9-1-2005

A Piecewise Linear Approximation D/A Converter for Small Format LCD Applications

Mark Reisiger

A Piecewise Linear Approximation D/A Converter for Small Format LCD Applications

By

Mark D. Reisiger

A Thesis Submitted in

Partial Fulfillment

of the

Requirements for the Degree of

MASTER OF SCIENCE

in

Electrical Engineering

Approved by:

PROF:_____
Dr. Robert Bowman (Thesis Advisor, Department Head)

PROF:_____
Dr. Syed Islam

PROF:_____
Dr. James Moon

DEPARTMENT OF ELECTRICAL ENGINEERING

COLLEGE OF ENGINEERING

ROCHESTER INSTITUTE OF TECHNOLOGY

ROCHESTER, NEW YORK

SEPTEMBER, 2005

Thesis Release Permission Form

Rochester Institute of Technology

Kate Gleason College of Engineering

A Piecewise Linear Approximation D/A Converter for Small Format LCD Applications

I, Mark Reisiger hereby grant the permission of the Wallace Library of the Rochester Institute of Technology to reproduce my thesis in whole or in part. Any reproduction will not be for commercial use or profit.

Author: _____

Mark Reisiger

Date: 9/28/05

Table of Contents

List of Figures

1.0 Introduction

Low power operation is a driving requirement for the advancement of portable consumer electronics. As products get smaller and have more functionality the device integration requirements get tighter. This is certainly true of small format LCD applications like PDAs and cell phones. Recent advances in LCD technology have allowed for advanced circuitry to be built on the glass. This allows for the unique opportunity to integrate the LCD column driver with other circuitry rather than the traditional flip chip mounting on the glass. The integration of these D/A converters with digital circuitry presents a new set of design considerations. These considerations allow for the exploration of non-traditional architectures and algorithms. This work will explore these design considerations in detail and present a novel algorithm for conversion as well as a system implementation of this algorithm. The system implementation is compared to a standard linear converter to weigh the relative advantages of each. A high performance dynamically biased amplifier is developed for use in the D/A converter. This amplifier has a high slew rate while consuming a small amount of quiescent power.

2.0 System Design Considerations

In order to appreciate all of the design considerations that apply to the D/A converters, the operation of LCD panels must first be analyzed. While several new technologies such as Low Temperature Polysilicon (LTPS) and Continuous Grain Silicon (CGS) have emerged, the core operating principle remains the same.

2.1 Liquid Crystal Behavior

Liquid crystals are molecules with some unique physical properties. These materials have properties which are somewhere between those of solids and liquids. Liquid crystals can take on several different types of phases; the *nematic* phase is useful for display technologies. In the nematic phase all of the crystals have a tendency to orient themselves in the same direction as the molecules next to them. Liquid crystals also possess the property of *birefringence*. Birefringence means that a material has two different indices of refraction. The way in which light is refracted depends on its orientation as it passes through the material. Liquid crystals are also a dielectric and they are polar molecules so that their orientation can be controlled with an electric field.

All of these properties can be manipulated to produce an electrically controlled optical filter, which is the basis of an LCD. The *twisted nematic* architecture is the most common way to manufacture a LCD panel. Since the molecules in the nematic phase tend to align themselves to their surroundings, their orientation can be controlled. Figure 1 shows this structure. Liquid crystals are sandwiched between two layers of glass. Each layer of glass is coated in a transparent conductive material such as Indium Tin Oxide. The inside of the glass is coated with a polymer that has been brushed in one direction. The nematic liquid crystals tend to align themselves with this polymer. The twisted nematic structure gets its name because the polymers of the top and bottom panels are brushed orthogonally to each other. The liquid crystals naturally follow this 90 degree shift in orientation.

This structure produced some interesting optical properties. The gradual twist in the liquid crystals causes any light passed through the panel to be refracted resulting in a 90 degree shift in polarization. In order to complete the panel, external polarizers are placed on the outsides of the panel. These polarizers are oriented in the same direction as the brushed polymer on their respective panels. At steady state, this allows light to pass though the panel since the liquid crystals perform the necessary refraction to allow light to pass through orthogonal polarizers.

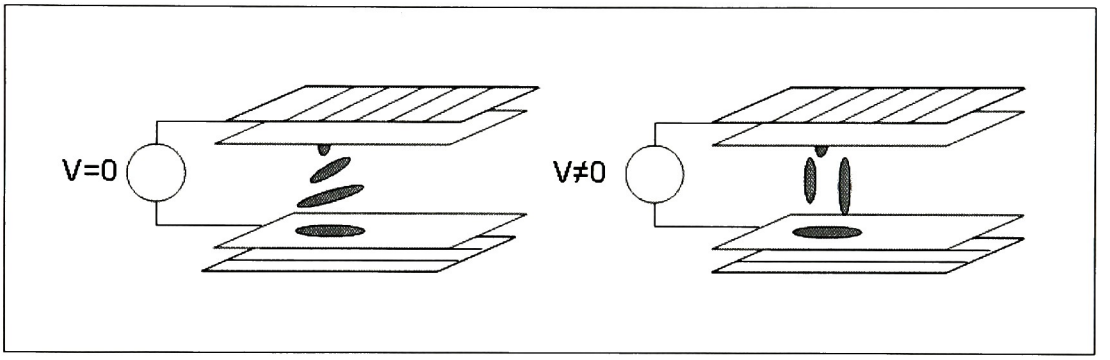

Figure 1 – Twisted Nematic LCD Panel

This behavior can be altered by applying an electric field between the glass plates as show in the right half of Figure 1. The electric field acts upon the polar molecules causing them to orient themselves in the direction of the field. This force is balanced by the nematic behavior of the liquid crystals. For low field strengths, the crystals stay more oriented in the twisted nematic structure. At high fields, the molecular polarization dominates the behavior. At these high fields, the liquid crystals do not refract the incoming light in a predicable manner. In this condition, the orthogonal polarizing filters block any light from passing though the panel. At intermediate field strengths, some of the liquid crystals provide the necessary refraction to allow light to pass through the

panel. These electric field strengths are the most important for displays since they allow for the adjustment of the optical transmission efficiency or *transmissivity* of the panel.

The transfer function between the electric field and the transmissivity of the panel is an important characteristic. Applied voltage will be considered rather than electric field since the dielectric width of the panel is fixed. Figure 2 shows the typical transmissivity transfer function of a panel. At low drive voltages the field does not exert much control over the liquid crystals. This manifests itself as a dead band in the transfer function of approximately 200 mV. Similarly, high voltages force all of the molecules to become aligned, completely blocking all light from passing. This can be seen as the saturation when the panel becomes overdriven. Another important characteristic to note is that the transmissivity does not depend on the polarity of the applied voltage. Practically, all panels are driven with alternating voltages to prevent the electroplating of ion impurities onto the glass due to DC bias. This is a major cause of stuck images.

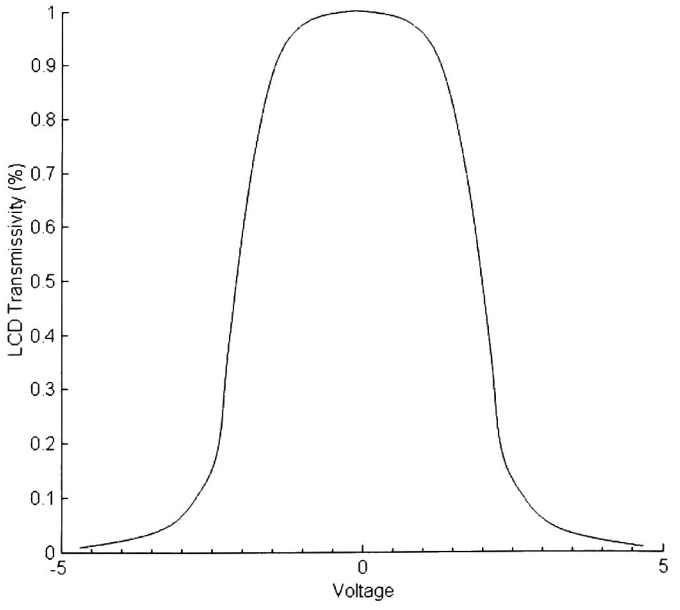

Figure 2 – LCD Transmissivity vs. Applied Voltage

2.2 LCD panel architecture

The ability of the twisted nematic architecture to modulate the light intensity must be exploited to create displays capable of producing images. Like other display technologies that exist, LCD panels are broken down into logical image cells called pixels. These pixels are composed of individual red, green, and blue cells. By modulating the relative amounts of these colors which are emitted, the whole spectrum of color is approximated as the light is mixed together in the human eye.

Since liquid crystal panels can only modulate light, color filtering is needed to produce the individual sub-pixels. A back light provides the complete optical spectrum at the rear of the panel. The liquid crystals modulate the amount of light that passes through the panel and individual color filters are masked on the top of the panel to select which portions of the spectrum pass through. Since the individual colors of every pixel needs to have independent control of the transmissivity, individual electrodes are needed. A system is needed to control each of these sub-pixels individually.

The *Active Matrix* (AM) architecture is used for high-resolution displays which are the focus in this work. In this scheme, the individual sub pixels are broken down into a matrix of rows and columns. A distinguishing feature of the active matrix display is that each pixel is only driven for a small fraction of the total display time. Capacitance present in the sub pixel holds the control voltage until it is updated during the next frame. Switches are needed on the panel in order to control the addressing of the sub-pixels. The *Thin-Film Transistor* (TFT) LCD panel is a method of constructing these switches.

In the TFT-LCD, a thin layer of silicon is deposited directly on the glass. This silicon is used to create MOSFETs which are arranged in the configuration shown in Figure 3.

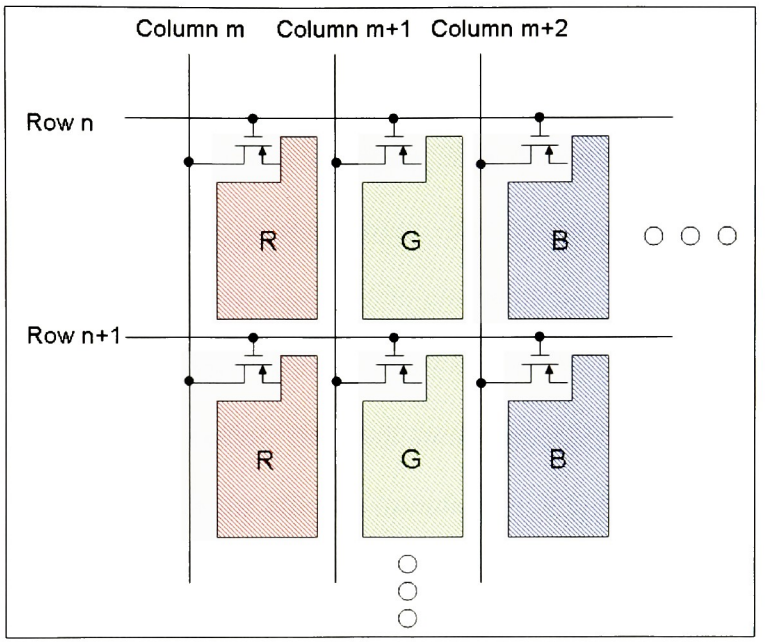

Figure 3 – TFT LCD Pixel Structure

The sub pixels in Figure 3 are addressed in the following manner. A gate drive signal is applied to the current row. This turns on the FET which connects the upper pixel electrode to its column. The appropriate voltage is applied to the column lines by the source driver. This voltage creates the field across the liquid crystals in that sub pixel region. The bottom electrode of the panel is common to all of the pixels.

The traditional TFT-LCD just described is typically an amorphous silicon display. The silicon that is deposited on the glass is ends up in an amorphous state. The silicon cannot be annealed using traditional electronics processing methods because the anneal temperature is beyond the 450 °C glass melting point. Transistors constructed from

amorphous silicon suffer from low carrier mobility. These devices are large and the gate drivers generally need to swing between +10 V to turn the device on and -5 V to turn the device off. The devices built in amorphous silicon are only useful as pixel switches.

The poor device characteristics in amorphous silicon panels force complicated packaging and drive solutions. All of the interconnect must be brought off of the glass panel to the row and column drivers. Typically, these drivers are mounted as a flip-chip package due to the large number of connections. The voltages required to drive the gates are much higher than modern standard CMOS processes use. This means that the gate driver chips must be independent from the other drivers and they will need an external DC-DC converter for the supplies. In small format displays, all of this complexity is extremely costly and prohibits the efficient manufacture of high resolution displays. For example, a quarter VGA (320x240) panel would require 960 connections (3 RGB * 320) to the source driver and 240 connections to the gate driver.

2.3 Low-Temperature Polysilicon

Enhancements in the processing of LCD panels reduce some of the complex systems integration problems with standard amorphous silicon. Low-Temperature Polysilicon (LTPS) is a process that has been developed by Phillips to produce transistors with better performance attributes on glass. A LPTS panel starts out like an amorphous silicon panel. After the layer of silicon is deposited the glass panel an excimer laser is used to anneal the silicon. The surface anneal creates polycrystalline silicon while remaining below the 450 °C glass limit. Transistors made from this polysilicon can achieve approximately 20% of the carrier mobility of those created in crystalline silicon.

Basic electronics can be created directly on the LTPS glass. Typical panel architectures are shown in Figure 4. The capacitors in this figure represent the capacitance of the individual pixel elements.

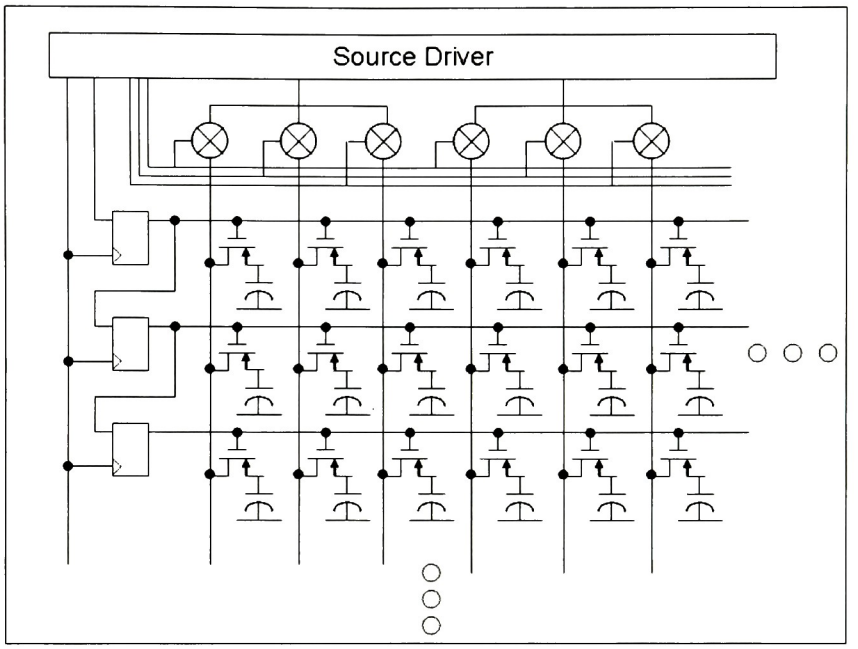

Figure 4 – LTPS Panel Architecture

In a LTPS panel, the gate driver has been eliminated. A series of shift registers are connected to the gates of the individual pixels. These shift registers are clocked at the line rate of the panel. A start token is driven into the first line register at the beginning of the frame. This token gets clocked down the panel, automatically selecting the proper row. In addition to the elimination of the gate driver, the interconnect requirements of the source driver have been reduced. It is possible to create low ratio (3:1) multiplexers in the source columns. The amount of external connections that are needed have been reduced by a factor of three. Since LPTS panels take extra processing steps to make, they cost more than amorphous silicon panels. This extra cost is somewhat compensated by

the elimination of the row drivers and other support circuitry. LTPS panels offer real financial benefits for small format displays.

2.4 Continuous Grain Silicon

Further improvements have been made with Continuous Grain Silicon (CGS) LCD panels. This proprietary process which has been pioneered by Sharp, claims to offer carrier mobility of 300 cm^2 / V s. This figure is 600 times faster than amorphous silicon and three times faster than LTPS. The improvements in transistors allow for even greater multiplexing ratios of 80:1. These multiplexers can now be controlled by on-glass addressing electronics; this means that a similar token passing structure is used for the source drivers as is used for the gate drivers. Figure 5 shows a typical CGS panel architecture.

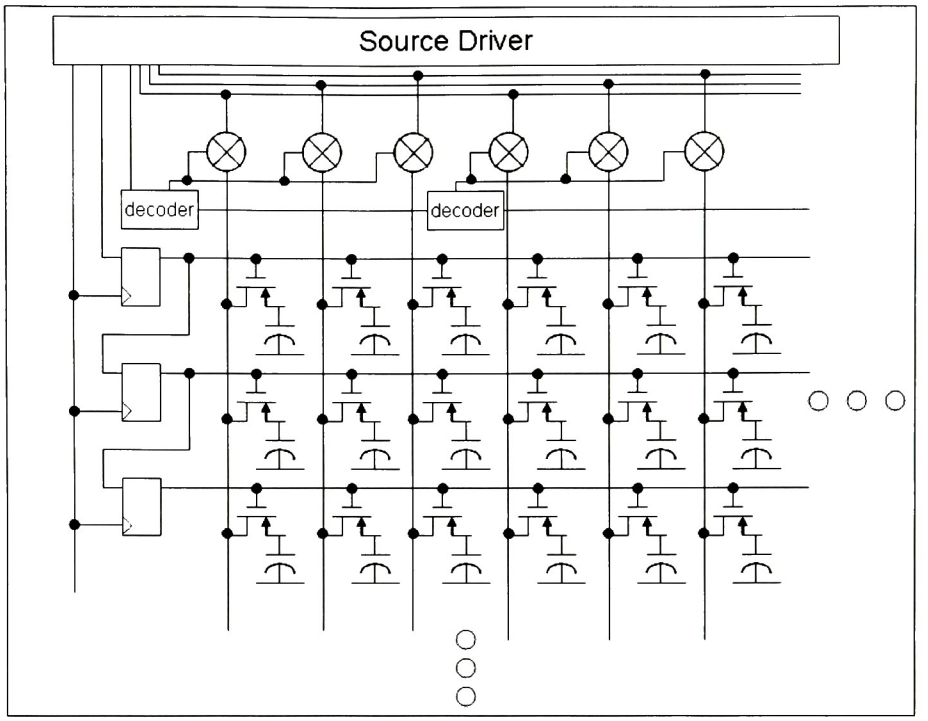

Figure 5 – CGS Panel Architecture

The benefits of CGS displays can be readily seen from a system design perspective. The 80:1 multiplexer ratios allow a QVGA display to be driven with four RGB (12 total) source drivers. The panel also needs just a few control signals. This means that a simple flex tape could be used to route interconnect and no flip-chip packaging is required. Since the source drivers are no longer required to be isolated, interesting packaging possibilities exist. A system on a chip approach encourages integrating all of the source driver D/A converters and output buffers in a single chip with the rest of the display processor.

2.5 *Timing and Drive Requirements*

The practical drive requirements placed upon the source drivers must be taken into account. The first requirement is the effective resolution that LCD panels support. Generally a panel can support 64 levels of transmissivity. The D/A converter needs more than 6-bits of precision due to the nonlinearity of the transfer characteristic. The intrinsic nonlinearity of the liquid crystals was discussed in Section 2.1. However, this is not the only nonlinearity in the system. Like all other video systems, gamma correction is needed in LCD panels. This curvature correction is needed to linearize the response of both the human eye and the video encoding method. Figure 6 shows a system level diagram with the gamma coefficients at each stage.

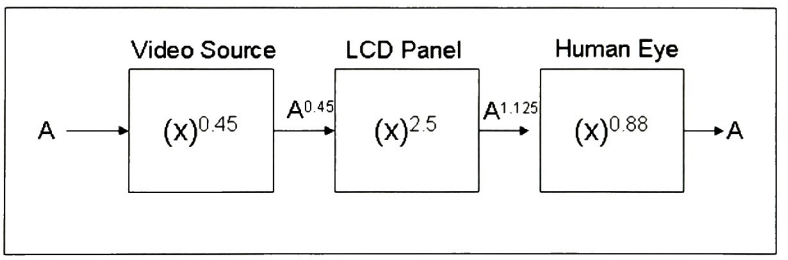

Figure 6 – Gamma Correction System

The source dependency of the gamma correction shows the need for an adaptable system. The LCD drivers should be able to generate any possible curve for the required system. In order to do this the worst case drive voltage must be found. Figure 7 shows a LCD transfer curve in which gamma correction has been applied. The panel transmissivity has been divided into 64 equal segments. The smallest voltage difference

between single steps of transmissivity is 18 mV. This step represents 0.36% of the 5 V full scale. Rounding up to the nearest bit implies that a 9-bit effective resolution linear converter is needed.

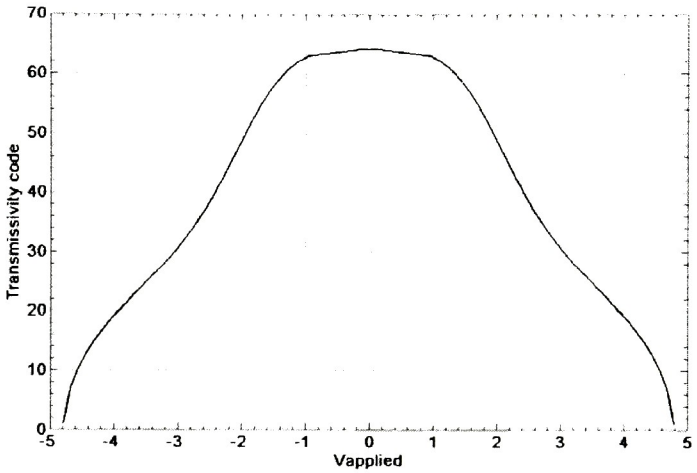

Figure 7 – LCD Transfer Function

The dynamic requirements of the LCD panel drive must also be considered. The first specification is the maximum conversion time for each pixel. This is a difficult requirement to determine since it is panel-specific. The horizontal and vertical blanking periods change slightly between manufactures. Worst case quantities are considered to determine an approximate pixel time. The line rate is calculated as follows:

$$line\_rate = \frac{1}{f_{rate}(Active\_lines + Blanking\_lines)} = \frac{1}{60(240+7)} = 67.5\mu s \qquad (1)$$

The assumptions made in (1) are as follows: The QVGA display is refreshed using 60 Hz video and the display requires 7 lines of video during the vertical refresh period.

The pixel time may now be determined from the line rate. Again assumptions about the horizontal blanking period are made. A worst case blanking interval of 3 μs is used. The system is also assumed to have four D/A converters per color.

$$pixel\_rate = \frac{(line\_rate - Blank\_time)}{\#Columns/\#DACs} = \frac{(67.5\mu s - 3\mu s)}{320/4} = 805ns \qquad (2)$$

Each D/A must be able to settle to 9-bit precision in 805 ns or have a conversion rate of 1.3 MS/s.

The LCD panel presents a load that looks primarily capacitive to the source driver. Each individual pixel has some small amount of capacitance, plus all of the metallization on the glass and in the flex cable presents additional parasitic capacitance. The multiplexer and pixel switches will present some on-resistance and the interconnect will have some small sheet resistance. These parasitics are difficult to model and the load that they present will change based upon the physical pixel which is addressed. A load of 100 pF is assumed as a design target for the output buffer.

3.0 D/A Architecture Selection

The D/A architecture of traditional column drivers is fairly straight forward. These 9-bit converters are implemented with a resistor string architecture. This structure offers several advantages such as simplicity in implementation and guaranteed monotonic behavior. Simple polysilicon resistors match well enough to be used to achieve the required 9-bit precision. This means that a standard CMOS process may be used. The major limitation of the R-string D/A is the power-speed tradeoff. The major time constant of the system is formed by the current limited charging of the parasitic capacitances of the amplifier and interconnect. Assuming that these parasitic capacitances are constant for any given system, the only way to make the converter faster is to burn more power in the R-string. This implies that there exists some minimum power consumption for any given conversion rate. A less important limitation of this architecture is that the output amplifier offset will limit the channel matching of the system.

To consider alternative architectures the system goals should be restated. Low power consumption is the primary design objective. Small die area and compatibility with digital CMOS processes are also important. In order to meet these goals a switched-capacitor implementation is developed. The use of switched-capacitors eliminates the need for resistor string bias current.

The major objective of this investigation not so much the architecture of the D/A, but the method of producing the LCD transmissivity curve. Typically, a 9-bit linear converter is used and a look-up table creates the nonlinearity. This work explores the

concept of using a piecewise linear approximation of the curve as a method of reducing the converter performance requirements.

3.1 Piecewise Linear Approximation Method

The goal of the piecewise linear approximation method is to reduce the system requirements of producing a nonlinear curve by the brute force lookup table method. In this method the precision of the converter is set by the minimum step size of the curve that is to be reproduced. In curves with large dynamic ranges, this approach can result in many unused bits.

The core concept behind the piecewise linear approach is to break the curve up into segments. Each segment can then be constructed using a low precision converter. Figure 8 shows the basic concept. The solid curve is the desired transfer function. The two linear approximations are fitted tangent to the desired curve. This method leaves some residual error as the linear approximations leave the curve near their junction. The precision of the application dictates whether this error is acceptable.

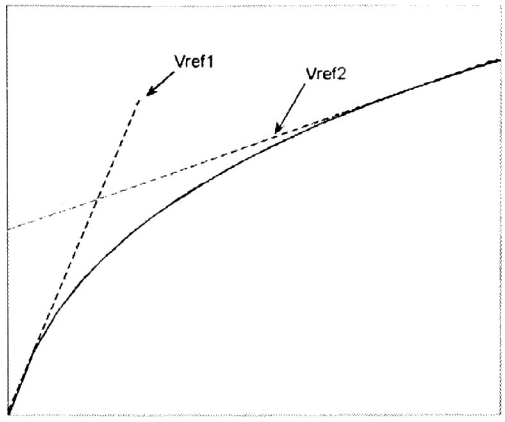

Figure 8 – Piecewise Linear Approximation

If the error is too great, another linear segment can be inserted between segments 1 and 2. This process can continue until the absolute error of each word is under some maximum error threshold. This limit of this approach occurs when the curve is so nonlinear that each word requires its own linear approximation. At this level the technique degenerates into the look-up table approach. Therefore it is appropriate for slightly nonlinear curves.

It should also be noted that there is room for optimization in the linear segment placement. Placing the segments tangent to the curve creates the worst possible absolute error at the transition points. Most applications will benefit from spreading the approximation error across all of the codes. An optimized algorithm would use a technique such as a least-squares regression line to minimize the mean square error in any segment. This technique has further problems since the length of any single segment is variable so many iterations would be necessary to determine the optimal solution for the

entire curve. The optimization of the approximation method is beyond the scope of this work.

A proof of concept algorithm is developed to apply this technique to any general transfer curve. The curve is set along with its limits of precision. This curve is then approximated with a single segment; the end points of which are on the limits of the curve. The absolute error for each code is calculated. Assuming that any errors are outside the limits of precision another break point is added at the code with the maximum error. New linear segments are created using these additional break points. The end points of the new segments are always placed on the desired curve. After each operation, the calculated values are rounded to the nearest bit of precision. This process iterates until all of the errors are below the preset threshold. For each curve the following information is generated: The digital word boundaries where the segment is valid, the gain required to generate the slope and the offset voltage required to align the curve.

A Matlab simulation was used to verify this technique for the LCD transfer curve. The test curve in Figure 9a is generated as the required inverse curve to linearize the LCD transmissivity curve in Figure 7. The limits of precision are set to 6-bit levels for the gain and offset coefficients. The curve error limits were set to ½ LSB of transmissivity at the 6-bit level. Figure 9b shows the curve that is generated using this technique. The approximation is done with sufficient accuracy in twelve segments. Figure 10 presents the absolute approximation error for each word along with the error thresholds.

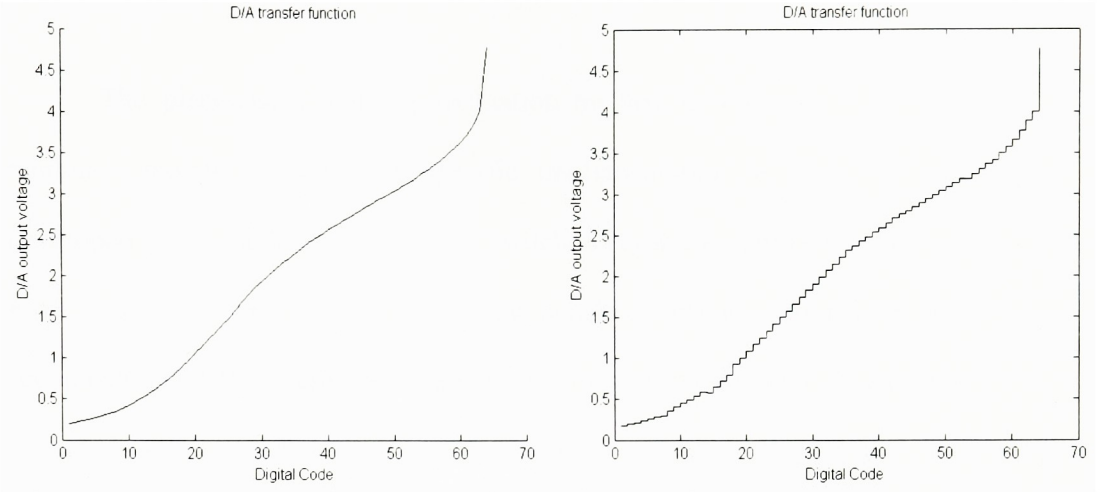

Figure 9 – (a) Sample Curve (b) Approximation

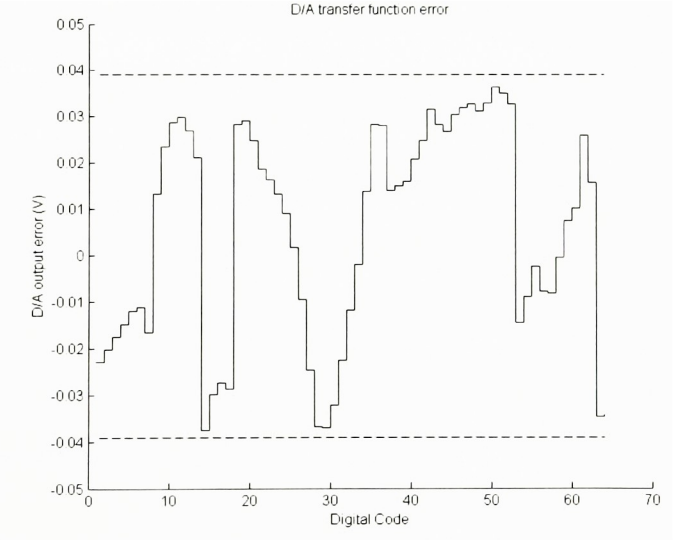

Figure 10 – Sample Curve Approximation Error

3.2 Piecewise Linear System Implementation

The piecewise linear approximation method is a generic process for creating nonlinear transfer curves. A specific implementation of the architecture is now developed. As stated in Section 3.0 a switched capacitor approach is used. The idea of producing nonlinear transfer curves using switched capacitor circuits has been looked at previously [1], [2]. There are several problems with the hard-wired approach. The first is that the LCD transfer curve and the required gamma correction may not be known. Secondly, considerable circuit effort is needed to generate the nonlinearities. For a system such as this one with twelve breakpoints, the number of amplifiers and comparators needed would ruin any efficiency that the technique might have over a single 9-bit converter.

The piecewise linear technique developed in Section 3.1 only requires a few components. In order to generate any one segment a reference voltage and an offset voltage are needed. The reference and offset voltages which are used for any arbitrary incoming word can be selected digitally since the break points were set on code boundaries. Thus the transfer function of the system looks like (3) where m is a selection index for each breakpoint region.

$$Vout = Vref[m] \cdot N + Vos[m] \tag{3}$$

In order to maximize the programmability of this system, D/A converters will be used to generate the V_{REF} and V_{OS} coefficients. This technique also allows for all of the

curve generation information to be stored digitally. This also means that the system must perform three 6-bit D/A conversions for every output word.

To reduce the area requirement of the system, one converter is used for all three conversions. The V_{REF} and V_{OS} voltages will be temporarily maintained in sample and hold circuits. Figure 11 shows the proposed architecture. The system operates in the following manner. A 6-bit word comes into the digital control block. A comparator circuit determines which linear segment the word belongs to. The result of this comparison selects the proper conversion codes for V_{REF} and V_{OS} from look up tables. V_{REF} is the first word to be converted. The timing controller selects the fixed reference voltage to precharge the D/A converter. This result is stored in the sample and hold shown at the bottom of Figure 11. The V_{OS} word is the next to be converted. Once again the timing controller precharges the D/A to the fixed reference voltage. The controller selects the offset sample and hold to store the result. Finally the timing controller precharges the D/A with the soft reference voltage present on the sample and hold. This conversion represents the first term in equation (3). The result of this conversion will be referred to as V_{WORD}. The result is stored in the final sample and hold. In order to minimize the timing requirements for the system, the D/A core and the output buffer act as a two stage pipeline. The D/A core converts pixel $N+1$ while the output buffer drives pixel N. The output driver sums together V_{OS} and V_{WORD} to form the final word. This buffer also provides any scaling needed from the core and provides the current drive necessary for large capacitive loads.

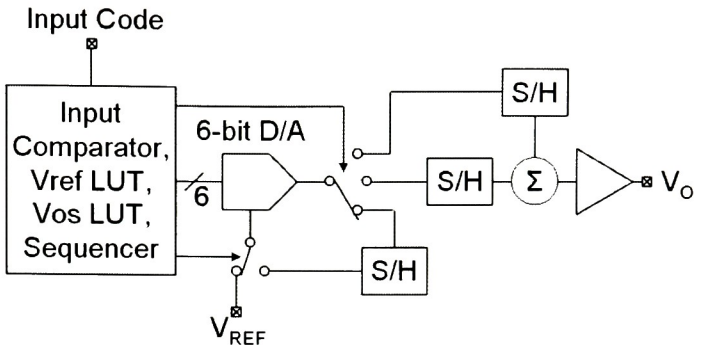

Figure 11 – System Block Diagram

4.0 Architecture Implementation Analysis

Many of the system requirements have been analyzed in Sections 2 & 3. These requirements are set by the physical properties of a LCD panel and by the needs of the piecewise linear approximation method. All of these system level requirements will be broken down into specific sections in order to accurately define the performance characteristics of each block.

4.1 D/A core

The D/A core must be the highest performing component of the overall system. This converter is required to generate the reference, offset, and word voltages during a single pixel cycle. The decision to use a switched capacitor approach has already been made, but there are several types of converters which fall within this domain. The benefits of each architecture will be analyzed; with a particular focus on conversion speed, power consumption and die area.

Architectures that use the charge redistribution principle [3] offer low power operation since they only utilize a single differential amplifier. Charge redistribution is a precharge-evaluate architecture so the amplifier must be capable of settling twice in a single conversion cycle. This architecture relies on a bit dependent number of capacitors. For the 6-bit converter, 64 unit cells of capacitance are needed in the feedback path. Likewise, 63 individual capacitors are needed in the input stage. These 127 capacitors need to be matched to greater than 9-bit precision. This resolution implies that matching better than 0.19% is needed. The matching requirements in [4], [5] place a limit to the minimum area and therefore the minimum capacitance of these unit cells. These minimum sized capacitors place additional drive requirements on the differential amplifier. This amplifier must be capable of settling these 64 unit capacitors in the feedback path.

Cyclic converter architectures offer several improvements when compared to charge redistribution converters [6], [7] and [8]. These converters operate serially; they convert a single bit at a time and hold the incremental result. This means they rely on a pair of matched capacitors. This considerably reduces the matching complexity of the system. The die area used by a cyclic converter is much less than a charge redistribution converter because of the reduction of capacitors. The slew rate requirements of the amplifier are reduced for the same reason. Unfortunately, the settling time requirements of the amplifier are drastically increased. Due to the serial nature of the architecture $2n+1$ clock cycles are needed for an n-bit converter. This requirement makes it unlikely to use a single D/A for all three words as shown in Figure 11. Since the pixel time is approximately 800 ns, the differential amplifier would have a settling time requirement of

61 ns for each 6-bit conversion. Anything faster than this would seem infeasible for a low-power design.

Pipeline converters combine the benefits of both of the previous architectures. This architecture is an extension of the cyclic architecture and is typically used in A/D converters to increase their conversion speed [9]. This architecture places one cyclic converter for each bit. These single bit converters pass their incremental results to the next converter. Pipeline converters still offer the advantages of the small area of a cyclic converter while eliminating the need to settle $2n+1$ times during a pixel. However these benefits are counteracted by the need to have n amplifiers. The large number of amplifiers does not lend itself well to low power designs.

Since the primary goal of this system is low-power operation, the charge redistribution architecture is chosen for its lenient settling time requirements and use of a single amplifier. Several refinements need to be made to the basic architecture in order to fit this application. First of all, this design is intended to operate within a digital environment. This means that it must have good rejection of all of the switching noise present on the chip. Due to the large number of capacitors that are needed, the smallest value that can be matched will be used. Small capacitors manifest charge injection errors as large voltage errors. In an attempt to reduce both of these effects a fully differential architecture will be used. The other non-ideality which must be contended with is offset voltage in MOS amplifiers. On a 5 V power supply, 9-bit precision implies approximately 10 mV as a LSB voltage. The offset voltage for untrimmed amplifiers will typically be in this range. Rather than resorting to trimming methods, an auto-

zeroing approach is taken. The gain amplifier in [10] samples the amplifier offset in the precharge clock phase and subtracts it during the conversion phase.

The complete D/A architecture is shown in Figure 12. The converter operates as follows. During the precharge phase (Φ_1), the programmable capacitor array is charged to either V_{REF} or signal ground[1] depending on the incoming digital word. The amplifier is placed in a unity gain configuration with no input signal except for the amplifier's offset voltage. This voltage is sampled on the feedback capacitors and the capacitor array. During the evaluate phase (Φ_2), the feedback capacitor is placed back in the amplifier feedback loop. The programmable capacitor array is driven to signal ground. This forces the voltage across the array to equal the amplifier offset voltage which effectively transfers the charge onto the feedback capacitor. The incomplete charge transfer due to offset voltage is counteracted by the offset error sampled during the precharge cycle.

[1] Signal ground is actually V_{CM} for this fully differential system

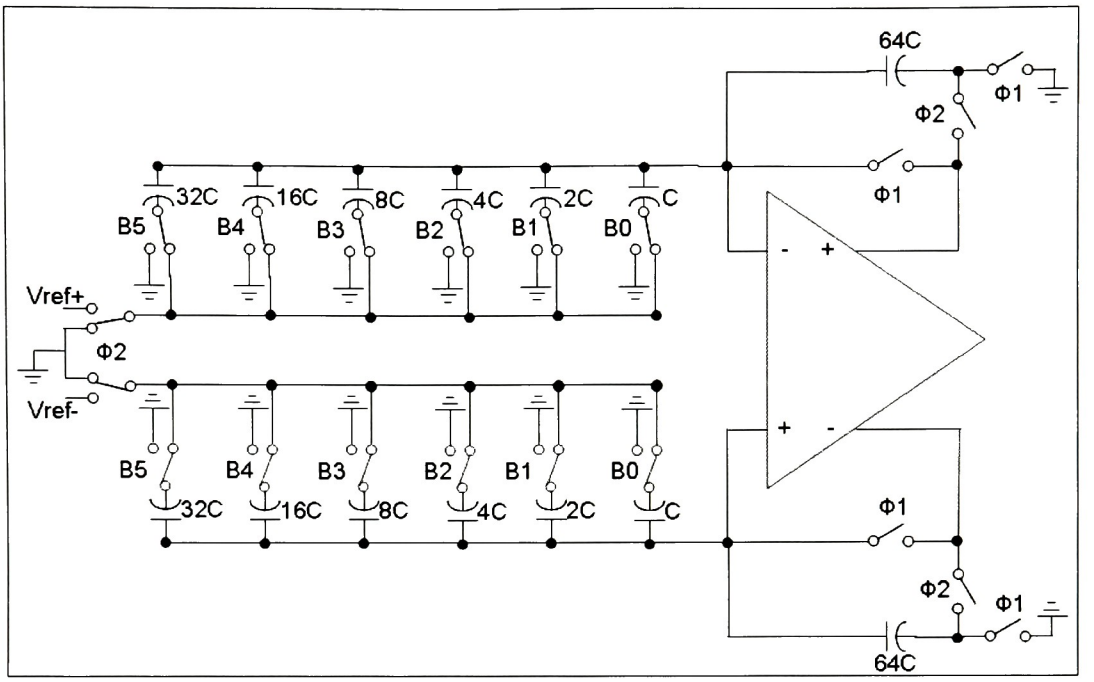

Figure 12 – D/A Architecture

4.2 Sample and Hold

The sample and hold circuits are required to store the intermediate voltages used in the construction of the final word. The requirements for each sample and hold are slightly different. The circuit that stores V_{REF} must also be capable of driving a significant capacitive load. This circuit must charge the programmable capacitor array in the D/A core during the soft V_{REF} conversion cycle. The sample and hold circuits for V_{OS} and V_{WORD} must simply provide the voltages to a summing junction in the output amplifier.

The V_{REF} sample and hold circuit has many of the same drive requirements as D/A core. This circuit has a similar capacitive load since it will be driving the programmable capacitor array. The settling time requirements must also be the same. All of the non idealities of the system are present too. Because of these issues and to maximize design reuse the same differential gain stage with auto-zeroing will be used for the sample and hold. Figure 13 shows this architecture. The issue of amplifier compensation needs to be addressed. These differential amplifiers are load compensated. The sample and hold will have varying capacitive loads based upon the switch states. The first variation is based on the incoming digital word. The number of unit capacitors charged can vary between 0 and 63. The amplifier also only drives its own feedback capacitance during the cycles when the load is disconnected. Depending on the specific design of the amplifier it may require switched dummy capacitors for stability.

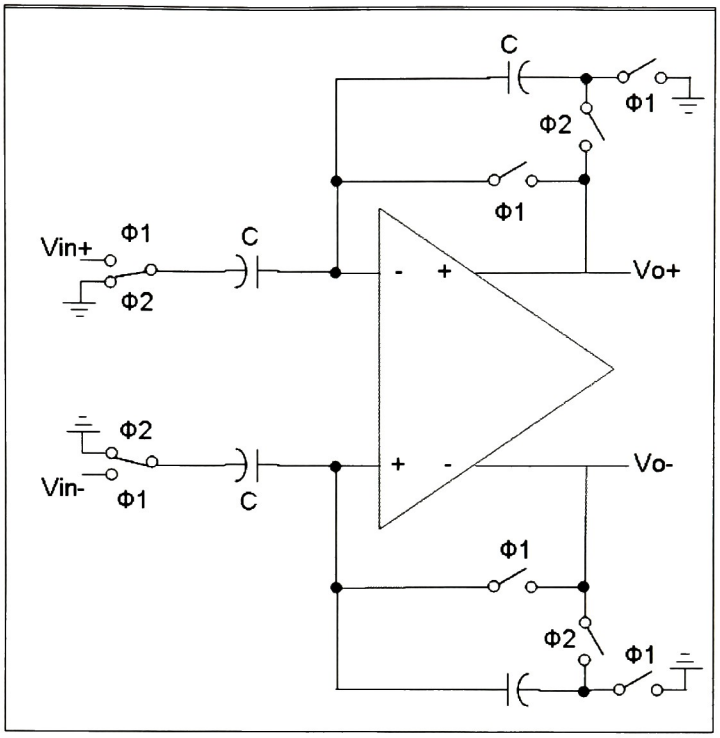

Figure 13 – V_{ref} Sample and Hold Architecture

The requirements for the output sample and hold circuits are quite different. Since these circuits only need to drive their charge into an amplifier virtual ground they can be purely passive devices. These passive devices have several unique requirements. The first of these issues is charge injection. This issue is dealt with in the same manner as before. Fully differential structures and transmission gates are utilized in order to minimize the charge injection errors. The other issue with the passive sample and hold circuits is junction leakage. The reverse biased source and drain diffusions will flow some small amount of current which affects the hold accuracy of these circuits. The amount of this junction leakage is largely temperature dependent and is the biggest issue at high temperatures. Both of these considerations will define the minimum size capacitor which may be used in order to stay between the error thresholds.

4.3 Output Driver

The output driver is responsible for several tasks. This amplifier performs the summing action of V_{OS} and V_{WORD}. The amplifier also performs the differential to single ended conversion and signal scaling needed to interface off chip. In order to accomplish this task mirrored SC gain stages are used. The negative differential input is used to charge the bottom capacitor array which is driven into ground[2] with a dummy feedback loop during the evaluate phase. This operation sets the amplifier input common mode voltage at the positive input. The clocks that are labeled Φ_1 are masked based upon their function. This precharge clock is used to zero the feedback capacitors during the beginning of the pixel cycle. All of the other Φ_1 clocks are loading controls from the D/A core. Depending on the core conversion state, this array could be charged with V_{WORD} or V_{OS}. In addition, even and odd pixels have their own holding capacitors. These redundant capacitors are not shown in Figure 14 for clarity. These extra capacitors are necessary since the system operates with a pipeline. The D/A core must hold word $N+1$ while the output driver is still settling word N. Another feature not shown in Figure 14 is the V_{OS} polarity switch. The Φ_1 mask for the V_{OS} hold capacitors queries the state of input bit B_0. This state determines whether the hold capacitor is charged with standard or inverted polarity. The polarity switch allows offsets to be subtracted from as well as added to V_{WORD}.

---

[2] This terminal is tied to V_{SS} while all of the other signal grounds are V_{CM}

An important performance limitation in the output driver needs to be mentioned. It does not employ any auto-zeroing circuitry like the core amplifiers because of settling time restrictions. This means that the random offset in this amplifier will dominate the absolute accuracy and the channel matching of the system.

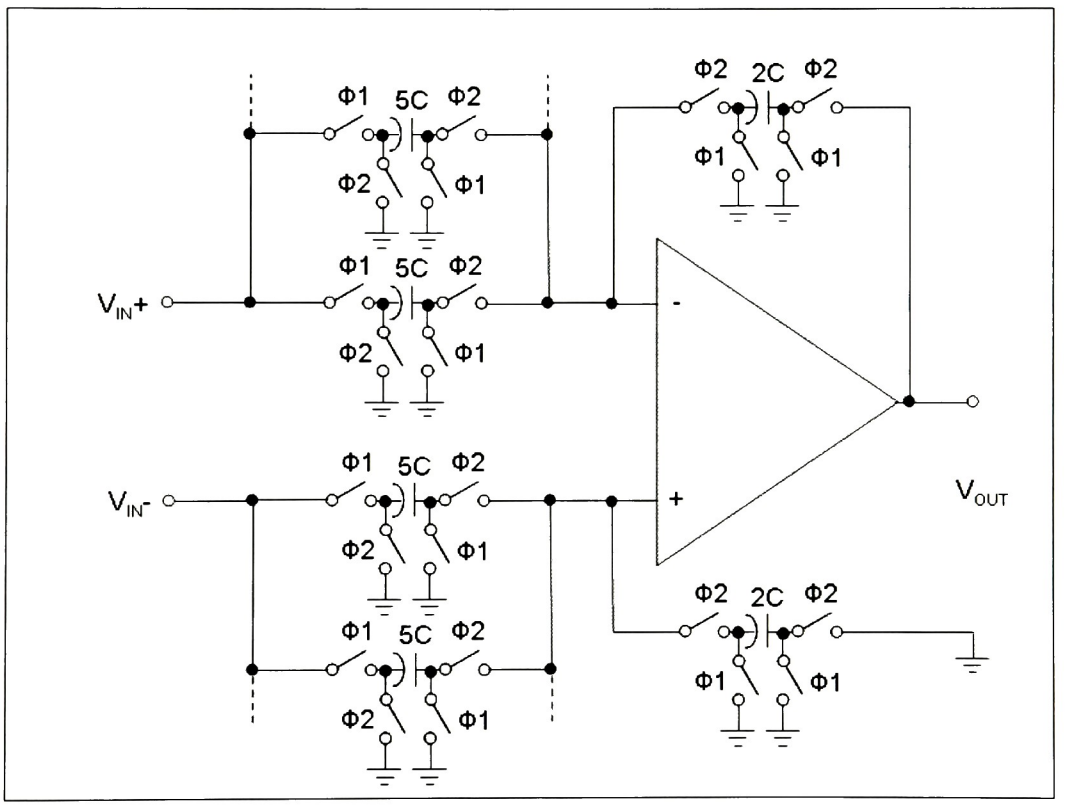

Figure 14 – Output Amplifier Block Diagram

5.0 Proof of Concept Design

The system level design of the piecewise linear D/A converter was described in Section 4. This established the performance criteria for each of the system components. Detailed component design and simulation will be described in this section. The physical implementation of all of the circuitry will also be included.

The fabrication costs of this proof of concept design were supported by a research grant though the MOSIS Educational Program. In order to provide a realistic test chip the TSMC035P2 process has been selected. This 0.35 μm process is the smallest feature size process that offers the 5 V support that is needed to directly drive the LCD panel. The 0.35 μm transistors offer enough performance that they could be utilized in the digital controller which would be integrated with this design. TSMC035P2 is a dual poly process which is necessary for precision capacitor designs. All simulations were performed using the Cadence Spectre simulator as part of the Cadence design suite.

5.1 Differential Amplifier

The fundamental component in this design is the differential amplifier. This amplifier is burdened with fast settling time requirements while maintaining low power operation. The amplifier must also have relatively high gain for converter accuracy and it must drive large capacitive loads.

This amplifier must have a certain level of static performance in order to meet the accuracy requirements of the system. Since the system is required to have 9 bit accuracy this amplifier must have much better precision than that. In order to get an estimate of the required gain, DC errors will be allowed to be ¼ LSB. The gain requirements for the worst case unity gain conversion are calculated from the standard feedback transfer function equation in (4).

$$F = \frac{A}{1 + A\beta} = 1 - \frac{1}{\frac{1}{4} \cdot 2^9} = 0.999517$$
$$A = \frac{F}{1 - F\beta} = \frac{0.999517}{1 - 1 \cdot 0.999517} = 2047 = 66dB \tag{4}$$

In order to understand the current drive and frequency response requirements of the amplifier the minimum sized capacitor must be determined. The matching capabilities of this process were obtained from the recommendations of a successful commercial 10-bit switched-capacitor design[3]. The recommendations state that there was no yield loss due to capacitor matching when 11 μm by 11 μm top plates were used and the corners were chamfered to maximize the area to perimeter ratio. The unit cell is derived in (5).

$$Cunit = \left[11\mu m \cdot 11\mu m - 4(\frac{1}{2} \cdot 1\mu m \cdot 1\mu m) \right] \cdot 870 \, {}^{aF}\!/_{\mu m^2} = 103.5 fF \tag{5}$$

[3] Thanks to Satoru Shingai of Analog Devices for this information

The TSMC035P2 process does not have a thin oxide between the two polysilicon layers which results in the small sheet capacitance value. Regardless, the D/A core amplifier in Figure 12 has 64 unit cells in its feedback path resulting in a 6.6 pF load plus any capacitance present in the sample and hold circuits that it drives.

The load capacitance will certainly be the primary pole of the amplifier. The bandwidth requirements of the system will determine the requirements for the output driver. In order to get a rough estimate of the bandwidth, single pole settling behavior is assumed. The amplifier will need to settle to ¼ LSB (0.048%) of the final value in 133 ns. The required bandwidth assuming a linear response is calculated as follows.

$$
\begin{aligned}
V &= 1 - e^{-\frac{t}{A/\tau}} \\
\frac{A}{\tau} &= {}^{t}\!/_{-\ln(0.000488)} = {}^{133\,ns}\!/_{7.6} = 17.5ns \\
GBW &= \frac{1}{2\pi\tau} = 9\,MHz
\end{aligned}
\tag{6}
$$

This bandwidth estimate is very liberal since it does not account for multiple pole settling behavior or nonlinear effects like slew rate limitations.

The slew rate requirement of the amplifier is needed to determine the amplifier bias. The slew rate is estimated in the following manner. The worst case swing will be approximately 2 V internally. The assumption is made that the amplifier will not slew for more than 10% of the settling time. This implies that a slew rate of 150 V/μs is needed. In order to slew the 6.6 pF load at this rate 1 mA of drive current is needed.

These large load currents eliminate Class A amplifiers as an option. In order to achieve a D/A power consumption of several hundred microwatts much smaller quiescent current consumption is needed. A power consumption goal of 100 μW per amplifier has been set. The internal amplifiers will use the 3.3 V power supplies in order to keep power consumption down and to utilize the thin gate oxide devices. This means that this amplifier must work with 30 μA of bias current. To achieve this, efficient class AB structures are needed with active-to-idle current ratios of more than thirty-three.

Traditional class AB structures are not likely to work in this application. These are generally output stages which are added onto a core amplifier. Most of the amplifier gain is created before these output stages and this large gain places a dominant pole at the internal node. This means that the pole created with the load capacitance would need to be pushed well beyond the gain bandwidth. A large amount of bias current is needed to create a low impedance output node. Assuming that the output stage is a source follower, the driving impedance is approximately *1/gm*. To place the secondary pole at 50 MHz where it would contribute 10 degrees of phase shift would require a transconductance of 2.2 mS. Using moderate size transistors of 50 μm / 350 nm and with k` for NMOS of 97 $\mu A/V^2$; 180 μA of bias current is needed to achieve this output impedance. Alternative Class AB structures are needed that will allow the load to be the primary pole and a high impedance gain node.

Dynamic Biasing (DB) is one solution to this problem. A dynamically biased amplifier works by changing the bias current of its output stage as a function of the input signal. This is typically done with a traditional OTA. The bias current is changed under large signal conditions. This approach allows for large slew rates with high static output

impedances for good DC gain. Dynamic biasing is typically accomplished with a transient detector which adjusts the tail current on the input pair. This technique creates weak positive feedback in the amplifier. Besides the stability concerns, the extra circuitry typically consumes a good deal of current to maintain a high bandwidth.

Pseudo-Differential (PD) amplifiers offer the potential to be used for dynamic biasing. These circuits are usually employed in low-voltage applications where head room can't be spared on a tail current source [11], [12]. Pseudo-differential pairs are of interest because they do not saturate like true differential pairs. The intrinsic square-law behavior of a MOSFET can be used to provide the amplifier bias and optimize g_m.

Figure 15 shows the schematic of the basic dynamically biased PD amplifier. Under small signal and idle conditions the bias current is set by the gate drive on the input pair *M1*. Signals applied to the input are mirrored to the output stage formed by *M3* and *M4*. Transistors *M5* and *M6* ideally create symmetric gate drive signal for the NMOS devices so that the n-channel and p-channel devices are driven equally. Gain is created due to the high channel impedance of these transistors. The gain of this amplifier is approximately $gm_1(ro_5||ro_3)$ assuming that the mirror ratios are 1:1.

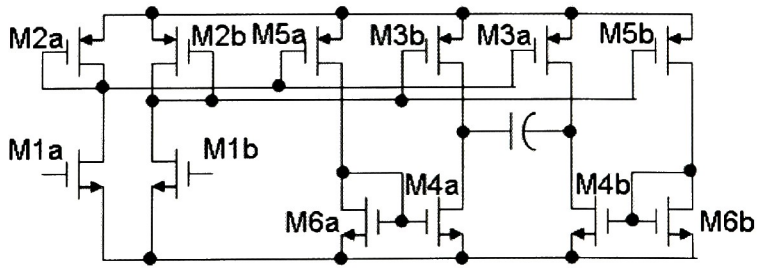

Figure 15 – Pseudo Differential Amplifier

The large signal operation is as follows: The signal step is assumed to be larger than the quiescent gate overdrive on the input pair. The input step is also symmetric about the DC gate drive voltage. This step turns *M1b* off since the gate drive is now below the threshold voltage. Without supply current the *b*-side mirrors shut down and the output mirrors become high impedance. Transistor *M1a* sinks much more current, ideally following square-law behavior. This current is mirrored throughout the *a*-side transistors to the output bridge. *M3a* pulls up on one side of the load while *M4a* pulls down on the other side of the load. Excluding mirror losses the current ratio of the amplifier is proportional to the square of the input gate drive. If the gate drive is biased near the threshold voltage then small absolute voltage steps can result in large current ratios since $(V_{GS} - V_T)$ will be large. Ideally, even bigger current ratios could be obtained with subthreshold conduction due to the exponential voltage-current relationship.

Unfortunately, this precision biasing is one of the largest practical problems in PD amplifiers. Process variations in V_T won't allow tight enough control of the quiescent current. Another factor is that the source is tied to the negative rail. This means that power supply noise will couple into the amplifier. Dynamic biasing is a non traditional use for the PD amplifier. Large static overdrives are used with small signal swings to maximize amplifier linearity in low-voltage applications [11].

A solution to the process problem has been described in [13]. This source follower circuit acts as a voltage sink with a V_T dependency. When a pseudo-differential pair is connected to the low impedance node the process sensitivity is reduced. This circuit is shown in Figure 16.

Transistor *M3* functions as a source follower in this circuit. The feedback loop created utilizes *M2* as a current sink to enhance the follower behavior of *M3*. Assume that a disturbance current is applied to the tail node of this circuit. This additional current would create a small error voltage at the tail node due to the channel impedance of *M2*. Transistor *M3* appears like a common-base amplifier to the disturbance. The impedance transformation between the source side and drain side of *M3* amplifies the disturbance signal. This amplified voltage is applied as additional gate drive on current sink *M2*. The additional gate drive allows the channel to conduct the disturbance without increasing V_{DS2}. This feedback loop effectively creates a low-impedance node at node 1.

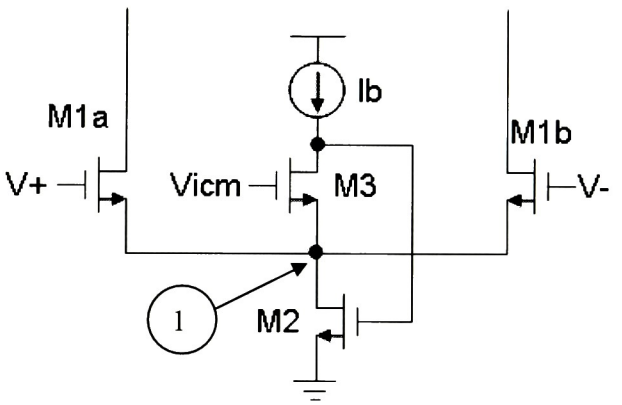

Figure 16 – Pseudo-Differntial Pair with Reduced V_T Sensitivity

Now the input pair transistors *M1* from Figure 14 are connected to this voltage node instead of the negative rail. If the follower transistor *M3* is matched to the input pair *M1* then the quiescent biasing is controlled by the relative voltage applied. This is more easily seen with a large signal analysis applied to Figure 16. The voltage source is considered as a degenerating element to the input pair modeled as a non linear resistor.

This means that any signal V_{IN} is split between V_{GSI} and V_{DEG}. The follow assumptions are made. A channel length modulation parameter λ_3 is used to model the channel conductance of *M3* since the current flowing through *M3* is always a constant I_b. Infinite channel resistance is assumed for *M2* since the impedance looking into the source of *M3* is much less than the drain of *M2*. It is also assumed that the current source I_b has infinite channel resistance. The equation which results offers valuable insight:

$$V_{in} - V_{cm} = \sqrt{\frac{I_1}{\beta_1}} - \sqrt{\frac{I_b}{\beta_3\left[1+\lambda_3\left(\sqrt{\frac{I_1+I_B}{\beta_2}}+V_T\right)\right]}} \tag{7}$$

The first term in (7) is the standard result for the current produced by a certain gate overdrive. The second term is the reduction in drain current because the voltage source isn't a perfect hard node. One thing to note is that the current output does not depend linearly on the threshold voltage as it usually does. The threshold voltage dependence from *M1* and *M3* cancel each other. This means that the current output is dependent on the differential voltage: $(V_{in}-V_{cm})$. From (7) it is seen that the degeneration term has a double square root dependence on current. This term can be neglected for drain currents of reasonable value. The input pair drain current then reduces to (8). The contribution of the voltage source can be considered a DC offset that only affects the quiescent bias.

$$I_1 \approx \beta_1 \left(V_{in} - V_{cm} + \sqrt{\frac{I_B}{\beta_3}} \right)^2 \tag{8}$$

At this point several circuit issues remain. From (7) it can be seen that the characteristics of transistor *M2* have little to do with the circuit operation. The feedback loop will apply as much gate drive as necessary to sink the disturbance current. The gate drive voltage will eventually force the transistors which are sourcing *Ib* into the triode region. This means that there is a minimum size for *M2* which is determined by the maximum amount of current that must be tolerated. *M2* should not to be larger than necessary since it forms the primary pole in the voltage source feedback system system. The high impedance node at the drain of *M3* must charge the capacitance C_{GS2}. For low power operation, it is desirable to make *Ib* small, which lowers the pole frequency. The voltage source feedback loop needs to have at least as much bandwidth as the amplifier. If the voltage source feedback loop cannot respond as quickly as the amplifier then the tail node will have a transient high impedance condition. Slow loop responses result in amplifier slewing and transient input pair transconductance degeneration.

The V_T tracking voltage source also has input common mode sensitivity. The tail voltage created from *M3* defines V_{DS2}. V_{DS2} must be greater than the saturation voltage of *M2*. If this voltage is too high, the gate drive on *M2* will need to be very low in order to maintain a small quiescent current. This limits the upper range because *M3* will enter the triode region, which effectively kills the loop gain, driving the *M2* gate.

This circuit must be used in a system with good input common mode voltage control since this defines the quiescent current consumption. The common mode

feedback (CMFB) circuit faces several challenges. The CMFB circuit must not present any resistive loading to the amplifier. The output stage of this amplifier is a high impedance node and any resistive common mode detector will reduce the DC gain. Capacitive loading is acceptable since this amplifier must drive several picofarads of capacitance. This would suggest the use of the traditional differential pair common mode detector circuit. Unfortunately this circuit only works well for small output swing amplifiers due to the saturation effect.

The input differential voltage can be extended in two ways. The first is to increase the tail current, which is unacceptable in low-power designs. The second is to degenerate the input pair with either long channels or tail resistors. In order to obtain the ±1 V minimum swings the frequency response of this pair is unacceptable. The CMFB circuit must have similar response time to the amplifier itself so that the common mode will get reset every conversion cycle. Another CMFB circuit is the triode mode mirror degeneration circuit. This circuit operates as variable resistive degeneration in a current mirror. One device is driven with V_{CM} and a drain coupled pair is driven with V_{O+} and V_{O-}. The output voltage is averaged in the drain coupled pair and any difference between the average and V_{CM} results in a change in resistance. This variable degeneration adjusts the mirror ratio until the feedback loop stabilizes. The advantage of this circuit is there are no differential voltage limitations. The loop operates until the output voltage approaches the threshold voltage of one device. In order to be effective at low currents, small devices are needed to support a wide swing in degeneration. This causes a problem with the dynamically biased amplifier. When the devices are tuned for good quiescent performance they present a substantial amount of resistance. During the large step

transition, the current in the mirror increases a couple orders of magnitude. The voltage drop across the degeneration saturates the amplifier. The CMFB circuit which is ultimately used is presented in [14]. This circuit is presented in Figure 17. This circuit has been modified slightly from the authors' work.

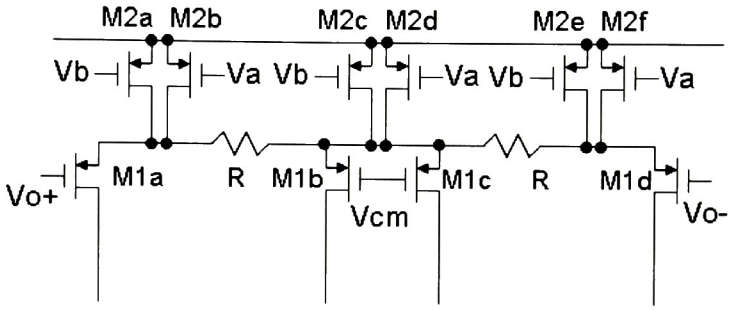

Figure 17 – CMFB Schematic

This circuit operates as a current steering network. Transistors *M2* are matched current sources. *M1* acts like a source follower; if all of the applied voltages are equal to V_{CM} then no current flows through resistor *R*. Under differential signaling, assume that V_{O+} increases and V_{O-} decreases. The source of *M1a* will increase attempting to follow the gate. Likewise the source voltage of *M1d* will decrease. Because of the voltage difference, some current will flow through R from the source of *M1a* to the sources of *M1b* and *M1c*. However an equal amount of current flows from this node into the source of *M1d* since the voltage difference is the same. Thus the current out of the drains of *M1b* and *M1c* is the same as it was quiescently. However if a common mode signal is present these currents through resistance *R* do not cancel. This results in an increased or decreased drain current in *M1b* and *M1c* depending on the direction of the common mode

signal. This feedback current signal is applied to the drains of transistors *M6* in Figure 15. The drain currents of *M1a* and *M1d* from the CMFB circuit are fed into dummy diodes matched to *M6* to minimize the V_{DS} mismatch between these devices.

Another modification that was made to the CMFB circuit is in the current sources. The common mode control that this circuit exhibits is defined by the relative common mode current applied by Figure 17 vs. the input signal current applied by the pseudo-differential input pair. If the CMFB circuit were statically biased, it would lose control of the common mode during large signal events. In order to correct this, drain coupled pairs *M2* are used as current sources. These transistors derive their gate drive from the *M2* diodes in Figure 15. The pair of transistors is needed so that the CMFB circuit increases its bias during steps in both directions since one of the diodes always shuts down under large signal events.

The improved version of the pseudo-differential amplifier still suffers from other problems. This amplifier has a very limited input common mode range. The input pairs must be driven in a symmetric manner around V_{CM}. This amplifier relies on the CMFB circuit for common mode rejection. High bandwidth CMFB circuits consume large amounts of current. An Input Balancing Stage is proposed which allows the V_T tracking PD to operate with improved input range and without the requirement for high bandwidth CMFB. Figure 18 is a functional representation of this stage. It produces an internal V_{CM} and drives the input signals symmetrically about this voltage to allow the V_T tracking PD to operate properly.

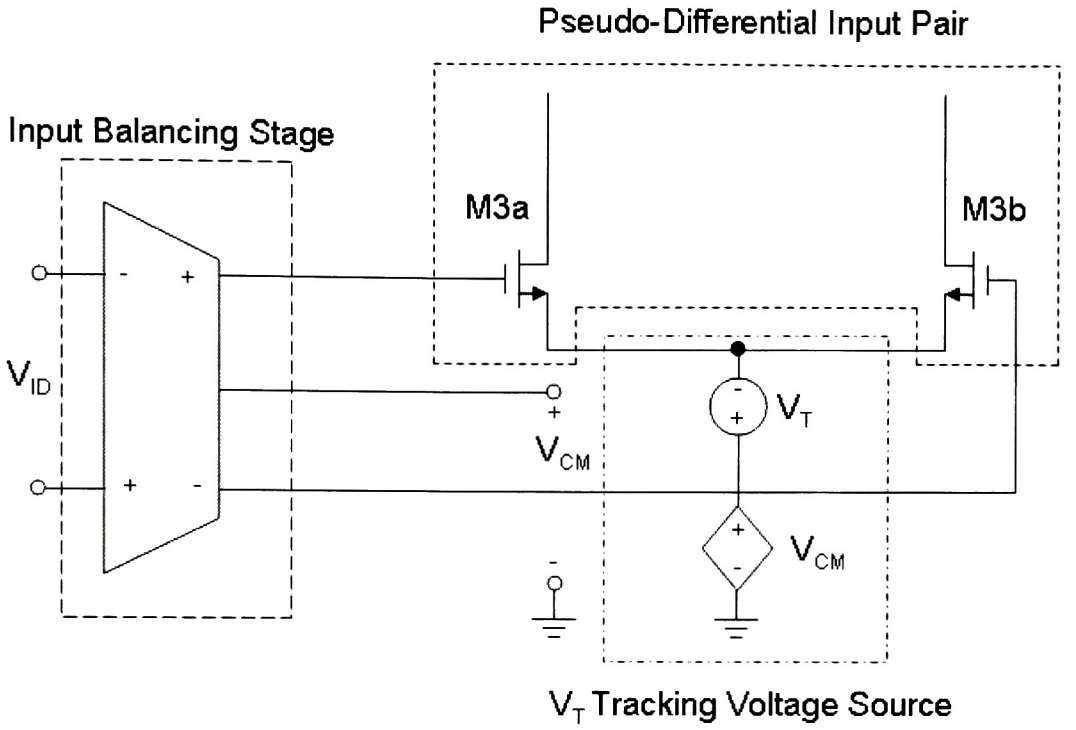

Figure 18 – Functional Diagram of the Proposed Amplifier

The circuit developed to do this is shown in Figure 19. The balancing action is achieved by using a traditional differential pair. This pair steers the current across matched loads *M2c* and *M2d*. In order to define the common mode, an identical differential pair is driven with the input signal. The drain currents are summed together and dropped across a diode load. This diode provides the gate drive signal for the other load pair. If all of the *M2* transistors are matched then their drain voltages will be the same. The dummy differential pair is used to provide V_{DS} matching for all of the loads. Since the voltage across *M2a* and *M2b* doesn't change with the input signal it is used to define the common mode for the V_T tracking voltage source.

The input signal range of this circuit is greatly extended beyond that of the pseudo-differential amplifier. V_{CM} is defined by a diode connected device with a fixed drive current. As long as the diode devices are sized such that their overdrive is greater than V_{DSSAT} of the V_T tracking voltage source *M3* then the circuit will operate in its active region. The balancer itself has a large input common mode range. It can be driven within a V_T plus V_{DSSAT} of the current source to the upper rail. The negative common mode limit occurs when the diode devices lose their overdrive. This actual value depends on the overdrive needed to operate the V_T tracking voltage source.

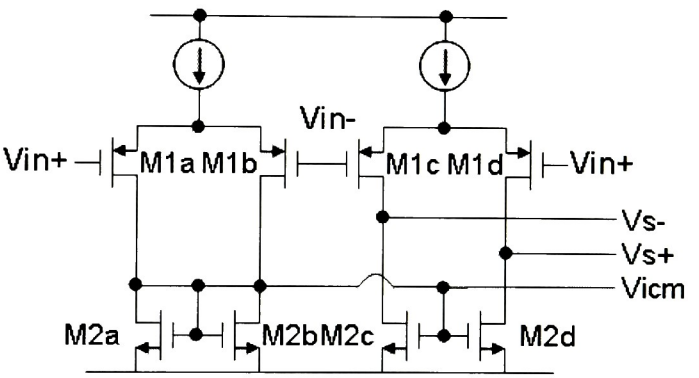

Figure 19 – Input Balancing Stage

The final amplifier architecture is shown in Figure 20 without the CMFB circuit. There are three primary poles in the signal path. The dominant pole is formed with the load capacitance and the channel impedance of the output pairs. The secondary pole is formed due to the balancer. The primary parasitic capacitances at this node are C_{GS3}

from the pseudo-differential inputs and C_{GD1} from the input pair. The limiting channel impedance is that of the *M2* load pair in parallel with the input pair. The tertiary pole is formed at the PMOS mirror node at the gates of *M6, M7,* and *M9*. C_{GS} of these transistors plus the Miller effect enhanced C_{GD7} dominate the node capacitance. The low $1/gm_6$ impedance at this node is what makes it the tertiary pole. The open circuit time constant approximations of the three poles are as follows:

$$
\begin{aligned}
p1 &= \frac{1}{C_{LOAD}(r_{on7} \| r_{on8})} \\
p2 &= \frac{1}{(C_{GS3} + C_{GD2})(r_{on2} \| r_{on1})} \\
p3 &= \frac{g_{m6}}{(C_{GS6} + C_{GS9} + [1 + A_{V7} \cdot C_{GS7}] + C_{GD3})}
\end{aligned}
\tag{8}
$$

The design tradeoffs for this amplifier will now be considered. The gain-bandwidth of this architecture is basically fixed by the load capacitance. This can be adjusted slightly by changing the gain in the input balancing stage but this cannot be done efficiently. This amplifier generally suffers from too much bandwidth rather than too little. Additional load capacitance must be added in order to compensate the amplifier.

The output stage is the primary voltage gain generating stage so the quiescent currents here are kept as low as possible in order to maintain high channel impedance. These bias currents will have minimum levels set by other amplifier requirements.

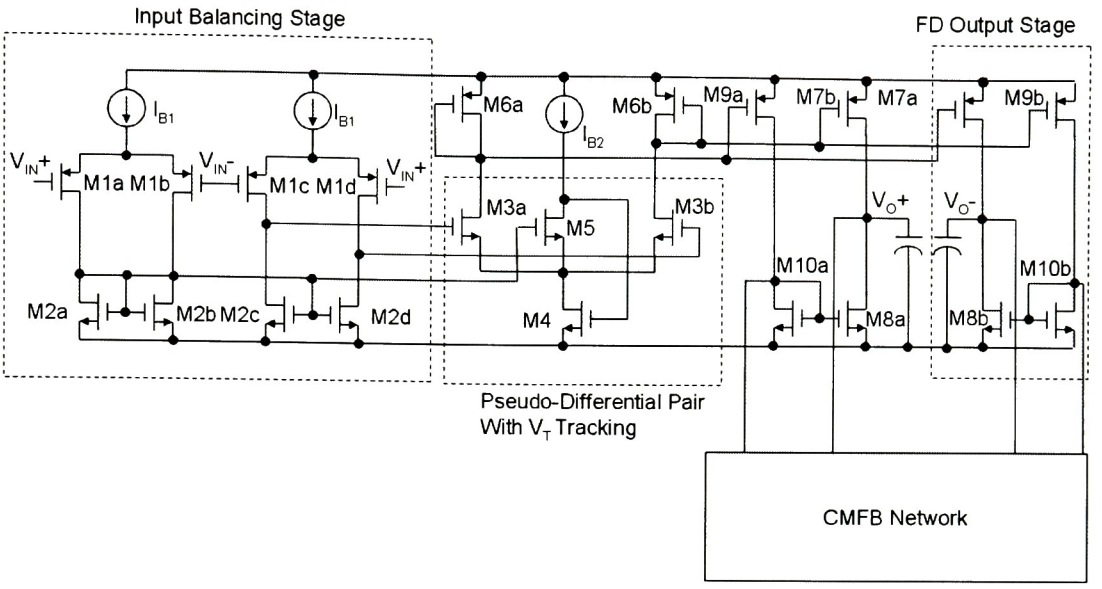

Figure 20 – Dynamically Biased Amplifier

The locations of the secondary and tertiary poles are set by the settling requirements of the system. If it is assumed that 50 degrees of phase margin is acceptable then each of these poles must be at least 2.2 times the unity gain crossover frequency.

The secondary pole node has several effects which dictate the use of bias current. The parasitic capacitances at this node cannot be reduced since they are created by the input transistors. *M1* is the primary offset generating mechanism of the amplifier. The matching requirements of the system dictate the minimum size of the input transistors. Likewise the matching of *M3* and *M5* control the quiescent bias point of the amplifier. Assuming that the capacitance is fixed at this node by matching issues the only way move this pole is to increase the bias current so the channel impedance of *M1* and *M2* is reduced. The required location of this pole sets a limit to how much gain is produced in

this stage. The pole location sets the channel impedance, and any attempts to recover the gain by increasing g_{m1} with either an aspect ratio increase or a bias current increase will be counteracted by linear increases in parasitic capacitances and linear reductions in channel impedance, respectively.

The location of the tertiary pole dictates how much quiescent current must be used in the pseudo-differential stage. The transconductance of *M6* cannot be adjusted effectively by changing the W/L ratio. The transconductance increases with the root of this ratio while the parasitic capacitance increases linearly. This, coupled with the need to maintain the mirror ratio with *M7* and *M9* results in an inverse relationship between W/L and the tertiary pole location. *M6* has a minimum acceptable g_m due to the dynamic bias conditions. The voltage drop across *M6* will force *M3* into the triode region at its limits. The following condition must be satisfied:

$$V_{DD} - V_{GSMAX6} \leq V_{DSSAT3} + V_{DSSAT4} \tag{9}$$

Where V_{GSMAX6} is given by:

$$V_{GSMAX6} = V_{TP} + \sqrt{\frac{I_{MAX}}{k'\frac{W}{L}_6}} \tag{10}$$

I_{MAX} is the dynamic current required to ensure that the amplifier never slews against the load capacitance. This value must be greater than the maximum slope of the linear settling behavior. If single pole settling is assumed then the maximum slope occurs at $t=0$ with:

$$\frac{\partial V}{\partial t}\Big|_{t=0} = \frac{A}{\tau}e^{\frac{-t}{\tau}} = \frac{A}{\tau}V/\mu s \tag{11}$$

The output dynamic range, A, amplifier time constant, τ and load capacitance can then be used to determine I_{MAX}. This value is then used to determine the aspect ratio of *M6*. Finally the quiescent current of the amplifier is used to place *p3* at a frequency with acceptable phase shift.

At this point the transistor sizes for *M7* and *M9* can also be set. The transistors could be used in mirror ratios other than 1:1 but without many benefits. The idea would be that the internal amplifier could run at smaller bias currents since the pseudo-differential input current would have an additional K:1 current boost. However, an increased mirror ratio cannot reduce the required current through *M6* due to the tertiary pole. The extra parasitic capacitance due to the larger output transistors does not allow for a reduction in current. In fact, more current is needed to compensate for the g_{m6} reduction due to the W/L_6 size reduction.

There is an interesting power optimization problem that has come from this amplifier. Large load capacitances are needed to compensate an amplifier unwilling to spend quiescent current. However these load capacitances will need to be charged dynamically. The amount of current that is needed is dependant on the switching frequency of the switched-capacitor circuit and the amount of voltage swing that the load capacitor travels. For systems that switch full scale loads often, spending quiescent current to build in more bandwidth than needed could result in a power reduction since the size of the load capacitor can be reduced.

5.2 *Output Amplifier*

The same amplifier architecture is used for the output amplifier. One of the output mirrors is removed along with the CMFB circuit. Figure 21 shows the output amplifier schematic. Transistor sizes are the key difference between these two amplifiers. This amplifier must also drive a 5 V full scale so thick oxide (150 Å) devices are used. From the requirements defined in Section 2.5 this amplifier must settle 100 pF in approximately 800 ns. This requirement suggests that the amplifier needs to be capable of delivering several milliamps of peak current to the load. The less stringent settling time requirements imply that several MHz of bandwidth will be sufficient.

Another important design consideration for this amplifier is the variability of the load. The driven capacitance can vary from about 10 pF to 100 pF. Since this architecture is load compensated it must be designed to be stable under all of the load conditions.

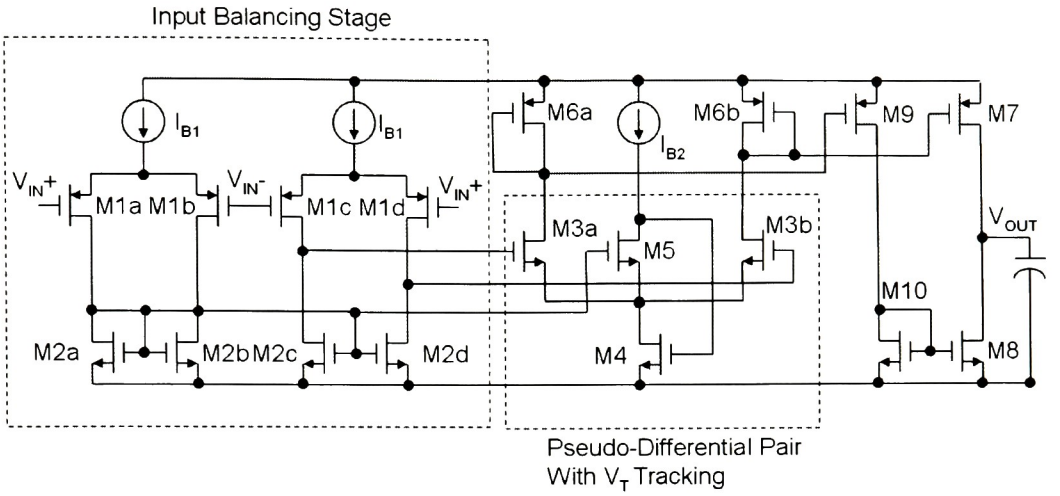

Figure 21 – Output Amplifier

Lastly the problem of amplifier offset voltage will be considered. The assumption will be made that the offset voltage is dominated by the V_T mismatch between the input pairs. The V_T variance per unit area for PMOS devices is:

$$\sigma_{0VT}^2 = \sqrt{1200mV \cdot t_{OX}} \tag{12}$$

So the V_T variance for the input pair is:

$$\sigma_{VT}^2 = \frac{\sigma_{0VT}^2}{WL} \tag{13}$$

Assuming that the maximum input pair aspect ratio is 40 µm / 1 µm due to bias current limitations, the standard deviation is:

$$\sigma_{VT} = \sqrt{\frac{(1200mV \cdot 150\dot{A})^2}{40\mu m \cdot 1\mu m}} = 2.84mV \tag{14}$$

This offset standard deviation is made even worse since this amplifier is used in a gain of 2 ½. This will cause a significant yield reduction since the 9-bit precision level is around 9.7 mV on a 5 V full scale.

For this test chip an offset trim mechanism is built into the output amplifier. The *M2* loads are degenerated with triode devices. The gate drives on *M2a, b, c* are sourced from the 3.3 V rail. The gate drive on *M2d* is brought out to a pad. Adjusting this voltage relative to 3.3 V allows the offset to be trimmed for testing.

5.3 Amplifier Simulation Results

The test benches that are used to simulate these amplifiers are constructed to mimic the operation of the actual application. These test benches are switched-capacitor gain circuits that use the same transmission gates as the rest of the circuit but only have a single bank of capacitors to reduce simulation times. The bank capacitance is adjusted to represent that actual net capacitance that is seen.

The frequency response of the fully differential (FD) amplifier used in the D/A core and sample and hold is shown in Figure 22.

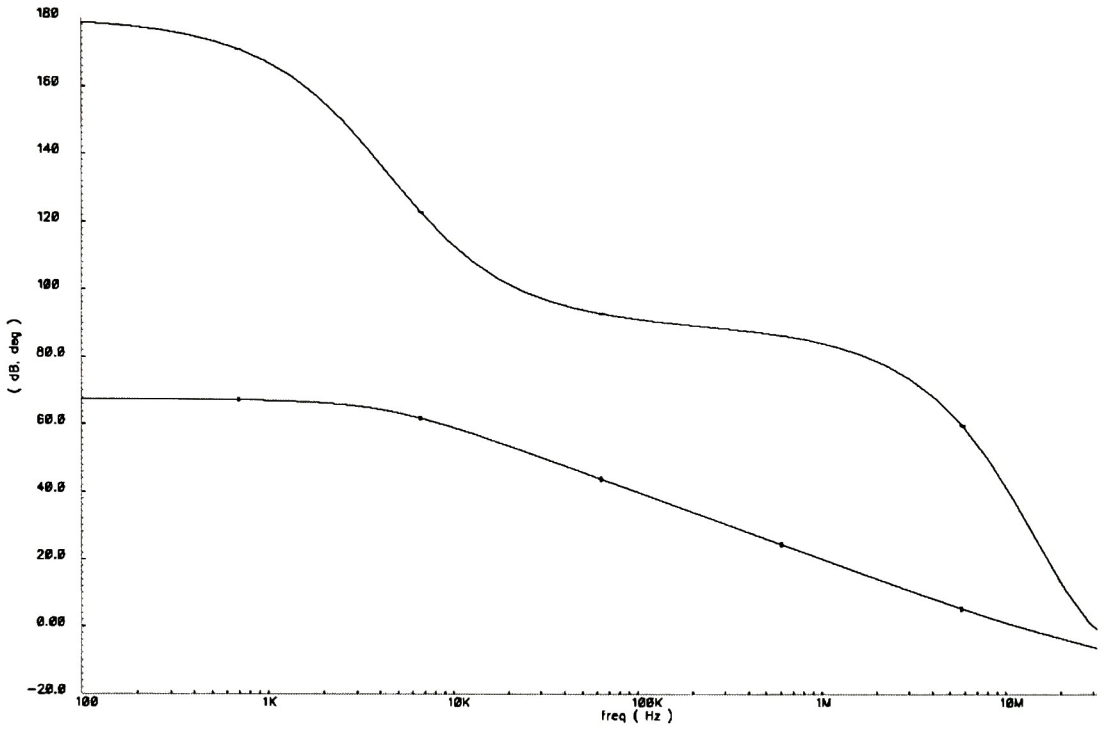

Figure 22 – FD Amplifier Frequency Response

This amplifier has a DC gain of 67 dB. The Unity Gain Frequency (UGF) is 12 MHz and the Phase Margin (PM) is 34 °. The dominant pole *p1* is at 4.5 kHz. The output mirror pole *p3* is at 35 MHz. The input balancing stage pole *p2* is at 34 MHz. The input balancing stage produces 23 dB of gain.

The input balancing stage has bias source I_{B1} set at 6.5 μA. The quiescent drain currents of the pseudo-differential pair are 2.3 μA. All of the signal mirrors in this amplifier are 1:1. V_{DS} mismatch causes the output stage bias current to be 2.8 μA. V_T tracking bias source I_{B1} is set at 1 μA. All of these bias sources are derived from a single 1 μA current source. The total amplifier bias including the CMFB circuit is 40 μA.

The frequency response of the Single Ended (SE) amplifier used in the output driver is shown in Figure 23.

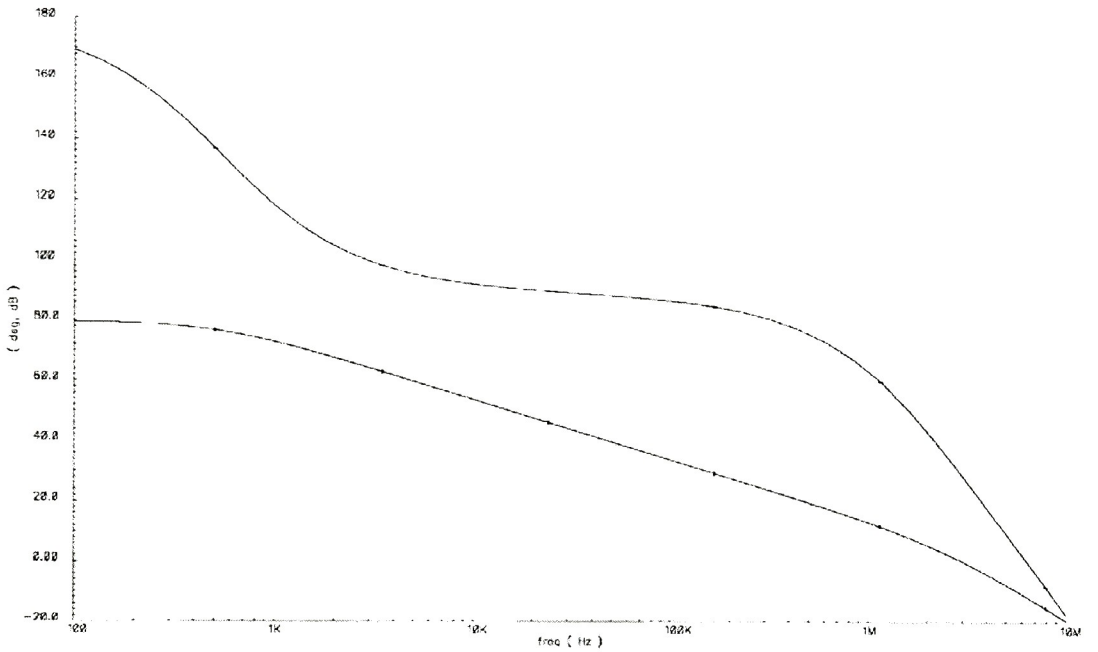

Figure 23 – SE Amplifier Frequency Response

This amplifier has a DC gain of 79 dB. The UGF is 3 MHz and the Phase Margin (PM) is 29 °. The dominant pole *p1* is at 540 Hz. The output mirror pole *p3* is at 26 MHz. The input balancing stage pole *p2* is at 2.5 MHz. The input balancing stage produces 27 dB of gain.

The input balancing stage has bias source I_{B1} set at 6.8 μA. The quiescent drain currents of the pseudo-differential pair are 3.5 μA. All of the signal mirrors in this

amplifier are 1:1. V_{DS} mismatch causes the output stage bias current to be 3 μA. V_T tracking bias source I_{B1} is set at 1 μA. The total amplifier bias is 28 μA.

The step response of both amplifiers is shown in Figure 24. This response corresponds to full scale steps for each of the amplifiers. The FD core amplifier operates on a 2 V differential full scale centered about a common mode voltage of 1.3 V. The initial negative differential voltage in Figure 23 is an artifact of the initial conditions in the simulation. The load for this amplifier is 14 pF. Half of this capacitance is in the feedback loop and the other half is dummy PMOS compensation capacitors. This amplifier settles to 1% in 81 ns and 0.1% in 113 ns.

The term slew rate is still used here to represent the maximum voltage slope of these amplifiers. Unlike the traditional slew rate this current limitation is not fixed and is based upon the input signal rather than the amplifier bias current. The slew rate of the FD amplifier is 42 V/μs.

The SE output amplifier must have linear drive from 200 mV to 4.8 V for this application. The SE amplifier can drive within 12 mV of each rail although not to 9-bit precision. This amplifier drives 800 fF in the feedback loop and 100 pF load referenced to ground. The load is somewhat isolated from the amplifier by a 200 Ω ESD diffusion resistor. The SE amplifier settles to 1% in 374 ns and 0.1% 517 ns.

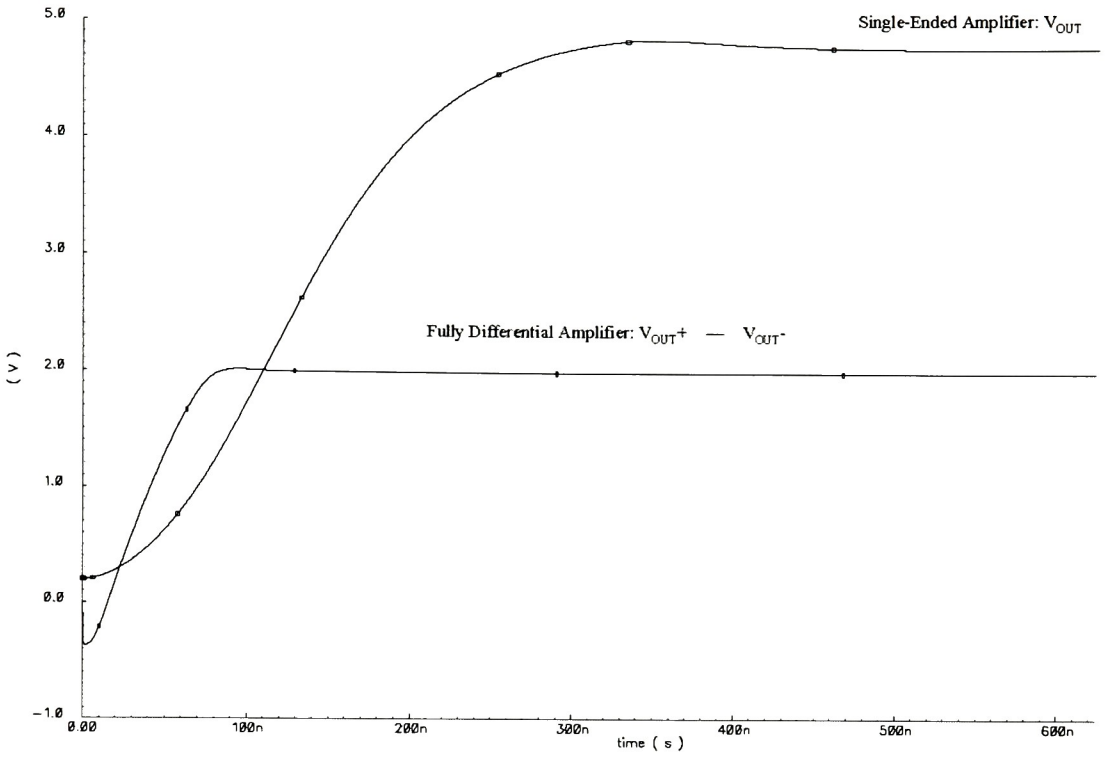

Figure 24 – FD & SE Amplifier Step Response

An important measurement in a dynamically biased amplifier is efficiency of bias current use. Since dynamic biasing increases the bias current of the entire amplifier, a significant portion of current is wasted on internal paths. A performance measurement called Dynamic Current Efficiency (DCE) is used here. DCE is the ratio of the current supplied to the load to the total amplifier bias. Figure 25 shows the supply current and load current of the SE amplifier. The peak DCE for this amplifier is 56%. The DCE stays above 25% until after the amplifier has settled beyond 1%. The DCE for the FD amplifier only achieve a peak of 8%. The FD amplifier has extra mirror paths which waste bias current. The CMFB circuit is also has three dynamically biased paths which ruin its efficiency.

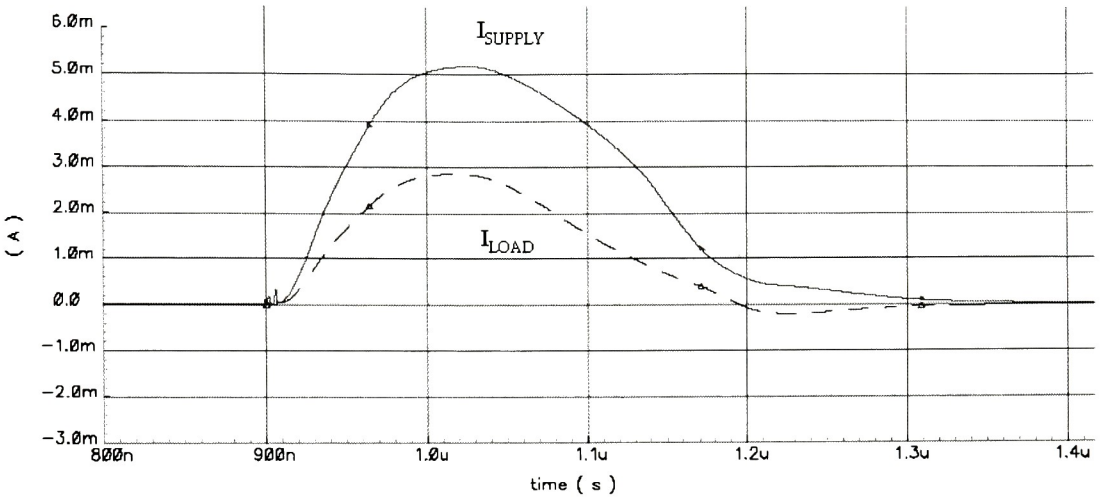

Figure 25 – Dynamic Current Efficiency

Figure 26 presents a performance table comparing the dynamically biased buffers developed here and other current work with similar power consumption and loading. The slew rates and settling times for this work are a notable improvement.

PARAMETER	Peng [15]	Giustolisi [16]	Lu [17]	Yao[18]	This work (FD)	This work (SE)
Slew Rate	1.03 V/µs	0.54 V/µs	5 V/µs	0.2 V/µs	42 V/µs	27 V/µs
Load	150 pF	20 pF & 1 kΩ	600 pF	18 pF	14 pF	100 pF
Settling time	2.8 µs (1%)	1.6 µs (1%)	2.7 µs	2.7 µs (rise)	81 ns (1%)	374 ns (1%)
GBW	2.85 MHz	2.2 MHz		1.2 MHz	12 MHz	3 MHz
Phase Margin		56 °		60 °	34 °	29 °
DC Gain	>100 dB			52 dB	67 dB	79 dB
Quiescent Power	45 uW	150 µW	23.1 µW	8 µW	132 µW	140 µW
Supply Voltage	1.5V	1.2 V	3.3 V	0.8 V	3.3 V	5 V
Die Area	0.02 mm^2	0.02 mm^2	0.026 mm^2	0.006 mm^2	0.003 mm^2	0.0063 mm^2
Technology	0.35 µm	1.2 µm	0.35 µm	0.25 µm	0.35 µm	0.35 µm

Figure 26 – Performance Comparison Table

5.4 System Design

The system level design is has been mostly defined by the architecture specification in Section 4. The remaining considerations are primarily switch optimizations. Transmission gates are used for switching in order to reduce the charge injection errors. The channel size of the NMOS and PMOS transistors is the same so that their injected charges cancel to first order. The minimum unit cell is defined in Equation 5. This means that the switches must be sized such that their channel resistance does not form a dominant time constant in the system. Once this is achieved, each switch is kept as small as possible to keep the charge injection and junction leakage errors small.

Figure 27 shows the sample and hold circuit output differential voltage operating on a full scale transition. The precharge clock has been placed on the plot for a frame of reference. During the precharge phase the input storage capacitors are charged to the value of V_{REF}. This value is driven by the D/A core. The sample and hold precharge clock coincides with the evaluate clock during the V_{REF} word conversion. The bottom inset in Figure 26 shows the output voltage during the precharge phase. This simulation includes a 30 mV input offset voltage. This offset is sampled during this clock phase. During the evaluate clock, the output voltage is driven to its final value, which is 2 V in this case. During the settling of the sample and hold a glitch can be seen. This glitch corresponds to the D/A core precharge phase in which the programmable capacitor array is charged to the soft V_{REF}. As this array is charged the output voltage settles back to its final value. The upper inset in Figure 27 shows the final settling behavior. The sample and hold settles to 2.7 mV (0.13%) of its final value which is within 9-bit precision. This shows the effectiveness of the auto-zeroing circuit which canceled out 30 mV of offset.

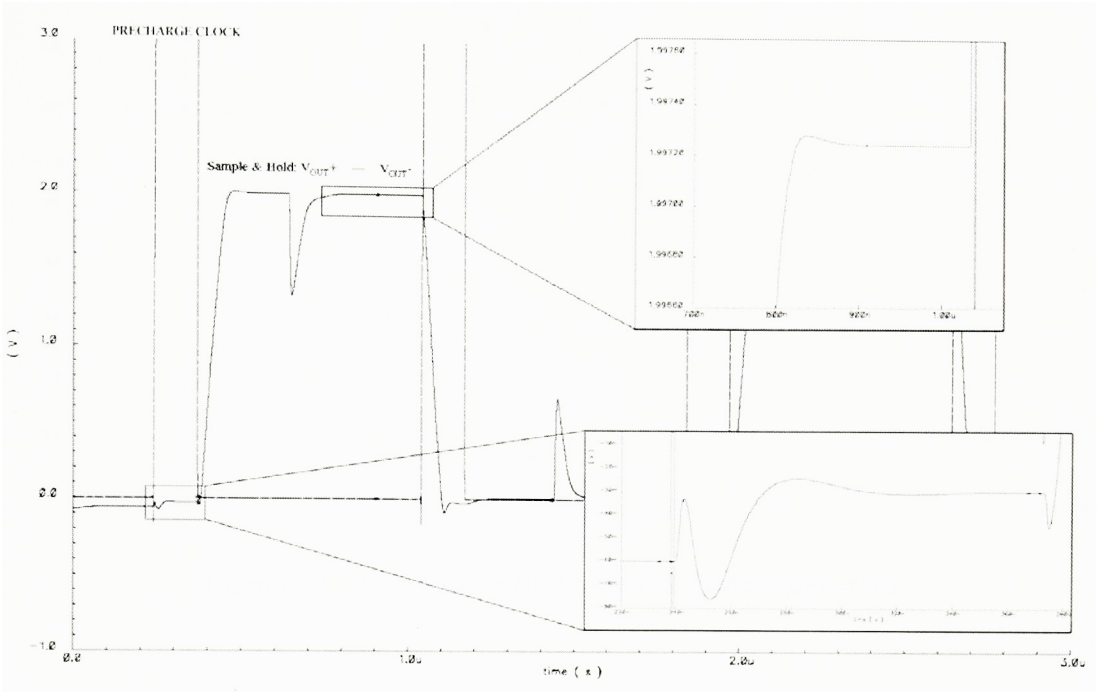

Figure 27 – Full Scale Sample & Hold Operation

Figure 28 shows the interaction with the D/A core and the sample and hold circuit. The top graph is the differential output voltage of the D/A core overlaid with its precharge clock. The D/A converts incoming words in the order: V_{REF}, V_{OS}, V_{WORD}. The bottom graph shows the sample and hold output differential voltage overlaid with its precharge clock. This figure shows the interaction between the D/A core clocking and the sample and hold clocking.

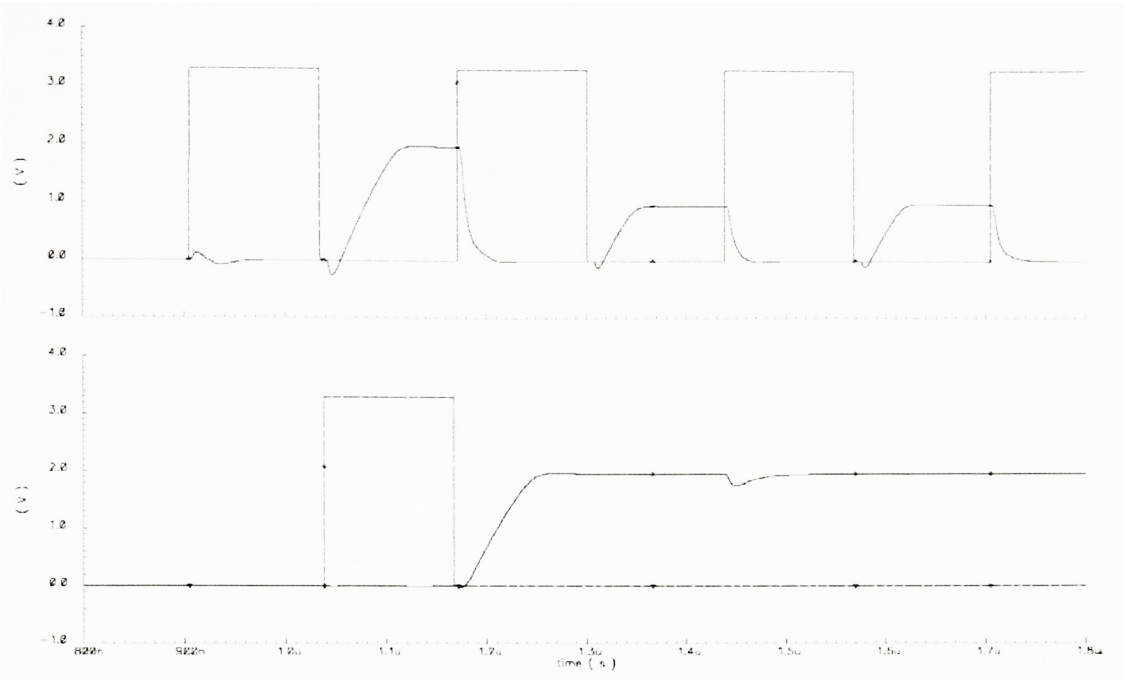

Figure 28 – D/A Core with Sample & Hold Operation

Figure 29 shows the interaction of the D/A core with the output driver hold circuits. The top graph is the differential output voltage of the D/A core overlaid with its precharge clock. The middle graph is the V_{OS} hold differential voltage overlaid with the output amplifier load offset clock. The bottom graph is the V_{WORD} hold differential voltage overlaid with the output amplifier load word clock. These two clocks are two of the four possible Φ_1 masked clocks used in Figure 14. The two clocks used in Figure 29 are the even pixel clocks. This means that the output driver is settling the odd pixel that was previously converted during this time.

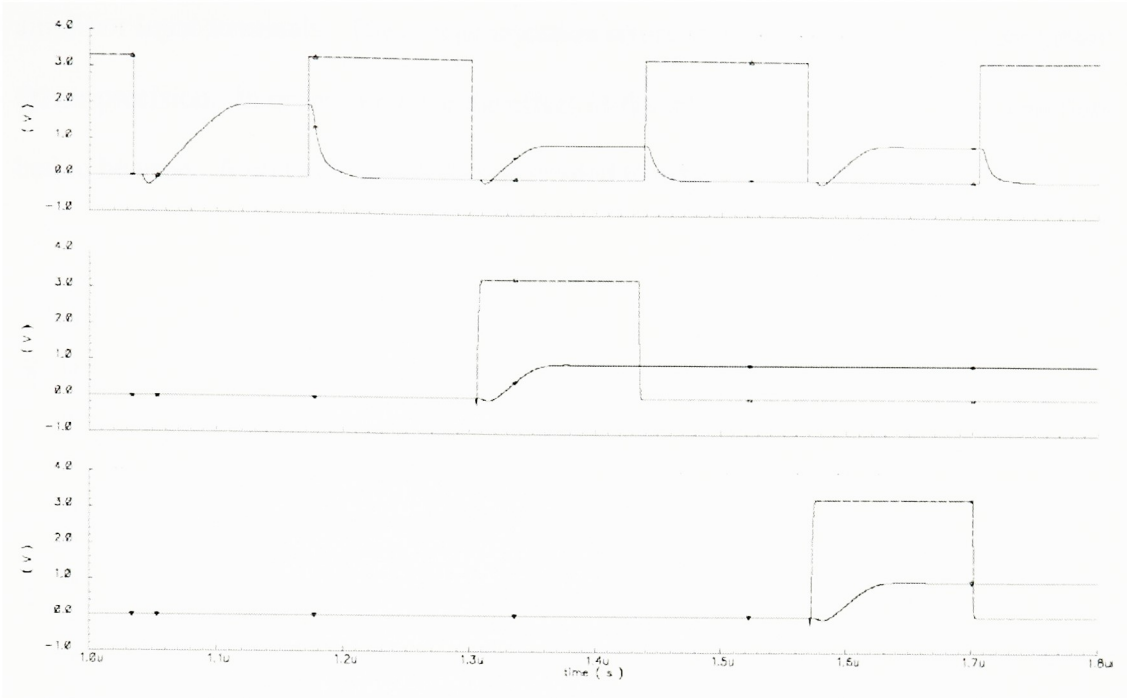

Figure 29 – D/A Core with Output Amplifier Hold Operation

Figure 30 shows the output driver clock interaction. The middle and bottom graphs are the same as in Figure 29. The top graph is the output amplifier output voltage. This voltage is measured on the external 100 pF load. The output driver precharge clock is overlaid on this signal. The V_{OS} and V_{WORD} hold capacitors maintain their current value until the Φ_2 drive clocks connect these capacitors to the input terminals of output amplifier. The charge stored on the hold capacitors is transferred to the feedback capacitor.

The output driver does not operate as effectively as the D/A core or sample and hold circuits. The fundamental difference is the switch locations. The output driver must precharge and hold values at the same time as it is driving previously held values. Transmission gates must be placed in the signal path between the capacitors and

amplifier input terminals. The charge injection errors at this location dominate the output driver precision. In order to reduce the effect of the injected charge, large unit capacitors had to be used. A unit cell of 400 fF was needed to achieve 9-bit precision.

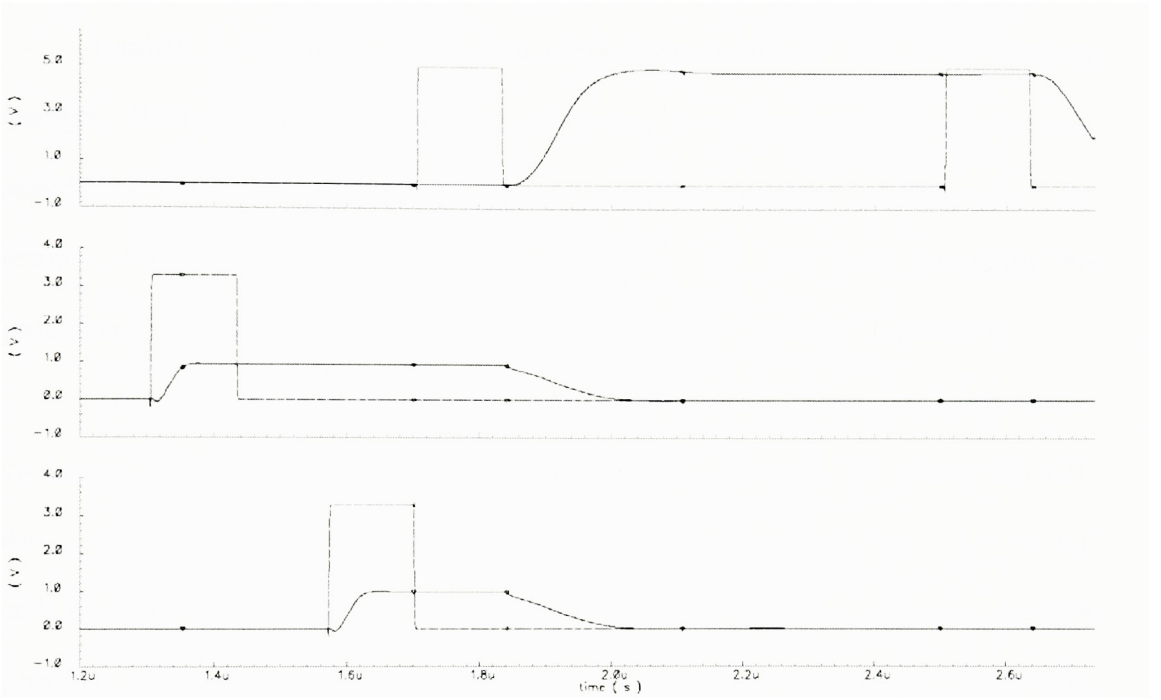

Figure 30 – Ouput Amplifer Hold and Drive Operation

5.4 Linear D/A Comparison

As the system was developed to implement the piecewise linear approximation converter the complexity of this method became apparent. This begs the question as to whether this technique is actually better than a standard 9-bit converter. This basic system architecture can implement 9-bit linear conversions with a few changes. Rather

than perform three 6-bit conversions to convert the linear segment coefficients, the system could be used to perform three 3-bit conversions. Each of these conversions is recycled using the core sample and hold forming a complete 9-bit conversion. This technique is a hybrid cyclic-charge redistribution D/A converter.

Limited by design time constraints the basic system is used with clocking changes. The first change is to unmask the precharge and evaluate clocks which drive the sample and hold. This circuit now captures every conversion of the D/A core. The second change is reference voltage switch. The D/A core must precharge the programmable capacitor array to the soft V_{REF} for the second and third conversions.

The D/A core must be converted into a 3-bit system. The core must always convert at least 1 3-bit LSB so that a new reference is generated. The three 6-bit LSB drivers are connected to logic high so that they will always provide 7 (1 + 2 + 4) unit cells of capacitance for the conversion. The eighth unit capacitor is a dummy capacitor present in the array to give an even 128 devices. If a complete redesign were performed on this system, only 16 capacitors would be needed per terminal rather than 128. This would reduce the layout size and complexity.

The output driver remains the same but the V_{OS} hold capacitors are disabled. This system component has the potential to offer good system improvements if a complete redesign were done. The output driver only needs to sample a single conversion value rather than two. This sample operation could be synchronized with the output driver precharge phase. This would eliminate the need for the even and odd pixel system. A single switched capacitor gain stage similar to the sample and hold circuit could then be used. This would reduce the number and size of the unit cell capacitors for the output

amplifier. Figure 31 shows the operation of the 9-bit hybrid converter. A LSB is converted each cycle to generate a single 9-bit LSB. The first conversion should be one eighth of the 2 V hard reference voltage. The actual converted value is 250.2 mV. The second conversion should be one eighth of the 250 mV soft voltage reference. The actual converted value is 31.26 mV. The last conversion should again be an eighth of the soft reference voltage. The actual converted value is 3.83 mV compared the ideal 9-bit LSB of 3.9 mV.

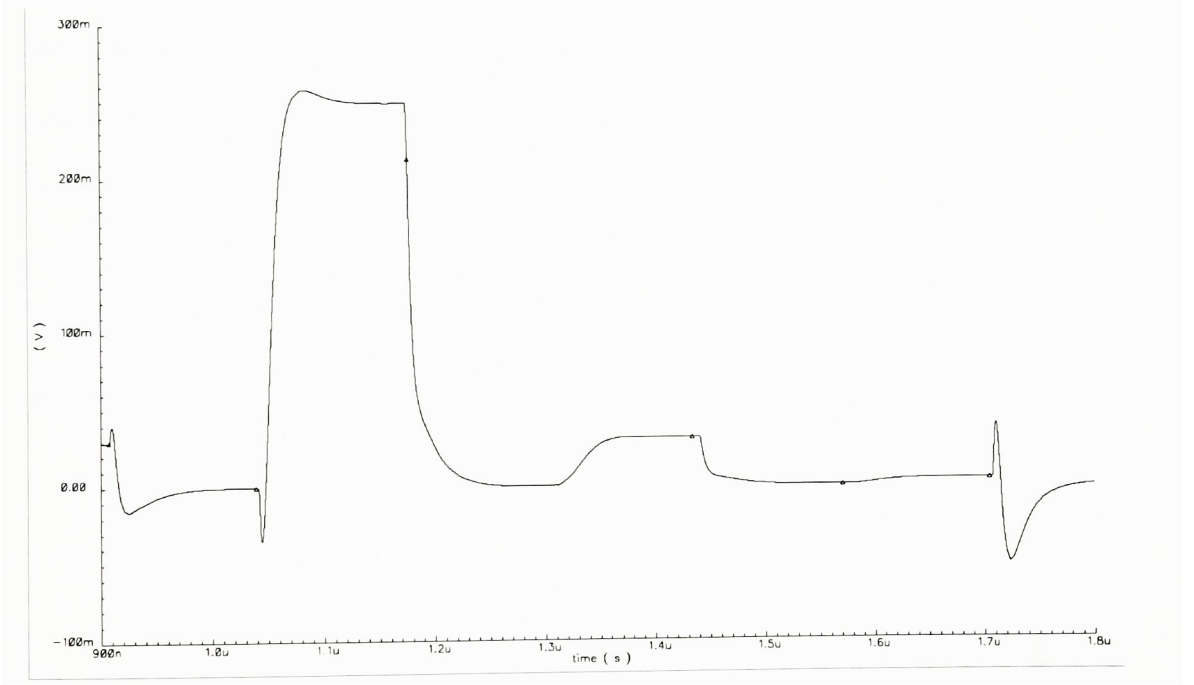

Figure 31 – 9-bit D/A Core Operation

5.5 Physical Implementation

A test chip has been implemented with all of the circuitry previously described. The test chip is to be fabricated on the TSMC035_P2 process through MOSIS.

A total of six D/A converters are on the test chip. They are divided into RGB channels with two converters per channel. A 2:1 multiplexer supplies input data to each of the converters from three designated 6-bit input buses. Three of the D/A converters are configured using the piecewise linear approximation method. The remaining three are configured as hybrid cyclic charge redistribution converters. A single clock generator creates the non overlapping clocks and the 16 necessary masked clocks for both types of D/A converter. Global buffers distribute these clocks to all of the converters.

In addition to the D/A converters two stand-alone dynamically biased amplifiers are placed on the die for additional testing. A single-ended and a fully differential buffer are fabricated. Figure 32 shows the layout of the test chip.

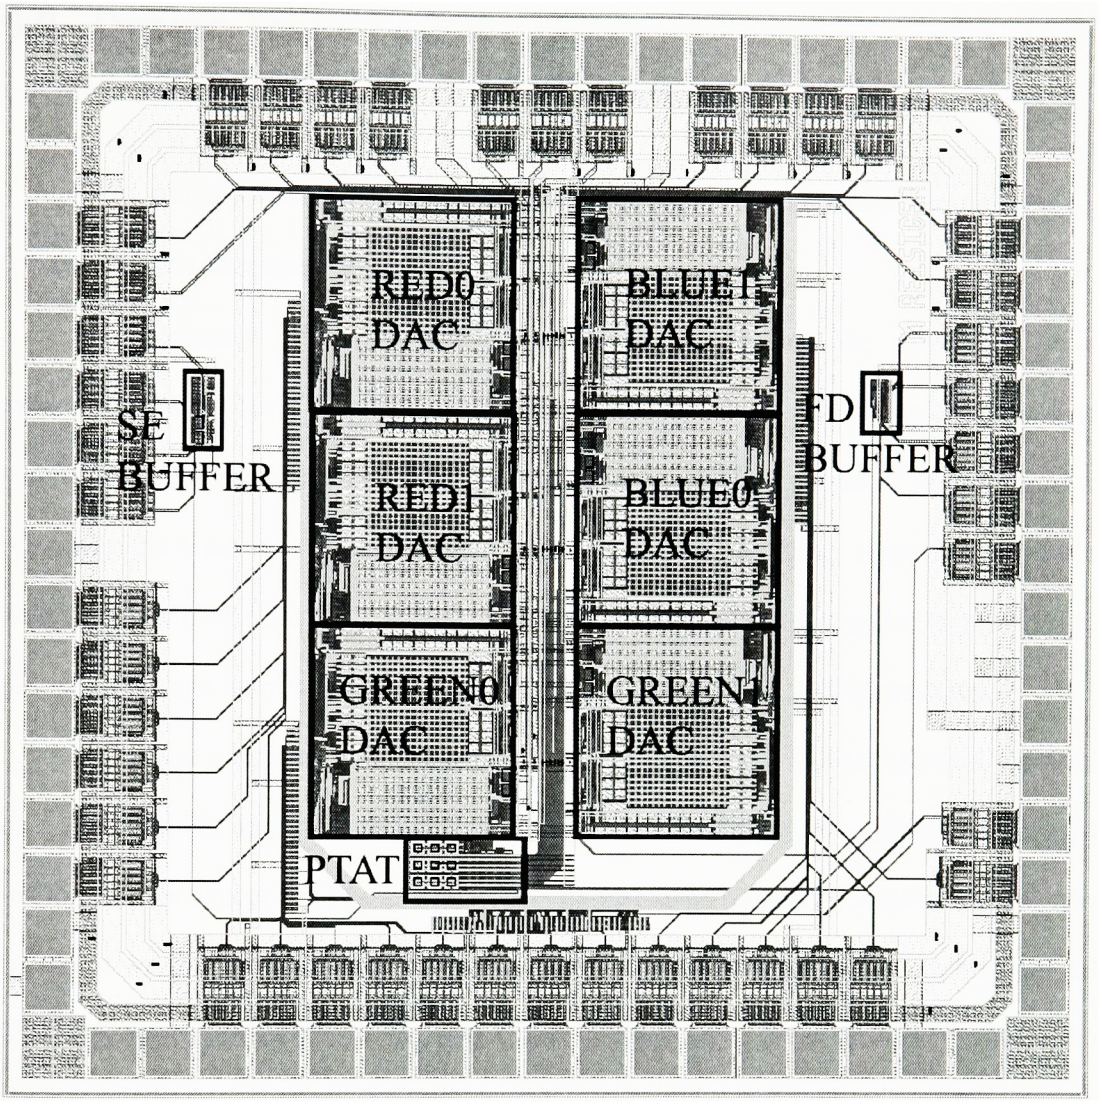

Figure 32 – Test Chip Layout

6.0 Design Test

The test silicon is not available in time to be included with this work. However a test plan will be discussed briefly.

The primary goal of this test chip is to demonstrate the piecewise linear approximation method. In order to do this several subsystems need to be tested. The first component that must be verified is the bias generator. This circuit is responsible for providing the 1 μA bias current to each of the amplifiers. A PTAT current source built with substrate PNP transistors is used. One of the bias currents is supplied to an external pin. This pin can be measured to verify the startup of the PTAT and its accuracy.

The next components to be verified are the amplifiers. Uncompensated versions of the SE and FD amplifiers are placed on the die. All of their connections are brought out to test pins except for V_{CM} and I_{REF}. The performance of both of these amplifiers can be tested without the constraints imposed by the rest of the system. Measurements of step response, frequency response, and offset voltage will be taken as these are the most important considerations for this system. The amplifier performance over different load conditions will be tested.

A test bench will be needed to verify the amplifier performance. Several voltage supplies are needed. 3.3 V and 5 V supplies are needed in addition to approximately 1.3 V for V_{CM}. The latter voltage should be adjustable to find the optimal common mode voltage. A feedback loop will be needed to test the transient performance. This loop will need to need to be either switched capacitors or a high impedance resistive feedback system.

Once the functionality of the amplifiers is confirmed, the D/A converters need to be tested. The test bench that is used for the amplifiers needs to be expanded in order to test the converters. Two more voltage supplies are needed to generate V_{REF+} and V_{REF-}. These supplies charge all six programmable capacitor arrays during the precharge clock phase. This means that they need good high frequency load regulation. Designated local bypass capacitors should be placed on these supplies. A data source is also needed. This source needs to be able to drive three 6-bit buses. Each data bus is multiplexed for two D/A converters. A single programmable data source such as a PIC microcontroller can be used to generate these signals. This microcontroller can also source the system clock and the reset signal to the test chip. Each of the output drivers also has an offset trim line which is centered about 3.3 V. Trim potentiometers referenced to the 5 V supply will drive these trim signals. The test system also needs to provide the load capacitance for each converter.

The conversion accuracy of each converter needs to be tested. For the approximation converters, this will involve converting input ramps of V_{WORD} at different soft V_{REF} values. During this test, V_{OS} will be kept at zero. Next, V_{OS} will be tested. The offset voltage will be ramped for a constant V_{WORD}. These tests enable the measurement of the linearity of the converter. The D/A core has an intrinsic converter linearity due to the programmable capacitor array. The linearity of the summing operation at the output amplifier is also important. Treating the V_{WORD} and V_{OS} ramps separately will allow the independent extraction of these two terms.

Lastly, an actual CGS display needs to be tested. In order to do this an actual video source is needed. The test bench will need a DVI receiver chip such as the Analog

Devices AD9887a. This chip provides a CMOS logic pixel output of incoming video signals. This pixel stream will need to be fed into the microcontroller. The microcontroller will perform the look up table actions necessary to generate the V_{REF}, V_{WORD} and V_{OS} words for the D/A converter. This test bench can be used to verify the image quality performance of the piecewise linear approximation algorithm.

7.0 Conclusion

This work has presented an algorithm for digital to analog conversion of nonlinear transfer functions. This algorithm was targeted for the transmissivity curve of a LCD panel. A switched-capacitor based system is developed to implement this algorithm. The same switched-capacitor system is also used to implement a 9-bit linear converter. These converters are used to compare the relative advantages of each system. The 9-bit linear converter offers a significant reduction in circuit complexity. This suggests that the proposed piecewise linear system may not prove worthwhile.

During this investigation, a novel dynamically biased amplifier architecture is developed. This architecture is the most valuable contribution of this work. This dynamically biased amplifier offers higher performance than other types of class AB amplifiers. The proposed amplifier can be used in more general applications than traditional pseudo-differential amplifiers because of the front-end balancing circuit. Work on this amplifier has been submitted to the 2006 ISSCC conference.

8.0 Future Work

There is a good deal of future work that can be done in this area. The system is completely untested in hardware. All of the test procedures discussed in Section 6 need to be performed.

The fact that this particular implementation of the piecewise linear algorithm is very complex has already been mentioned. The algorithm may still prove to be useful in another form. Different types of converters can be explored. Other means of coefficient generation should also be considered as opposed to using D/A converters to implement the references.

The dynamically biased buffer needs to be analyzed in more detail. One of the areas where analysis is needed is the input balancing stage gain. This gain dominates the amplifier biasing behavior in a feedback loop. A feedback loop tends to drive the input differential voltage to zero. When this voltage approaches zero, the amplifier loses much of its bias current. The output voltage may still be changing with a significant slope at this time. The lack of input voltage creates a current starvation condition. This current starvation manifests itself as another time constant in the settling behavior.

9.0 Acknowledgments

I would like to acknowledge Art Kalb and Satoru Shingai of Analog Devices for their valuable advice regarding the matching of capacitors and precision circuitry in general.

10.0 Appendices

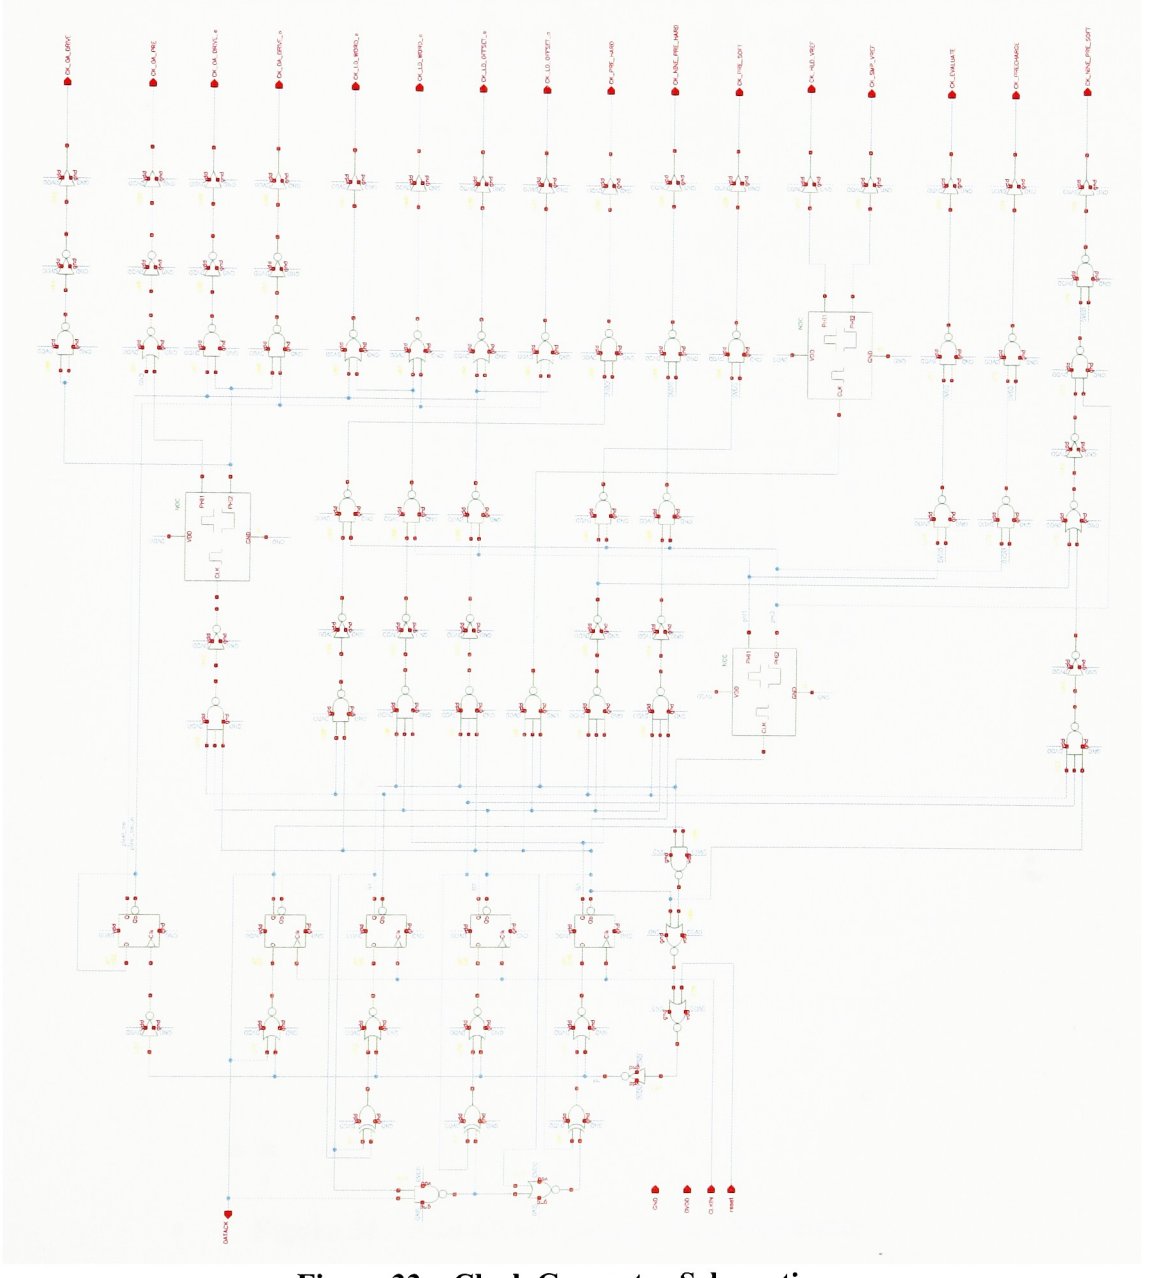

Figure 33 – Clock Generator Schematic

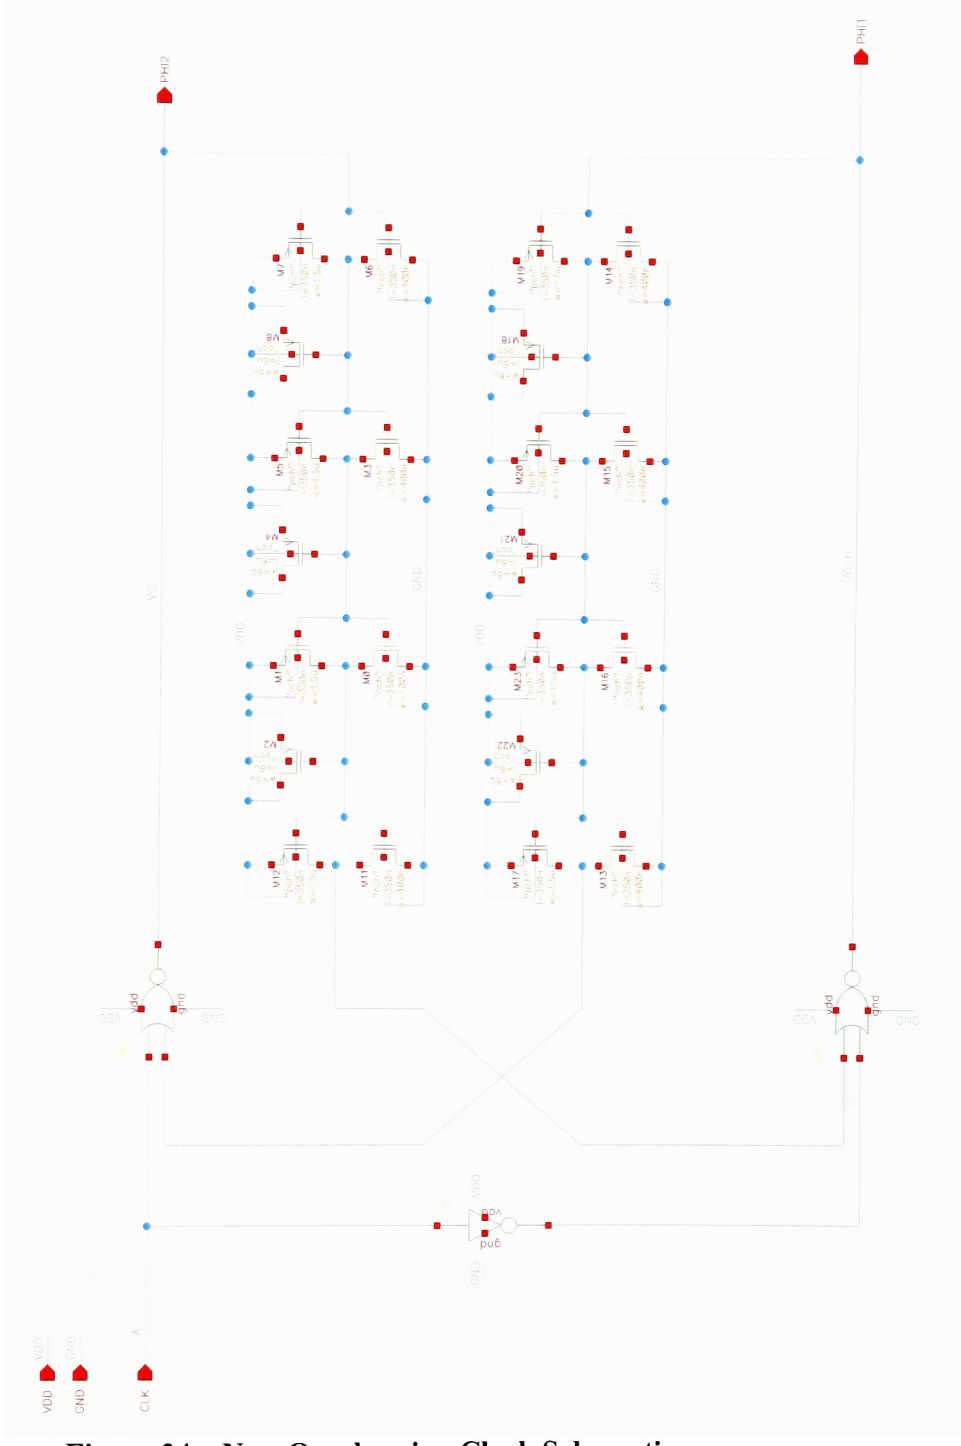

Figure 34 – Non-Overlapping Clock Schematic

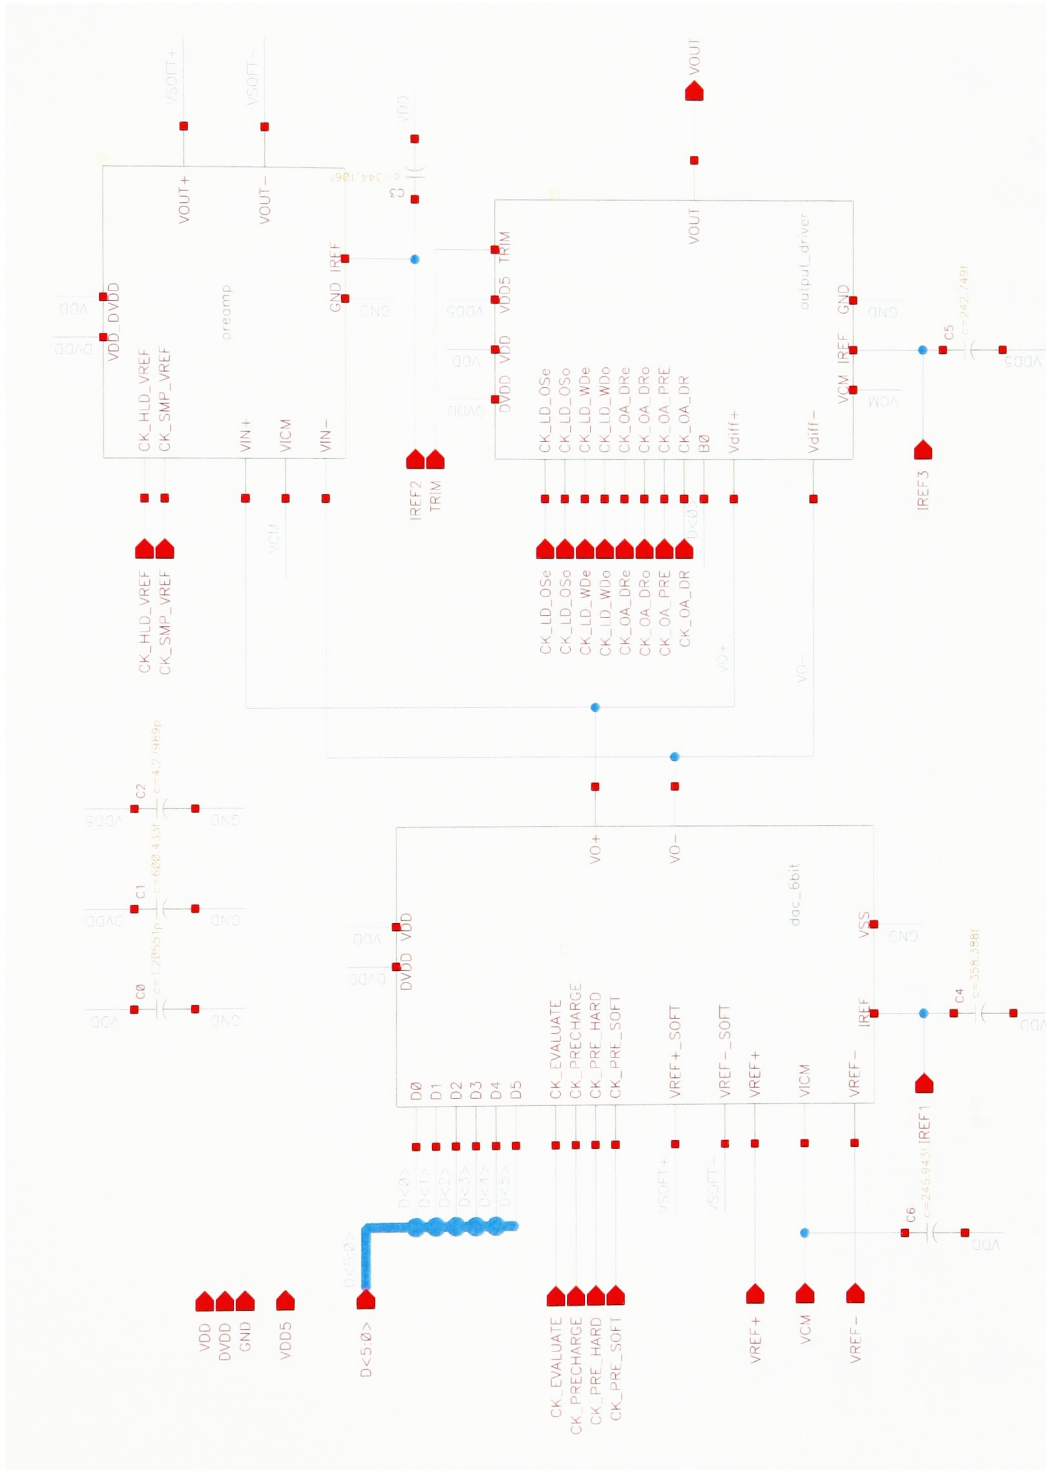

Figure 35 – 6-bit D/A Top Schematic

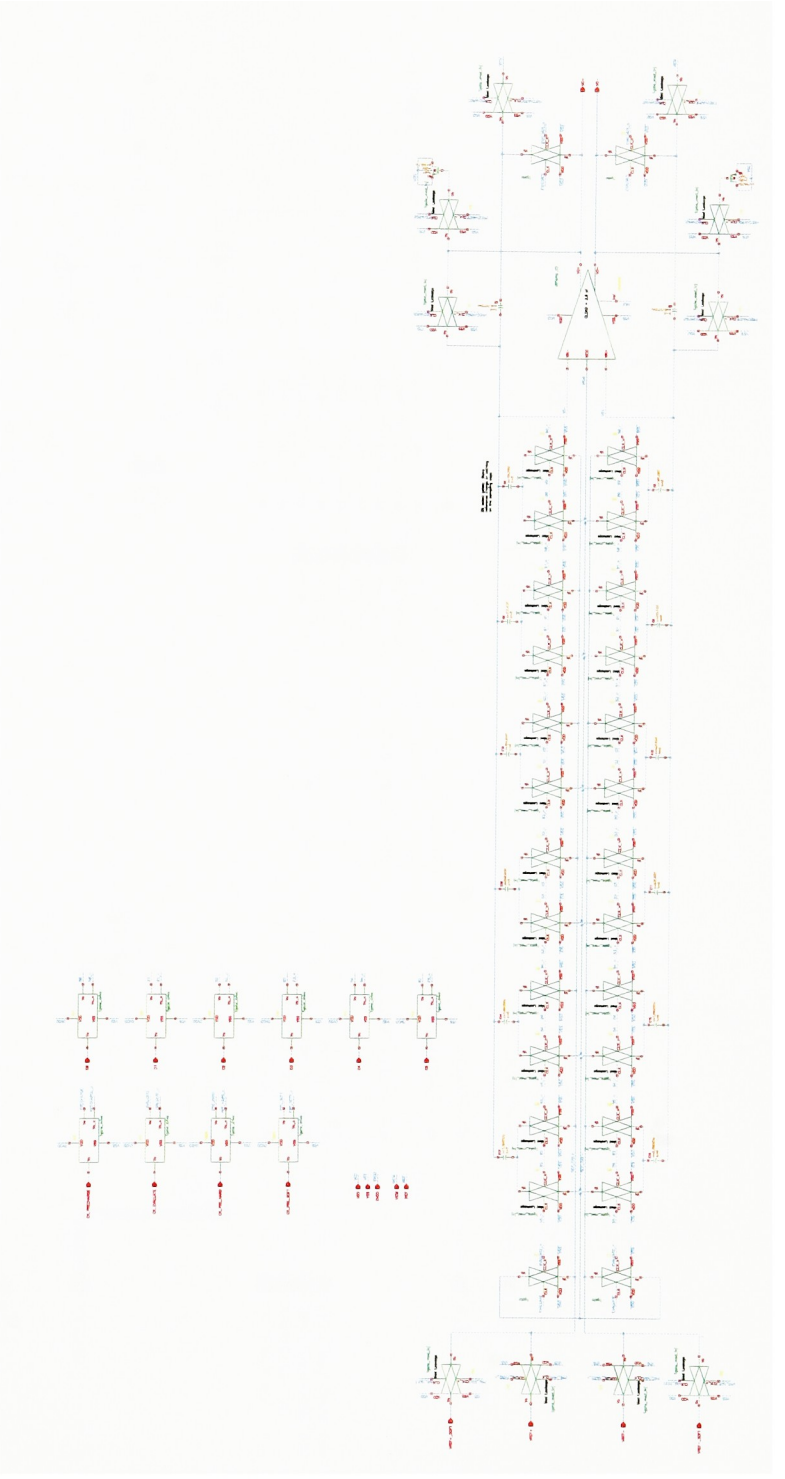

Figure 36 – 6-bit D/A Core Schematic

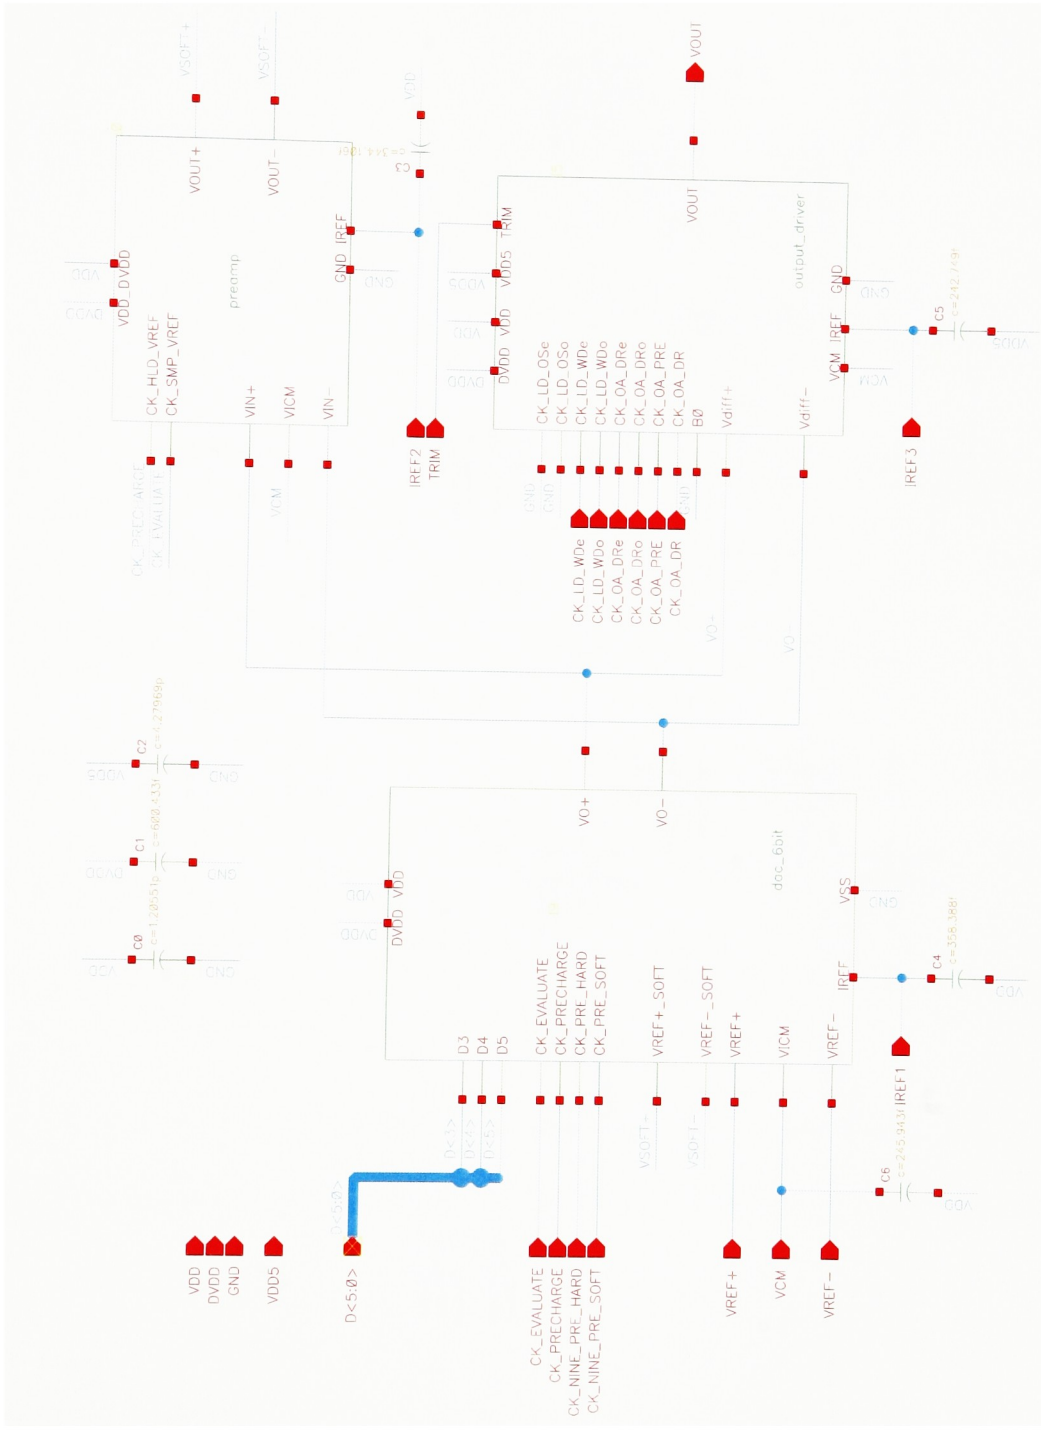

Figure 37 – 9-bit D/A Top Schematic

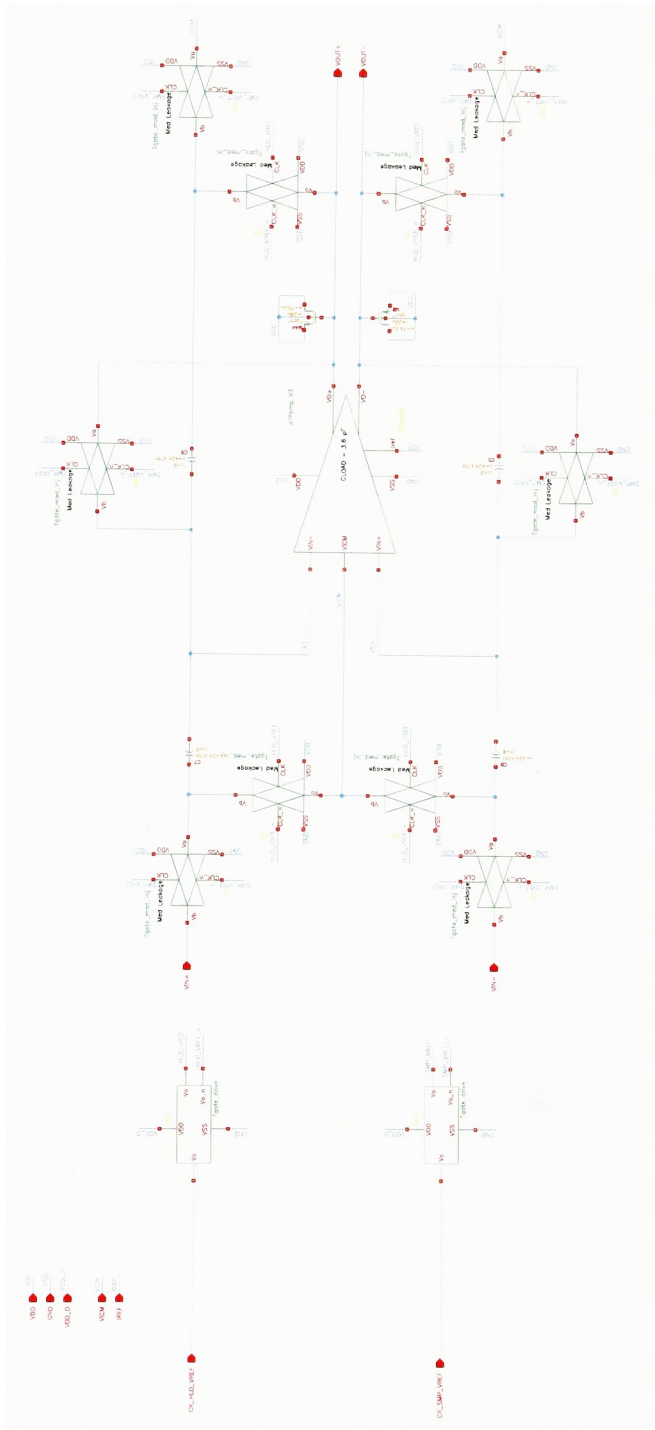

Figure 38 – V_{REF} Sample and Hold Schematic

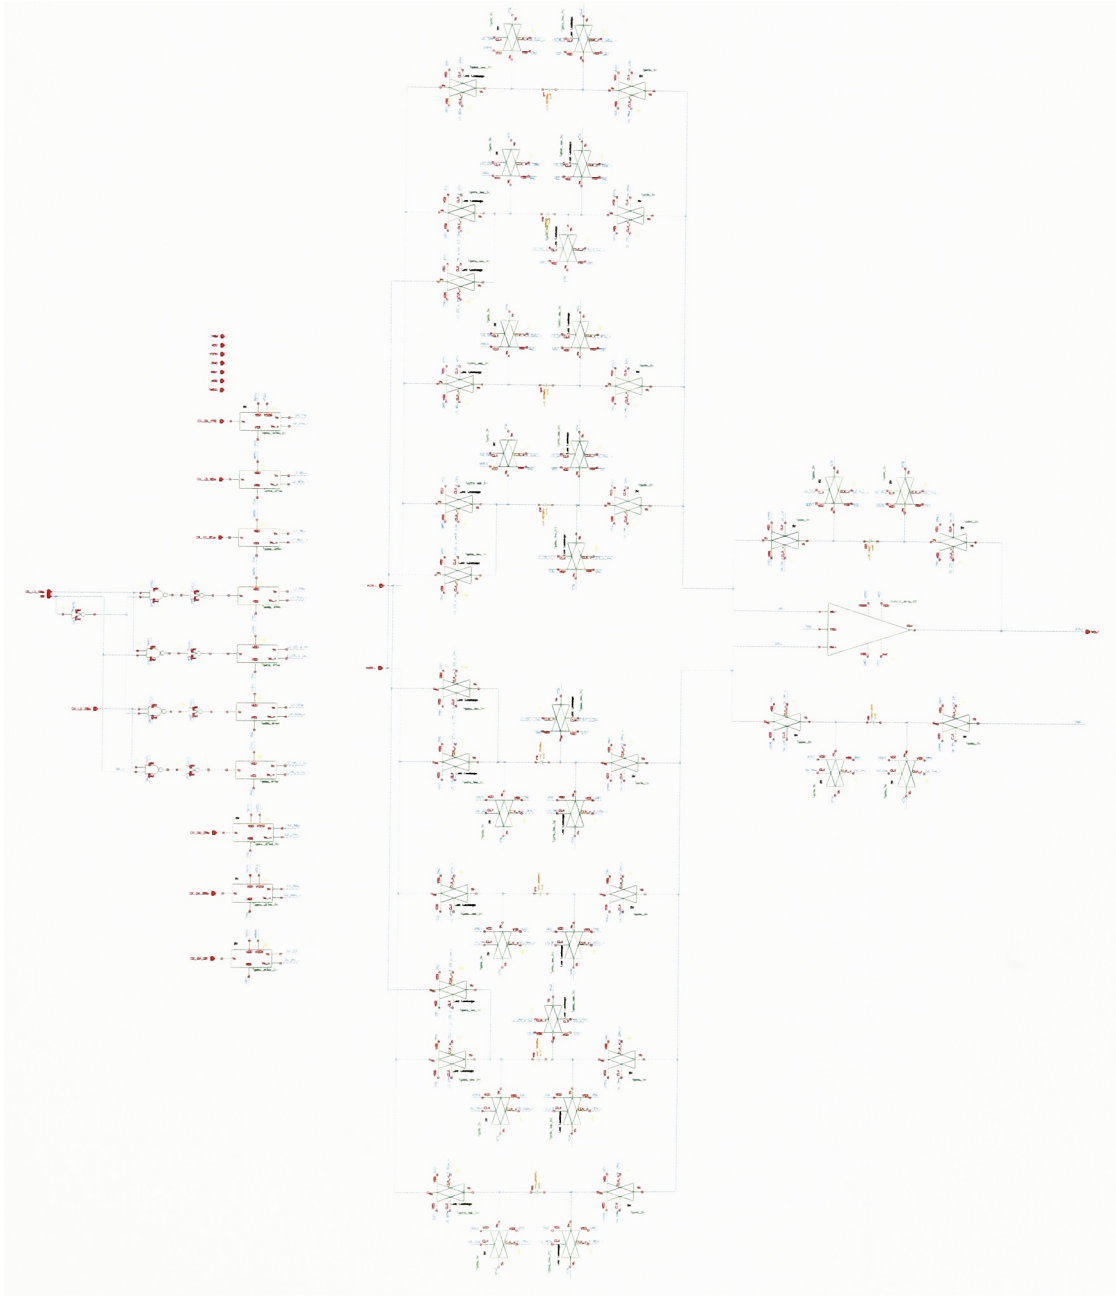

Figure 39 – Output Driver Schematic

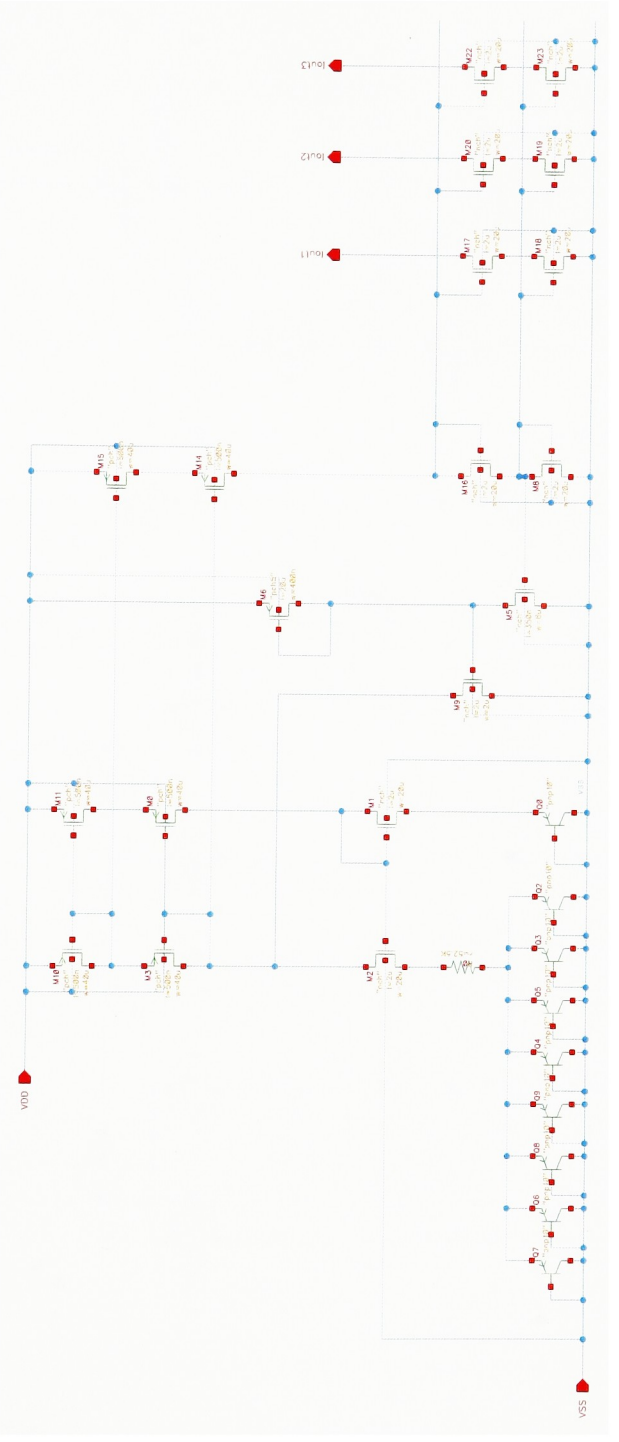

Figure 40 – PTAT Bias Generator Schematic

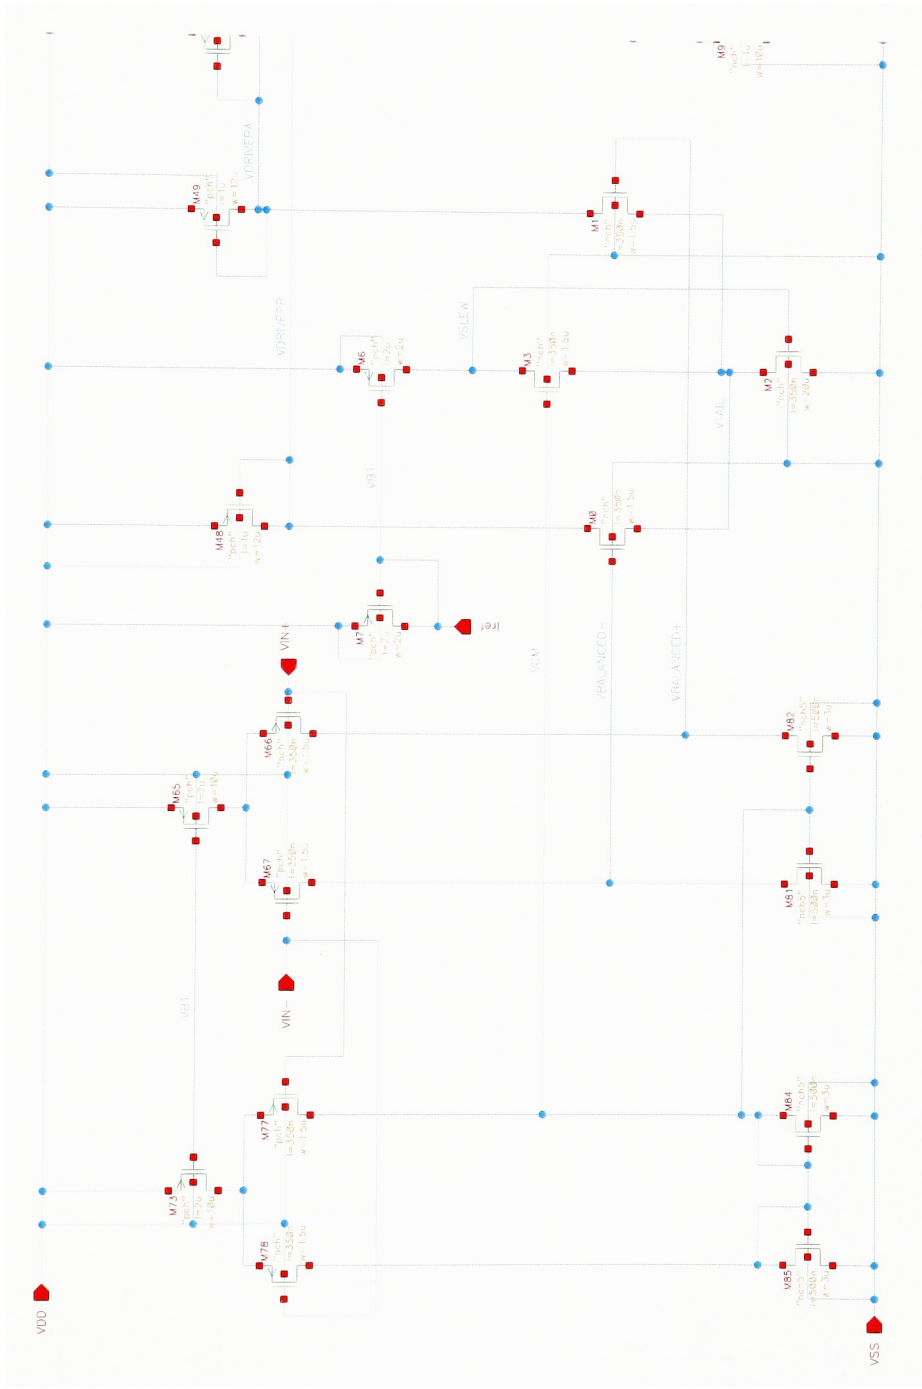

Figure 41 – FD Amplifier (Left Side) Schematic

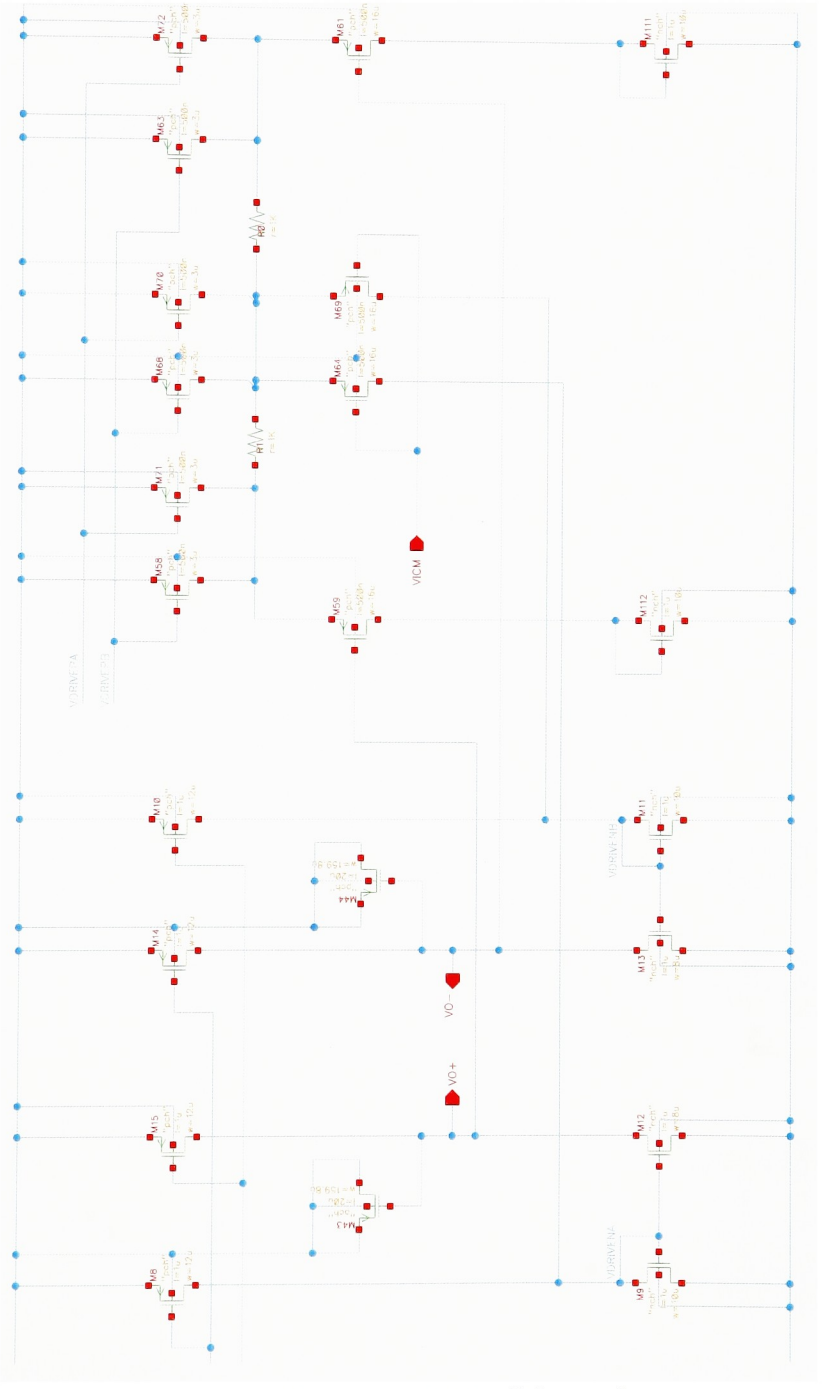

Figure 42 – FD Amplifier (Right Side) Schematic

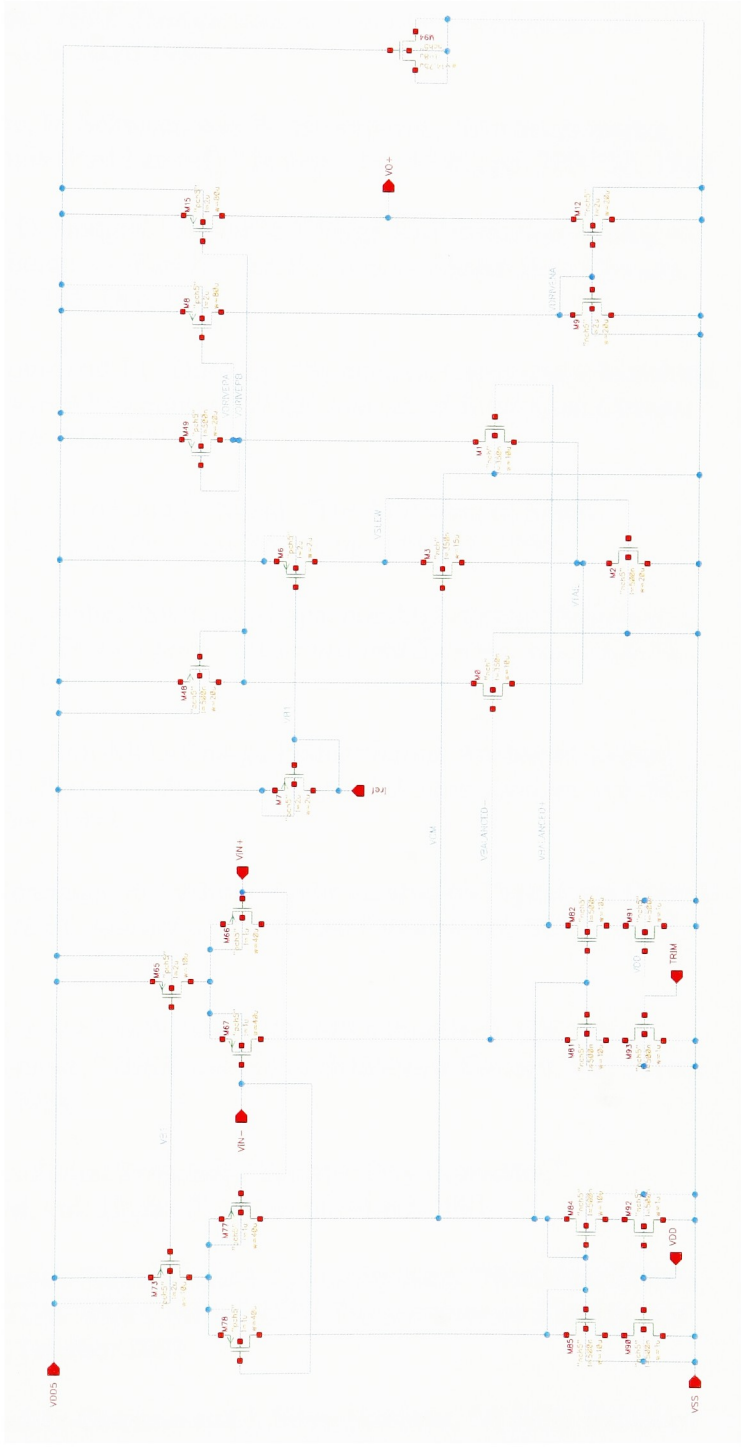

Figure 43 – SE Amplifier Schematic

11.0 References

[1] J. Huertas, *et al.* "Nonlinear Switched-Capacitor Networks: Basic Principles and Piecewise-Linear Design," *IEEE Transactions on Circuits and Systems*, vol. CAS-32, no. 4, pp. 305-319, April 1985.

[2] B. J. Hosticka, U. Kleine, R. Schweer, and W. Brockherde, "Nonlinear analog switched-capacitor circuits (Part I and II)," *in Proc. ISCAS-82*, pp. 729-732, 1982.

[3] R. Suarez, P. Gray, and D. Hodges, "All-MOS Charge Redistribution Analog-to-Digital Conversion Techniques—Part II," *IEEE Journal of Solid-State Circuits,* vol. SC-10, no. 6, pp 379-385, Dec. 1975.

[4] M.J. McNutt, S. LeMarquis and J.L. Dunkley, "Systematic Capacitance Matching Errors and Corrective Layout Procedures," *IEEE Journal of Solid-State Circuits,* vol. 29, no.10, pp 611-616, May 1994.

[5] B.A. Minch, C. Diorio, P. Hasler and C. Mead, "The Matching of Small Capacitors for Analog VLSI," *in Proc. ISCAS-96,* pp. 239-241, 1996

[6] H. Matsumoto and K. Watanabe, "Switched-Capacitor Algorithmic Digital-to-Analog Converters," *IEEE Transactions on Circuits and Systems,* vol. CAS-33, no. 7, pp. 721-724, July 1986.

[7] J. McCreary and P. Gray, "All-MOS Charge Redistribution Analog-to-Digital Conversion Techniques—Part I," *IEEE Journal of Solid-State Circuits,* vol. SC-10, no. 6, pp 371-379, Dec. 1975.

[8] C. Shih and P. Gray, "Reference Refreshing Cyclic Analog-to-Digital and Digital-to-Analog Converters," *IEEE Journal of Solid-State Circuits,* vol. SC-21, no. 4 pp. 544-554, Aug. 1986.

[9] J. Li and U. Moon, "A 1.8-V 67-mW 10-bit 100-MS/s Pipelined ADC Using Time-Shifted CDS Technique," *IEEE Journal of Solid-State Circuits,* vol. 39, no. 9, pp. 1468-1476, Sept. 2004.

[10] R. Gregorian, "High-Resolution Switched-Capacitor D/A Converter," *Microelectronics Journal*, vol. 12, no. 2, pp.10-13, March 1981.

[11] A. N. Mohieldin, E. Sanchez-Sinencio and J. Silva-Martinez, "Nonlinear Effects in Pseudo Differential OTAs With CMFB," *IEEE Transactions on Circuits and Systems,* vol. 50, no. 10, October 2003.

[12] A. N. Mohieldin, E. Sanchez-Sinencio and J. Silva-Martinez, "A Fully Balanced Pseudo-Differential OTA With Common-Mode Feedforward and Inherent

Common-Mode Feedback Detector," *IEEE Journal of Solid-State Circuits,* vol. 38, no. 4, April 2003.

[13] A.J. Lopez-Martin, *et al,* "Low-Voltage Super Class AB CMOS OTA Cells With Very High Slew Rate and Power Efficiency," *IEEE Journal of Solid-State Circuits,* vol. 40, no. 5, May 2005.

[14] D. Hernandez-Garduno and J. Silva-Martinez, "Continuous-Time Common-Mode Feedback for High-Speed Switched-Capacitor Networks," *IEEE Journal of Solid-State Circuits,* vol. 40, no. 8, August 2005.

[15] X. Peng and W. Sansen, "Transconductance With Capacitances Feedback Compensation for Multistage Amplifiers" *IEEE Journal of Solid-State Circuits,* vol 40. no 5. pp. 1514-1520, July 2005.

[16] G. Giustolisi, *et al.,* "1.2-V CMOS Op-Amp with a Dynamically Biased Output Stage," *IEEE Journal of Solid-State Circuits,* vol. 39. no. 4 pp. 632-636. April 2004.

[17] C. Lu, "High-Speed Driving Scheme and Compact High-Speed Low-Power Rail-to-Rail Class-B Buffer Amplifier for LCD Applications," *IEEE Journal of Solid-State Circuits,* vol. 39. no. 11 pp. 1938-1947. Nov. 2004.

[18] L. Yao, M. Steyaert and W. Sansen, "A 0.8V, 8-μW, CMOS OTA with 50-dB Gain and 1.2-MHz GBW in 18-pF Load," *Proc. in ESSCIRC '03,* pp.297-300, Sept. 2003.

Filename: thesis_final.doc
Directory: F:
Template: C:\Documents and Settings\FASTKiosk_XRDS\Application Data\Microsoft\Templates\Normal.dot
Title: N
Subject:
Author: Mark Reisiger
Keywords:
Comments:
Creation Date: 11/1/2005 6:53:00 PM
Change Number: 17
Last Saved On: 11/4/2005 11:34:00 AM
Last Saved By: Patricia Vicari
Total Editing Time: 157 Minutes
Last Printed On: 11/4/2005 11:56:00 AM
As of Last Complete Printing
Number of Pages: 84
Number of Words: 14,065 (approx.)
Number of Characters: 80,173 (approx.)