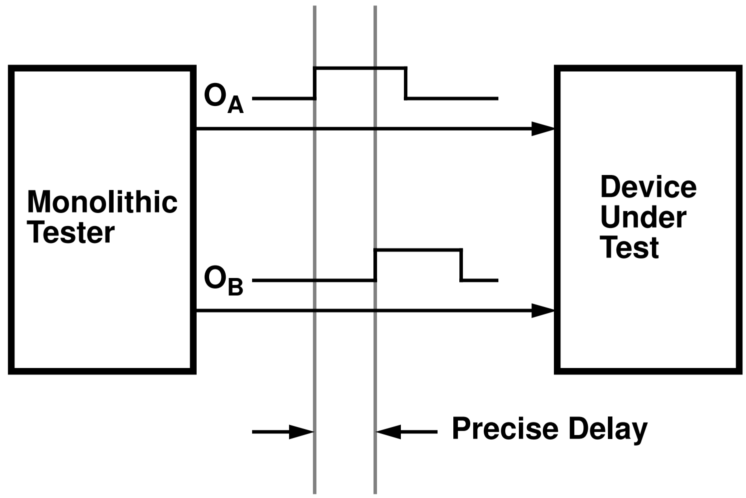

Figure 1-3: Monolithic testing of digital integrated circuits.

This thesis presents new and innovative techniques for using ring oscillators to overcome this resolution limit. These techniques are based on combining several ring oscillators together to form a series of coupled ring oscillators. Coupled oscillator techniques allow ring-like structures to produce a delay resolution equal to a fraction of a buffer delay with a delay precision equivalent to that of a ring oscillator. More importantly, these techniques can be applied in a variety of different delay generation applications.

The concepts and techniques pertaining to the design and operation of coupled oscillators are described in detail in Chapter 2. The chapter also presents two new structures, an array oscillator and a delay line oscillator, and discusses the theory related to their design and operation. The generation of precise delays requires more than high resolution coupled oscillator architectures. It also requires low noise constituent buffer stages to prevent an effective loss of precision due to jitter in the output signals. Chapter 3 discusses the general issues related to differential and single-ended buffer designs for high supply noise immunity and low supply voltage operation. A number of other performance critical implementation issues must be considered in the implementation of coupled oscillator delay generators, including the method used to read the delayed signals generated by the structures and the techniques used to reference and synchronize the delay intervals with the overall system. Chapter 4 examines these issues and describes the actual implementation of an array oscillator delay generator. It also presents test results from an experimental implementation based on differential buffers, confirming the ability of an array oscillator to produce precise delays with high resolution and low output jitter. Lastly, some concluding remarks on the overall effectiveness, significance, and practical applications of coupled oscillator delay generators are presented in Chapter 5.

Chapter 2

Coupled Oscillator Architectures

This chapter discusses new types of oscillator structures for generating delays with a resolution equal to a precise fraction of a buffer delay. They are called coupled oscillators because they are based on a series of identical interconnected ring oscillators that function as a single unit. In a coupled oscillator, adjacent pairs of ring oscillators in a linear array are connected together through one or more coupling inputs as illustrated in Figure 2-1. These coupling inputs allow each ring to affect one another such that they become interdependent. The resulting interaction between ring oscillators gives rise to a number of useful properties. Because the rings are all identical, they will tend to phase lock and oscillate at the same frequency with some fixed phase relationship between their outputs. With identical coupling between each adjacent pair of rings, the phase shift between each set of outputs from adjacent rings will be identical. The total phase shift across all sets of outputs from all rings can then be constrained to equal some easily generated delay unit such as a buffer delay or an external delay. This boundary constraint can be imposed with a phase-locked loop or by utilizing the symmetry within the coupled ring oscillators themselves as will be discussed in this chapter. With the boundary constraint in place, the phase shift between each set of outputs from adjacently coupled rings can then be made a precise fraction of a buffer delay or some external delay reference. The result is that these structures can achieve a significant increase in delay resolution without sacrificing the high precision offered by the ring oscillators.

The coupling inputs for each ring oscillator are formed by adding an additional input to one or more of the constituent buffers. This chapter will begin with a description of the dual-input buffer used to form the coupling inputs. It will then show how ring oscillators coupled together with one or more of these dual-input buffers can form a variety of structures for generating precise delays with high resolution. These structures include the array oscillator and the delay line oscillator.

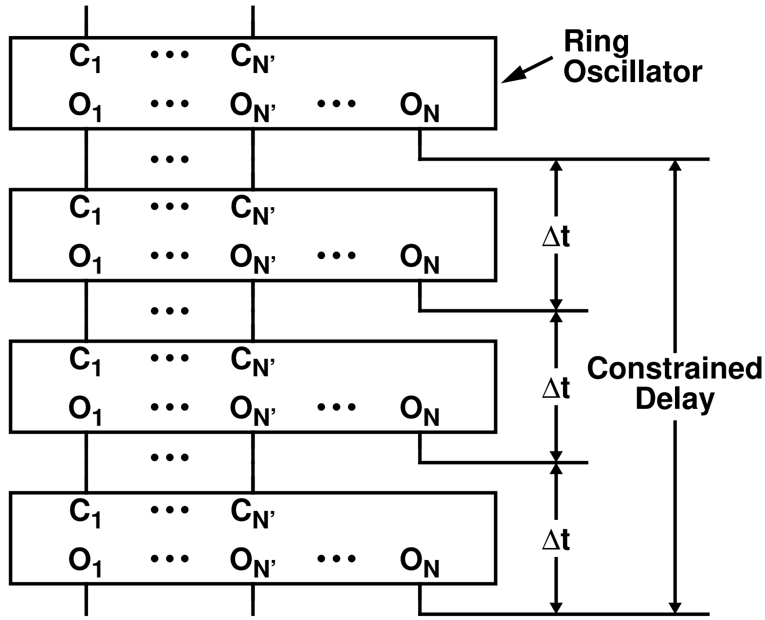

Figure 2-1: Coupled oscillator structure.

2.1 Dual-input inverting buffer

To couple rings together, coupled oscillators require an inverting buffer stage with two inputs. While one input is connected to the output of a buffer stage in the same ring, the other input is connected to the output of a buffer stage in another ring. Accordingly, the first input is referred to as the ring input and the second input is referred to as the coupling input. Although this dual-input inverting buffer has two inputs of the same polarity, it is very similar to a single-input inverting buffer. An example of such a dual-input buffer is shown in Figure 2-2. It is constructed from a static CMOS inverter by shunting the outputs of two half-sized static CMOS inverters together.[1] For this buffer to function properly, the delay between the ring and coupling input transitions must be small, so that the transitions overlap to some extent. Although neither transition in isolation may be able to cause a complete transition at the output, the overlapping input transitions will generate a continuous and complete output transition. When the input transitions overlap, both the ring and coupling inputs will affect the time of the output transition. Early coupling input

1. 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.

transitions will reduce the buffer delay, defined from ring input to buffer output, while late coupling input transitions will increase the buffer delay.

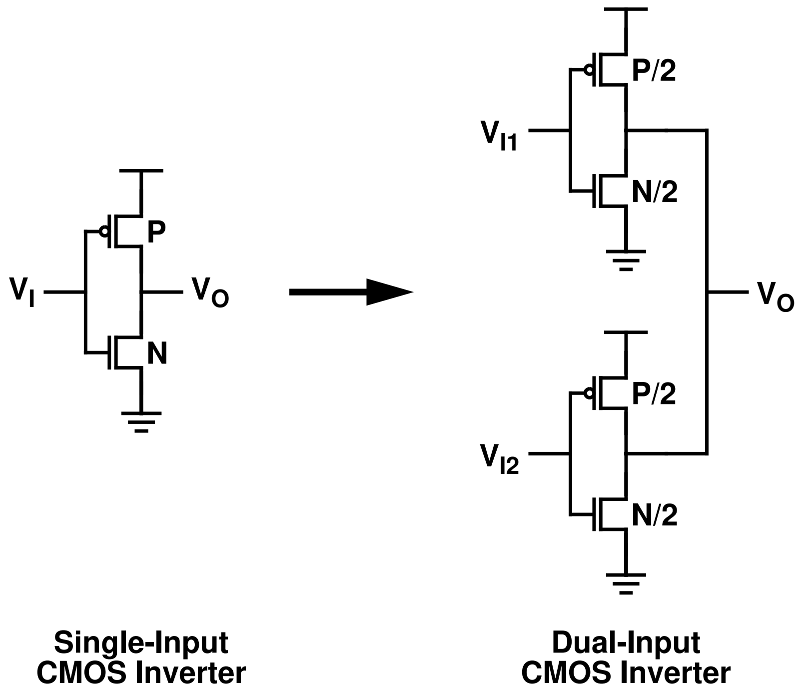

Figure 2-2: Example of a dual-input inverting buffer created by shunting the outputs of two half-sized CMOS inverters.

Figure 2-3 shows the simulated delay characteristics for the static CMOS dual-input buffer. The buffer delay from the ring input to the buffer output is plotted as a function of the delay between the ring and coupling inputs. The scales are normalized to a CMOS inverter delay. Because the output transition time is roughly twice the inverter delay, the input transitions will overlap for a wide range of delays between the ring and coupling inputs. As the delay between the ring and coupling inputs varies from -1 to +1 inverter delays, the buffer delay varies from about 0.5 to 1.5 inverter delays. As the delay between the two inputs increases significantly beyond one inverter delay, the operation breaks down as the input transitions no longer overlap. While the buffer delay has a fairly linear dependence on the input delay, the linearity is not required for the operation of coupled oscillators. It will be shown later in the chapter that a monotonic dependence alone leads to the important properties of coupled oscillators.

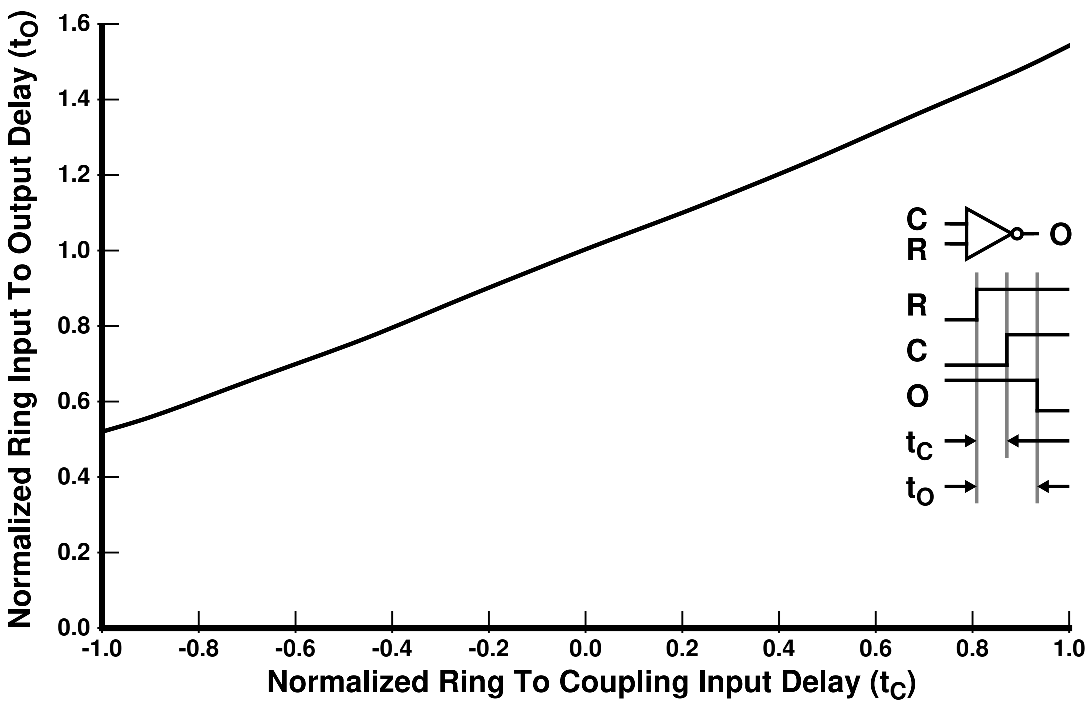

Figure 2-3: Simulated delay characteristics of a static CMOS dual-input buffer.

2.2 Array oscillator

When ring oscillators are formed based on this dual-input buffer, the coupling inputs provide a means for connecting rings together. For example, the outputs of one ring can be simply connected to the coupling inputs of the next ring. By coupling several rings together, a two-dimensional array of dual-input inverting buffers can be formed. An array oscillator, shown in Figure 2-4, is such a structure. Rings extend 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.

An array oscillator can generate precise delays with a resolution equal to a buffer delay divided by the number of rings. The basic idea is to force several rings oscillating at the same frequency to be uniformly offset in phase by a precise fraction of a buffer delay. Corresponding outputs from each ring will then divide a buffer delay into several equal delay intervals. The coupling between the rings that generates this uniform spacing is the key to the design of the array oscillator.

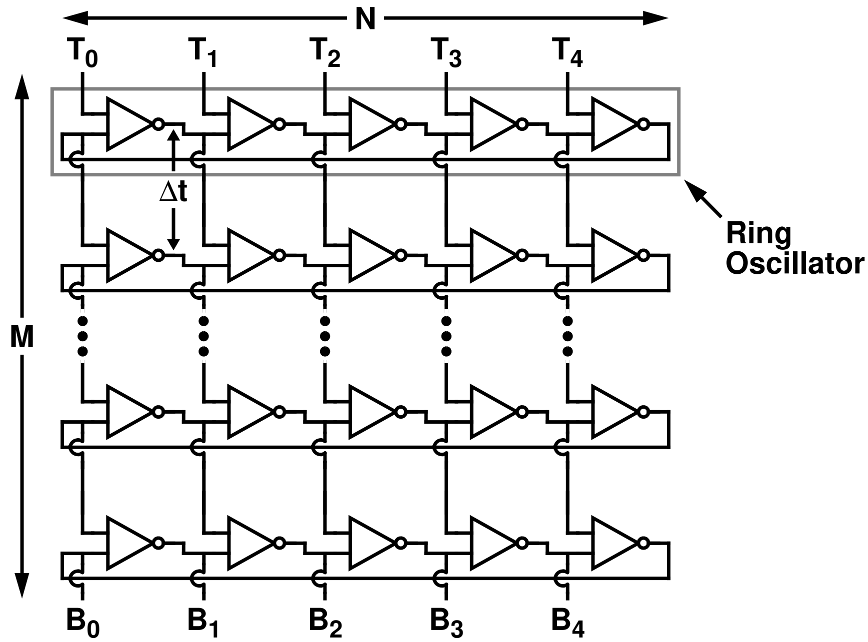

Figure 2-4: Array oscillator structure.

2.2.1 Array operation

The closing connections in the array oscillator, which connect the top array inputs to the bottom array outputs, impose boundary constraints that constrain its operation to specific consistent modes of oscillation. Before this section considers how the closing connections affect the array operation, it will first consider consistent states of oscillation for a related structure with no boundary constraints. This structure, shown in Figure 2-5, is an infinite array oscillator which is composed of an infinite series of coupled rings.

To determine the fundamental state for this structure, 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 that each ring will oscillate at the same frequency. With all rings oscillating at the same frequency, the phase relationship among all nodes in the array will remain constant. 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 a possible mode of oscillation, 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.

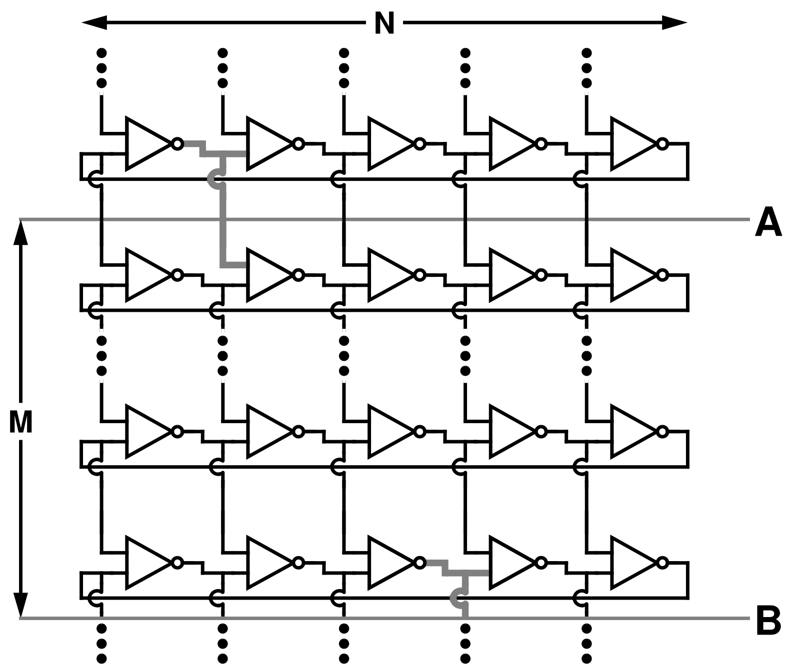

Figure 2-5: Infinite array oscillator.

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 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 in the infinite array oscillator could identically match those of a previous ring, but not at the same ring positions. Suppose that the phase of the ring outputs along line A in Figure 2-5 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 infinite array would behave the same as a finite array with 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 Figure 2-4 is closed as described above, where top array nodes T_i connect to bottom array nodes B_{i+2}. These connections force the top array nodes to lag two buffer delays in phase behind the corresponding bottom array nodes. 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 Figure 2-6 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.

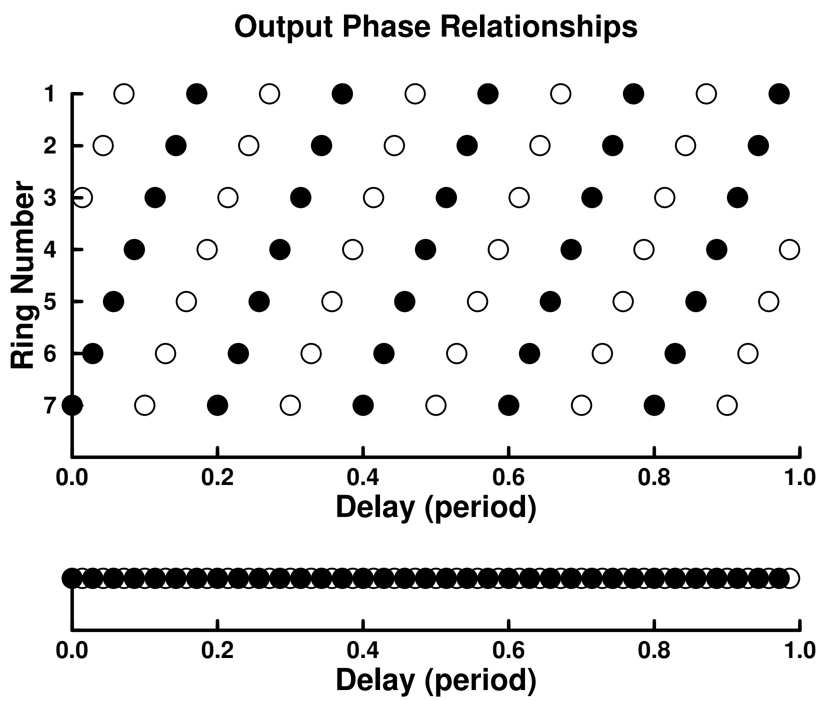

Figure 2-6: Phase relationship among array oscillator outputs for an array with seven rings, each with five buffers.

2.2.2 Array dynamics

The previous section implicitly assumed that the consistent states for the array oscillator are also stable states. Consistent states are stable states only when the phase imbalances that arise from small perturbations drive the array back to the original consistent state. For the states described in the previous section, these imbalances will always be restorative provided that the dual-input buffer delay increases monotonically with increasing delay between the ring and coupling inputs.

The restorative action can be illustrated by again considering an infinite array oscillator. Suppose that initially one ring has a smaller phase difference between the ring and coupling inputs of its buffers than for all previous rings which have an equal phase difference. Further assume that the coupling inputs lag in phase behind the ring inputs. Because of the monotonic functional relationship of the dual-input buffers, the delays of the buffers in this ring will be smaller, causing the ring to oscillate at a higher frequency. With this ring oscillating at a higher frequency, its ring inputs will arrive earlier in phase each cycle relative to its coupling inputs. As this process continues, the oscillation frequency of this ring will decrease as the delay between the ring and coupling inputs increases. Eventually the phase difference between the ring and coupling inputs of this ring will match that of the previous rings. With this equality in phase difference, the oscillation frequency of this ring will match that of the previous rings so that the phase difference between the ring and coupling inputs of this ring will remain fixed, leading to a stable state. Any deviation in the phase difference between the ring and coupling inputs of this ring will be restored in the same manner. The recovery of the phase difference can be modeled as a first order system where the time derivative of the phase difference is proportional to the change in the dual-input buffer delay that results from the error in the phase difference. Thus, the linearity of the functional relationship of the dual-input buffers only affects the rate at which equilibrium is established in the phase relationship among the rings in the array.

To see how this restorative action keeps the consistent states stable in the closed array oscillator, consider a closed array oscillator that has been perturbed from a consistent state. Suppose that the phase difference between the ring and coupling inputs of all buffers are initially equal except for those in one ring which have a smaller phase difference. The ring with a smaller phase difference will initially oscillate at a higher frequency than the other rings. With this ring oscillating at a higher frequency, its ring inputs will arrive earlier in phase each cycle relative to its coupling inputs. As this process continues, the oscillation frequency of this ring will decrease as the delay between the ring and coupling inputs increases. In addition, each following ring will begin oscillating at a higher frequency due to the decreasing phase difference between its ring and coupling inputs. The result is that the smaller phase difference between the ring and coupling inputs of the first ring will disperse with each cycle to the following rings. The dispersion will continue until all rings oscillate at the same frequency with an equal phase difference between the ring and coupling inputs of all buffers. Because all of the ring frequencies are equal in steady-state, the phase relationship among all ring outputs in the array will remain fixed, leading to a stable state. Any deviation from this fixed relationship will be restored in the same manner. As the next section will show, an array will typically have more than one consistent state. Thus, the restorative action will drive the array into the consistent state closest to the initial conditions.

2.2.3 Modes of oscillation

The previous sections 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. However, phase is periodic. If the phase shift between rings is larger, so that the magnitude of 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 discuss the possible phase relationships among the buffer outputs in the array and 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 that the array in Figure 2-4 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 hence a buffer delay $D = \frac{T}{2N}$, then

$$M\Delta t = kD + xT \tag{2-1}$$

where Δt is the delay between corresponding ring nodes in adjacent rings as indicated in Figure 2-4, 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,

$$\Delta t = \frac{xT + k\frac{T}{2N}}{M} \tag{2-2}$$

so that

$$\frac{\Delta t}{T} = \frac{C}{2NM} \tag{2-3}$$

where

$$C = k + x2N \tag{2-4}$$

C is defined to be 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 factor C and will exhibit a different period fraction $\Delta t/T$ as the delay between corresponding ring nodes in adjacent rings. The magnitude of $\Delta t/T$ is not the resultant delay resolution since, for larger values of C, 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 oscillation frequency of the array changes with the mode of oscillation due to changes in the buffer delay of all buffers. For overlapping ring and coupling input transitions, 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 the ring input to the buffer output, will be equal to the delay of a buffer with simultaneous ring and coupling input transitions, less one half of the time the coupling input transition occurs before the ring input transition. This linear approximation agrees well with the simulated delay characteristics shown in Figure 2-3. Thus

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

so that

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

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

The nature of the boundary conditions suggests that an array oscillator will support an infinite number of modes of oscillation periodically spaced by $2N$ in both the positive and negative directions. In actuality, the number of modes that the array will support is quite limited. Oscillations in the array will not occur in modes C far from zero. With large positive modes of oscillation, oscillations are inhibited by a bandwidth limitation in the buffers. This bandwidth limitation prevents them from providing adequate gain to support an increased oscillation frequency as determined by the buffer delay, from the ring input to the buffer output. With large negative modes of oscillation, insight into the limiting effect can be obtained from the equation specifying the dual-input buffer delay as a function of the mode of oscillation for equally sized ring and coupling inputs. As the mode C approaches $-2M$ where the delay between the ring and coupling input transitions is two buffer delays, the equation, as plotted in Figure 2-7, shows that the buffer delay becomes infinite. This result occurs because the array structure and the dual-input buffer operation each impose a relationship between the buffer delay, from the ring input to the buffer output, and the delay between the ring and coupling input transitions. The array forces the delay between the two input transitions to be a precise fraction of the buffer delay as defined by the mode C. For a dual-input buffer, the output transitions are on average delayed by about one single-input buffer delay from the midpoint between the ring and coupling input transitions.[2] Thus, the buffer delay will always be slightly larger than one half the delay between the two input transitions, no matter how large this input delay becomes, making it impossible to achieve a delay between the ring and coupling input transitions that is greater than or equal to two buffer delays. Therefore, as the array forces the delay between the ring and coupling input transitions to approach two buffer delays, the dual-input buffer operation forces the buffer delay to approach infinity. These results suggest that the number of stable modes with C less than zero is an integer near $\frac{M}{N}$ with equally sized ring and coupling inputs.

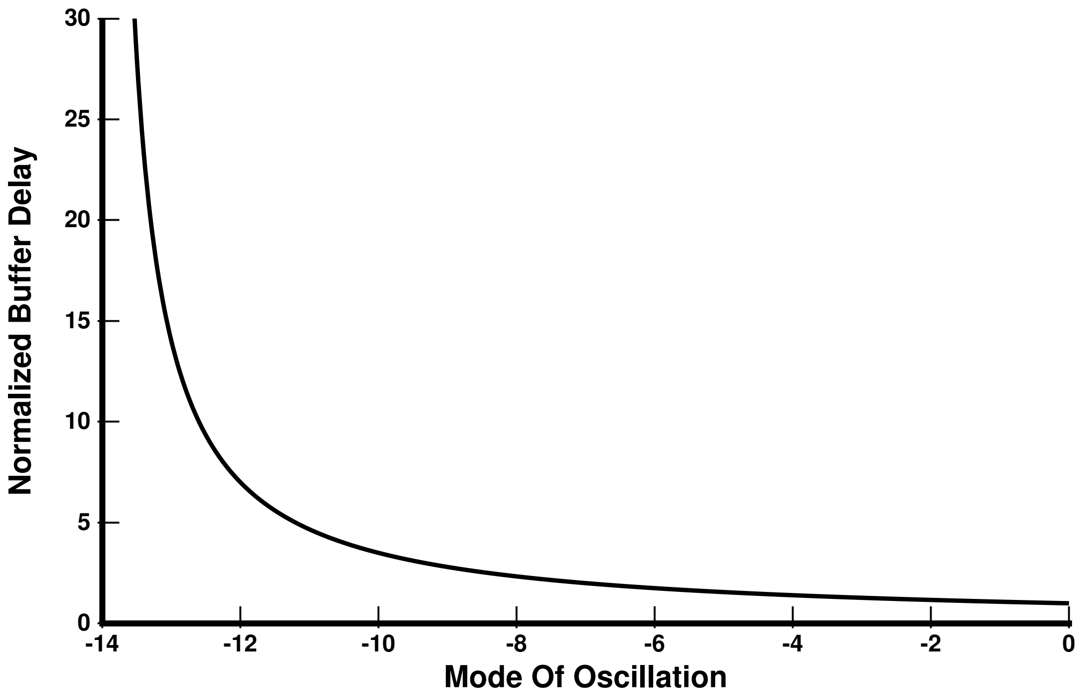

Figure 2-7: Relationship between the dual-input buffer delay and the mode of oscillation for an array with seven rings.

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 mode of oscillation 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

2. SPICE simulations of dual-input buffers confirm that the output transitions are on average delayed by about one single-input buffer delay from the midpoint between the ring and coupling input transitions. They show that the buffers exhibit two extreme types of behavior that depends on the proximity of the buffer input switching voltage to the center of the buffer output voltage swing. If the input switching voltage is at the center of the output voltage swing, all signal transition times are reduced in proportion to the delay between the ring and coupling input transitions so that the output transitions are delayed by about one single-input buffer delay from the midpoint between the two input transitions. This behavior is always exhibited by dual-input differential stages since the switching point is always at the center of the differential voltage swing. Alternatively, if the input switching voltage is away from the center of the output voltage swing, all signals are identically duty cycle shifted by half the delay between the ring and coupling inputs so that the rising and falling output transitions are each delayed by about one single-input buffer delay from transitions on opposite inputs. With the input switching voltage below the center of the voltage swing, the rising transitions are delayed by one half the delay between the ring and coupling inputs. Thus, the rising output transitions occur one single-input buffer delay from the falling coupling input transitions, while the falling output transitions occur one single-input buffer delay from the rising ring input transition.

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

$$\frac{2NM}{GCD(C, M)} \tag{2-7}$$

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

$$\frac{NM}{GCD(\frac{C}{2}, M)} \tag{2-8}$$

where C is even.

Because the oscillation frequency and phase ordering of array outputs change with each mode of oscillation, 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 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

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

for the array oscillator to enter the mode C after the reset operation. The reset operation is accomplished by disabling the coupling in one of the rings so the array of coupled rings no longer forms a closed loop. The desired boundary conditions can then be forced on the array by adjusting the operating frequency of the first ring. Modes with C close to zero are readily achieved since it is usually easy to force all rings in an open array to oscillate in phase. Chapter 4 describes in detail a few methods for disabling the coupling between rings.

2.2.4 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 of the array core occur as a result of random device and capacitance mismatches in the buffers which cause random variations in the buffer delays. To explain the effect of these errors, this section will begin by considering them in the context of a ring oscillator.

If a ring oscillator is composed of buffers with equal delays, all of its output phases will equally divide the oscillation period. However, random variations in the buffer delays will cause the output phases to deviate from their uniformly distributed ideal values. The deviations in the delays of the output phases from their ideal values result in delay errors. Insight into the overall delay precision of a ring oscillator can be obtained by considering the root-mean-square (RMS) error in the delays of its output phases.

For a ring oscillator, the RMS error in the delays of the output phases from their ideal delays can be derived from the RMS variation in the buffer delays from their average value by considering a ring of buffers, each with independent random delays [13]. The result of such a derivation shows that

$$\Delta P = \frac{\sqrt{3}}{6}\sqrt{N - \frac{1}{N}}\Delta D \tag{2-10}$$

where ΔP is the absolute RMS error in the delays of the output phases from their ideal delays, ΔD is the absolute RMS variation in the buffer delays from their average value, and N is the number of buffers in the ring. Thus, the absolute RMS error in the delays of the output phases of a ring oscillator approximately scales with the square root of the number of buffers. However, since the oscillation period is proportional to the number of buffers, this RMS error expressed relative to the period scales inversely with the square root of the number of buffers. It is also important to note that the RMS variation in the buffer delays is typically a constant fraction of their average value. As such, both the absolute and period relative RMS errors will scale inversely with the square root of the number of buffers if the oscillation frequency is held constant by reducing the average buffer delay as the number of buffers is increased.

For an array oscillator, an exact expression for this RMS error cannot be directly derived. Random-process statistical simulations with a fixed absolute RMS variation in the buffer delays show that the absolute RMS error in the delays of the output phases of an array oscillator approximately scales with the square root of the number of buffers in the array, as is the case for a ring oscillator. They also show that this RMS error is minimized when the coupling between buffers is maximized with equally sized ring and coupling inputs. Thus, the absolute accuracy of the array core scales similarly to that of a simple ring oscillator. With a constant number of buffers per ring, both the absolute and period relative RMS errors will approximately scale with the square root of the number of rings. However, with an array that contains a number of rings closely equal to the number of buffers per ring, the period relative RMS error will be roughly constant, independent of the actual size of the array. If the oscillation frequency is held constant, both the absolute and period relative RMS errors will also be independent of the actual size of the array for an array that contains a number of rings closely equal to the number of buffers per ring. The actual accuracy of the complete array oscillator may also be limited by the output channel, as will be discussed in Chapter 4.

2.2.5 Layout issues

Like a simple ring oscillator, the array oscillator's operation depends on all of the buffer delays being identical. In an array oscillator, however, the required matching is more stringent since the sub-buffer delay resolution requires extremely well matched buffer delays.

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 all buffer output nodes must be carefully balanced. Achieving balanced interconnect capacitance can be challenging due to the boundary connections in rows and columns and the shift by k buffers in closing connections.

In a ring oscillator, balanced interconnect capacitance is accomplished by interleaving the buffers in a single row as illustrated in Figure 2-8. Not only can the wires at the ends be made the same length as the wires in the middle, they can also be made to pass over the same underlying material guaranteeing identical interconnect capacitance. This same interleaved strategy can be applied to an array oscillator as illustrated in Figure 2-9. The buffers in each ring can be interleaved to balance the interconnect capacitance due to the horizontal connections between buffers. In addition, the rings in the array can also be

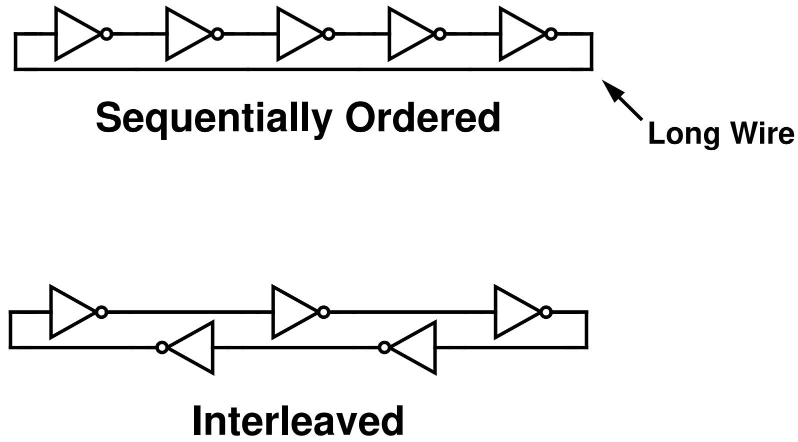

Figure 2-8: Possible floor plans for a ring oscillator.

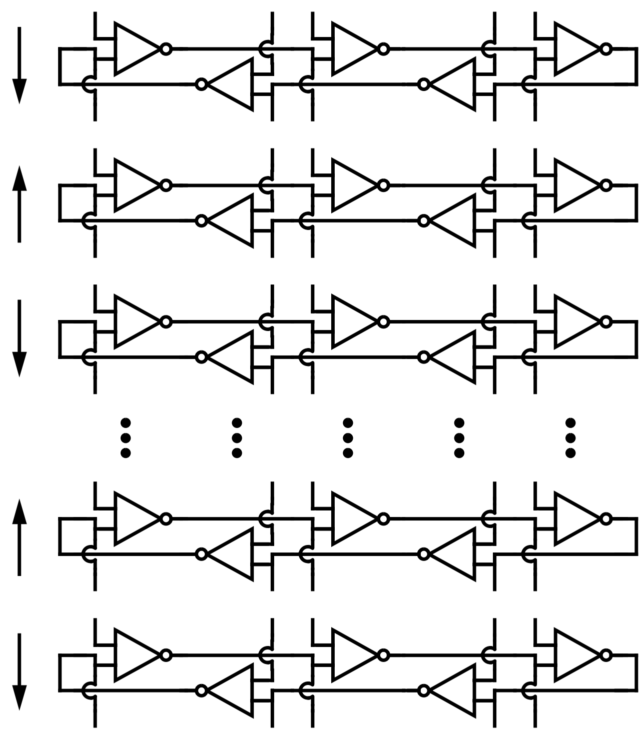

Figure 2-9: Floor plan of the array core illustrating the interleaved buffers in both horizontal and vertical directions.

interleaved to balance the interconnect capacitance due to the vertical connections between rings.

The array closing connections are more difficult to address. The array closing connections require that the wires between two of the rings connect buffers in different ring positions in order to achieve the k buffer shift. These wires need to be of the same length and have the same interconnect capacitance as all the other output wires in the array. It would be advantageous to be able somehow to close the array without connecting buffers in different ring positions. To see how this can be accomplished, consider the non-interleaved version of the array shown in Figure 2-10. Each ring in the array is situated with a single buffer shift to the left with respect to the previous ring. If the number of rings is properly constrained, a net shift of k buffers can be obtained in the closing connections simply by connecting the top array nodes to the bottom array nodes aligned with the columns as shown in Figure 2-10. 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, also illustrated in Figure 2-10. A shift by the number of buffers per ring is equivalent to no shift since the nodes at the boundary of the array will be at the same buffer positions as with the original unshifted rings. Therefore, in order to achieve a net shift of k buffers through M rings, the number of rings M must be constrained so that

$$M = yN - k \tag{2-11}$$

for some positive integer y.

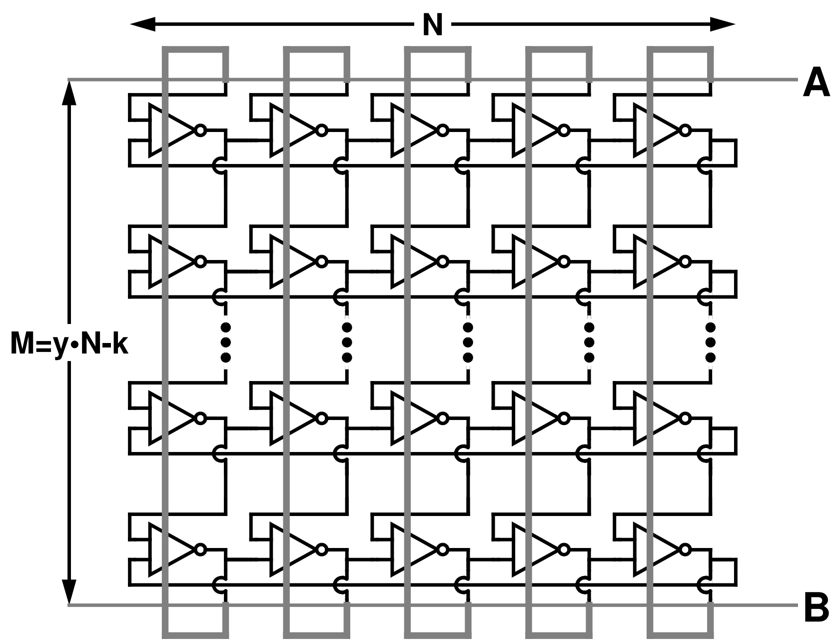

Figure 2-10: Closed array oscillator with single-buffer shift.

The combination of interleaving and single buffer shifts will allow all of the buffers to have identical interconnect capacitance at their outputs. Figure 2-11 illustrates the interleaved buffers in both rows and columns and the single buffer shift in each ring. Adjacently connected buffers are separated by a single buffer and adjacently connected rings are separated by a single ring. All connecting wires travel only in the horizontal and vertical directions between adjacent interleaved buffers in all rows and columns of the array. In an array with seven rings and five buffers per ring, there will be a net shift of two buffers establishing the array boundary conditions.

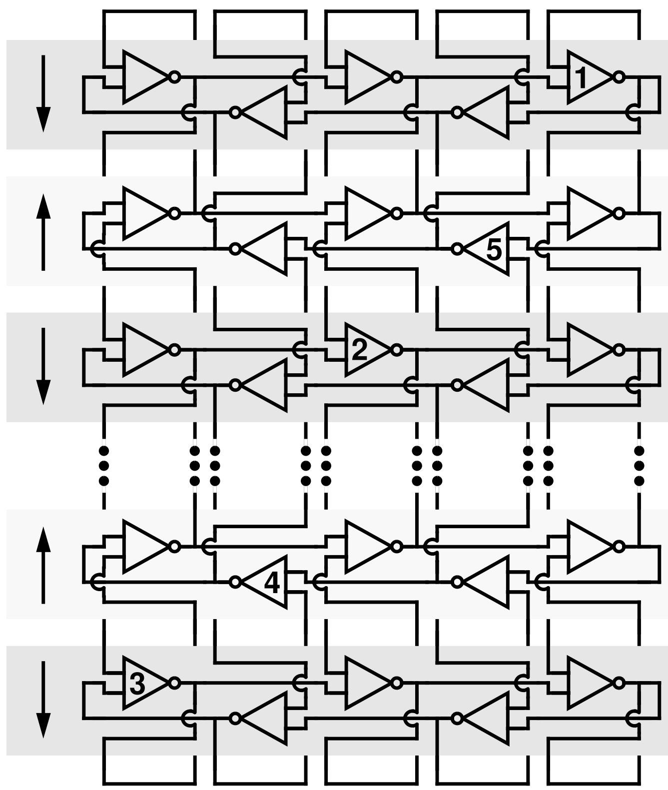

Figure 2-11: Complete floor plan of the array core, illustrating the interleaved buffers in both horizontal and vertical directions and the single buffer shift in every ring. The numbered buffers indicate consecutive buffers along a single logical column.

2.2.6 Summary

In contrast to a simple ring oscillator, the array oscillator can achieve 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 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 scales similarly to that of a simple ring oscillator.

An array oscillator is desirable when an oscillator with delays referenced to the period is required. If a specific delay reference exists or if the use of an oscillator is undesirable, then another structure analogous to a delay line based on coupled oscillators can be used for generating precise delays with high resolution. This structure is described next.

2.3 Delay line oscillator

Another application of coupled oscillator techniques is found in a delay line oscillator. This structure is analogous to a delay line based on a series of buffer stages. However, instead of using buffer stages as delay elements, it uses stages based on ring oscillators called ring stages as shown in Figure 2-12. A delay element can be made from a ring oscillator by replacing one of its single-input buffers with a dual-input buffer. The coupling input of the dual-input buffer is the input to the delay element while the ring input of the dual-input buffer is the output. The delay element has the property that it can delay a periodic input signal over a positive and negative range.

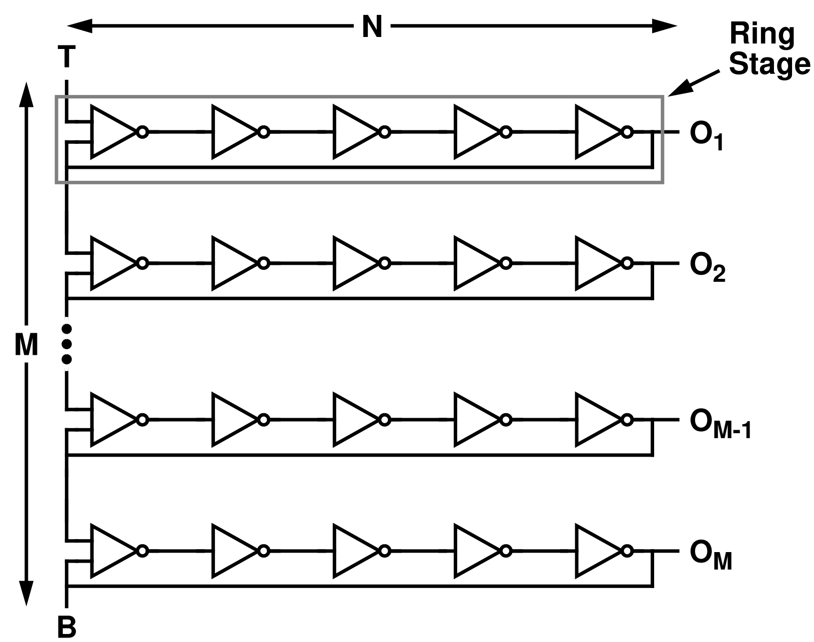

Figure 2-12: Delay line oscillator structure.

A delay line oscillator is different from a buffer delay line in that it can delay periodic signals by arbitrarily small equal delay increments. The delay increment can span a range of over plus or minus one buffer delay. However, unlike a buffer delay line, the delay line oscillator can only operate on periodic signals. Its main application is to subdivide arbitrarily small delay intervals into equal delay increments. A phase-locked loop is used to adjust the ring stage delay so the delay across the delay line oscillator is equal to some periodic delay reference.

2.3.1 Delay line operation

Each ring stage in a delay line oscillator produces an equal delay that can vary over a positive and negative range when provided with an input signal that is identical in wave shape to the signal at the ring stage output. The operation of a ring stage is best understood by first considering a ring stage with its input and output connected together. Since the input and output of the ring stage are the coupling input and ring input of the dual-input buffer respectively, this connection shorts the ring and coupling inputs of the dual-input buffer together and converts the ring stage into a simple ring oscillator. This ring oscillator will oscillate at some characteristic frequency that is referred to as the natural frequency for the ring stage. Because the input and output of the ring stage are connected together, they will be at the same phase. If the input of the ring stage is instead connected to a periodic signal with a frequency equal to the natural frequency, the input and output of the ring stage will again be at the same phase.

A ring stage will oscillate at the same frequency as the periodic input signal provided that it is close to the natural frequency. This frequency tracking is made possible by the dual-input buffer. The dynamics of the frequency tracking are the subject of the next section. In the frequency range where the ring stage tracks the input signal, a constant phase difference will be generated between the coupling and ring inputs of the dual-input buffer or, equivalently, the input and output of the ring stage. This phase difference corresponds to the amount that the coupling input transitions must lead or lag the ring input transitions in order to adjust the dual-input buffer delay so that the ring stage period matches the period of the input signal. It is equivalent to a delay for the ring stage measured from the ring stage input to the ring stage output. If the frequency of the input signal is greater than the natural frequency, then the ring stage will have a positive delay. Similarly, if the frequency of the input signal is less than the natural frequency, then the ring stage will have a negative delay. Since the output of one ring stage is the input to the next ring stage in the delay line, all ring stages will be driven by periodic signals with the same frequency and thus will also produce identical delays.

2.3.2 Delay line dynamics

The operation of the delay line depends on the oscillation frequency of the ring stages tracking the frequency of the input signal so that the phase difference between the ring stage inputs and outputs will be constant. As with the array oscillator, the dual-input buffers in each ring stage of the delay line oscillator establish a restorative action that drives all ring stages to oscillate at the same frequency as that of the input signal and to provide fixed and equal delays. This action will always be restorative provided that the dual-input buffer delay increases monotonically with increasing delay between the ring and coupling inputs.

To see how the restorative action causes the oscillation frequency of a ring stage to track the frequency of its input signal, suppose that the oscillation frequency of the ring stage is higher than the input frequency. With the ring stage oscillating at a higher frequency, its output will arrive earlier in phase each cycle relative to its input so that the phase difference between the ring and coupling inputs will increase. Accordingly, the dual-input buffer delay will also increase, causing a reduction in the oscillation frequency. Since the oscillation frequency of the ring stage is determined in part by the dual-input buffer delay, the increasing dual-input buffer delay will reduce the oscillation frequency. This process will continue until the oscillation frequency matches the input frequency. The opposite occurs if the oscillation frequency of the ring stage is less than the input frequency. The result is that the ring stage will be forced to phase lock to its input signal and will establish a particular delay for the ring stage.

The monotonicity of the functional relationship of the dual-input buffers results in a unique ring stage delay for each input frequency. Thus, the frequency of the input signal effectively provides control over the delay of the ring stages. However, there is a limit to how far the frequency at the input of the ring stage can deviate from the natural frequency. This frequency range of operation corresponds to some minimum and maximum ring stage delay. When the transitions at the coupling input of the dual-input buffer occur significantly before or after transitions at the ring input, they will have little effect on the time of the output transitions. In this situation, the oscillation frequency of the ring stage will no longer track the frequency of the input signal.

2.3.3 Phase-locked operation

With the aid of a phase-locked loop, the delay line oscillator can be used to accurately divide arbitrarily small delay intervals into equal delay increments. A simplified block diagram of a phase-locked delay line oscillator is shown in Figure 2-13. Two periodic input signals define the input delay interval. The first signal is used as input to the delay line oscillator, while the second is compared to the output of the delay line oscillator. The feedback in the PLL adjusts the delay of the ring stages by changing their natural frequency so that the output of the delay line oscillator and the second input signal are at the same phase. The ring stage outputs will then precisely divide the delay interval established by the input signals. Output multiplexers can be used to selectively access the intermediate delay intervals. This PLL configuration is identical to that which could be used for a buffer delay line [1]. This section describes some differences from a simple buffer delay line. These include special treatment of the waveforms supplied to the delay line oscillator and phase comparator, the operating characteristics of the delay line oscillator outside the frequency range of operation, and the reduced possibility of incorrectly phase locking with a delay that spans additional multiples of the oscillation period.

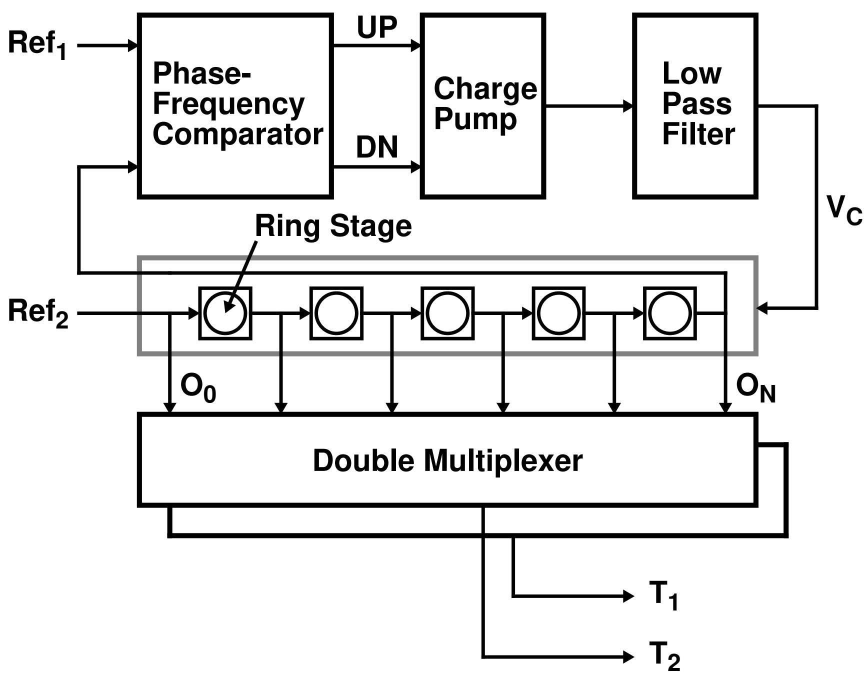

Figure 2-13: Simplified block diagram of a phase-locked delay line oscillator.

When generating delays with high resolution and equivalently high precision, the shape of the input waveforms to the delay line and phase comparator are very important. For the delays generated by all ring stages to be identical, the input waveform to each ring stage must be identical. In addition, to accurately compare the phase of the output of the delay line oscillator with the second input signal to the PLL, the waveforms to the phase comparator must also be identical. Such requirements suggest that the two input signals to the PLL should be buffered by one or two single-input buffers followed by a ring stage before entering the delay line oscillator and phase comparator. The load presented by the phase comparator inputs on the ring stages should be identical to that of a ring stage. Figure 2-14 shows a block diagram of a PLL with the additional buffering. The input waveforms to the complete phase-locked delay line oscillator must also be identical for accurate phase comparison. However, this requirement must be addressed by the overall system.

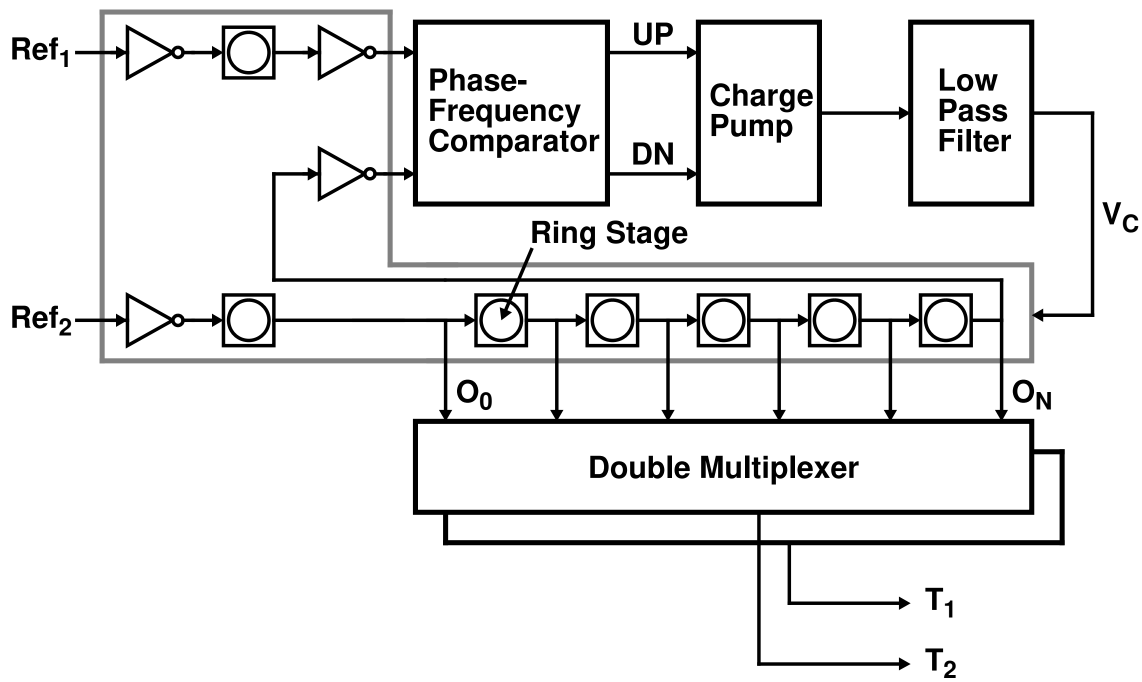

Figure 2-14: Refined block diagram of a phase-locked delay line oscillator.

Inside the frequency range of operation, the delay line oscillator behaves similar to a buffer delay line. With a buffer stage that has delay directly proportional to control voltage, the frequency domain transfer function from control voltage input to delay output is simply equal to a constant. However, outside the frequency range of operation, the delay line oscillator behaves like an oscillator rather than a delay line. The frequency range of operation is analogous to the capture range for a PLL based on an oscillator, since inside these ranges both loops respond like linear systems. Because of this oscillator-like behavior, the gain of the PLL should be limited so that output frequency does not pass over the frequency range of operation in an oscillatory manner when phase locking. The use of a phase-frequency comparator rather than a simple phase comparator will not only allow for an in-phase lock which is needed for subdividing delay intervals, but will also reduce the time required for the PLL to reach the frequency range of operation and prevent the PLL from phase locking to a harmonic or subharmonic of the input signal, similar to a PLL based on an oscillator. The oscillator-like behavior at the boundary of the delay range of operation will always reset the phase-frequency comparator so that it will drive the delay spanned by the delay line oscillator back toward a point where it can phase lock. This behavior eliminates the possibility that the PLL can incorrectly drive the spanned delay toward the boundary of the delay range, which is a typical problem for buffer delay lines.

The delay line oscillator effectively eliminates the possibility of incorrectly phase locking with a delay that spans additional multiples of the oscillation period. This possibility can affect a buffer delay line when the delay range supported is too large. The delays required from a delay line are typically a specific fraction of the input period. With a buffer delay line, the delay range needed must be large enough to accommodate the desired input frequency range and any shifts in the delay range due to process skews. The maximum delay designed into the buffer stages must be large enough to accommodate delays at the lowest input frequency, yet small enough to prevent false locks with delays at the highest input frequency. These constraints will inevitably limit the input frequency range obtainable with a buffer delay line. If a wide range of input frequencies must be supported, then elaborate loop reset mechanisms must be employed with special constraints on when and how the input frequency can change. However, with the delay line oscillator, the delay range over which it operates as a delay line rather than an oscillator is established as a specific and fixed fraction of the input period. The input frequency range, determined by the core buffer stage design, can be made as broad as desired, making process skews inconsequential, without affecting the relative delay range. The relative delay range must be broad enough only to account for the range of delays required with a given input frequency. Thus, with the possibility of a false lock occurring eliminated, the frequency range and absolute delay range achievable by a delay line oscillator can be significantly larger than that for a buffer delay line.

2.3.4 Alternate implementation

The ring stages used inside the delay line oscillator can be constructed using an alternate approach. Rather than using the ring input of the dual-input buffer as the ring stage output, the output of the dual-input buffer can be used instead. Figure 2-15 shows a delay line oscillator with these modified ring stages. These modified ring stages offer a number of possible advantages. The delay range is reduced by a factor of two and shifted more toward positive delays while still passing through zero. In addition, the sensitivity of the ring stage delay to the buffer delay control input is also reduced by a factor of two. This factor of two reduction in delay range and delay control sensitivity result because the delay from coupling input to the buffer output includes only half the delay between the ring and coupling input transitions. More importantly, the coupling between the ring stages is stronger which leads to a reduction in the jitter that results from independent variations in the buffer delays. The jitter that results from correlated variations in the buffer delays is reduced by a factor of two primarily by the reduced sensitivity of the ring stage delay.

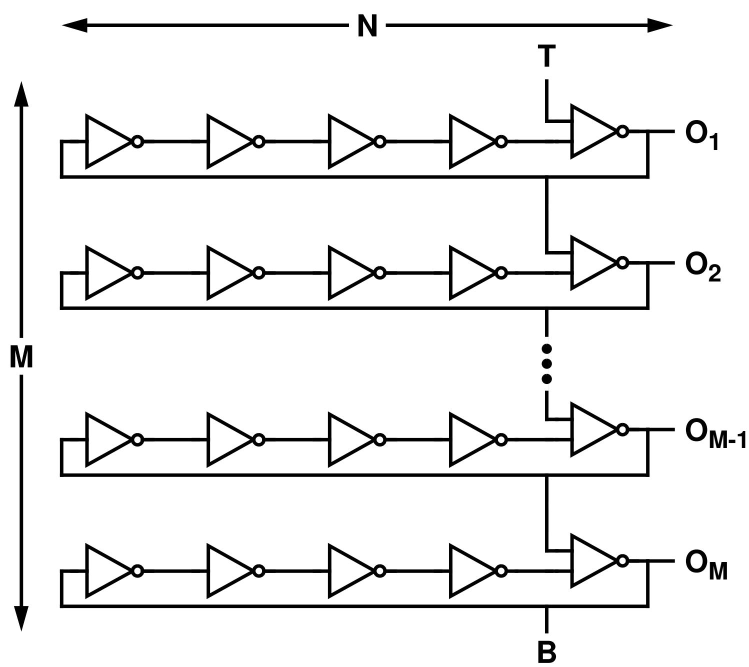

Figure 2-15: Alternate structure for a delay line oscillator.

2.3.5 Summary

A delay line oscillator is a another structure for producing precise delays with high resolution. This structure is analogous to a delay line based on a series of buffer stages with the buffer stages replaced by ring oscillator stages. Unlike a buffer delay line, it can delay periodic input signals by arbitrarily small equal delay increments. A delay line oscillator has some advantages over a buffer delay line in terms of its phase-locked operation. In particular, the input frequency and relative delay ranges can be independently set, allowing the delay line oscillator to support a much larger frequency range and absolute delay range than a buffer delay line without the danger of a false lock occurring. In addition, the property that the delay of a ring stage changes with input frequency makes it very useful in other applications.

Coupled oscillator structures such as the array oscillator and delay line oscillator all obtain their significantly enhanced delay resolution from the dual-input buffer. This buffer has the ability to generate a delay based upon two input signals. An arbitrarily small phase difference between the two input signals can be used as the basic delay unit. When replicated in an appropriately symmetric fashion, these buffers can then provide not only very high delay resolution but also very high precision.

The coupled oscillator structures address only some of the issues that must be resolved to make a precise delay generator with high resolution. The oscillator must also be able to operate over a large frequency range and provide for high supply noise immunity. These issues must be addressed in the buffer stage design and will be discussed in the following chapter.

Chapter 3

Buffer Designs

The buffer stage is the basic underlying element of a coupled oscillator. It must provide a delay that is well defined and identical for all buffer stages. In order to establish the exact value of the generated delay, its delay must be adjustable through some control voltage input. Ideally, the control voltage alone would control the buffer stage delay. However, the delay may be affected by other external factors such as the supply and substrate voltages. Changes in the supply and substrate voltages will typically cause changes in the buffer delays which leads to output jitter. The high resolution and precision offered by coupled oscillator structures impose severe constraints on the amount of allowable jitter produced by the underlying buffer stages since output jitter will decrease the effective precision of the delay generator. As a result, one of the primary considerations in the buffer stage design is to minimize the sensitivity of the buffer stage to supply and substrate noise in order to minimize the amount of output jitter. In the noisy environment of a digital integrated circuit, achieving low output jitter can be difficult due to the level of supply and substrate noise rejection required.

In addition to high noise immunity, there are other constraints on the buffer design. The buffer stage should have low supply voltage requirements to allow fabrication in low-voltage state-of-the-art digital IC technologies and to avoid imposing supply constraints with its use in low-voltage board level designs. It should support a broad frequency range to allow the delay generator to generate delays at a variety of frequencies. The buffer stage should also provide linear control over frequency for oscillators and linear control over delay for delay lines since the stability of the phase-locked loop depends on the loop gain.

This chapter will begin by discussing how noise can lead to output jitter in a buffer stage. It will then explore how output jitter, noise sensitivity, and requirements imposed by a phase-locked loop shape the design of a buffer stage using a simple differential buffer stage as an example. Following this discussion, it will present two particular buffer designs in detail. The first design is a differential buffer stage based on a source coupled pair using load elements with symmetric I-V characteristics. These load elements combined with the differential stage give it very high noise rejection. The second design is a single-ended buffer stage based on a diode-clamped common source device. While the expected jitter performance of the single-ended buffer stage is a little worse than the differential design, its reduced supply voltage requirements, lower power consumption, small size, and overall simplicity make it worth considering.

3.1 Output jitter and noise sensitivity

Supply and substrate noise are typically the primary cause of output jitter in a delay generator that is based on phase-locked ring oscillators or delay lines. The amount of output jitter that will result from supply and substrate noise depends on several factors that include the noise sensitivity of the constituent buffer stages, the architecture of the PLL, and the technique used to measure the jitter. This section will first consider the mechanisms that allow the noise to affect the delay of a buffer stage. It will then discuss the amount of jitter that results from the noise for a given PLL architecture and measurement technique.

The sensitivity of a buffer stage to supply and substrate noise fundamentally determines the amount of output jitter. Supply noise sensitivity can be analyzed by dividing its effect on the buffer stage into two components. One component results from the fact that the buffer delay will depend on the DC value of the supply voltage. Noise induced changes in the DC value of the supply voltage will cause a change in the buffer delay. Accordingly, this component is referred to as static supply noise sensitivity. Static supply noise sensitivity is measured as the percent difference in buffer delay that results between measurements made at two supply voltages divided by the supply voltage difference. The other component results from the fact that the buffer delay also will be affected by transient changes in the supply voltage. Transient changes in the supply voltage caused by noise will result in temporary changes in the buffer delay. This transient dependence is referred to as dynamic supply noise sensitivity. Dynamic supply noise sensitivity is measured as the peak percent change in delay that results immediately after an instantaneous unit step change in supply voltage with the change in delay resulting from the static supply noise sensitivity subtracted out. Substrate noise sensitivity also has static and dynamic components. Because the substrate voltage is a supply voltage for the buffer stage, the effect of substrate noise on the buffer delay will be included with static and dynamic supply noise sensitivity. However, any special considerations unique to substrate noise will be considered separately.

The amount of output jitter depends on whether the delay generator is based on a ring oscillator or delay line and on the jitter measurement technique employed. There are two different approaches for measuring output jitter. The first approach measures the jitter in the period at the output of the delay generator, referred to as period jitter. The second approach measures the jitter between the reference frequency input of the PLL and the output of the delay generator, referred to as absolute jitter. Absolute jitter is usually larger than period jitter. For a delay generator based on a delay line, the jitter of interest is between the reference and the output. If the reference is jitter free, then the jitter between the reference and the output of the delay line will just be the period jitter. For a delay generator based on a ring oscillator, either type of jitter can be of significance depending on the system organization. If the reference input is used as one of the outputs of the delay generator, then the absolute jitter must be considered. However, if the reference input is only used by the oscillator to establish the value of the oscillation period so that all outputs used originate from the oscillator, then only the period jitter is of interest. Because the period jitter for a ring oscillator and delay line are typically the same, and absolute jitter is distinct only for a ring oscillator, the following discussion will consider period jitter only for a delay line and absolute jitter only for a ring oscillator.

Dynamic supply noise sensitivity will cause the delay of buffer stages to be affected by a transient change in supply voltage for the extent of the transient. For a delay line, a change in the buffer delay will only result in a change in the phase difference between its input and output during the transient. A supply voltage transient will then result in a transient change in the phase difference between the input and output approximately equal in magnitude to the change in the buffer delay multiplied by the number of buffer stages. The resultant period jitter, equal to this change in the phase difference, will persist only for the extent of the supply voltage transient. However, for an oscillator, the absolute jitter is more significant. The buffer delay only determines the oscillation period as the buffer stages are not connected directly to the reference input of the PLL. In the absence of supply noise, the feedback in the PLL will keep the oscillation period precisely equal to the period of the reference input. A supply voltage transient will result in a transient change in the oscillation period. This transient change will then cause the phase of the oscillator to deviate from the phase of the reference. Since this phase error is feed back to the input of the oscillator, a phase step will be observable between the reference and the oscillator output. The phase step will persist until it can be corrected by the feedback loop in the PLL at a rate limited by the loop bandwidth. The resultant absolute jitter will equal the phase step in magnitude.

Static supply noise sensitivity means that the delay of buffer stages varies with a static change in the supply voltage. For a delay line, a supply voltage change will result in a change in the delay across the delay line stages. The change in delay will persist until it can be corrected by the PLL at a rate limited by the loop bandwidth. The resultant period jitter will equal the peak change in the delay, which is equal to the supply voltage change times the static supply noise sensitivity divided by the operating frequency. For a ring oscillator, a supply voltage change will result in a change in the ring oscillation frequency or period. The effect on the absolute jitter can be quite significant. Such an oscillation period change will cause the phase difference between the reference and the oscillator output to accumulate over time to form a total phase error that increases with time. The PLL will compensate for such period changes at a rate limited by the loop bandwidth. The resultant absolute jitter will approximately equal the supply voltage change times the static supply noise sensitivity divided by the loop bandwidth. Since the loop bandwidth must typically be less than one tenth of the operating frequency, this phase accumulation is a serious issue for ring oscillators. As such, it makes static supply noise sensitivity much more significant than the dynamic supply noise sensitivity.

3.2 Buffer design considerations

The output jitter and noise sensitivity issues discussed in the previous section must be carefully addressed in a buffer stage design for it to have high supply and substrate noise immunity. They will each fundamentally shape the design of the buffer stage. In addition to the noise issues, the use of the buffer stage within a phase-locked loop imposes additional requirements on the buffer stage design. The buffer stage design must have adequate voltage gain, low minimum operating supply voltage, and appropriate delay range control.

This section will begin by considering the simple differential buffer stage shown in Figure 3-1. This buffer stage is based on an NMOS source coupled pair driving two resistive load elements and biased by a simple NMOS current source. The buffer stage is basically a differential amplifier operating in the large signal regime where the NMOS source coupled pair acts as a current switch. Its basic mode of operation is to invert and amplify the input signal and provide a unit delay. The output swing is determined by the resistive load elements and the simple NMOS current source. The buffer delay is primarily determined by the effective resistance of the load elements and the capacitive load at the outputs.

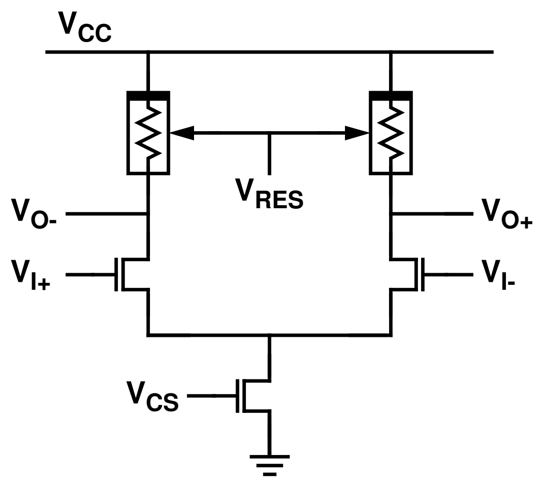

Figure 3-1: Schematic of a simple differential buffer stage.

3.2.1 Noise rejection

A basic requirement of the buffer stage is for the buffer delay to be adjustable. Since the buffer delay is determined by the resistance of the load elements, this resistance must also be adjustable. However, adjustable resistive loads made with real MOS devices will typically have non-linear I-V characteristics over the output voltage swing range. Non-linear loads have important implications on the noise rejection properties of the buffer stage.

One of the most important issues to be addressed in a buffer stage design is static supply noise rejection, since it will typically dominate the phase-locked jitter performance of both oscillators and delay lines. The differential buffer stage in Figure 3-1 would achieve very high static supply noise rejection if the load elements were linear resistors since the resistance of the load elements would be independent of supply voltage. However, since the load elements must provide an adjustable resistance, they will inevitably be non-linear. With non-linear loads, the effective load resistance is strongly determined by the differential pair bias current. Thus, with non-linear loads, the level of static supply noise rejection that can be achieved will be determined by the sensitivity of the differential pair bias current to supply voltage. The static supply noise sensitivity can also be limited by the supply voltage dependency of non-linear diffusion capacitances at the drains of the differential pair devices. However, the effect of these non-linear capacitances will typically be very small.

The differential buffer stage in Figure 3-1 can only achieve limited static supply noise rejection due to the limited output impedance of the simple NMOS current source. Since the output swings are referenced to the top supply, a change in supply voltage will cause an equal change in the drain voltage of the simple NMOS current source device. A change in drain voltage will cause a change in drain current due to the limited output impedance. A change in drain current will then lead to a change in the buffer delay. Thus, a current source with limited output impedance will result in a buffer delay dependency on supply voltage. The supply noise rejection of this buffer stage can be improved by using a current source with a higher output impedance. However, supply voltage limitations will impact the options available for accomplishing the current source isolation. The supply isolation can be accomplished by using cascode current [14] sources or supply filtering [15]. With supply filtering, an intermediate supply voltage is produced from a reference voltage that is independent of supply voltage. This approach is more conducive to buffer designs where delay control is achieved by adjusting the intermediate supply voltage. When inadequate supply voltage is available for cascoded current sources, simple current sources can be used with an actively adjusted bias voltage to compensate for their limited output impedance. Such an approach can give an effective isolation impedance as determined from the dependency of the output current on supply voltage that is equivalent to that obtainable with cascode current sources while requiring at least $V_{GS} - V_T$ less supply voltage. This dynamically-biased current source approach is discussed in detail for the buffer designs presented later in this chapter. The complexity of the buffer stage biasing circuits used is an important issue since the biasing approach can affect both the static and dynamic supply noise rejection of the buffers.

The next issue that needs to be addressed is substrate noise rejection, as it will have important implications on the topology of a buffer stage. To achieve high immunity from substrate noise, care must be taken regarding how the buffer is controlled and biased. For an N-well technology, variations in the substrate voltage will affect the operation of NMOS devices through variations in their threshold voltage. Suppose that the control voltage is used to bias the simple NMOS current source in the differential buffer stage shown in Figure 3-1. Since the control voltage is effectively a fixed bias voltage established by the feedback in the PLL, this current source will have an output current that is very sensitive to substrate noise. PMOS devices, however, are located in local wells and thus are isolated from variations in the substrate voltage. For this reason, in order to minimize the sensitivity of a buffer stage to substrate noise in an N-well technology, the control voltage should always be referenced to the positive supply and only connected to PMOS devices. In this case, the possibility is eliminated that substrate noise can directly modulate the control voltage through changes in the threshold voltage, as would occur if NMOS devices were connected to a control voltage referenced to the negative supply. In addition, this choice of reference allows the loop filter to be placed in an N-well, minimizing the capacitive coupling of changes in the supply voltage and in particular the coupling of substrate noise. In general, any devices requiring externally produced bias voltages should be PMOS devices, while all biased NMOS devices should have their bias voltages produced locally referenced to other NMOS devices through current mirrors or the equivalent, so that any variations in the threshold voltage are compensated for by changes in the bias voltages.

This observation about the control voltage combined with the biasing approach can have other important implications on the topology of a buffer stage. For example, a design that uses cascoded current sources can be based on either NMOS or PMOS current source devices. With the control voltage referenced to the top supply, both variations will use it as the simple current source bias voltage for PMOS cascode current sources. Designs using NMOS cascode current sources will obtain their bias voltages by mirroring the output of a PMOS cascode current source. Although the use of NMOS cascode current sources has the advantage of isolating the control voltage from the buffer stages through the current mirror, it can have the disadvantage of increasing the supply voltage requirements of the buffer stage design for the generation of the NMOS cascode current source bias voltages. In contrast, a design that uses dynamically-biased simple current sources will be limited to use only NMOS current source devices, where control voltage is used to produce the NMOS current source bias voltage and possibly to bias other PMOS devices. This constraint originates from the fact that the control voltage must be used either to establish some property of the output voltage swing which is referenced to the opposite supply to which the current sources are connected, or to directly bias devices of the opposite type to the current source devices. In either case, with the control voltage referenced to the top supply, the dynamically-biased current sources must use NMOS devices. The result is that different biasing approaches may favor complementary versions of a buffer stage design with the NMOS and PMOS device types interchanged.

With the topological and biasing issues relating to static supply and substrate noise rejection worked out, the issues concerning dynamic supply noise rejection can be addressed. For differential buffer designs like that in Figure 3-1, the effect of dynamic supply noise can be minimized if its load elements provide a differential-mode resistance, which determines the buffer delay, that is independent of the common-mode voltage carrying the supply noise. A common-mode shift in the output voltage of a differential buffer cannot affect its delay if the differential-mode resistance at the outputs remains constant. This differential-mode resistance is sometimes established directly by a resistive element between the differential outputs [10]. For single-ended buffer designs, the effect of the dynamic supply noise can only be limited by providing strongly restored output levels. In general, given that dynamic changes in the supply voltage will always disturb the output nodes of the buffer stages causing some variation in the buffer delays, the effect of the supply noise can be minimized by reducing the capacitive coupling to the buffer outputs and reducing the recovery time of the buffer output nodes. The capacitive coupling to the buffer outputs can be minimized by routing them to the greatest extent possible over either N-well or substrate regions depending on the supply to which the buffer output voltages are referenced.

Another important issue that can affect the noise rejection of a buffer stage involves capacitive coupling to the control voltage and the bias voltages. This issue will strongly impact the design of the bias generator. Capacitive coupling of supply voltage changes to the control voltage and bias voltages can significantly reduce the static and dynamic noise rejection of a buffer. Capacitive coupling to the control voltage is the most significant of the two since in a PLL the control voltage is only stored on capacitors in the loop filter. Any capacitance from the control voltage to the non-referenced supply will allow changes in the supply voltage to directly couple into the control voltage and will appear as increased static supply voltage sensitivity. Capacitive coupling to the bias voltages will only result in temporary disturbances since they are actively restored by the bias generator and, thus, will appear as increased dynamic supply noise sensitivity. Capacitive coupling to the control voltage can be reduced by minimizing the number of devices that connect to it. If the control voltage is used as a bias voltage in the buffer stage for the load elements or current source, it may be advantageous to buffer it. Capacitive coupling to the bias voltages can be reduced by minimizing the capacitive coupling caused by the interconnect capacitance.

3.2.2 Buffer operation

With the noise issues addressed, the voltage gain, output voltage swing, and delay range control for the buffer stage should be considered. They are all determined by the load elements, making the load elements central to the performance of the buffer stage.

The control range over which a ring of buffer stages will oscillate can be limited by the minimum voltage gain of the buffer stage. In order for oscillations to occur in a ring oscillator, the buffers must have a minimum gain at the DC bias point, the point where the input and output voltages are equal. Small signal analysis [16] will show that oscillations will only occur for

$$A \geq \sec^2 \frac{\pi}{N}, \qquad N \geq 3 \tag{3-1}$$

where A is the buffer gain at the DC bias point and N is the number of buffer stages in the ring. With five buffers per ring, this minimum gain corresponds to about 1.5. Although the voltage gain is determined by a number of factors, it is most significantly affected by changes in the I-V characteristics of the load elements over the control range. This minimum gain will typically limit the control range over which oscillations are possible for many buffer stage designs. A control range limitation will in turn limit the output frequency range.

The load elements provide control over the output voltage swing which is critical to both the static and dynamic noise performance of a buffer stage. Imposed limits on the output voltage swing will enable supply isolating current sources to always operate in their high impedance region allowing the attainment of high static supply noise rejection. The sensitivity of the output voltage levels to dynamic supply noise can strongly influence the buffer delay. The DC swing limits of a buffer are the same as those established by the output voltages when the buffer is configured as a latch. The transient swing limits of buffers in a ring oscillator may fall short of the DC swing limits. However, they must include the output swing range for which the gain exceeds the minimum gain for oscillations. The actual transient swing limits will depend on the number of stages and the gain as a function of the output voltage.

In order to maximize the stability of a PLL over its input operating range, linear control over delay or frequency for the constituent delay line or oscillator is desired. Because the loop gain is proportional to the gain of the delay line or oscillator, non-linearities in the delay or frequency versus control voltage will cause a range of possible loop gains. However, the loop gain of a PLL is very important in establishing its phase margin. The larger the range of possible loop gains, the more difficult the PLL will be to stabilize. Linear control over delay or frequency requires the buffer delay to change linearly with control voltage for delay lines and inverse linearly with control voltage for oscillators. Control over the delay of the buffer stages is also provided by the load elements. To provide delay control, the load elements must have an effective impedance over the buffer's output swing range that is controllable.

The next two sections present specific designs for a differential and a single-ended buffer stage. Both these designs incorporate all of the considerations discussed in this section with particular emphasis on high supply and substrate noise rejection and minimum supply voltage requirements.

3.3 Differential buffer design

Differential buffer stages are desirable for use in oscillators and delay lines because of their ability to achieve high dynamic supply noise rejection. In a CMOS IC technology, a differential buffer stage is typically based on a source coupled pair driving two load elements and biased by a high impedance current source, as illustrated in the previous section. However, considerable flexibility surrounds its design, mainly with respect to the load element structure and biasing approach. The load elements are chosen to provide control over the buffer delay, to limit the output voltage swing, and to provide high dynamic supply noise rejection. The biasing approach is selected, not only to obtain high isolation from one of the supplies for achieving high static supply noise rejection, but also to minimize the supply voltage requirements of the buffer design.

This section will describe a differential buffer stage that is designed to have high supply noise immunity while being able to operate at low supply voltages for use in coupled oscillators for precise delay generation. 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. The current source bias circuit enables the buffer stage to achieve high static supply rejection without cascoding through the use of self-biasing and replica-feedback techniques. The symmetric load elements used in the buffer stage provide for high dynamic supply rejection through a first order cancellation of noise coupling.

3.3.1 Differential buffer stage

The dual-input differential buffer stage is based on an NMOS source coupled pair with symmetric load elements and a dynamically-biased simple NMOS current source, as shown in Figure 3-2. 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. 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.

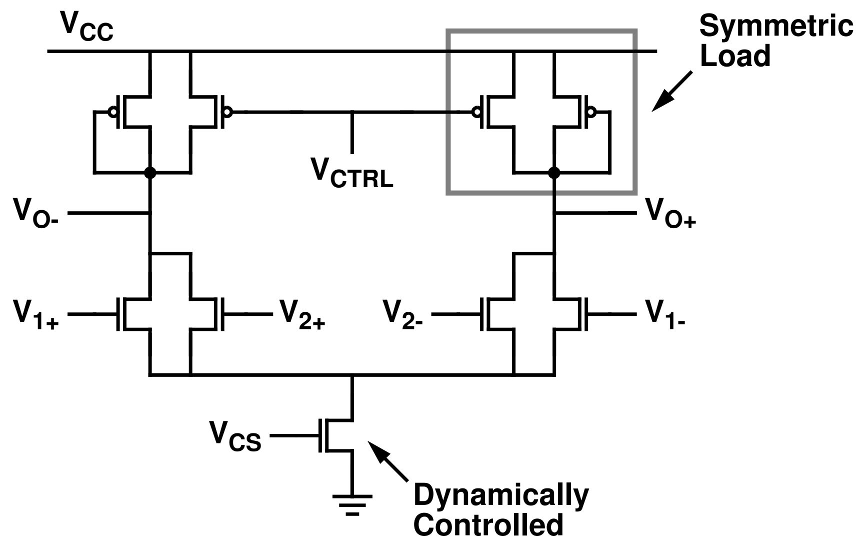

Figure 3-2: Schematic of the dual-input differential buffer stage, containing symmetric loads and a dynamically-biased current source.

The devices used in the buffer stage are sized in order to equalize the mobility-scaled current densities. This strategy provides a good trade-off between speed, required supply voltage, and device area. As such, the sum of the widths for the NMOS source coupled pair devices is set to equal the width of the NMOS current source device. The sum of the widths of all devices in both PMOS load elements is set to equal two times the width of the NMOS current source device. The ring and coupling input source coupled pair devices are typically of equal width. All devices are of minimum channel length.

3.3.2 Symmetric loads

For achieving high dynamic supply noise rejection, the load elements should ideally have linear I-V characteristics. Dynamic variations in the supply voltage of the buffer stages will cause disturbances in the buffer output voltages as a result of parasitic coupling capacitance between the output nodes and the non-referenced supply. Because these disturbances are common to both outputs, they disturb the common-mode output voltage of the buffers. Linear characteristics provide differential-mode resistance that is independent of common-mode voltage. Since the delay of the buffer depends only on the differential-mode resistance, it is not affected by the common-mode disturbances resulting from supply noise. Unfortunately, adjustable resistive loads made with real MOS devices cannot maintain linearity while generating a broad frequency range.

There is a class of non-linear loads that perform almost as well dynamically as linear resistive loads for small supply disturbances. These loads, called symmetric loads, have the property that their I-V characteristics are symmetric about the point at the center of the voltage swing for the extent of the voltage swing. The theory behind these loads is described in the following section.

3.3.3 Theory of symmetric loads

Symmetric loads obtain their high dynamic supply noise rejection as a direct result of their symmetric I-V characteristics. Although non-linear load resistances normally convert common-mode noise into differential-mode noise which affects the buffer delays, the symmetric I-V characteristics of symmetric loads inhibit this noise coupling. The output voltages of a differential buffer with symmetric loads during undisturbed operation are always symmetric about the center of the voltage swing which is also the common-mode voltage level. This symmetry results from the fact that the charging current at each differential output node is equal for complementary output voltages within the output voltage swing which leads to equal rates of change in the output voltages. Thus, at any point in time, the incremental resistances at each differential output node is nominally equal. Since the incremental resistance and capacitance at each output node is nominally equal, a dynamic supply variation will cause equal voltage shifts on the output nodes independent of the edge rate of the dynamic supply variation. With a small dynamic supply variation, the change in the incremental resistance will be very small. With no change in the incremental resistance, the equal voltage shifts will cause the charging current for the capacitance at each output node to change by an equal and opposite amount. As a result, the total differential-mode charging current for the two output nodes will remain constant and independent of the dynamic supply variation. Since the buffer delay depends only on the differential-mode voltage of the output nodes, the buffer delay will be unaffected by the dynamic supply variation.

Consider a differential buffer with non-symmetric non-linear loads for comparison. Because the I-V characteristics of the load elements are not symmetric about the center of the swing, the two differential output nodes with complementary output voltages will each experience different charging currents which will cause their voltages to change at different rates. With the output voltages changing at different rates, the output voltages will be asymmetric about the center of the voltage swing with no defined common-mode voltage level. Thus, at any point in time, the incremental resistance at each output node will typically be unequal. With unequal incremental resistance at each output node, a dynamic supply variation will cause unequal voltage shifts on the output nodes that depend on the edge rate of the dynamic supply variation. With unequal incremental resistances and voltage shifts, the charging current at each output node will be unequal in magnitude. As a result, the total differential-mode charging current for the two output nodes will change due to the dynamic supply variation which will cause the buffer delay to also change.

A mathematical representation of differential-mode output voltage will give more insight into the actual performance of a differential buffer with symmetric loads when exposed to dynamic supply variations. It is assumed that the output voltages will remain symmetric after a common-mode disturbance except for a common-mode offset. The change in the differential-mode output voltage is governed by the differential equation

$$I_{DIFF} = -C\frac{d}{dt}(2V_L) \tag{3-2}$$

where I_{DIFF} is the differential-mode output current that charges the output capacitance, V_L is one half the differential-mode output voltage, and C is the capacitance at each differential pair output node. Since I_{DIFF} is simply the difference between the two output currents, this equation can be expressed more completely by the equation

$$\left(f\left(\frac{V_S}{2} + V_L + V_{OS}\right) - I_O\right) - f\left(\frac{V_S}{2} - V_L + V_{OS}\right) = -2C\frac{dV_L}{dt} \tag{3-3}$$

where $I = f(V)$ is the I-V characteristic of load, V_S is the voltage swing, V_{OS} is the common-mode offset voltage generated by the dynamic supply variation, and I_O is the differential pair bias current. For this equation V_L has an initial value of $-V_S/2$ and a final value of $V_S/2$. Given that *f*(*V*) is symmetric about $V_S/2$ so that

$$f\left(\frac{V_S}{2} + V_L\right) = I_O - f\left(\frac{V_S}{2} - V_L\right) \tag{3-4}$$

a Taylor's Series expansion of *f*(*V*) about $V_S/2 + V_L$ will show that all of the odd derivative terms cancel out. With the higher-order terms beyond the second derivative ignored, the change in the differential-mode output voltage is given by

$$f\left(\frac{V_S}{2} + V_L\right) + f''\left(\frac{V_S}{2} + V_L\right)\frac{V_{OS}^2}{2} - I_O = -C\frac{dV_L}{dt} \tag{3-5}$$

This equation shows that the coupling of the common-mode offset voltage to the buffer delay is significantly reduced by the cancellation of the first order term which leads to a substantial reduction in the jitter caused by common-mode noise present on the supplies. Comparing the mathematical result to the qualitative description, it becomes clear that symmetric loads do not perform quite as well as linear resistive loads only because the incremental conductance of the output nodes will change a little after the dynamic supply variation, as expressed by the second derivative term. Thus, the dynamic supply noise rejection of a differential buffer with symmetric loads scales with the square of the noise level on the supply. It is minimized when the curvature of the I-V characteristics of the loads is minimized, the limit being linear resistive loads.

The amount of jitter that results from dynamic supply variations also relates to how quickly the common-mode voltage level is restored by the differential buffer stages. The delay of a buffer stage can only be affected by a dynamic supply variation during or shortly before its output transition. The time during which the delay can be affected before the output transition is determined by the common-mode recovery time of the buffer stage. Since only the delay to the logic threshold is of interest, the delay can only be affected during the time prior to the middle of the output transition. Symmetric loads provide a low common-mode resistance so that the common-mode voltage recovery time will be relatively small, allowing a small window of opportunity for dynamic noise to affect the buffer delay.

3.3.4 MOS circuit implementation of symmetric loads

A buffer stage based on a simple symmetric load element is shown in Figure 3-3. The load element contains a diode-connected PMOS device in shunt with an equally sized biased PMOS device to produce symmetric load characteristics.

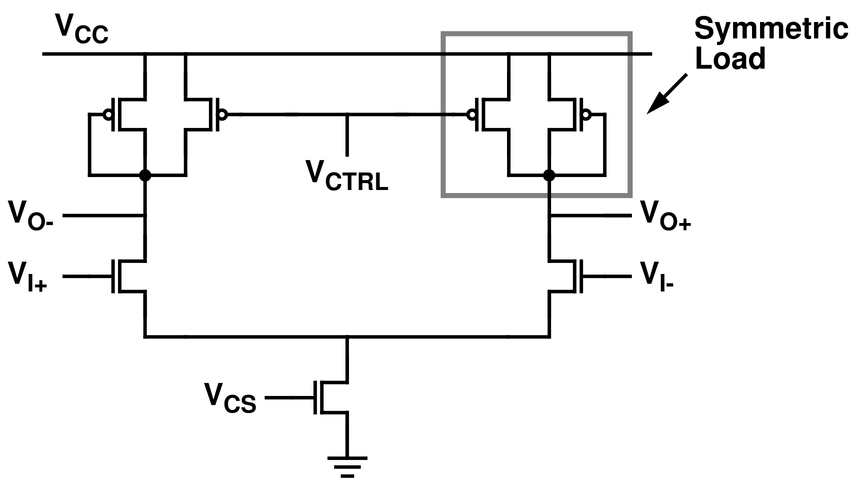

Figure 3-3: Schematic of a differential buffer with MOS symmetric load elements.

Figure 3-4 contains simulated 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 voltage for 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. In order to maintain the symmetric I-V characteristics of the loads, the current source bias circuit is designed to adjust the buffer bias current so that the output swings vary with the control voltage rather than being fixed.

By considering the standard quadratic model for MOS devices, it can be shown that the I-V characteristics of symmetric loads are completely symmetric about the center of the voltage swing as defined from the top supply to the bias voltage for the biased PMOS device. Ignoring the output conductance in the saturation region, the load current I_L can be very succinctly expressed as a function of the load voltage drop V_L when both are referenced to the point of symmetry. At the point of symmetry, $I_L = I_B/2$ and $V_L = V_B/2$, where V_B is the bias voltage of the biased PMOS device which is also equal to the voltage swing, $I_B = k(V_B - V_T)^2$ is the differential pair bias current, k is the device transconductance of both PMOS devices, and V_T is the device threshold voltage. The load current

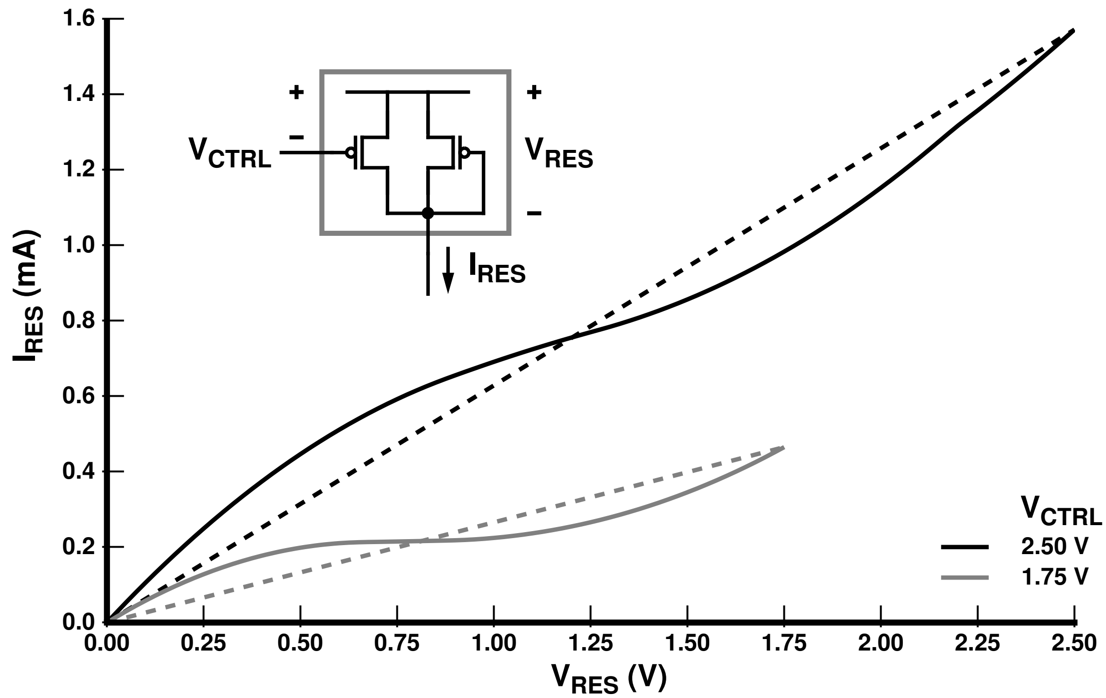

Figure 3-4: Simulated symmetric load I-V characteristics at low and mid-range bias voltages. The dashed lines show the effective resistance of the loads and highlight the symmetry of the I-V characteristics.

$I_L' = I_L - I_B/2$ expressed as a function of load voltage $V_L' = V_L - V_B/2$, both referenced to the point of symmetry, is given by

$$I_L' = -\frac{k}{2}(V_X - V_L')^2, \qquad V_L' < -|V_X| \tag{3-6}$$

$$I_L' = 0, \qquad |V_L'| < -V_X \tag{3-7}$$

$$I_L' = 2kV_XV_L', \qquad |V_L'| < V_X \tag{3-8}$$

$$I_L' = \frac{k}{2}(V_X + V_L')^2, \qquad V_L' > |V_X| \tag{3-9}$$

where $V_X = V_B/2 - V_T$ is the region crossover voltage. Output conductance in the saturation region can be included by modeling it as a shunt resistance. This resistance, however, will be large compared to the average resistance of the load element.

The PMOS bias voltage for the load element must be set so that the load current at the point of symmetry equals one half of the differential pair bias current. The PMOS bias voltage can be generated simply by connecting the bias voltage to the output of the load element biased by the differential pair bias current. This connection will establish the PMOS bias voltage as the lower DC voltage swing limit, the point where the load current equals the differential pair bias current. Alternatively, the differential pair bias current can be established for a given PMOS bias voltage using a replica-feedback bias circuit as described in the next section.

The MOS realization of symmetric loads has additional advantages beyond their high dynamic supply noise rejection. Because the higher gain region always occurs at the center of the voltage swing, the buffers will provide adequate gain for oscillations over a broad frequency range. Furthermore, because the load resistance of symmetric loads decreases towards 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 only contain two equally sized devices of similar size to the differential pair and current source devices used in the buffer, the IC layout for the differential buffer stage will be very compact.

Simulation results confirm that symmetric loads realized with a diode-connected PMOS device in shunt with an equally sized biased PMOS device can provide high dynamic supply noise rejection. SPICE simulations of a ring oscillator with all output interconnect capacitance connected to the negative supply and with the control voltage fixed relative to the top supply, show that an instantaneous 500-mV supply voltage step only results in a total phase error of 0.5% of an oscillation period with static effects factored out. Simulations conducted under the same conditions for the differential buffer described in [14], which utilize a different non-symmetric but resistive load element, show that the supply voltage step results in a total phase error of 0.7% of an oscillation period.

3.3.5 Current source bias circuit

Figure 3-5 shows a simplified schematic of the current source bias circuit which achieves two basic functions. First, it sets the current through the simple NMOS current source in the buffers in order to provide the correct symmetric load swing limits. Second, in order to counteract the effect of the finite output impedance of the simple NMOS current source to achieve high static supply noise rejection, it dynamically adjusts the NMOS current source bias so that this current is held constant and highly independent of supply voltage.

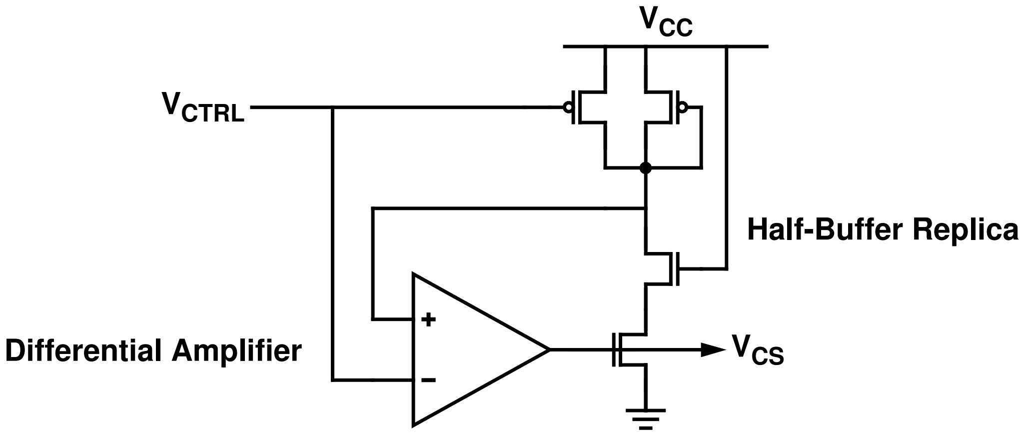

Figure 3-5: Simplified schematic of the self-biased replica-feedback current source bias circuit for the differential buffer stage.

The current source bias circuit is based on 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 device changes. However, the gate bias is adjusted by the amplifier to keep the output current constant, counteracting the effect of the finite output impedance.

The complete schematic of the current source bias circuit is shown in Figure 3-6. In order for the amplifier not to limit the supply voltage operating range of the buffers, it too must have low supply voltage requirements. For this reason, an amplifier based on a self-biased PMOS source coupled pair is used. In order for the PMOS current source device in the amplifier to remain in saturation, the current densities in the PMOS source coupled pair devices and the PMOS current source device must be one quarter that in the PMOS symmetric load devices. 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. Because the amplifier is self biased, the bias circuit is stable with the amplifier unbiased and an NMOS current source bias at the negative supply. As a result, an initialization circuit is needed to bias the

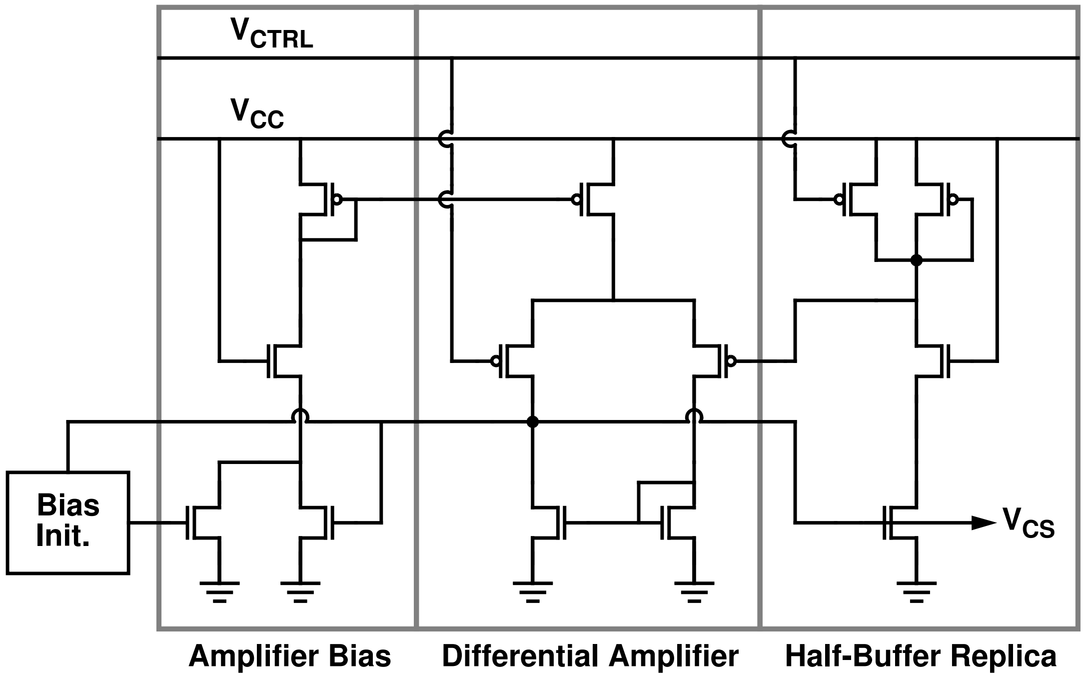

Figure 3-6: Complete schematic of the self-biased replica-feedback current source bias circuit for the differential buffer stage.

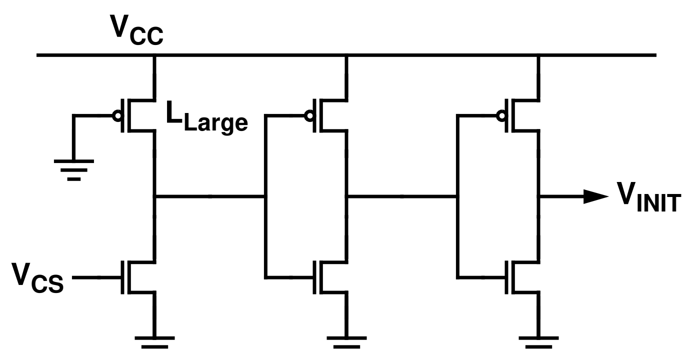

Figure 3-7: Schematic of the initialization circuit for the current source bias circuit.

amplifier at power-up. This initialization circuit, shown in Figure 3-7, prevents the NMOS current source bias from completely turning off the current sources.

Because this current source bias circuit contains a feedback loop with two gain stages and with two significant poles, some frequency compensation is required. Since the pole at the amplifier output will dominate with a much higher output impedance than for the pole at the half-buffer replica output, it can be moved down in frequency to increase the phase margin of the bias circuit simply through the capacitive output load of the simple NMOS current source gates in the buffer stages. 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.

With no required swing reference voltage, the only external 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 with identical current densities.

For a large structure such as an array oscillator, the control voltage must be routed to a large number of buffer stages. As a result, the control voltage input capacitance will be relatively large, inevitably with some unavoidable interconnect capacitance to the negative supply. As discussed earlier in this chapter, control voltage coupling to the negative supply can present a significant jitter problem. In addition, a large control voltage input capacitance may require the loop filter capacitors to be larger than desired. In such a situation it may be advantageous to buffer the control voltage. Rather than attempting to buffer the control voltage external to the bias circuit and possibly introducing other noise coupling sources, an addition to the bias circuit can be made to accommodate this function. A diode-connected load element in a half-buffer replica, where the gate of the biased PMOS device is connected to the output of the load element, can be used to generate an internal control voltage from the NMOS current source bias to be used as the PMOS bias voltage. The current source bias circuit with control voltage buffering is shown in Figure 3-8.

3.4 Single-ended buffer design

Single-ended buffer stages, like differential buffers, can be used to make oscillators and delay lines. While the dynamic noise performance of single-ended buffers is not as good as that of differential buffers, they have advantages in terms of size and power consumption. Because single-ended buffers are inherently smaller with fewer devices and bias voltages and do not require differential symmetry, they have less interconnect capacitance at their outputs, allowing them to be scaled down further in size for an equivalent frequency range at a fraction of the power consumption. In addition, without the need to bias a source coupled pair, single-ended buffers typically require less supply voltage. Thus,

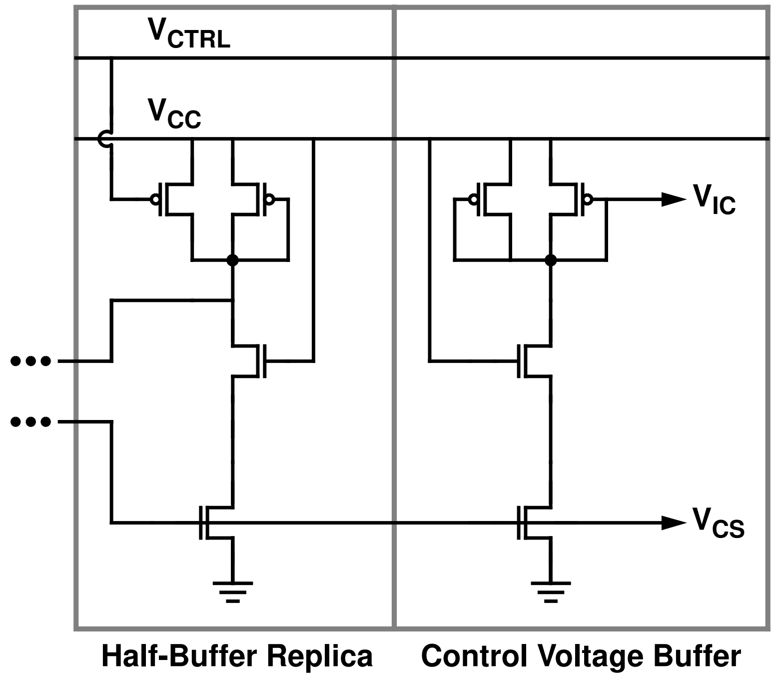

Figure 3-8: Schematic of the current source bias circuit with control voltage buffering.

single-ended buffers can be advantageously used instead of differential buffers as delay stages.

A single-ended buffer stage has single-wire input and output signals. It can be based on a variety of structures, including source coupled pairs, but more typically is based on common source devices. The simplest design for a single-ended buffer in a CMOS IC technology is a static CMOS inverter, which utilizes complementary common source devices. Unfortunately, the static CMOS inverter has very poor static supply noise rejection properties, making it only usable in supply filtered designs where the local supply is used to control the delay. More promising designs are based on a single common source device. A load element connected to the same supply provides control over the buffer delay and limits the output voltage swing. Both the common source device and the load element are biased by a high impedance current source to achieve high static supply noise rejection.

This section will describe a single-ended buffer stage that is based on a PMOS common source device, a PMOS diode-connected load element, and a dynamically-biased NMOS current source. The current source bias circuit is similar to the one used for the differential buffer stage described in the previous section. The single-ended buffer stage is designed to have high supply noise immunity while being able to operate at low supply voltages.

3.4.1 Single-ended buffer stage

The dual-input single-ended buffer stage, shown in Figure 3-9, is a PMOS common source amplifier with a PMOS diode clamp. The PMOS common source device is biased by a dynamically-biased simple NMOS current source. It provides the inversion and gain required for a ring oscillator. The PMOS diode-connected device serves to clamp the swing in order to keep the current source device in saturation and to provide sufficiently low output resistance to minimize the settling time of the output high voltage level. The I-V characteristics of the PMOS diode-connected device combined with the NMOS current source bias are almost identical to those of an NMOS common source device with its source voltage at the output low level and its gate voltage fixed. As with the differential buffer stage described in the previous section, the bias voltage for 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.

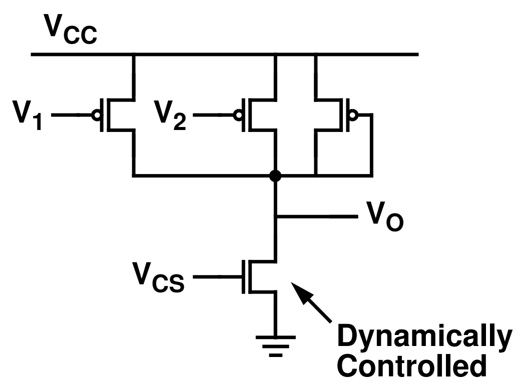

Figure 3-9: Schematic of the dual-input single-ended buffer stage, containing a dynamically-biased current source.

In an N-well technology, a PMOS common source amplifier is chosen over an NMOS common source amplifier as the buffer stage for several reasons related to minimizing the sensitivity of the buffer stage delay to substrate noise, as discussed earlier in this chapter. An NMOS common source amplifier would have an input switching threshold, closely related to the NMOS device threshold, that is very sensitive to the substrate voltage through the body effect. This input switching threshold sensitivity provides a means for substrate noise to couple to the buffer delay. A PMOS common source amplifier avoids this problem since PMOS devices, located in isolated wells, are not affected by substrate noise. In addition, with the control voltage referenced to the positive supply, the dynamically-biased simple current source must use an NMOS device. With an NMOS current source, a PMOS common source amplifier must be used. In general, this constraint arises because the bias circuit for the dynamically-biased simple current source must generate the current source bias voltage by referencing an external voltage, possibly the control voltage, to some macroscopic property of the buffer stage, such as the output voltage swing or the output voltage levels. The output voltage swing is always referenced to the opposite supply to which the current sources are connected. Since in an N-well technology the control voltage and other external bias voltages must be referenced to the positive supply in order to minimize sensitivity to substrate noise as discussed earlier in this chapter, the dynamically-biased current source must use an NMOS device.

The buffer device sizes are determined from simple constraints. The size of the current source device should be large enough so that it remains in saturation when the output is at its low level. The size of the diode-connected device should be set to establish the output low level. The size of the common source device should be set to provide enough gain for oscillations to occur at the point where the input and output voltages are equal. The gain at this point is equal to the ratio of the sizes of the common source device and the diode-connected device. This simple relationship results from the fact that with equal gate voltages, the ratio of the small-signal transconductances is equal to the ratio of the device sizes. The output high level, determined by the size of the common source device and bias current, should not constrain the device sizes.

The diode-connected MOS device allows the buffer delay to be controlled by the current source bias current. The voltage swing at the buffer output increases less than linearly with increasing bias current due to the square root dependence in the I-V characteristics of the load device. An increase in bias current decreases the rise time of the buffer output because more charging current is available without a proportional increase in the voltage swing. An increase in bias current also decreases the fall time of the buffer output, because the fall time is determined primarily by the output resistance of the common source device which is set by the output high level. When the bias current decreases, the opposite happens. The result is that the buffer delay varies with bias current while the output voltage swing is only marginally affected.

With differential buffers, linear resistive loads are most desirable for achieving high dynamic supply noise rejection since they provide a differential-mode resistance, and thus a buffer delay, that is independent of common-mode voltage. Because single-ended buffers do not have differentially swinging nodes, linear resistive loads will not prevent supply noise from affecting the delays of the buffers and, therefore, provide no advantage to single-ended buffers. As a result, the dynamic supply noise rejection of differential buffer stages can always be made lower than that of single-ended buffer stages. However, the dynamic supply noise rejection of single-ended stages can be made low enough so that it is not the main source of jitter in a PLL. Unavoidable capacitive coupling of supply voltage changes to the control voltage and bias voltages and the static supply noise sensitivity of buffer stages will typically dominate the output jitter with most types of buffer stages.

Even though the dynamic supply noise rejection of single-ended buffer stages can never be made as high as that of differential buffer stages, the dynamic supply noise rejection of this buffer stage is still relatively good. Although disturbed output voltages resulting from capacitive coupling to dynamic variations in the supply voltage will affect the delay of any buffers with outputs in transition, they are quickly restored. A disturbance will not affect the buffer delay beyond a single buffer delay after the supply variation due to the strong level restoration of the buffers. The output high level occurs when the common source device completely turns off, leaving the diode-connected device to establish the output high level with a relatively small time constant. The output low level, although determined by the output high level, occurs with the common source NMOS device in the linear region, again with a relatively small time constant. As a result, the accumulated phase error resulting from a disturbance will be relatively small.

SPICE simulations of a ring oscillator with all output interconnect capacitance connected to the negative supply 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 2% of an oscillation period with static effects factored out. This result is four times larger than that of the differential buffer stage described in the previous section.

3.4.2 Current source bias circuit

The current source bias circuit, shown in Figure 3-10 is similar in design to the one used for the differential buffer stage described in the previous section. However, instead of a half-buffer replica of the differential buffer stage, it uses the PMOS common source amplifier with its input connected to the control voltage. It is designed to achieve two functions analogous to the previous design. First, it sets the current through the simple NMOS current source in the buffers so that the voltage at the level where the input and output voltages are the same is equal to the control voltage. In the previous design, the current was set so that the output low level was equal to the control voltage. Second, it dynamically adjusts the NMOS current source bias so this current is held constant and highly independent of supply voltage. It operates identically to the previous design in that the amplifier adjusts the simple NMOS current source bias voltage in order to make the output voltage of the buffer replica equal to the control voltage which is connected to the input of the buffer replica. Since the bias current required by the buffer replica to establish the output voltage level does not depend on supply voltage, high static supply rejection is achieved.

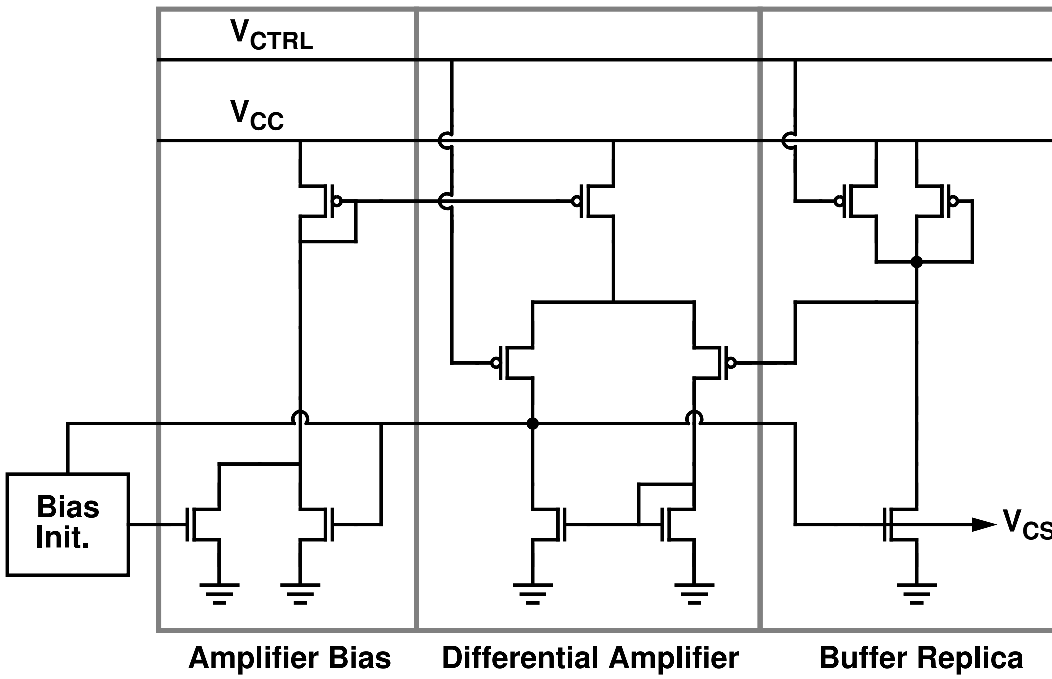

Figure 3-10: Schematic of the self-biased replica-feedback current source bias circuit for the single-ended buffer stage.

Although its design and operation are similar to the previous design, there are some important differences worth noting. First, the control voltage is only used in the bias circuit by the amplifier so that buffering the control voltage because of noise coupling or capacitive loading is not an issue. Second, the supply voltage requirements for the buffer stage and bias circuit are about a series NMOS and PMOS diode voltage drop with identical current densities minus the NMOS threshold voltage. These supply voltage requirements are less than the previous design by about the NMOS threshold voltage.

3.5 Summary

This chapter has examined both a differential and a single-ended buffer design. Both have particular merits making them each viable design alternatives. Differential buffer stages are desirable because they provide complementary outputs and have the potential to achieve very high dynamic supply noise rejection. Single-ended designs are desirable because they require less area, supply voltage, and power consumption while achieving a frequency range similar to that of differential buffer stages.

In attempting to minimize the sensitivity of the buffer stage to substrate noise by referencing the control voltage to the positive supply, an interesting trade-off develops with respect to the biasing approach for the two designs. This trade-off arises from the fact that the use of an NMOS source coupled pair or common source device will usually provide a significant speed advantage for the buffer stage over PMOS devices. When using dynamically-biased simple current sources, NMOS current source devices must be used in order for the control voltage to be referenced to the top supply. As a result, the differential buffer design is able to use an NMOS source coupled pair, while the single-ended buffer design must use a PMOS common source device. However, when using cascode current sources, both NMOS and PMOS current source devices can be used. As such, the differential buffer design can still use an NMOS source coupled pair, while the single-ended buffer design is now able to use an NMOS common source device. Thus, from the perspective of maximizing the operating frequency, the differential buffer design is ideally suited to both biasing approaches while the single-ended buffer design is better suited to using cascode current sources.

In digital integrated circuit applications, supply and substrate noise can pose a serious challenge to the jitter performance of a buffer stage. Although a variety of noise coupling sources can affect the operation of buffer stages, by far the most significant are static supply noise sensitivity and, possibly more important, unavoidable capacitive coupling of supply voltage changes to the control voltage and bias voltages. However, if dynamic supply and substrate noise sensitivities are not properly addressed in a buffer stage design, they can become limiting factors in the performance of the buffer stage. The following chapter will discuss these noise issues further in the context of a complete PLL design.

Chapter 4

Array Oscillator Delay Generator

This chapter describes the actual implementation of an array oscillator delay generator based on the array oscillator architecture presented in Chapter 2 and the buffer stage designs with high noise immunity presented in Chapter 3. It will examine a variety of issues pertaining to the implementation that include the output channel structure, reset mechanism, biasing issues, layout strategy, and design techniques for achieving low jitter. This chapter will also present experimental results demonstrating the ability of an array oscillator delay generator to produce precise delays with a resolution equal to one seventh of a buffer delay.

An adjustable time interval generator is an ideal use for an array oscillator since an array oscillator can produce very high delay resolution. With the addition of two or more output channels to provide independent access to the internal array nodes, one serving as a delay reference and the others addressing the desired delays, an array oscillator is able to produce adjustable periodic delay intervals. In order to maintain the high precision provided by the array core, each 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 channels.

The chapter will focus primarily on the implementation of an array oscillator delay generator using the differential buffer stage with symmetric loads and replica-feedback biasing presented in Chapter 3. However, it will also discuss implementation differences for a single-ended buffer stage implementation.

4.1 Output channel

Since an array oscillator is a two-dimensional structure, it is ideally suited to utilize row and column addressing when only supporting a limited number of output channels. As a result, the output channels contain word lines and bit lines with an organization very similar to that found in a multi-ported memory array. A simplified block diagram of a single output channel is shown in Figure 4-1. The column and column output multiplexers are distributed structures and are represented in the figure as a set of buffers with output enabling inputs and with their outputs shunted together. The word lines select a row of buffer cells containing a single buffer from each ring with the rings oriented along the columns. The array buffer in each buffer cell is isolated from the column multiplexers for each output channel with an additional buffer. The outputs from the selected buffers are multiplexed on to the bit lines through the column multiplexers. After an additional buffer to increase the signal swings, the bit lines from the selected column are multiplexed on to a single pair of output wires through the column output multiplexer. Although this additional buffer is implemented as a simple buffer stage, it can be implemented instead as a single-input multiplexer with an output enabling input originating from the decoder of the column output multiplexer in order to reduce power dissipation. These output wires are followed by a buffer, a conditionally inverting multiplexer that can swap the differential signals, and a final output buffer.

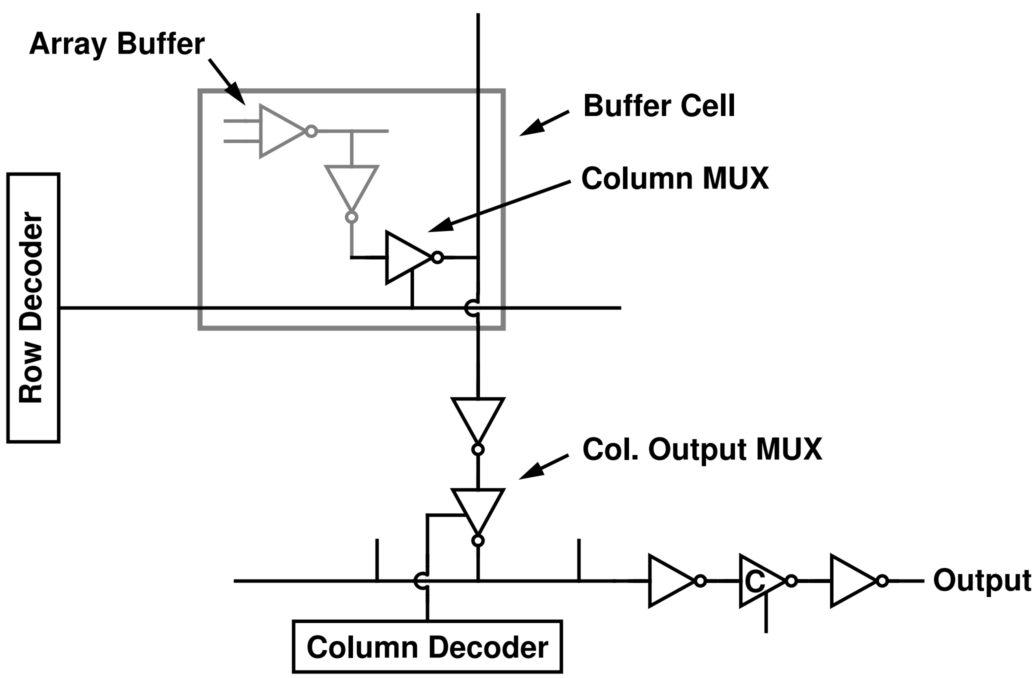

Figure 4-1: Simplified block diagram of a single differential output channel. The column and column output multiplexers are distributed structures and are represented as buffers with output enables. All wires in the signal path represent differential signals.

An isolating buffer is required between each array buffer and corresponding column output multiplexers in order to minimize the output load of the array buffer and to prevent address-dependent loading changes from affecting the array buffer delays and thus the array operation. The possible address-dependent loading changes originate from a column multiplexer input capacitance that depends on whether the input is selected. The delay through the output channels can also be affected by address-dependent changes in the output load of the isolation buffer. With several column multiplexers connected to the output of the same isolation buffer, delay of the isolation buffer can depend on the number of column multiplexers that have the same row address. When only two output channels are provided, both channels will be equally affected by any changes in the output load of the isolation buffer. However, such an effect could result in apparent jitter when the row addresses for the two channels are changed to and from the same value. To prevent this effect, an independent buffer can be used at the input of each column multiplexer. Although address-dependent changes in the output load of the isolation buffer should generally be considered, the multiplexers described next provide enough input isolation to make independent buffering unnecessary.

All buffers and multiplexers used in the output channels are based on single-differential-input versions of the dual-input differential buffer stage. The schematic of a multiplexer is shown in Figure 4-2. 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 multiplexer is a distributed structure 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 on to a pair of shared output wires with a single pair of shared load elements. As the number of inputs is increased, the bandwidth of the multiplexer is reduced due to the increased load on the shared output wires. The conditionally inverting multiplexer is a special case of the multiplexer shown in Figure 4-2. It 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.

The delays through all output channels are made equal from any internal array node by balancing and equalizing the lumped interconnect capacitance. The balancing of lumped interconnect capacitance within a single output channel is achieved by using identical layouts for all buffer cells, bit lines, and column output multiplexers. The balancing of lumped interconnect capacitance between output channels is achieved by using identical wire lengths and symmetrical layout for each output channel within the buffer cells, bit lines, and column output multiplexers. The balancing of lumped interconnect capacitance

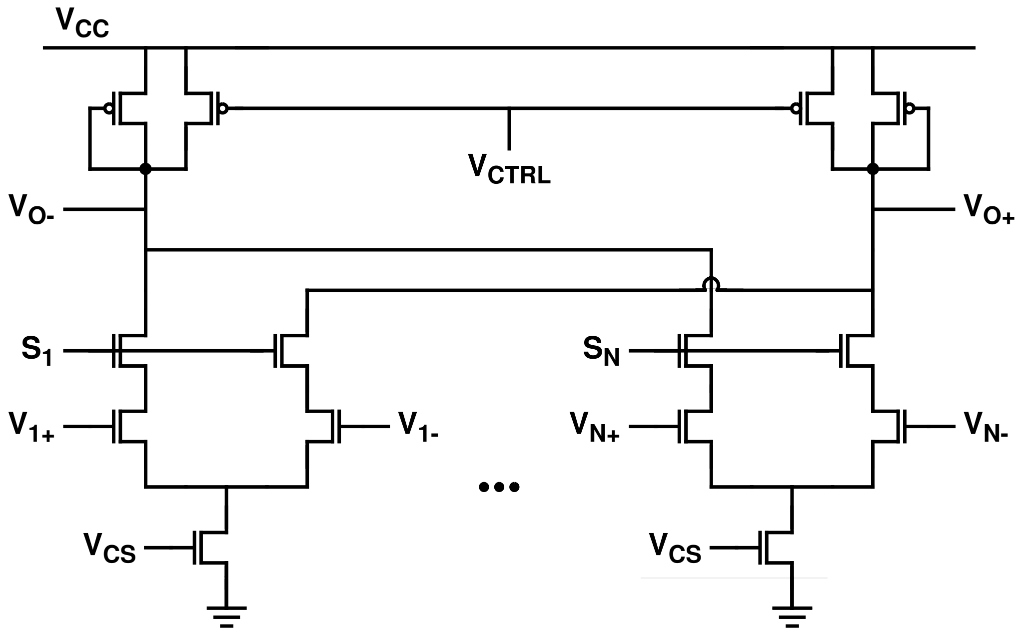

Figure 4-2: Schematic of the differential output channel multiplexer.

between output channels typically requires some wire length to be wasted beyond the necessary connections.

4.1.1 Bandwidth limitations and device mismatches

There are a number of potential problems that can affect the output channel. As the number of buffers per ring is reduced, the bandwidth required 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 due to the increased output load. More importantly, as the number of rings is increased, the precision required from the output channel will increase due to the increased delay resolution. A bandwidth limitation in the output channel can cause delay errors greater than the array resolution by amplifying the effect of small device mismatches.

To understand how the delay errors arise, consider the outputs of the column and column output multiplexers. When the multiplexers are bandwidth limited, the voltage swings at their outputs will be less than their static limits. With incomplete voltage swings, random device mismatches in the multiplexers result in address-dependent differential-mode offsets.[1] As the bandwidth decreases, these differential-mode offsets will become a larger fraction of the decreasing voltage swing. In addition, the differential-mode offsets are amplified by the DC gain of the multiplexers, unlike the output signals, which experience little if any amplification. This amplification effect, illustrated in Figure 4-3, suggests that differential-mode offsets originating early in the output channel will result in larger differential-mode 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 in the signals. 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.

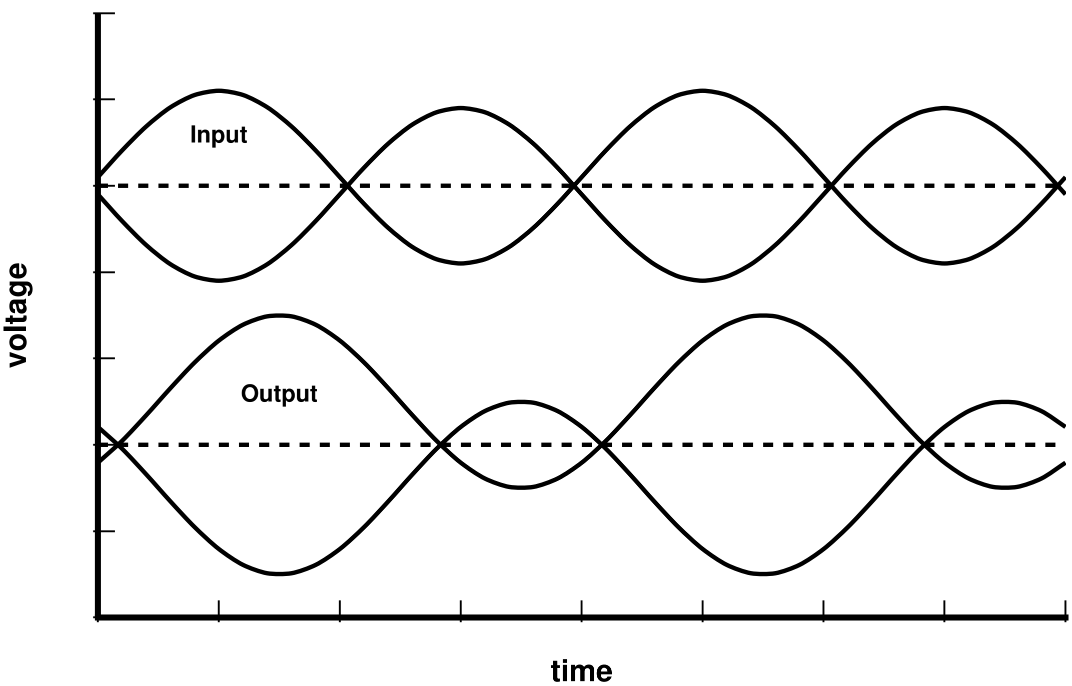

Figure 4-3: Amplification of differential-mode offsets through bandwidth limited multiplexers.

4.1.2 Offset cancellation circuit

As long as the output channels are linear, no information about the delays of the output phases will be lost due to address-dependent offsets. The output channels 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 each column output multiplexer to cancel out the random offsets and prevent

1. For single-ended designs, the random device mismatches instead result in address-dependent offsets from the switching point

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 Figure 4-4, while the schematic is shown in Figure 4-5. 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. Both buffers use the same size devices as the column output multiplexer and the value of each capacitor is about 10pF. The operation of the circuit is most easily analyzed [16] with the feedback path broken at the input to the first buffer as indicated in Figure 4-4, so that only the output of the second buffer connects to the output of the column output multiplexer. Since 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. Since the second buffer is the same size as the column output multiplexer, its gain will be one half that of the first buffer. Thus, this circuit reduces the differential-mode offset at the output of the column output multiplexer by the square of the buffer gain. It should be noted that this circuit is not suitable for applications where fast settling in the output channel is required. These applications require higher bandwidth multiplexers.

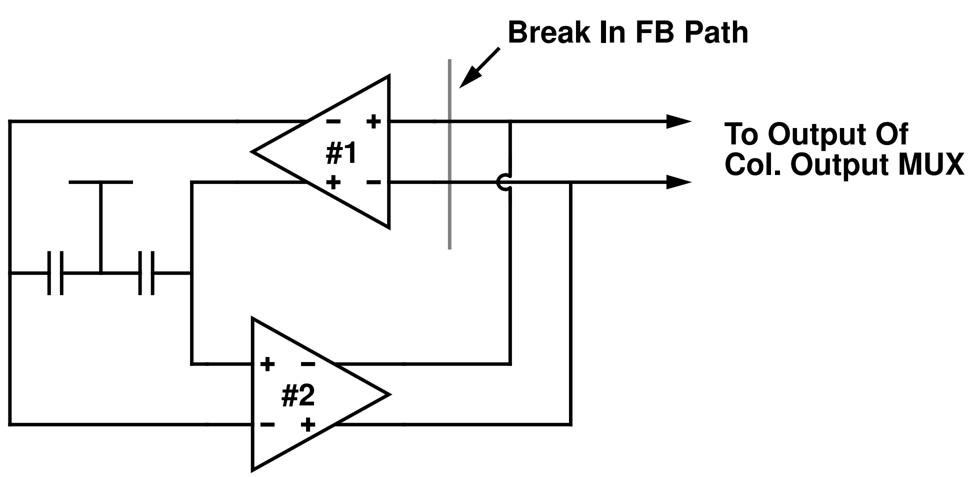

Figure 4-4: Block diagram of the differential offset cancellation circuit.

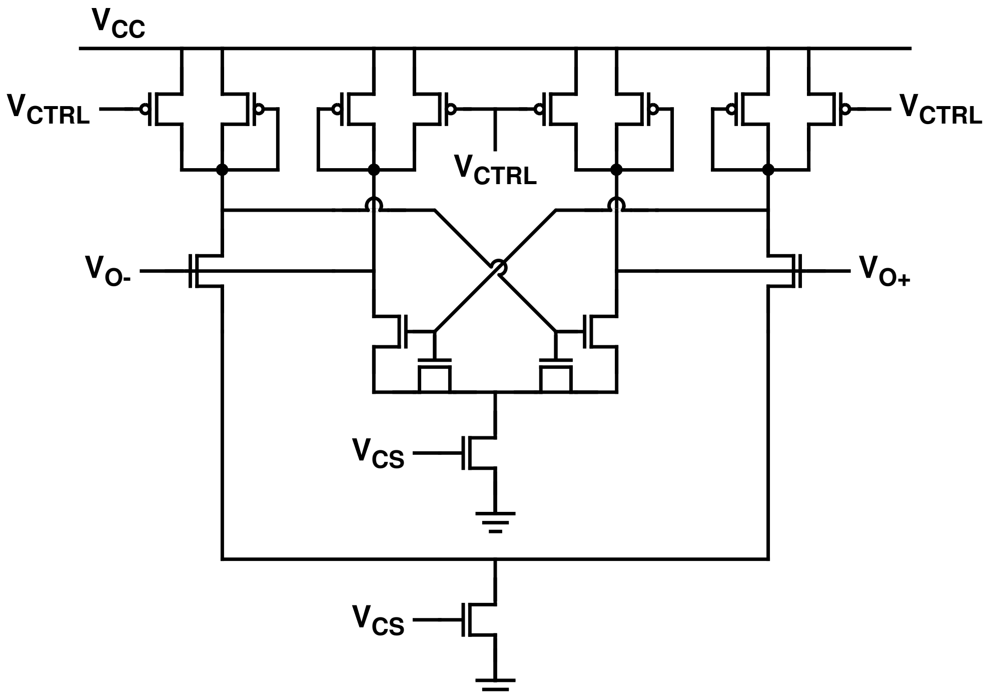

Figure 4-5: Schematic of the differential offset cancellation circuit.

4.1.3 Device sizing

The sizing of the buffer stages and multiplexers in the output channel are based on device matching and bandwidth considerations. The smallest devices in the array core and output channel will have the poorest matching of electrical characteristics. With this consideration, all buffer stages and multiplexers would ideally be of the same size for a given power dissipation level. However, since it is desirable to minimize the loading on the array buffers, the smallest devices are constrained to be one half those in the array buffers. The bandwidth of the column multiplexer and the column output multiplexer are the most constrained of all stages in the output channel. As a result, the multiplexer device sizes must be set so that device capacitances rather than interconnect capacitances dominate the bandwidth of the multiplexers. Based on these considerations, the sizes of the buffer stages and multiplexers are established. With the array buffers of unit size, the isolation buffers and the column multiplexers are of one half unit size, and all buffers and multiplexers that follow the column multiplexers are of unit size.

4.1.4 Single addressable output channel

The number of output channels required from the array depends on system requirements. When only a single adjustable delay is needed, only one of two output channels needs to be completely addressable. In this case, a straightforward organization would provide two identical output channels in order to achieve equal channel delays, even though only one output channel is needed for addressing. However, a more optimal organization would split the row and column addressing function between two output channels. This organization provides a significant savings in the size of the two output channels while still achieving equal channel delays.

A block diagram of an array with two split addressed output channels is shown in Figure 4-6. This organization requires only one set of column bit lines. However, two column output multiplexers are still required to match the output channel delays. The word lines from the row decoder only drive the first column. All following columns have their row address fixed to a single row. One channel output is derived from a multiplexer that selects between the bit lines from all columns. The second output is derived from an identical multiplexer that only selects the bit lines from the first column. The input to the row decoder and the first column output multiplexer will select the desired delay. Since the first column has only one set of bit lines, this organization has the limitation that output delays equal to an integral number of buffer delays are not accessible.

If the system can tolerate a fixed delay offset between the two output channels, then an even more simplified output channel design can be used. One output can be multiplexed directly from the array buffer outputs along a row at the perimeter of the array, and the other can be multiplexed directly from the array buffer outputs along a column at the perimeter of the array without using bit lines. An isolation buffer, however, is still required at the output of every array buffer. This simplified organization allows all output delays to be accessed.

4.1.5 Externally multiplexed output channels

The optimal organization of the output channels changes when the array is required to support a large number output channels. When the number of output channels is greater than the number of buffers per ring, it is more advantageous to multiplex the array buffer outputs external to the array core. In this case, with the column multiplexers inside the array core, the required number of bit lines will exceed the number of dedicated bit lines required to channel all array buffer outputs from the array core. In addition, the area

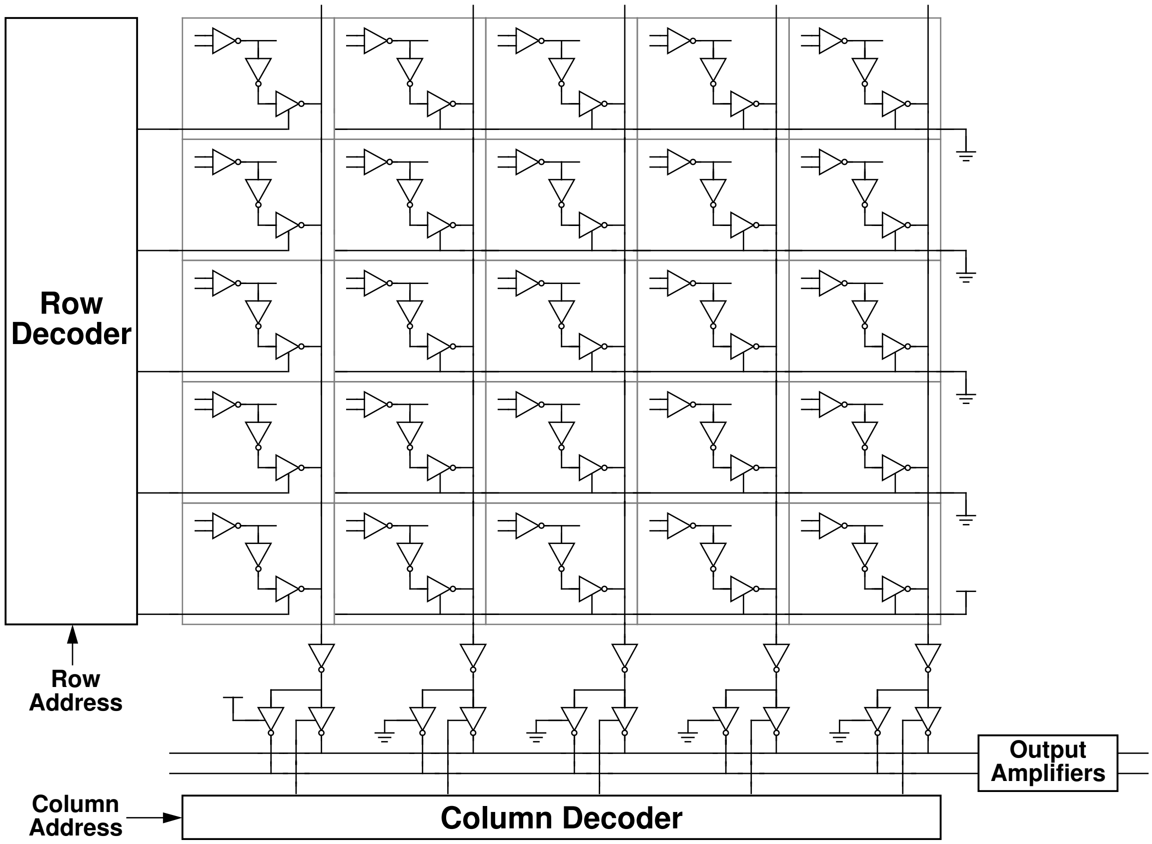

Figure 4-6: Block diagram of the array with two output channels using split row and column addressing.

overhead for the column multiplexers in each buffer cell will result in excessive interconnect capacitance at the array buffer outputs which will cause the oscillation frequency range to be reduced. With the array buffer output multiplexed external to the array core, the array core can be kept small and independent of the number of output channels. In addition, the layout for the output channels can be independently tiled adjacent to the array core.

A block diagram of the array with a single externally multiplexed output channel is shown in Figure 4-7. Each array buffer output is isolated through an additional buffer following which it is directed to the periphery of the array through separate dedicated bit lines. With the rings oriented vertically, each column will contain a number of dedicated bit lines equal to the number of buffers per ring. In order to equalize the lumped interconnect capacitance for each of the dedicated bit lines, they all must be of equal length and, thus, all must completely traverse the columns. With the rings oriented vertically, the interconnect capacitance can be minimized since the vertical dimension of the array will typically be the smallest. The multiplexers for each output channel are located at the periphery of the array, each in a separate row. The dedicated bit lines, in addition to traversing the length columns of the array core, must also traverse across each of the output multiplexers. Each output multiplexer is typically partitioned into two levels of multiplexers for increased bandwidth. Buffering of the inputs to each output multiplexer is usually required for isolation from the other output multiplexers.

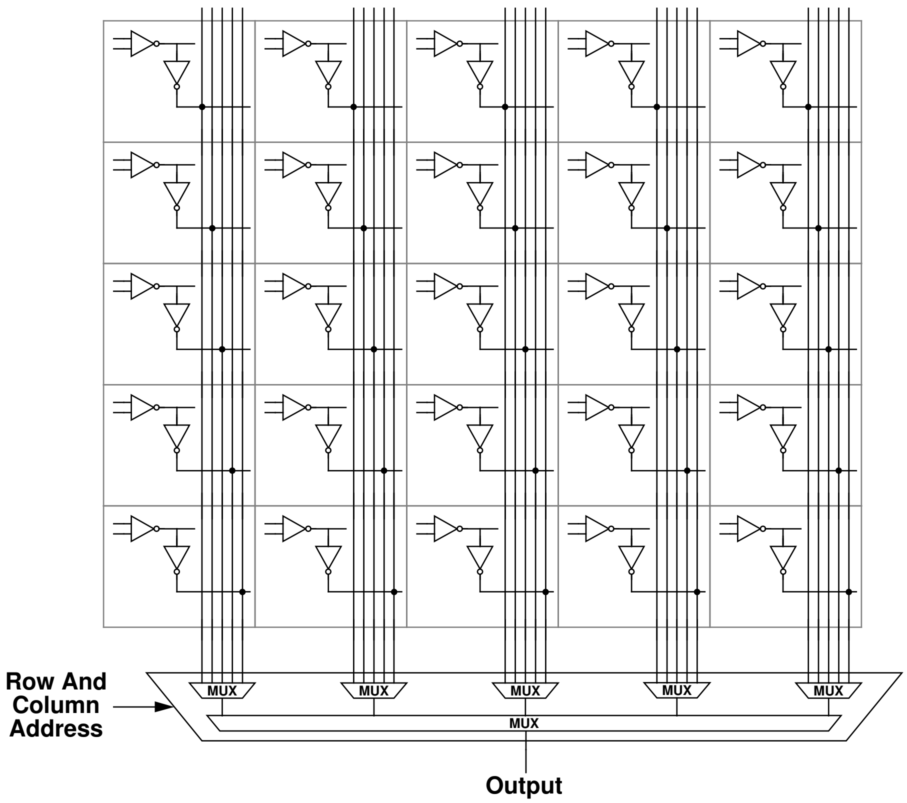

Figure 4-7: Block diagram of the array with a single output channel using the externally multiplexed organization.

4.2 Reset operation

Another issue that the delay generator must address is how to assure that the array oscillator operates in the desired mode of oscillation. As discussed in Chapter 2, an array oscillator will support several modes of oscillation. Each mode will have a different operating frequency range and a different output phase ordering. Most applications will require that the array operate in the mode with the highest oscillation frequency range. To satisfy this requirement a reset mechanism is needed.

As explained in Chapter 2, the array is selectively reset into a particular mode C by initializing the phase relationship among the nodes in the array so that it is closer to the phase relationship for this particular mode than for any other mode. The reset operation is most easily accomplished by disabling the coupling between two of the rings so that the array of coupled rings no longer forms a closed loop. Disabling the coupling by 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. Instead, the coupling can be disabled by switching off the bias voltages in one or more of the rings in order to deactivate the coupling inputs of the first following ring. The desired boundary conditions can then be forced on the array by adjusting the oscillation frequency of this first ring through its bias voltages.

With only five buffers per ring, modes with C close to zero will provide the highest operating frequency range. These modes 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 provided that all coupling input devices in the first ring are off. The inactive coupling inputs on the first ring will not change the oscillation frequency of the first ring 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.

4.2.1 Bias voltage switching

For the differential buffer stage based on symmetric loads, it is not possible to turn off a single ring oscillator by switching its bias voltages in a manner that deactivates the coupling input differential pair devices of the following ring in order to achieve modes C close to zero. If the PMOS bias of one ring is switched to the top supply in order to turn off the biased PMOS devices, the diode-connected PMOS devices in the load elements will still allow the ring to oscillate and the array will still be closed. Alternatively, if the NMOS bias of one ring is switched to the bottom supply in order to turn off the NMOS current source, its outputs will be at the upper swing limit which will slow down the next ring so that the following rings will not oscillate with their ring and coupling inputs at the same phase which is required for modes close to zero.

By switching off selected bias voltages on two of the rings, the outputs of the second ring can be at the lower swing limit which will allow the next ring to oscillate at the frequency of a ring with no coupling inputs so that the following rings will oscillate with their ring and coupling inputs at the same phase. The NMOS current source bias voltage for the first ring is switched to the bottom supply in order to turn off the NMOS current sources. All the outputs of the first ring will then be fixed at the upper swing limit. In addition, the PMOS bias voltage for the second ring is switched to the top supply in order to turn off the bias PMOS devices. With the coupling inputs all at the upper swing limit, this second ring will have inadequate gain to oscillate. Since the biased PMOS devices are off, only the diode-connected PMOS devices remain active in the load elements of the second ring. These load elements now only require half of the differential pair bias current to achieve an output voltage at the lower swing limit. With the differential pair bias current equally split in half between each side of the differential pair, the output voltages of the second ring will be at the lower swing limit. When the bias voltages are switched back to their nominal levels, the array oscillator will assume the desired mode where C is close to zero.

With the outputs from each ring bias circuit shunted together, the bias voltage switching must be accomplished through voltage switches between the outputs of the bias circuits and the ring bias inputs. The circuits used for switching the NMOS and PMOS bias voltages are shown in Figure 4-8 and Figure 4-9 respectively. The diode-connected device provides the initial switching current, while the complementary long channel device establishes the final gate voltage of the pass device at the non-referenced supply through high resistance to minimize the supply noise coupling. In order to make the impedances of the bias voltage nodes within each ring oscillator identical, these switches must be placed on both the NMOS and PMOS bias voltage inputs of each ring, even though only one switch on two rings will ever be turned off.

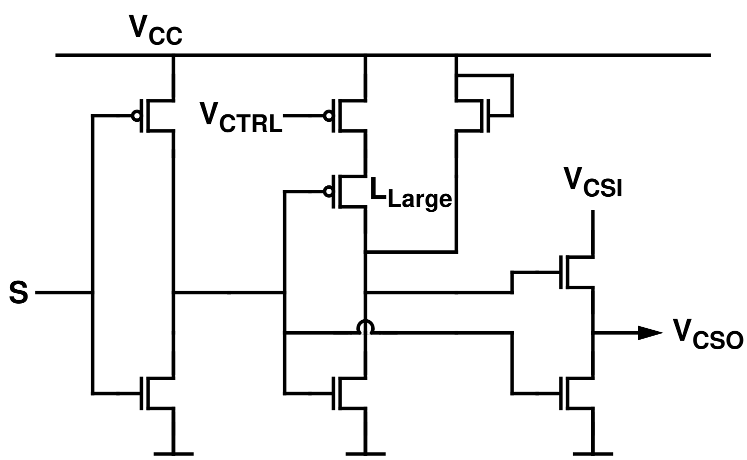

Figure 4-8: Schematic of the NMOS bias voltage switch.

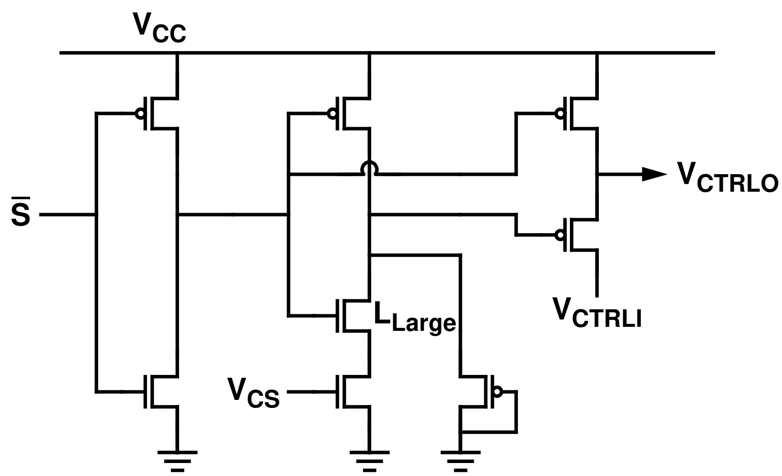

Figure 4-9: Schematic of the PMOS bias voltage switch.

4.2.2 Noise coupling

A disadvantage of this bias voltage switching technique is that it allows supply noise to capacitively couple to the bias voltages from the non-referenced supply for both the NMOS and PMOS bias voltages. The coupling path to each bias voltage is through the gate capacitance of the pass device and the channel of the device establishing the gate voltage of the pass device at the non-referenced supply. Another disadvantage of this bias voltage switching technique is that it constrains the minimum supply voltage to be the NMOS bias voltage plus a NMOS threshold, or the PMOS bias voltage plus a PMOS threshold, below which the bias voltage switches never turn on. Fortunately, this supply voltage limit is usually always below the series NMOS and PMOS diode voltage drop established by the buffer stage and bias circuit representing the minimum supply voltage where high static supply noise rejection can be achieved.

An alternative bias voltage switching technique that was not implemented is shown in Figure 4-10. This technique avoids the supply noise coupling problem and the minimum supply voltage constraint by pushing the switches into the bias circuit with control voltage buffering. The NMOS bias voltage is turned off by supplementing the differential pair bias current for the replicated load element in feedback with the amplifier which causes the amplifier to respond by turning off the current source to restore equilibrium. The bias current for the replicated load element used to generate the PMOS bias voltage also has its bias current supplemented to guarantee that the buffer outputs will be pulled all the way up to the top supply when the NMOS bias voltage is turned off. The PMOS bias voltage is turned off by turning off the replicated NMOS source coupled pair device used to generate the PMOS bias voltage and pulling the PMOS bias voltage up to the top supply. With this technique, however, the outputs of all ring bias circuits cannot be shunted together nor can the bias circuits share any internal nodes. These constraints pose a slight disadvantage for the bias circuit layout and could make device mismatches between the bias circuits an issue.

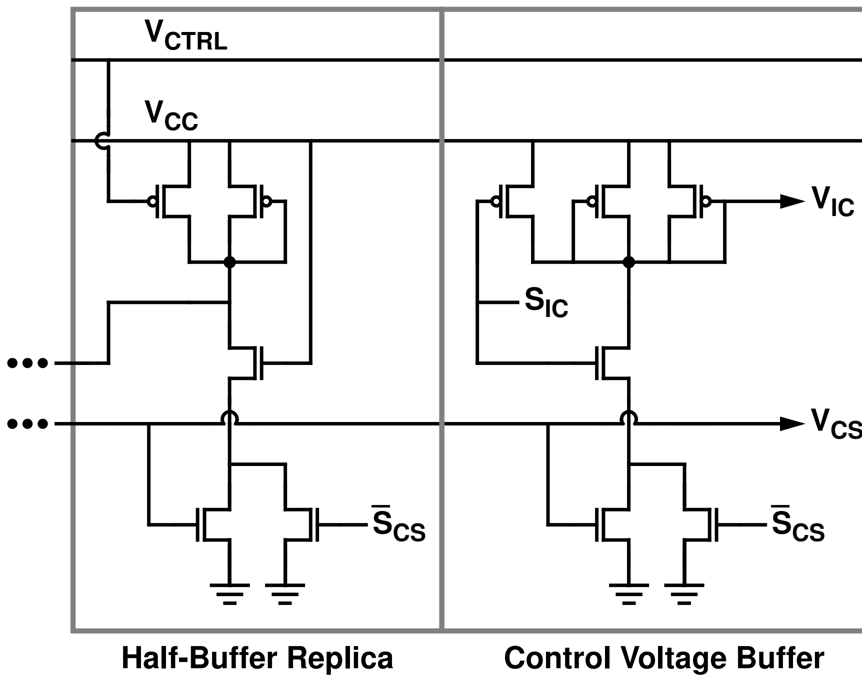

Figure 4-10: Schematic of the current source bias circuit with integrated bias voltage switches for the differential buffer stage.

4.2.3 Modes of oscillation

Although some choices of array dimensions will support modes of oscillation with higher operating frequency ranges than others, the reset mechanism may not be able to realize them. This limitation will in turn constrain the array dimensions when the highest operating frequency range is required. During initialization, the buffers in the first ring experience a reduction in total transconductance with the deactivation of the coupling input differential pair devices that causes a small reduction in the oscillation frequency from a ring with no coupling inputs. As a result, this reset mechanism places the array in a temporary mode with C slightly less than zero, as the corresponding ring outputs from successive rings are slightly earlier in phase rather than being at the same phase. This starting point makes it difficult to achieve a +2 mode after initialization for all operating frequencies, especially when the number of rings in the array is large. For this reason, the array dimensions are instead constrained to make a -2 mode possible which can be reliably achieved for all operating frequencies. By delaying the activation of the NMOS bias voltage in the first ring, the next slower mode of oscillation can be usually achieved. Upon activation of the control voltage, before initialization, the array will typically start oscillating in the slowest mode of oscillation. For a high operating frequency range, five stage ring oscillators are desirable. In order to support the -2 mode, the array dimensions are constrained to be 5×7, 5×17, or, in general, $5 \times (2 + 5k)$ for any odd integer k. Even values of k result in redundant complementary phases.

4.3 Biasing issues

The implementation of the array oscillator delay generator uses the bias circuit with control voltage buffering. This circuit is used primarily to relax the requirements for a minimal amount of capacitive coupling to the control voltage within the array core. It also reduces the capacitive load on the control voltage and thus allows smaller capacitors to be used in the loop filter. A separate bias circuit is used for each ring in the array so that the bias circuits can be arrayed with the rings to allow the array to be scaled easily. The outputs of each bias circuit are shunted together to minimize any device mismatching problems between the bias circuits. Each bias circuit contains devices of the same size as the corresponding devices in the array buffer stages. With five buffers per ring, the bias circuit will have an equivalent output load of 16.5 array buffer stages.

In the delay generator implementation, the output channels can conveniently share the outputs of the bias circuits used by the array buffers. However, the bandwidth of the output channels can be improved relative to the array oscillation frequency by using a separate bias circuit for the output channels in order to provide higher current densities in the buffer stages of the output channels. The disadvantages of this technique are that another set of bias circuits are needed and that the minimum operating supply voltage will be determined by the buffer stages in the output channels with their higher current densities.

4.4 Layout issues

A block level floor plan of the array oscillator delay generator is shown in Figure 4-11. The rings extend vertically and are arrayed horizontally through the array core. The bias circuits for each ring are located at the top, the row decoders are located at the left, and the column output decoders are located at the bottom of the structure. The complete layout is centered around the layout of the buffer cell. The buffer cell contains the array buffer, the isolation buffer, two column multiplexer differential pairs, and all required interconnect. All other cells in the design must be pitch matched to this cell. These cells include the bias circuits, row decoders, column output multiplexers, and column decoders. To the extent that the horizontal and vertical dimensions of the buffer cell are minimized, the oscillation frequency range possible will be maximized. The layout for the buffer cells are typically mirrored about the horizontal axis in each column and about the vertical axis in each row.

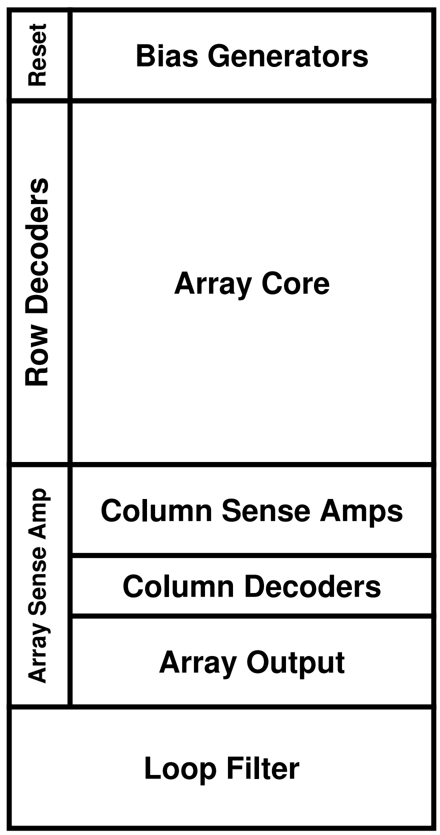

Figure 4-11: Block level floor plan of the array oscillator delay generator.

The need for matching devices imposes additional requirements on the layout. Almost all devices must have matching electrical properties with other devices. As a result, their layout must be insensitive to horizontal and vertical mirroring and always be at the same rotational orientation. To accomplish this matching symmetry, all devices must be folded at least once and their diffusion contacts must cover the whole width of the source and drain regions. However, the differential circuits require additional symmetry. These circuits must be completely balanced in terms of device characteristics and interconnect capacitance on each side of the differential pair.

4.5 Single-ended implementation

The array oscillator delay generator can also be implemented using the single-ended buffer stage with replica-feedback biasing. This section will discuss the implementation details particular to the single-ended implementation.

4.5.1 Output channel

The output channel design for the single-ended implementation is very similar to that for the differential implementation. A simplified block diagram of a single output channel is shown in Figure 4-12. A single buffer stage is used in each buffer cell to isolate the array buffer from the column multiplexers. Also, an additional buffer stage is added between each pair of bit lines, which are the outputs of the column multiplexers, and the column output multiplexers to increase the signal swing.

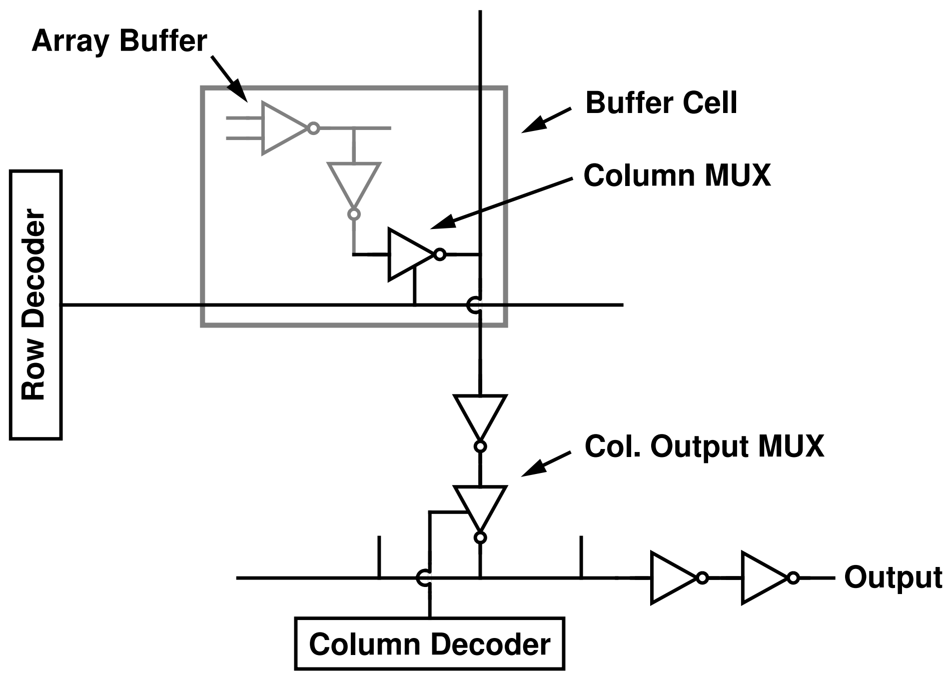

Figure 4-12: Simplified block diagram of a single-ended output channel.

All buffers and multiplexers used in the output channel are based on single-input versions of the dual-input single-ended buffer stage. The schematic of a multiplexer is shown in Figure 4-13. The multiplexer operates in the same manner as the multiplexer described for the differential implementation. The sizing of the buffer stages and multiplexers in the output channel are identical to those for the differential implementation. With the array buffers of unit size, the isolation buffers and the column multiplexers are of one half unit size, and all buffers and multiplexers that follow the column multiplexers are of unit size.

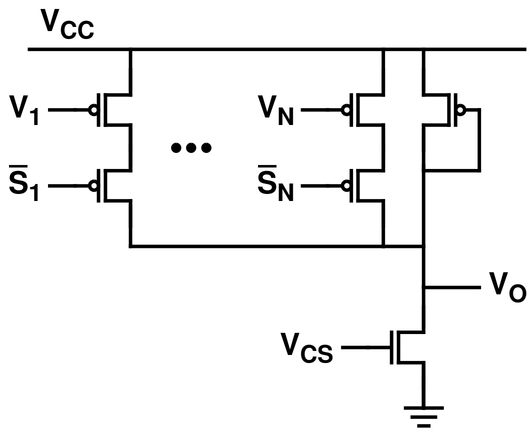

Figure 4-13: Schematic of the single-ended output channel multiplexer.

Similar to the differential implementation, a single-ended offset cancellation circuit can be added to the output of each column output multiplexer. The schematic for the single-ended offset cancellation circuit is shown in Figure 4-14. It contains a single-stage differential amplifier, two single-ended buffer stages, one of which has its input and output shorted to provide a switching point voltage reference, and a PMOS capacitor to remove the AC signal components. The buffer stages and the amplifier use the same size devices as the column output multiplexer. The value of the capacitor is about 10pF. Similar to the differential offset cancellation circuit, this circuit reduces the offset at the output of the column output multiplexer by the product of the gains of the buffer stage and the amplifier.

4.5.2 Reset operation

For the single-ended current switched buffer stage, modes with C close to zero can be achieved by turning off the NMOS bias voltage on one of the ring oscillators. By switching off the NMOS bias voltage on one of the rings, its outputs will be at the upper swing limit which will allow the next ring to oscillate at the frequency of a ring with no coupling inputs so that the following rings will oscillate with their ring and coupling inputs at the same phase. When the NMOS bias voltage is switched back to its nominal level, the array oscillator will assume the desired mode where C is close to zero. The circuit used to switch the NMOS bias voltage is identical to that described for the differential implementation. In addition, the same modification to the bias circuit discussed for the differential

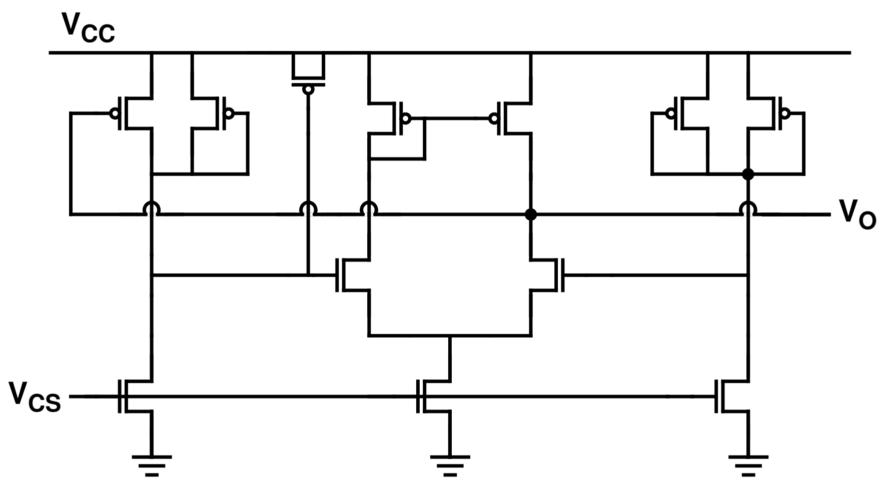

Figure 4-14: Schematic of the single-ended offset cancellation circuit.

implementation can be used in the single-ended implementation to accomplish the bias voltage switching.

Similar to the differential implementation, the reset mechanism will constrain the array dimensions when the highest operating frequency range is required. However, for the single-ended implementation, the array is placed in a temporary mode with C a little greater than zero during initialization. This non-zero mode results because, during initialization, the buffers in the first ring experience a reduction in effective output resistance that cause a small increase in the oscillation frequency from a ring with no coupling inputs. This reduction in effective output resistance results from a reduction in the output high voltage level with the deactivation of the coupling input common source devices. This starting point instead favors the reliable achievement of a +2 mode after initialization. In order to support the +2 mode with five stage ring oscillators, the array sizes are constrained to be 5×8, 5×13, 5×18, or, in general, $5 \times (3 + 5k)$ for any integer k.

4.5.3 Other issues

Like the differential implementation, a separate bias circuit is used for each ring in the array with their outputs shunted together. Each bias circuit contains devices of the same size as the corresponding devices in the array buffer stages. With five buffers per ring, the bias circuit will have an equivalent output load of 10.5 array buffer stages. The output channel shares the outputs of the bias circuit with the array buffers. However, the bandwidth of the output channel can be improved relative to the array oscillation frequency by using a separate bias circuit for the output channels.

The same layout issues that applied to the differential implementation apply to the single-ended implementation with one exception. Since there are no differential circuits, no differential symmetry is needed. This exception significantly relaxes constraints in the layout and typically results in an area savings.

4.6 Phase-locked loop

The array oscillator must be phase locked to some frequency reference in order to produce precise delays of known duration. In order to produce these delays with low output jitter, it is important for the array oscillator to have low noise sensitivity. However, a carefully designed phase-locked loop is also essential. Not only must the PLL provide low noise sensitivity, it must also provide high stability, fast settling, low peaking, and minimal spurious outputs. Because PLLs based on oscillators are at least second order systems, they are more difficult to design than PLLs based on delay lines which are typically first order systems [1]. An analysis of charge pump PLLs is presented in [17, 18, 19]. This section will focus on the design of a monolithic PLL for the array oscillator delay generator with particular emphasis on achieving low output jitter.

4.6.1 Circuit implementation

The array oscillator delay generator uses the charge pump PLL design shown in Figure 4-15. It is composed of a phase-frequency comparator, a charge pump, and a low pass loop filter in addition to the oscillator. Resistor R_1 in series with capacitor C_1 provides a zero in the open loop response which improves the phase margin and, thus, the overall stability of the loop. Capacitor C_2 in shunt with the control voltage arises in part from the input capacitance of the oscillator. It provides higher-order roll off for reducing the spurious outputs, but it can adversely affect the overall stability of the loop.

The phase-frequency comparator, shown in Figure 4-16 is similar in design to the phase-frequency comparator found in [20]. This type of circuit can have a dead band region in its operating characteristics, where the phase-frequency comparator produces no output for input signals with small but finite phase differences. The dead band region is caused by a rapid reset path through the phase-frequency comparator that resets the outputs before they ever become asserted for input signals with small but finite phase differences. With a dead band region in the operating characteristics, the phase of the oscillator output can depart from the phase of the reference signal to the extent of the dead band region without any correction from the PLL which leads to higher output jitter. In this circuit the dead

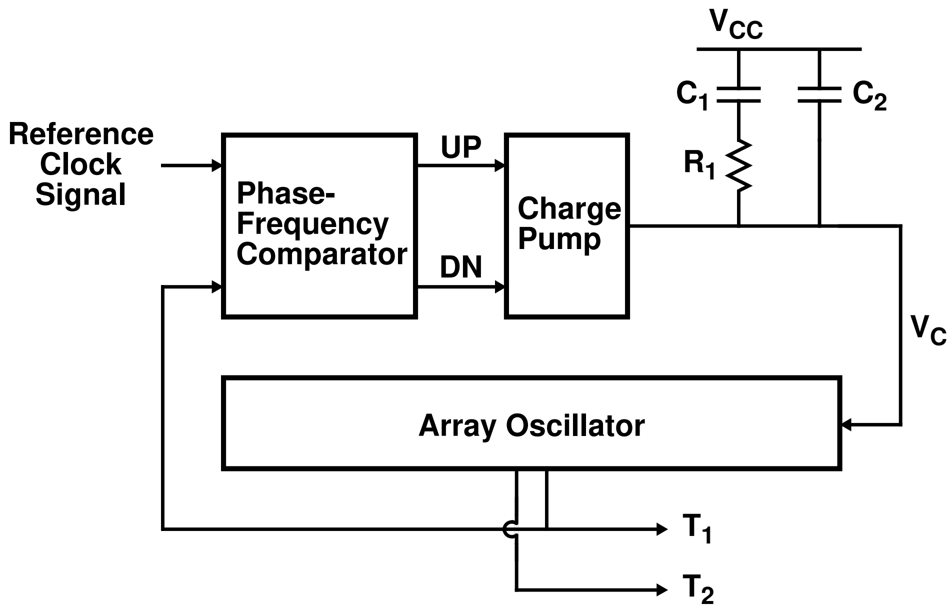

Figure 4-15: Block diagram of the charge pump PLL typically used for phase-locking an oscillator.

band region has been eliminated by adding extra delay in the reset path to insure that when both input signals are at the same phase, there will be equal and non-zero up and down pulses at the outputs. The size of these pulses for in-phase input signals should be kept at a minimum since they will reduce the effective impedance and cause extra ripple at the charge pump output. The former can lead to an input phase offset when the charge pump currents are mismatched, while the latter can result in reference frequency sidebands at the output of PLLs performing frequency multiplication. The existence of the output pulses for in-phase input signals will not, however, cause higher gain for input signals with small but finite phase differences than that for input signals with large phase differences. The constant gain results because the extra delay added to the reset path will always cause both up and down signals to be asserted simultaneously for an equal time interval while the outputs are reset independently of the input phase difference. Thus, the elimination of the dead band region results in overall linear operating characteristics for the phase-frequency comparator, especially for input signals with small but finite phase differences. The phase-frequency comparator implementation is based on single-ended static CMOS circuits. The differential oscillator output and differential reference input are converted to the single-ended signals used as input to the phase-frequency comparator with the differential-to-single-ended converter circuit shown in Figure 4-17.

The charge pump circuit is shown in Figure 4-18. An undesirable feature of charge pumps is the charge injection produced by the overlap capacitance of the switch devices and by

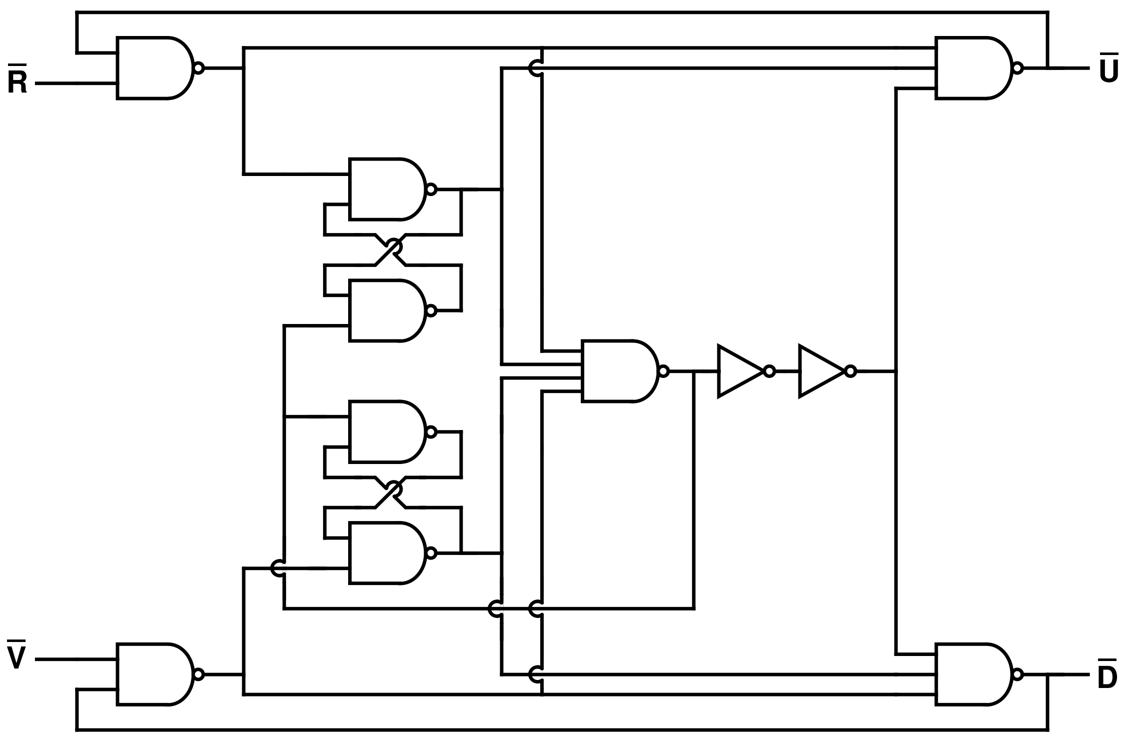

Figure 4-16: Logic diagram of the phase-frequency comparator.

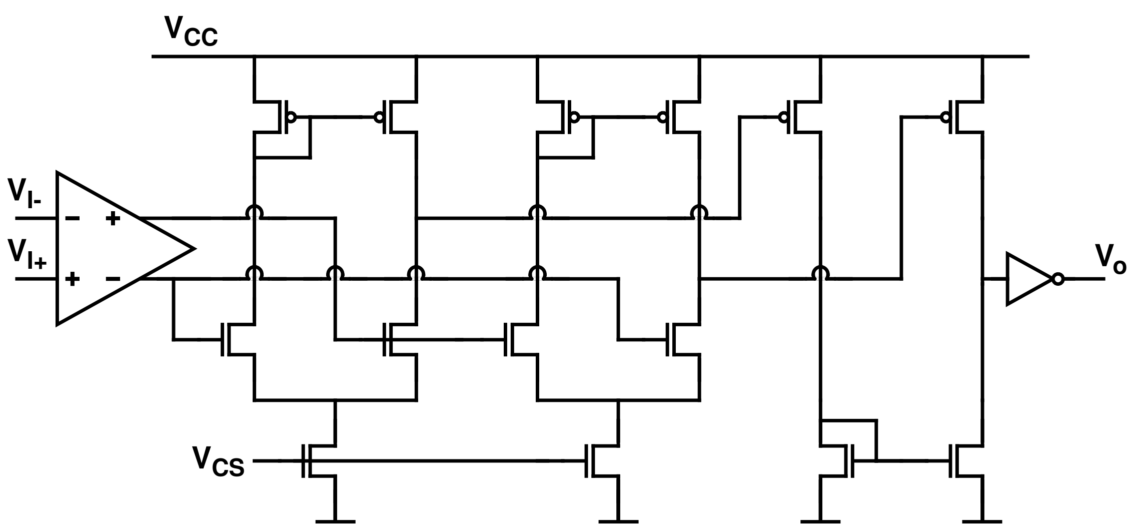

Figure 4-17: Schematic of the differential-to-single-ended converter circuit.

the capacitances at the intermediate node between the current source and switch devices. This charge injection will result in a phase offset at the input of the phase-frequency comparator when the PLL is in lock. This phase offset will increase as the charge pump current is reduced. When the ordering of the current source and switch devices are reversed so that the switch devices connect directly to the output node as shown in Figure 4-19 [1, 14], the output voltage is directly affected by the switching noise resulting from the overlap capacitance of the switch devices. In addition, the intermediate nodes between the current source and switch devices will charge towards the supplies while the switch devices are off. When the switch devices turn on, these intermediate nodes must charge toward the output voltage, removing charge from the output node in the process and resulting in a phase offset that depends on the output voltage. To combat this charge injection problem, the intermediate nodes can be switched to the output of the amplifier which buffers the output voltage while the other switches are off, but this approach can introduce an output voltage dependent phase offset of its own through an amplifier input offset. However, when the current source devices are connected to the output node as shown in Figure 4-18, the output voltage is isolated from the switching noise resulting from the overlap capacitance of the switch devices. In addition, the intermediate nodes between the current source and switch devices will charge toward the output voltage only by the gate overdrive of the current source devices, $V_{GS} - V_T$, an amount independent of the output voltage. Moreover, since both the NMOS and PMOS current sources always turn on each cycle, any charge injection will cancel out to first order with equal current source device sizes. The layout for the current source and switch devices contain no diffusion contact between the two gates to minimize the diffusion capacitance. The two devices are folded together for matching and layout constraints such that there are two unconnected intermediate nodes. The result is that the input offset can be minimized without employing a more elaborate solution.

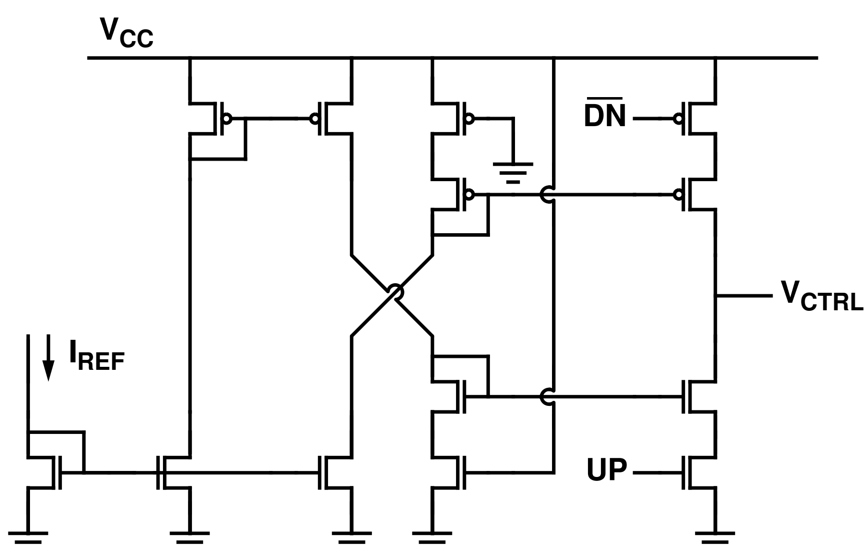

Figure 4-18: Schematic of the charge pump.

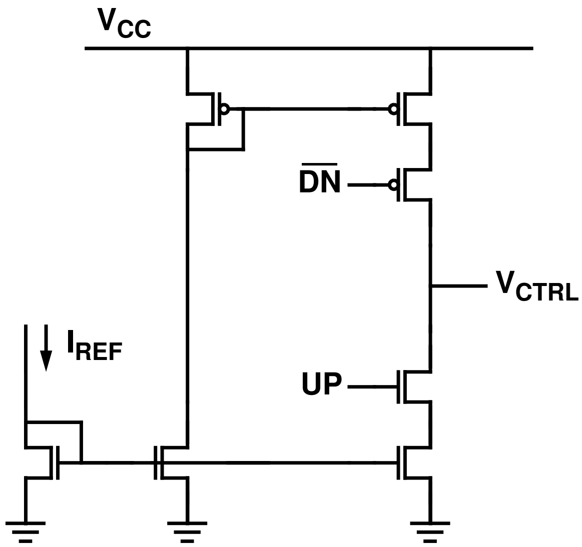

Figure 4-19: Schematic of the charge pump with charge injection offsets.

The low pass filter connects to the output of the charge pump and filters the control voltage to the oscillator. The capacitors are implemented using the gate capacitance of PMOS devices. The resistor is implemented with polysilicon. The low pass filter is completely contained in an N-well for supply and substrate noise isolation. Capacitor C_2 in shunt with the control voltage, which includes the input capacitance of the oscillator, is added to provide greater control over the value of the total capacitance on the control voltage. Ideally, the value of this added capacitance should represent at least 50% of the total capacitance in order to fix the total value at +/-50% if the input capacitance of the oscillator could vary from 0% to 200% of its nominal value. The capacitor in series with the resistor should have a value 15 to 20 times larger than the value of the total capacitance on the control voltage in order to provide adequate phase margin. With the charge pump current set at a minimum value of about 10μA, the value of the resistor is adjusted to maximize the overall phase margin of the PLL. The minimum allowable value for the charge pump is typically used in order to provide adequate separation between the natural frequency of the PLL and the minimum operating frequency of the oscillator.

4.6.2 Minimizing output jitter

To maintain the low jitter characteristics of the array buffer stages, the PLL must be specifically designed for low noise sensitivity. This section will address noise coupling sources affecting the complete PLL design for the array oscillator delay generator. It will also examine the dependence of the output jitter on the PLL loop parameters.

Capacitive coupling from the supply and substrate to the control voltage can be minimized in several ways. In order to eliminate the substrate noise sensitivity of the devices connecting to it, the control voltage is referenced to the positive supply and connected only to PMOS devices. This supply reference also allows the loop filter to be contained completely in a separate N-well, isolating it not only from capacitive coupling to the negative supply through the substrate, but also from substrate noise. All capacitors should be made from PMOS devices with their sources and drains connected to the positive supply. Connecting the capacitors to the negative supply would allow a direct coupling path for changes in the supply voltage to the control voltage. The control voltage should be carefully routed over N-well or polysilicon regions connected to the positive supply to avoid noise coupling through the interconnect. Care should be taken regarding how the positive supply that connects to the loop filter capacitors is routed. To avoid any supply currents modulating the control voltage through the resistance and inductance of the supply wires, the positive supply should be routed directly from the oscillator to the loop filter with nothing else attached.

Like the control voltage, all oscillator bias voltages should be routed over N-well, substrate, or polysilicon regions which are biased to the supplies to which the bias voltages are referenced. In addition, the bias circuits that generate the bias voltages should be made large enough so that any disturbances of the bias voltages are quickly restored. This is particularly important if the value of the bias voltage relative to its reference supply depends on the supply voltage of the oscillator, as is the case with the bias voltage for dynamically-biased NMOS current sources in the differential buffer stages. For such bias voltages, the coupling is implicit in the bias circuit. Adding capacitance to the bias voltages which does not depend on the supply voltage can be done to a limited extent. The capacitance will minimize the peak phase error at the expense of increasing its duration. The poles created by the capacitance added to bias voltages must be well beyond the loop bandwidth or they will decrease the phase margin of the loop.

The loop parameters also have a significant impact on the expected peak jitter of the PLL. The two key parameters that affect the amount of jitter are the loop bandwidth and the phase margin. All recovery and settling times scale with loop bandwidth. As a result, if jitter between the reference input and oscillator outputs is important, then the loop bandwidth should be made as large as possible, even if it means sacrificing some higher order roll off for filtering of spurious outputs caused by reference frequency sidebands in PLLs performing frequency multiplication. If only the period jitter is important, then a trade-off must be made between the recovery time of the PLL and the amount of reference frequency filtering desired. Since the bandwidth does not scale with the operating frequency, the bandwidth must be maximized with respect to the lowest operating frequency.

The phase margin of the PLL determines the amount of peaking or overshoot in the response of the PLL to a disturbance. For most disturbances, especially for a control voltage disturbance or a frequency change due to static supply voltage sensitivity, both of which cause phase ramps, the peak overshoot directly relates to the peak phase error observed. The peak overshoot will decrease as the phase margin is increased. However, once the system becomes overdamped, the settling time will be greatly increased by a slow loop time constant, even though the peak overshoot continues to decrease with increasing phase margin.

4.7 Experimental Results

The differential array oscillator delay generator discussed in this chapter has been fabricated in a 1.2-μm N-well CMOS technology. The array oscillator contains seven rings of five buffers per ring with nodes T_i connected to nodes B_{i+2}, as described in Chapter 2. 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 bias switches for the reset operation are implemented external to the ring bias circuits. The implementation includes two completely addressable output channels with offset cancellation circuits. The output channels share the outputs of the bias circuits used by the array core. These bias circuits provide control voltage buffering and have a total control voltage input capacitance of about 2pF. The loop filter of the PLL contains a 3kΩ resistor, a 212pF capacitor, and a 11pF capacitor. Simulations show that a charge pump current of 12μA will maximize the phase margin at 63 degrees and provide a unity gain frequency of 1MHz. A micrograph of the fabricated array oscillator delay generator with a superimposed floor plan is shown in Figure 4-20.

4.7.1 Measurement system

Successful testing of the array oscillator delay generator requires a large variety of measurements to be performed. These include measurements on the output phase linearity, the frequency sensitivity to supply voltage and control voltage, the time domain responses of the PLL to step changes in phase and supply voltage, and the noise frequency spectrum of the output jitter. Each of these measurements must be conducted on a number of different chips under a variety of operating conditions which include the value of the control

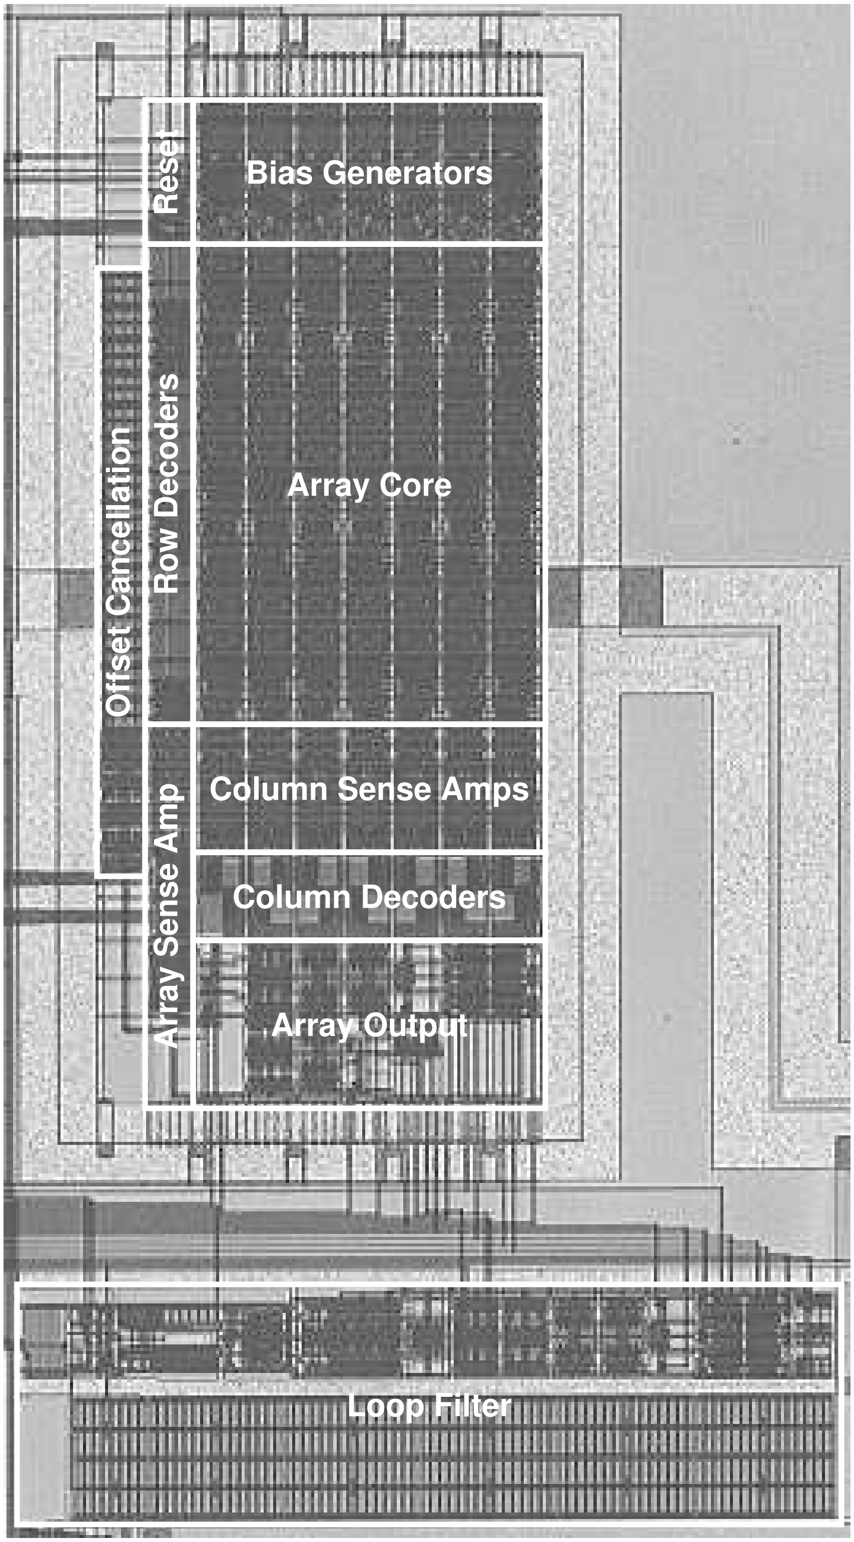

Figure 4-20: Die micrograph of the 5 x 7 differential array oscillator.

voltage and the frequency of the reference input signal. Some of these measurements are very complex in that they require careful control of the experimental setup and many supporting measurements. Other measurements require careful processing of large amounts of data. Many of these measurements require more than a day to complete.

In order to satisfy these measurement requirements, a completely automated computer controlled measurement system was developed specifically for testing this delay generator. It enables all the necessary measurements to be performed while unattended. It also supplies all of the data processing tools required for expeditiously analyzing the acquired data. The measurement system is centered around a computer controllable digital sampling oscilloscope which is used to perform most of the measurements under the supervisory control of a computer workstation. Other elements of the measurement system include computer controllable pulse generators, signal generators, multimeters, power supplies, and digital pattern generators.

4.7.2 Output accuracy

The measured output accuracy of the 5×7 array oscillator is summarized in Figure 4-21. 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 3ns and an LSB of 43ps, the peak error is about one LSB. In terms of differential non-linearity, the peak is a little less than one and a half LSB.

4.7.3 Oscillator performance

The measured frequency as a function of static supply voltage for several control voltages referenced to top supply is shown in Figure 4-22. 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 addition, the frequency varies linearly with control voltage except at high control voltages where thermal effects once again dominate.

Table 4-1 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 1.2-µm 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 [14].

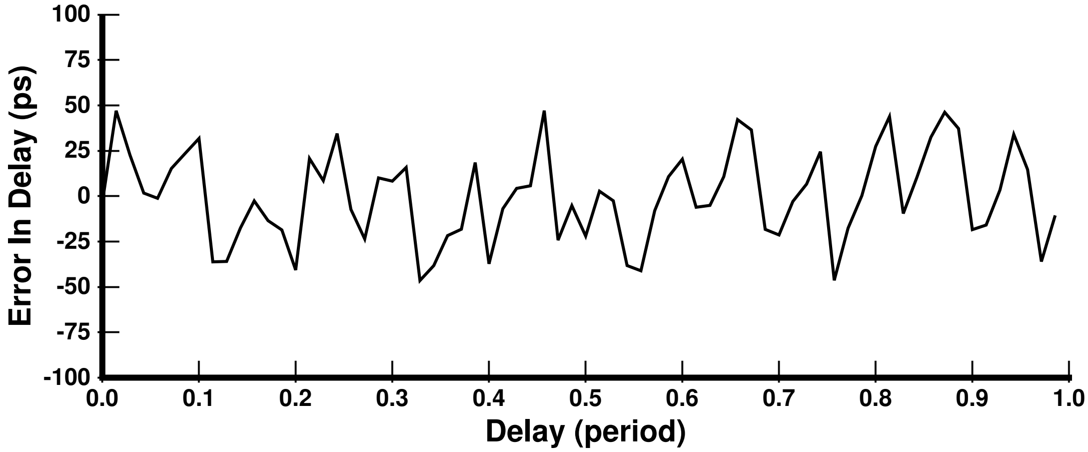

Vctrl	**3.25 V**	**LSB**	**43 ps**
Freq	**331 MHz**	**Peak INL**	**47 ps**
Period	**3.02 ns**	**Peak DNL**	**66 ps**

Figure 4-21: Measured output accuracy of the 5 x 7 differential array oscillator. The plot shows the error in delays for the 70 output phases as a function of the percentage of the period.

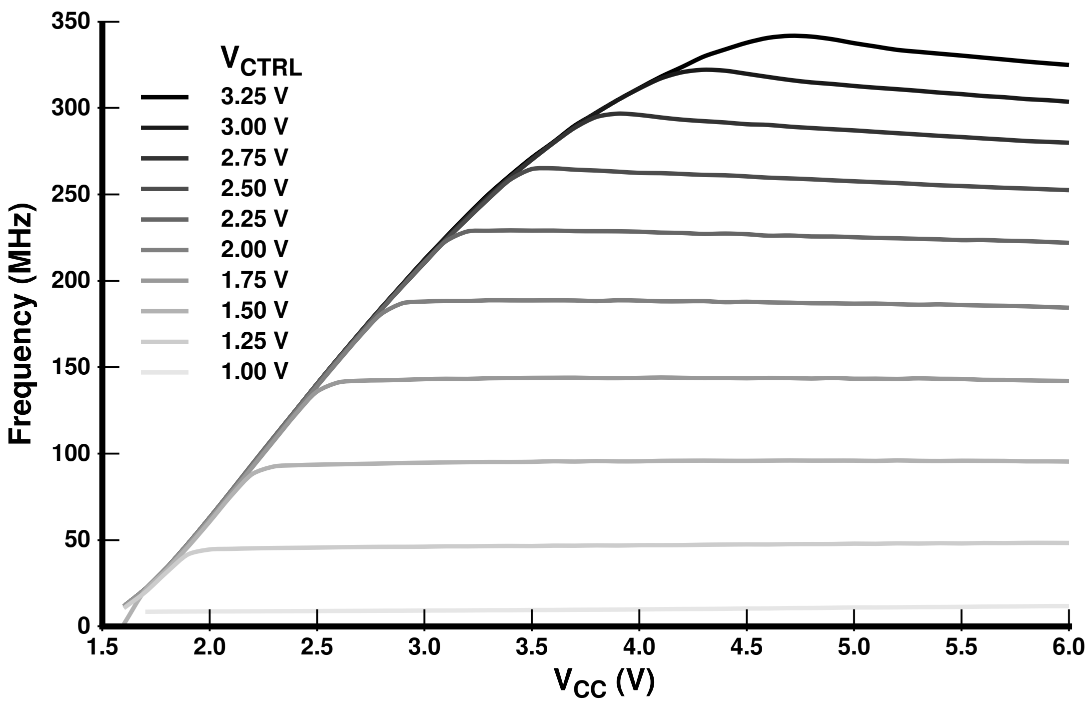

Figure 4-22: Measured frequency as a function of static supply voltage for the 5 x 7 differential array oscillator. At higher control voltages, thermal effects dominate the results due to increased power dissipation.

Frequency range, sensitivity:	5-355MHz, 141MHz/V
Static supply sensitivity:	0.25%/V @ 143MHz
Minimum supply voltage:	2.5V @ 143MHz
Power dissipation:	65mA @ 143MHz
Die area including PLL:	2.24mm^2
Technology:	1.2µm N-well CMOS

Table 4-1: Overall specifications of the array oscillator as a voltage-controlled oscillator.

4.7.4 PLL performance

The time domain response of the PLL to a step change in phase and a step change in supply voltage were characterized. In this design, the oscillator input to the phase-frequency comparator is multiplexed between the two output channels of the array oscillator. By toggling between the two array oscillator output channels at a submultiple of the reference frequency input, an oscilloscope that is triggered by the toggle signal can be used to explore the time domain response of the PLL to step changes in phase. A similar technique can also be used to inject step changes in the supply voltage. Figure 4-23 shows the phase response of the PLL operating at 100MHz to a half-cycle step change in phase at a charge pump current of 12µA. This response corresponds to a phase margin of about 60 degrees and a unity gain frequency of 1MHz. Figure 4-24 shows the phase response of the PLL operating at 100MHz to a 250-mV step change in supply voltage at a charge pump current of 12µA. The peak change in phase is about 7.5% of a cycle. This magnitude of phase change is about five times larger than that which would be expected from just static supply noise sensitivity. It results from supply noise coupling to the bias voltages in the array oscillator as explained below.

To further characterize the jitter performance of the PLL, a variety of peak jitter amplitude measurements were made between the reference frequency input and the oscillator output of the PLL in response to a 250-mV peak-to-peak sine wave signal on the supply voltage. These measurements were performed using the histogram function of the digital sampling oscilloscope under the control of an adaptive computer algorithm which controlled the

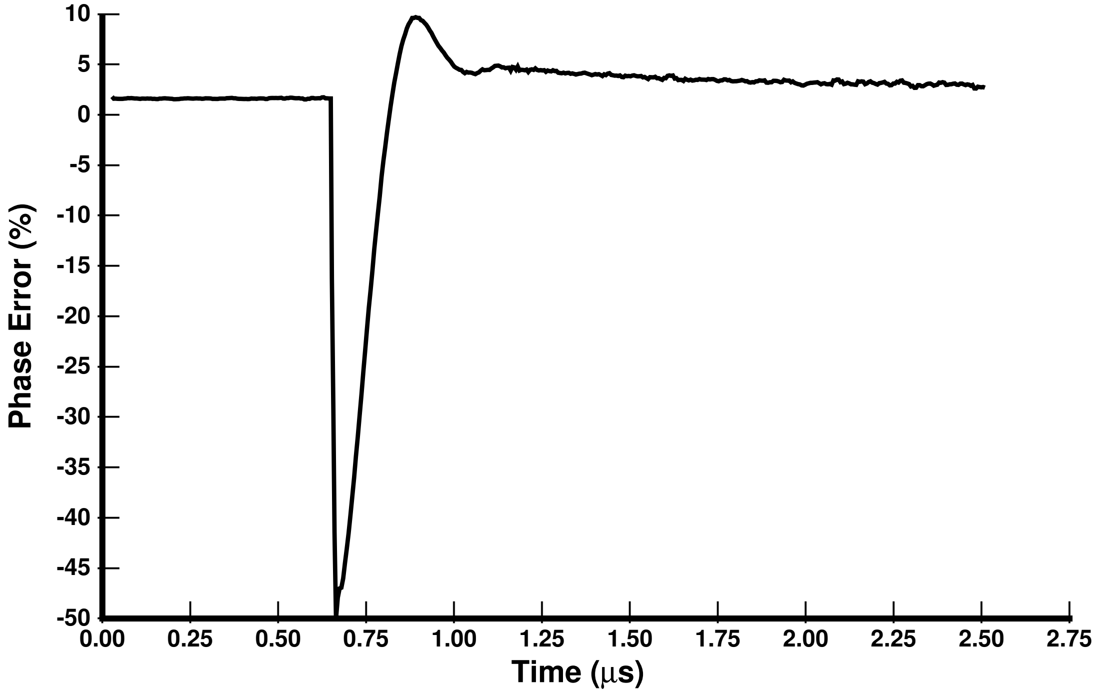

Figure 4-23: Response of the PLL to a half-cycle step change in phase.

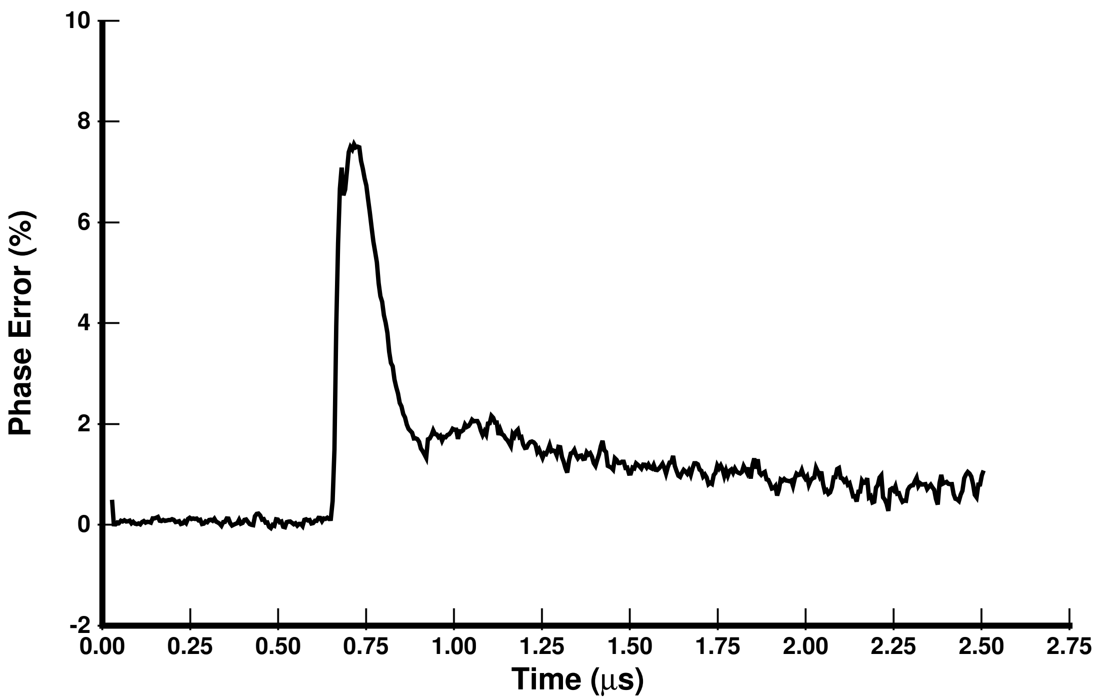

Figure 4-24: Response of the PLL to a 250-mV step change in supply voltage.

size and location of the histogram window. Figure 4-25 shows the peak jitter amplitude of the PLL with a charge pump current of 12µA as a function of sine wave frequency for several operating frequencies. All the operating frequencies indicate a maximum peak jitter amplitude with a sine wave frequency of about 3MHz. Figure 4-26 and Figure 4-27 show respectively the maximum peak jitter amplitude as a function of operating frequency at a charge pump current of 12µA in units of time and normalized to the operating period. These results show that the output jitter as a fraction of the operating period is roughly constant. This observation indicates that the recovery time of the bias voltages in the array oscillator, which is a constant fraction of the operating period, is a dominant factor in the measured output jitter.

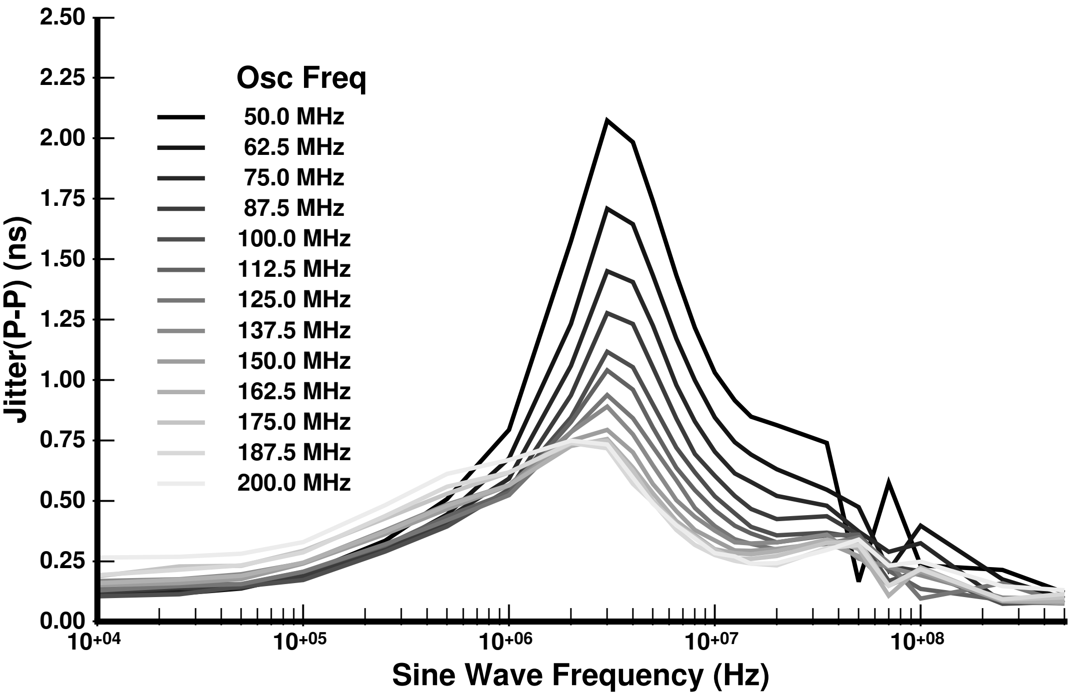

Figure 4-25: Peak jitter amplitude resulting from a 250-mV sine wave on the supply voltage for several operating frequencies.

Table 4-2 summarizes the overall specifications of the PLL. The peak jitter is larger than desired as a result of supply noise coupling to the bias voltages in the array oscillator through the bias voltage switches and interconnect capacitance. This supply noise coupling could have been removed from the existing layout, but in the interest of time, this option was not pursued. If the coupling had been removed, the jitter performance would have been significantly improved.

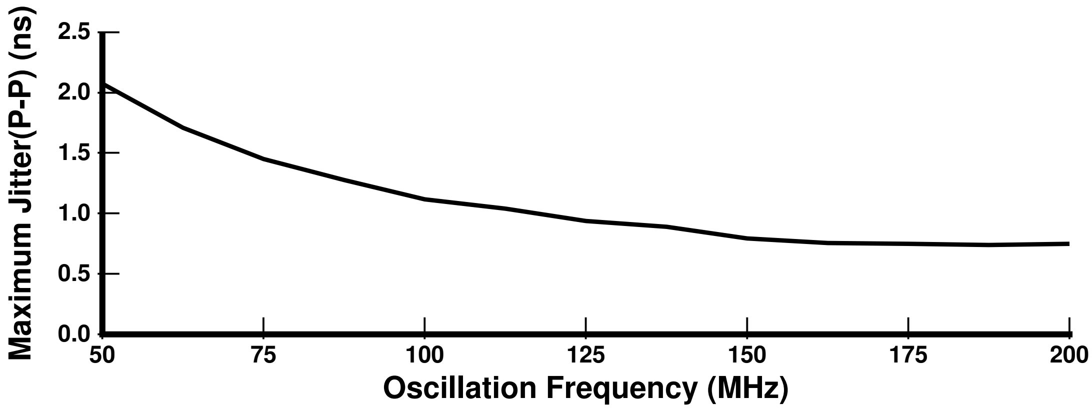

Figure 4-26: Maximum peak jitter amplitude resulting from a 250-mV sine wave on the supply voltage as a function of operating frequency.

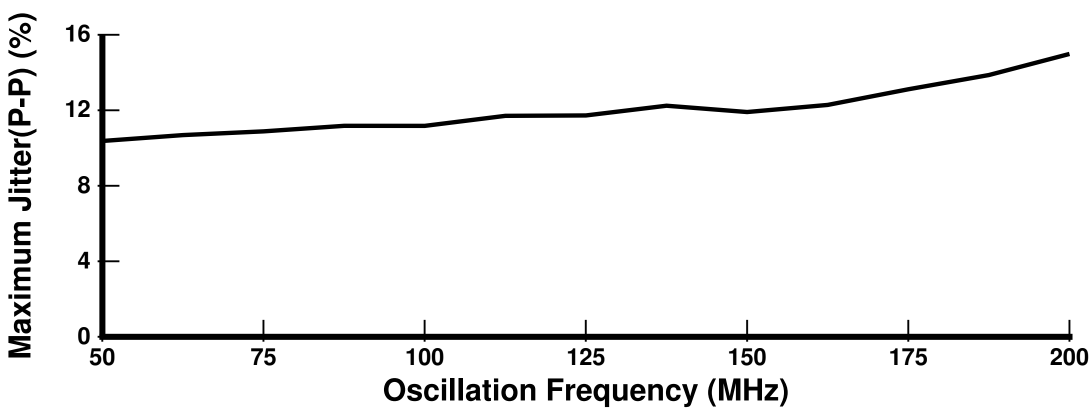

Figure 4-27: Maximum peak jitter amplitude normalized to period resulting from a 250-mV sine wave on the supply voltage as a function of operating frequency.

4.8 Summary

The experimental implementation of the array oscillator delay generator clearly demonstrates the ability of the array oscillator to produce precise delays with high resolution. The implementation also demonstrates the ability of the differential buffer stage with symmetric loads and replica-feedback biasing to achieve high static and dynamic supply noise rejection. While the desired jitter performance was not fully realized due to supply noise coupling to the oscillator bias voltages through the bias voltage switches and interconnect capacitance, such noise coupling is readily removable through a straightforward modification of the existing layout.

Charge pump current:	12µA
Unity gain frequency:	1MHz
Phase margin:	60deg
Phase comparator offset:	600ps
Operating frequency:	100MHz
Peak jitter (250-mV step):	1.51ps
Peak jitter (250-mV sine):	1.12ns

Table 4-2: Overall specifications of the PLL.

The array oscillator delay generator implementation discussed in this chapter is designed to work at the highest possible operating frequency. In maximizing the speed, the device sizes and thus the area and power consumption are increased to reduce the fraction of the array buffer output capacitance resulting from the interconnect. This design objective and the need to maintain high precision, pose severe constraints on the design of the output channel. However, the maximum operating frequency obtained is higher than can be used in the implementing IC technology by at least a factor of two.

The implementation of an array oscillator for a non-experimental application will typically have a maximum operating frequency that is significantly reduced from that of this implementation. The ring oscillators will have seven or more buffer stages each containing smaller devices. In addition, the output channel will have its own bias circuit allowing for higher current densities. These combined changes will permit the implementation of an array oscillator with minimal impact on resolution, yet relax the bandwidth constraints on the output channel design so that even higher precision can be obtained with less area and power than was achieved with the initial experimental implementation.

Chapter 5

Conclusions

Coupled oscillators can be effectively used to produce precise delays with high resolution. The array oscillator is the most interesting and innovative application of the coupled oscillator techniques described in this thesis. Based on a series of coupled ring oscillators, the array oscillator can produce delays with resolution equal to a buffer delay divided by the number of rings. By coupling several ring oscillators together, the dependence of the oscillation frequency on the number of buffers in a ring oscillator can be eliminated. Its high symmetry inherently provides very high linearity and precision. The delay line oscillator is another interesting structure. Its utilization of coupled ring stages enables it to precisely subdivide arbitrarily small delay intervals and to avoid false locks while operating over broad frequency ranges. The dual-input buffer stages and coupled ring stages can be applied to other applications involving precise delay generation. However, the precision actually realizable by coupled oscillators can be reduced if care is not taken to address such important external issues as jitter performance and delay errors due to bandwidth limitations in the output channel.

Buffer stages with high supply noise immunity and the ability to operate at low supply voltages can be designed using differential amplifiers based on symmetric loads and by employing replica-feedback biasing. 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. However, achieving high supply noise immunity goes beyond buffer stage design. The control voltage, all the bias lines, and the loop filter of the PLL must be carefully isolated from supply noise. The bias generator must be designed with noise immunity and fast dynamic recovery as important objectives. In addition, the loop parameters chosen for the PLL must be made with careful regard to supply noise.

Oscillators based on differential buffer stages have superior dynamic supply noise rejection properties over those based on single-ended buffer stages. However, the dynamic supply noise rejection of the oscillator itself is usually not the most significant factor in determining the overall output jitter from the complete phase locked delay generator. Single-ended oscillators have advantages to offer in other areas such as, reduced supply voltage requirements, lower power consumption, small size, and overall simplicity. With a well structured single-ended oscillator design, overall system performance can exceed that obtainable with differential oscillators.

The complexity of a PLL design can be easily managed provided that important tradeoffs are kept in proper perspective. The most important issue for delay generation is output jitter. When a system requires a PLL, its configuration should be established with the tradeoffs of the PLL in mind. Such tradeoffs include the possibility of establishing all of the system timing completely from the oscillator in order to avoid the jitter between the oscillator and reference which is always significantly larger than the jitter in the oscillation period.

An oscillator and PLL with high supply noise immunity still are not sufficient to achieve high precision. Great care must be taken in transmitting the high precision delays out of the delay generator without distortion due to mismatches in devices and interconnect capacitance. Maximizing the voltage swings is a desirable objective, but cannot always be accomplished. Offset cancellation structures can be added to recover distorted output delays resulting from address-dependent voltage offsets and duty cycle variations in the output signals.

A delay generator incorporating the design concepts described in this thesis was fabricated in a 1.2-μm N-well CMOS technology. Experimental results were obtained indicating that the delay generator can achieve an output delay resolution of 43ps while operating at 331MHz with a peak delay error of 47ps. These results conclusively demonstrate the ability of coupled oscillator delay generators to produce delays with very high resolution and precision. They also confirm that carefully implemented coupled oscillator delay generators can be successfully used in single-chip testers to test chips designed with integrated circuit technologies of higher speed than the testers themselves. More importantly, the experimental results obtained and the design concepts presented in this thesis have advanced the field of delay generation by creating a new class of delay generation structures and implementing techniques based on coupled ring oscillators which can produce precise delays with sub-gate delay resolution for a broad array of integrated circuit applications.

References

[1] M. Johnson, and E. Hudson, “A Variable Delay Line PLL for CPU-Coprocessor Synchronization,” IEEE J. Solid-State Circuits, vol. SC-23, no. 5, pp. 1218-1223, Oct. 1988.

[2] M. Loinaz, and B. Wooley, “A BiCMOS Time Interval Digitizer for High-Energy Physics Instrumentation,” Proc. of the IEEE 1993 CICC, May 1993, pp. 28.6.1-28.6.4.

[3] B. Kim, et al., “A 30-MHz Hybrid Analog/Digital Clock Recovery Circuit in 2-um CMOS,” IEEE J. Solid-State Circuits, vol. 25, no. 6, pp. 1385-1394, Dec. 1990.

[4] M. Bazes, “A Novel Precision MOS Synchronous Delay Line,” IEEE J. Solid-State Circuits, vol. SC-20, no. 6, pp. 1265-1271, Dec. 1985.

[5] F. Gardner, “Phaselock Techniques,” 2nd ed. New York, NY: John Wiley, 1979.

[6] Motorola MECL Device Data, 3rd ed. Motorola, Inc., pp. 6-27—6-42, 1979.

[7] D. Wolaver, “Phase-Locked Loop Circuit Design.” Englewood Cliffs, NJ: Prentice Hall, 1991.

[8] R. Best, “Phase-Locked Loops: Theory, Design and Applications.” New York, NY: McGraw-Hill, 1984.

[9] Hewlett-Packard 1.2um N-well CMOS Process.

[10] M. Horowitz, et al., “PLL Design for a 500MB/s Interface,” ISSCC 1993 Dig. Tech. Papers, pp. 160-161, Feb. 1993.

[11] T. Knotts, et al., “A 500MHz Time Digitizer IC with 15.625ps Resolution,” ISSCC 1994 Dig. Tech. Papers, pp. 58-59, Feb. 1994.

[12] J. Gasbarro, and M. Horowitz, “A Single-Chip, Functional Tester for VLSI Circuits,” ISSCC 1990 Dig. Tech. Papers, pp. 84-85, Feb. 1990.

[13] L. Alberto, “Probability and Random Processes for Electrical Engineering.” New York, NY: Addison-Wesley, 1989.

[14] I. Young, et al., “A PLL Clock Generator with 5 to 110 MHz of Lock Range for Microprocessors,” IEEE J. Solid-State Circuits, vol. 27, no. 11, pp. 1599-1607, Nov. 1992.

[15] K. Ware, et al., “A 200-MHz CMOS Phase-Locked Loop with Dual Phase Detectors,” IEEE J. Solid-State Circuits, vol. SC-24, no. 6, pp. 1560-1568, Dec. 1989.

[16] P. Gray, and M. Meyer, “Analysis and Design of Analog Integrated Circuits,” 2nd ed. New York, NY: John Wiley, 1984.

[17] F. Gardner, “Charge-Pump Phase-Lock Loops,” IEEE Trans. Communications, vol. COM-28, no. 11, pp. 1849-1858, Nov. 1980.

[18] D. Jeong, et al., “Design of PLL-Based Clock Generation Circuits,” IEEE J. Solid-State Circuits, vol. SC-22, no. 2, pp. 255-261, Apr. 1987.

[19] F. Gardner, “Phase Accuracy of Charge Pump PLL’s,” IEEE Trans. Communications, vol. COM-30, no. 10, pp. 2362-2363, Oct. 1982.

[20] Motorola MECL Device Data, 3rd ed. Motorola, Inc., pp. 6-81—6-84, 1979.

[21] J. Maneatis, and M. Horowitz, “Precise Delay Generation Using Coupled Oscillators”, ISSCC 1993 Dig. Tech. Papers, pp. 118-119, Feb. 1993.

[22] J. Maneatis, and M. Horowitz, “Precise Delay Generation Using Coupled Oscillators”, IEEE J. Solid-State Circuits, vol. 28, no. 12, pp. 1273-1282, Dec. 1993.

[23] W. Edson, “Frequency Memory in Multi-Mode Oscillators,” IRE Trans. Circuit Theory, pp. 58-66, Mar. 1955.

[24] A. Abidi, and R. Meyer, “Noise in Relaxation Oscillators,” IEEE J. Solid-State Circuits, vol. SC-18, no. 6, pp. 794-802, Dec. 1983.

[25] M. Wakayama, and A. Abidi, “A 30-MHz Low-Jitter High-Linearity CMOS Voltage-Controlled Oscillator,” IEEE J. Solid-State Circuits, vol. SC-22, no. 6, pp. 1074-1081, Dec. 1987.

[26] W. Liewellyn, et al., “A 33Mb/s Data Synchronizing Phase-Locked-Loop Circuit,” ISSCC 1988 Dig. Tech. Papers, pp. 12-13, Feb. 1988.

[27] R. Leonowich, and J. Steininger, “A 45-MHz CMOS Phase/Frequency-Locked Loop Timing Recovery Circuit,” ISSCC 1988 Dig. Tech. Papers, pp. 14-15, Feb. 1988.

[28] M. Banu, “100KHz - 1GHz NMOS Variable-Frequency Oscillator with “Analog and Digital Control,” ISSCC 1988 Dig. Tech. Papers, pp. 20-21, Feb. 1988.

[29] T. Liu, and R. Meyer, “A 250MHz Monolithic Voltage-Controlled Oscillator,” ISSCC 1988 Dig. Tech. Papers, pp. 22-23, Feb. 1988.

[30] K. Kato, et al., “A Low-Power 128-MHz VCO for Monolithic PLL IC’s,” IEEE J. Solid-State Circuits, vol. 23, pp. 474-479, Apr. 1988.

[31] M. Banu, “MOS Oscillators with Multi-Decade Tuning Range and Gigahertz Maximum Speed,” IEEE J. Solid-State Circuits, vol. 23, pp. 1386-1393, Dec. 1988.

[32] S. Sun, “An Analog PLL-Based Clock and Data Recovery Circuit with High Input Jitter Tolerance,” IEEE J. Solid-State Circuits, vol. 24, no. 2, pp. 325-330, Apr. 1989.

[33] M. Soyuer, and R. Meyer, “High-Frequency Phase-Locked Loops in Monolithic Bipolar Technology,” IEEE J. Solid-State Circuits, vol. 24, no. 3, pp. 787-795, Jun. 1989.

[34] T. Otsuji, and N. Narumi, “A 10-ps Resolution Process-Insensitive Timing Generator IC,” IEEE J. Solid-State Circuits, vol. 24, no. 10, pp. 1412-1418, Oct. 1989.

[35] S. Enam, and A. Abidi, “A 300-MHz CMOS Voltage-Controlled Ring Oscillator,” IEEE J. Solid-State Circuits, vol. 25, no. 1, pp. 312-315, Feb. 1990.

[36] P. Wallace, et al., “A GaAs 1.5Gb/s Clock Recovery and Data Retiming Circuit,” ISSCC 1990 Dig. Tech. Papers, pp. 192-193, Feb. 1990.

[37] J. Sonntag, and R. Leonowitch, “A Monolithic CMOS 10MHz DPLL for Burst-Mode Data Retiming,” ISSCC 1990 Dig. Tech. Papers, pp. 194-195, Feb. 1990.

[38] G. Moon, et al., “An Enhancement-Mode Voltage-Controlled Linear Resistor with Large Dynamic Range,” IEEE Trans. Circuits & Systems, p. 1284, Oct. 1990.

[39] S. Miyazawa, et al., “A BiCMOS PLL-Based Data Separator Circuit with High Stability and Accuracy,” IEEE J. Solid-State Circuits, vol. 26, no. 2, pp. 116-121, Feb. 1991.

[40] L. DeVito, et al., “A 52MHz and 155MHz Clock-Recovery PLL,” ISSCC 1991 Dig. Tech. Papers, pp. 142-143, Feb. 1991.

[41] B. Lai, and B. Walker, “A Monolithic 622Mb/s Clock Extraction Data Retiming Circuit,” ISSCC 1991 Dig. Tech. Papers, pp. 144-145, Feb. 1991.

[42] K. Kurita, et al., “PLL-Based BiCMOS On-Chip Clock Generator for Very High-Speed Microprocessor,” IEEE J. Solid-State Circuits, vol. 26, no. 4, pp. 585-589, Apr. 1991.

[43] T. Otsuji, and N. Narumi, “A 3-ns Rnage, 8-ps Resolution, Timing Generator LSI Utilizing Si Bipolar Gate Array,” IEEE J. Solid-State Circuits, vol. 26, no. 5, pp. 806-811, May 1991.

[44] H. Ransijn, and P. O’Connor, “A PLL-Based 2.5-Gb/s GaAs Clock and Data Regenerator IC,” IEEE J. Solid-State Circuits, vol. 26, no. 10, pp. 1345-1353, Oct. 1991.

[45] S. Enam, and A. Abidi, “MOS Decision and Clock-Recovery Circuits for Gb/s Optical-Fiber Receivers,” ISSCC 1992 Dig. Tech. Papers, pp. 96-97, Feb. 1992.

[46] N. Nguyen, and R. Meyer, “A 1.8-GHz Monolithic LC Voltage-Controlled Oscillator,” IEEE J. Solid-State Circuits, vol. 27, no. 3, pp. 444-450, Mar. 1992.

[47] M. Soyuer, and J. Warnock, “Multigigahertz Voltage-Controlled Oscillators in Advanced Silicon Bipolar Technology,” IEEE J. Solid-State Circuits, vol. 27, no. 4, pp. 668-670, Apr. 1992.

[48] N. Nguyen, and R. Meyer, “Start-up and Frequency Stability in High-Frequency Oscillators,” IEEE J. Solid-State Circuits, vol. 27, no. 5, pp. 810-820, May. 1992.

[49] L. Schmidt, and H. Rein, “Continuously Variable Gigahertz Phase-Shifter IC Covering more than One Frequency Decade,” IEEE J. Solid-State Circuits, vol. 27, no. 6, pp. 854-862, Jun. 1992.

[50] M. Flynn, and S. Lidholm, “A 1.2-um CMOS Current-Controlled Oscillator,” IEEE J. Solid-State Circuits, vol. 27, no. 7, pp. 982-987, Jul. 1992.

[51] C. Verhoeven, “A High-Frequency Electronically Tunable Quadrature Oscillator,” IEEE J. Solid-State Circuits, vol. 27, no. 7, pp. 1097-1100, Jul. 1992.

[52] T. Lee, and J. Bulzacchelli, “A 155-MHz Clock Recovery Delay- and Phase-Locked Loop,” IEEE J. Solid-State Circuits, vol. 27, no. 12, pp. 1736-1746, Dec. 1992.

[53] A. Buchwald, et. al., “A 6-GHz Integrated Phase-Locked Loop Using AlGaAs/GaAs Heterojunction Bipolar Transistors,” IEEE J. Solid-State Circuits, vol. 27, no. 12, pp. 1752-1762, Dec. 1992.

[54] M. Bagheri, et al., “11.6-GHz 1:4 Regenerating Demultiplexer with Bit-Rotation Control and 6.1-GHz Auto-Latching Phase-Aligner IC’s Using AlGaAs/GaAs HBT Technology,” IEEE J. Solid-State Circuits, vol. 27, no. 12, pp. 1787-1793, Dec. 1992.

[55] T. Hu, and P. Gray, “A Monolithic 480Mb/s Parallel AGC/Decision/Clock-Recovery Circuit in 1.2um CMOS,” ISSCC 1993 Dig. Tech. Papers, pp. 98-99, Feb. 1993.

[56] M. Banu, and A. Dunlop, “A 660Mb/s Clock Recovery Circuit with Instantaneous Locking for NRZ Data and Burst-Mode Transmission,” ISSCC 1993 Dig. Tech. Papers, pp. 102-103, Feb. 1993.

[57] M. Soyuer, and H. Ainspan, “A Monolithic 2.3Gb/s 100mW Clock and Data Recovery Circuit,” ISSCC 1993 Dig. Tech. Papers, pp. 102-103, Feb. 1993.

[58] A. Rothermel, and F. Dell’ova, “Analog Phase Measuring Circuit for Digital CMOS IC’s,” IEEE J. Solid-State Circuits, vol. 28, no. 7, pp. 853-856, Jul. 1993.

[59] T. Rahkonen, and J. Kostamovaara, “The Use of Stabilized CMOS Delay Lines for the Digitization of Short Time Intervals,” IEEE J. Solid-State Circuits, vol. 28, no. 8, pp. 887-894, Aug. 1993.

[60] T. Otsuji, “A Picosecond-Accuracy, 700-MHz Range, Si-Bipolar Time Interval Counter LSI,” IEEE J. Solid-State Circuits, vol. 28, no. 9, pp. 941-947, Sep. 1993.

[61] D. Reynolds, “A 320MHz CMOS Triple 8b DAC with On-Chip PLL and Hardware Cursor,” ISSCC 1994 Dig. Tech. Papers, pp. 50-51, Feb. 1994.

[62] R. Blauschild, “An Integrated Time Reference,” ISSCC 1994 Dig. Tech. Papers, pp. 56-57, Feb. 1994.

[63] K. Nogami, and A. El Gamal, “A CMOS 160Mb/s Phase Modulation I/O Interface Circuit,” ISSCC 1994 Dig. Tech. Papers, pp. 108-109, Feb. 1994.

[64] N. Ishihara, and Y. Akazawa, “A Monolithic 156Mb/s Clock and Data-Recovery PLL Circuit using the Sample and Hold Technique,” ISSCC 1994 Dig. Tech. Papers, pp. 110-111, Feb. 1994.

[65] B. Razavi, and J. Sung, “A 6GHz 60mW BiCMOS Phase-Locked Loop with 2V Supply,” ISSCC 1994 Dig. Tech. Papers, pp. 114-115, Feb. 1994.

[66] A. Pottbacker, and U. Langmann, “An 8GHz Silicon Bipolar Clock-Recovery and Data-Regenerator IC,” ISSCC 1994 Dig. Tech. Papers, pp. 116-117, Feb. 1994.

[67] Z. Wang, et al., “19GHz Monolithic Integrated Clock Recovery Using PLL and 0.3um Gate-Length Quantum-Well HEMTs,” ISSCC 1994 Dig. Tech. Papers, pp. 118-119, Feb. 1994.

[68] K. Nakamura, et al., “A 220MHz Pipelined 16Mb BiCMOS SRAM with PLL Proportional Self-Timing Generator,” ISSCC 1994 Dig. Tech. Papers, pp. 258-259, Feb. 1994.

[69] A. Waizman, “A Delay Line Loop for Frequency Synthesis of De-Skewed Clock,” ISSCC 1994 Dig. Tech. Papers, pp. 298-299, Feb. 1994.

[70] T. Lee, et al., “A 2.5V Delay-Locked Loop for an 18Mb 500MB/s DRAM,” ISSCC 1994 Dig. Tech. Papers, pp. 300-301, Feb. 1994.

[71] A. Efendovich, et al., "Multifrequency Zero-Jitter Delay-Locked Loop," IEEE J. Solid-State Circuits, vol. 29, no. 1, pp. 67-70, Jan. 1994.

[72] D. Mijuskovic, et al., "Cell-Based Fully Integrated CMOS Frequency Synthesizers," IEEE J. Solid-State Circuits, vol. 29, no. 3, pp. 271-279, Mar. 1994.