A Highly Integrated Adiabatic Energy Recovery Digital to Analog Converter

by

M. Josephine Ammer

B.S. Computer Science and Engineering
Massachusetts Institute of Technology

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 1999

Author .....
Department of Electrical Engineering and Computer Science
December 17, 1998

Certified by .....
Thomas F. Knight Jr.
Senior Research Scientist
Thesis Supervisor

Accepted by .....
Arthur C. Smith
Chairman, Department Committee on Graduate Students

A Highly Integrated Adiabatic Energy Recovery Digital to Analog Converter

by

M. Josephine Ammer

B.S. Computer Science and Engineering

Massachusetts Institute of Technology

Submitted to the Department of Electrical Engineering and Computer Science
on December 17, 1998, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

A low-power high-resolution digital input display architecture has been developed which is scaleable to larger display resolutions and grayscale resolutions. The architecture uses parallel sampled ramp DACs in a one-per-column configuration. Both adiabatic and energy recovery operation modes are incorporated for low power. This architecture allows the implementation of optimal drive schemes which are too power intensive when implemented on conventional displays. A 160 x 120 pixel 6-bit implementation of the architecture has been fabricated and tested; it consumes 7.5 mW (worst case) with a 4 MHz pixel clock. A theoretical power consumption model is validated by the working implementation and used to predict the power for other display resolutions, grayscale resolutions, and driving schemes.

Thesis Supervisor: Thomas F. Knight Jr.
Title: Senior Research Scientist

Acknowledgments

I would like to thank my thesis advisor, Tom Knight, for supporting me in this endeavor and for all that he has taught me. I aspire to approach problems as creatively and elegantly as he.

Carlin Vieri offered invaluable input, guidance and support and taught me to treasure "the right solution" once found. To him, and everyone in the Pendulum group, I owe my thanks, especially also Norm Margoulis.

Special thanks to everyone at MicroDisplay for everything that made this work possible: a wonderful learning environment, an excellent infrastructure for practicing engineering, and the friendly social atmosphere. Thanks especially to Phillip Alvelda, David Huffman, Chris Mullin, Rand Stadtman, Jean-Jacques Drolet, Ramon Olivier, and Todd Stiers.

Thanks to Jim Goodman for his input on the project, and for his painstaking editing. Also thanks to Raj Amirtharajah for his ideas on ramp generation circuits.

Several friends at MIT offered support throughout my thesis; without them, I would truly have gone insane. They include: Paul Lagace, Mark Foltz, Jeremy Brown, Anthony Zolnik, Tony Gray, Christine Alvarado, and Mike Oltmans.

I owe a world of gratitude to Mike Bolotski. He continually raised the bar, insisted I jump over it, and was the first person to cheer when I cleared it. He is my most brutal critic and my most avid supporter.

Finally, I would like to thank my family for all the years of love and support.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The development of low-power, high-resolution displays will improve the usability of portable devices by allowing them to effectively display the increasing bandwidth of information available to them. For instance, email is regularly forwarded to pagers and digital cellular phones, but the ability to read email is limited because of the small screen resolution on most of these devices. Bandwidth is currently available to transmit text-based web pages onto cellular phones, however, viewing these web pages is difficult due to the small screen resolution. Larger resolution displays are not incorporated into these devices because of their high power consumption. Laptop computers are one type of portable device which incorporate high resolution displays. However, the display is one of the dominant power consumers in laptop computers leading to large heavy batteries or short battery life. Email, laptop computer display information, and more commonly, cellular phones are all dominated by digital communication and storage. Therefore developing displays which interface with a digital input is crucial to the integration of displays with portable applications. The goal of this work is to develop a low-power, high-resolution display that interfaces with a digital input.

A promising technology for low power information display is liquid crystal on silicon (LCOS) which enables same-die integration of control circuitry and display for substantial overall system power reduction. However, information display in the form of images on these devices is an analog problem; the display brightness at any

position is determined by the voltage applied at that position. Therefore, the problem of converting stored or transmitted digital data to analog pixel values for display is a crucial step in developing low power display systems. Further, the power consumption of high-speed digital to analog converters (DACs) will become a main component of a portable device's power budget. Therefore, including low power on-chip digital to analog conversion for LCOS displays will decrease their size and decrease the overall display system power which presently includes a high-power off-chip monolithic DAC.

One feature of information displays that can be exploited to lower power is its inherent parallelism. A large, fast, high-power monolithic DAC can be used to convert one pixel at a time, or several slower, lower-power DACs in parallel can be used. Parallelism yields benefits in the form of scalability and allows for the application of several low power techniques which exploit parallelism to reduce power. One such common power reduction technique is voltage scaling. By parallelizing a system to reduce the speed requirements of each parallel component, the system supply voltage can be reduced. Since power consumption in traditional CMOS circuits falls off roughly as the square of the supply voltage, running N parallel components at $1/N$ the speed reduces power consumption by a factor of N [2]. For miniature liquid crystal displays, voltage scaling is not an option because the liquid-crystal (LC) material requires digital to analog conversions that span over a five volt range. It is rare to find an integrated DAC that runs over 5 volts and even more rare to find one that is optimized for low power.

There are potential disadvantages of using parallelism in display drivers. Process variation, temperature variations, and timing differences across the CMOS die can cause parallel elements to have different behavior which is noticeable in the displayed image. Parallel digital to analog converters must be designed carefully to avoid these variations. Another disadvantage is that parallelism can often require the duplication of logic, which can significantly increase the die size.

Two techniques that can reduce power, yet operate at high voltages are adiabatic charging and energy recovery. These techniques are particularly suitable for driving highly capacitive loads such as the pixel array of an LCOS display. Adiabatic charging

and energy recovery take advantage of parallelism with the same linear reduction in power as voltage scaling, yet at maximum voltages.

The goal of this work is to design a low power display driver architecture which includes digital to analog conversion that works at 5V and with the large pixel array capacitances. The resulting design must not introduce image artifacts due to converter variation, must integrate well with the pixel array layout, and must be scalable to larger display resolutions. Parallelism combined with adiabatic charging and energy recovery is the chosen low power method for the high voltage, high capacitance pixel array. Other, more conventional, low power methods are employed where possible. The resulting architecture is called **ACERDAC** which stands for **A**diabatic **C**harging **E**nergy **R**ecovery **D**igital to **A**nalog **C**onversion. A physical implementation of the ACERDAC architecture is required for proof of functionality and utility. The ability to predict the power for this architecture on other display resolutions and configurations is desired.

By improving on the low-power DAC problem for displays, this work extends the field of low power electronics by including environments of high voltage and high capacitance in the domain of low power digital to analog conversion. This work also shows that combining conventional low power circuit techniques with adiabatic charging and energy recovery can be practicable and practical.

Several fields of study are combined in this work, and the necessary background on each is given in Chapter 2. This includes background on LCOS display structure and driving methods, low-power circuitry design including an introduction to adiabatic charging and energy recovery, the basics of digital to analog conversion, and digital to analog converter methods used in other displays. The ACERDAC architecture is described in Chapter 3, and the specific implementation which was fabricated is discussed in Chapter 4. The ACERDAC implementation is tested and used to validate a theoretical power consumption model which predicts the power of different ACERDAC implementations with different display resolutions, grayscale resolutions, and driving schemes. The test results are given in Chapter 5. The lessons learned from the first ACERDAC implementation are used to propose a second implementation.

The theoretical model is used to predict the power of this second implementation and the results are compared to a conventional LCOS display in Chapter 6.

Chapter 2

Background and Previous Work

Because several fields of study are combined in this work, this chapter presents background on liquid-crystal (LC) materials, LCOS and other display technologies, low power design, adiabatic charging and energy recovery, and the basics of digital to analog conversion. Low power digital to analog conversion is the goal of this work and liquid-crystal-on-silicon (LCOS) displays are the test bed for this architecture. Two low power techniques, adiabatic charging and energy recovery, are at the primary low power methods used in this research, but other, more conventional low power techniques are also used. The chosen DAC method is the sampled ramp DAC because it is suited to parallelism and and the application of adiabatic charging and energy recovery.

2.1 LCOS Displays

LCOS displays achieve lower power than flat panel displays and CRTs because of a small physical structure and lower voltage requirements. However, the LCOS display power needs to be lowered further in order to be considered for integration into power-limited devices such as cellular phones and pagers. To develop an architecture for a low-power digital-input display, it is crucial to examine the physical structure of LCOS displays, the drive method for analog input LCOS displays, and the tradeoffs of placing the DAC at different points in the display architecture.

2.1.1 LCOS Physical Structure

LCOS displays consist of a silicon backplane, a liquid crystal layer, and a clear coverglass (Figure 2-1). The top most metalization layer on the silicon and a clear conductive film (made of indium tin-oxide or ITO) on the inside of the coverglass form the pixel electrodes. The display uses a twisted-nematic (TN) LC in which the LC molecules twist between the two electrodes, rotating incident light to pass through a laminated polarizer. When an electric field is applied, the molecules orient vertically, destroying the twist, not rotating the light, and resulting in a dark pixel [1]. Since the ITO electrode is held at a constant voltage, the pixel brightness is determined by the voltage applied to the top metalization layer of the silicon backplane.

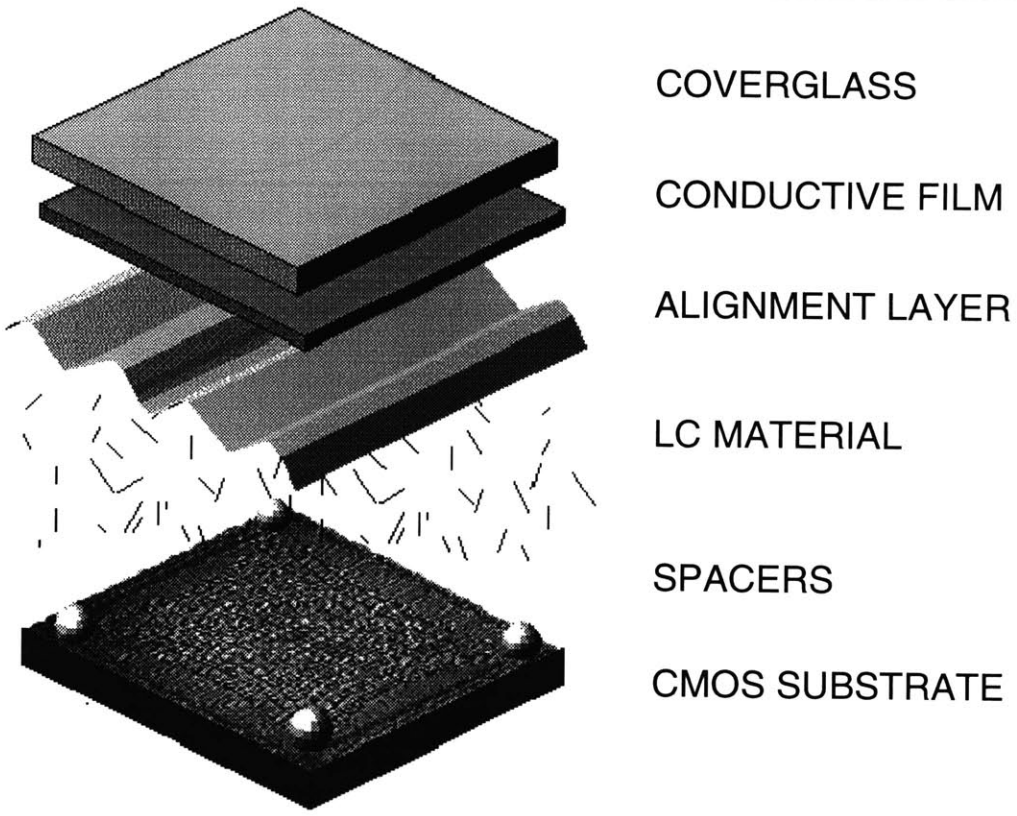

Figure 2-1: LCOS Display Layers. The top-most metalization layer on the CMOS substrate and a clear conductive film on the inside of the coverglass form the pixel electrodes.

2.1.2 Driving Liquid Crystal Materials

There are many different LC modes; the one used in this research is normally white, and achieves maximum darkness when 2 volts is applied across it. A typical LC transfer curve of input voltage versus brightness is non-linear (Figure 2-2). A pre-distortion is applied to the pixel voltages to compensate.

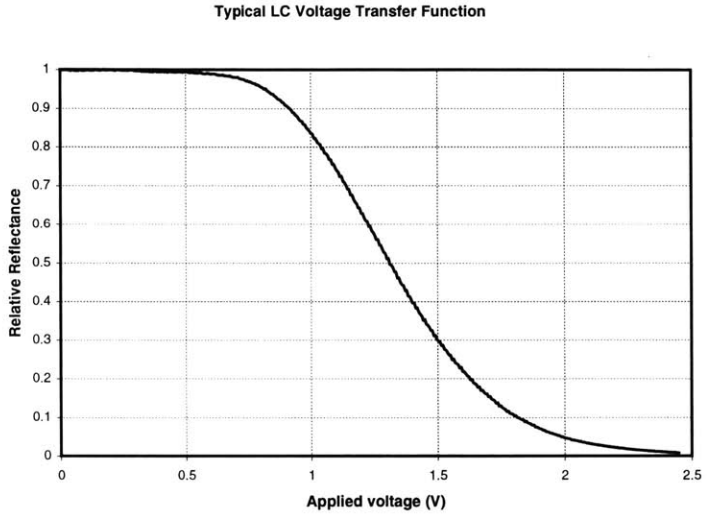

Figure 2-2: Liquid Crystal transfer curve. LC responds to the RMS voltage, so a positive or negative field can be applied. A correction is applied to the pixel voltages to compensate for the nonlinearity of this curve.

LC drive schemes are further complicated by the requirement of maintaining a net zero DC field across the LC material. LC materials contain ionic impurities that drift to the electrodes under a DC electric field. When sufficient impurities collect at an electrode, they nullify the charge on the electrode, preventing correct operation. Fortunately, LC materials respond to the RMS value of the applied field rather than the instantaneous value, and can be driven with alternating polarity signals to effectively eliminate the DC field. This method is called inversion, and can be applied at the pixel, line, or frame granularities [5].

Experiments have shown that in terms of image quality, pixel inversion is superior to line inversion which is in turn superior to frame inversion [5]. Unfortunately, pixel

inversion has the highest power cost for most display architectures, since successive pixel values of identical intensities require opposite voltages to be applied to the pixel electrode. Since images are frequently correlated both horizontally and vertically, pixel inversion can lead to the maximum number of voltage transitions with some display driver architectures. One advantage of the proposed architecture in this work is the low overhead of the pixel-inversion scheme compared to the frame-inversion scheme.

2.1.3 Analog LCOS Display Driving Method

Figure 2-3 shows a typical analog display architecture in which the video signal drives a horizontal wire connected to each pixel column via an analog switch. The switches are controlled by a shift register containing a single high bit that connects each column in turn to the video wire. A similar vertical shift register connects all pixels in the active row to their column wires. As each analog pixel value is placed on the video wire, it is sampled by the column wire and the currently active pixel. After all pixels in a row have been loaded, the vertical shift register advances and the process repeats [5].

The pixel array is similar to a DRAM; voltages are trapped as charge on the pixel electrode capacitances. As in DRAMs, the voltage on the pixel electrode will degrade over time which causes the displayed image also to degrade. Therefore, even when displaying a constant image, the pixel array must be refreshed periodically. In addition to the refresh requirement, the screen must be repainted periodically with the alternating polarity of signal so that a net zero DC field can be maintained. Repainting the screen causes the image to flicker. Experiments have shown that updating the screen at a rate slower than 60 Hz causes unacceptable flicker whereas a faster update rate causes proportionally less flicker.

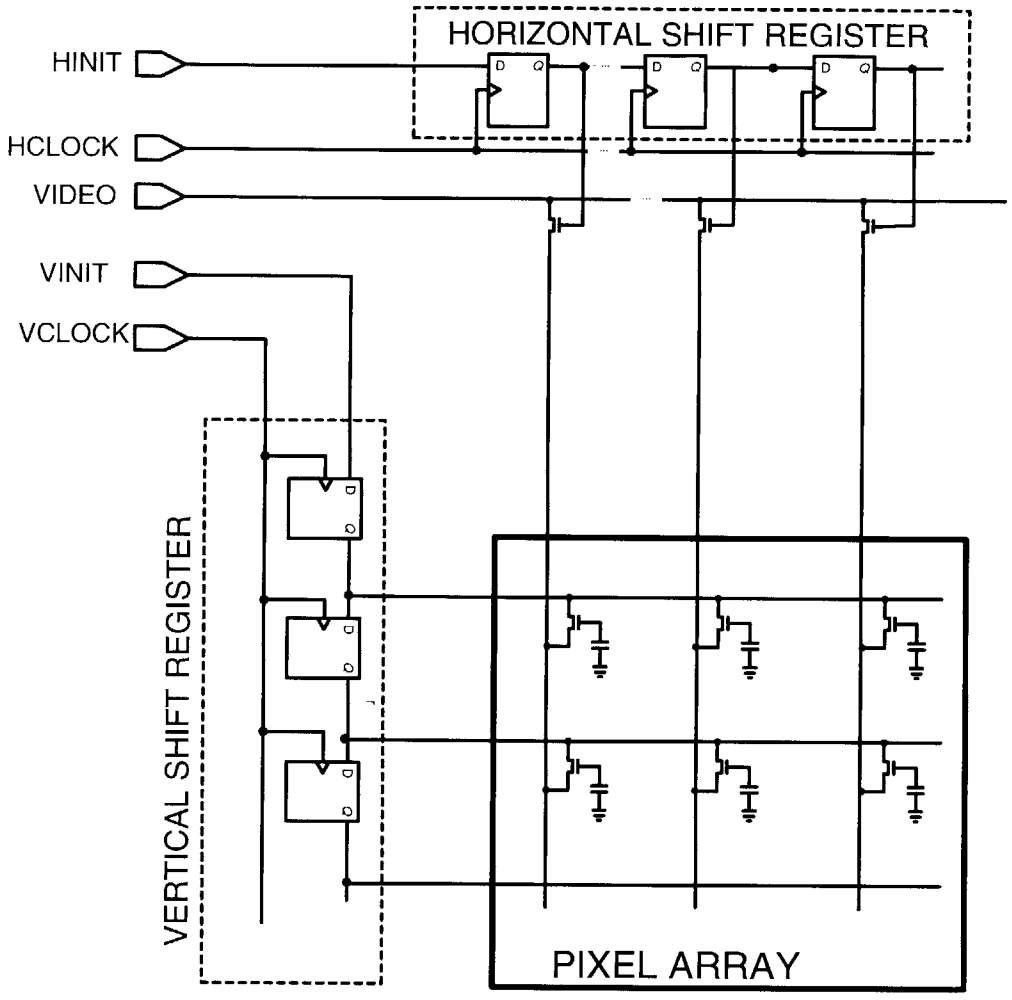

Figure 2-3: Typical analog display architecture. Horizontal shift registers connect a column of pixels to the video wire, and vertical shift registers connect a row of pixel to their respective column wires.

2.1.4 Operation Frequencies of LCOS Displays

There are three basic frequencies applicable when discussing displays: frame rate (F_f), row rate (F_r), and pixel rate (F_p). The frame rate corresponds to the frequency at which the entire screen is repainted. Typically, F_f is 60Hz, whereas F_r and F_p are a function of F_f and the display resolution.

$$F_r = N_r \cdot F_f$$

$$F_p = N_c \cdot F_r = N_c \cdot N_r \cdot F_f,$$

where N_r and N_c are respectively the number of rows and number of columns in the display. Table 2.1 gives F_r and F_p for two typical display resolutions.

Resolution	Frame Rate F_f (Hz)	Row Rate F_r (KHz)	Pixel Rate F_p (MHz)
160 x 120	60	7.2	1.152
640 x 480	60	28.8	18.432

Table 2.1: Row and pixel frequencies. Given for two display resolutions at a frame frequency of 60 Hz. F_r is F_f times the number of rows, and F_p is F_r times number of columns.

2.1.5 DAC Placement in LCOS Displays

The overall requirement for a digital-input LCOS display driver is to receive incoming digital data at F_p and write a fresh analog voltage to each pixel of the display at F_f. There are three major phases to this procedure: distributing the digital input data to the DAC, performing the conversion, and distributing the resulting analog value to the pixels. The DAC placement within the data distribution path is a key design parameter in the display driver design.

One architectural extreme is to have a DAC right at the chip input (Figure 2-4).

The digital distribution system is reduced to nothing but digital input pads. But, the DAC needs to convert at F_p, and the analog distribution system needs to distribute analog voltages at F_p to the columns. While this architecture requires the circuits to operate at the highest frequency, F_p, this is the most common architecture for contemporary LCOS displays because it is the easiest way to convert an analog input display to a digital input display.

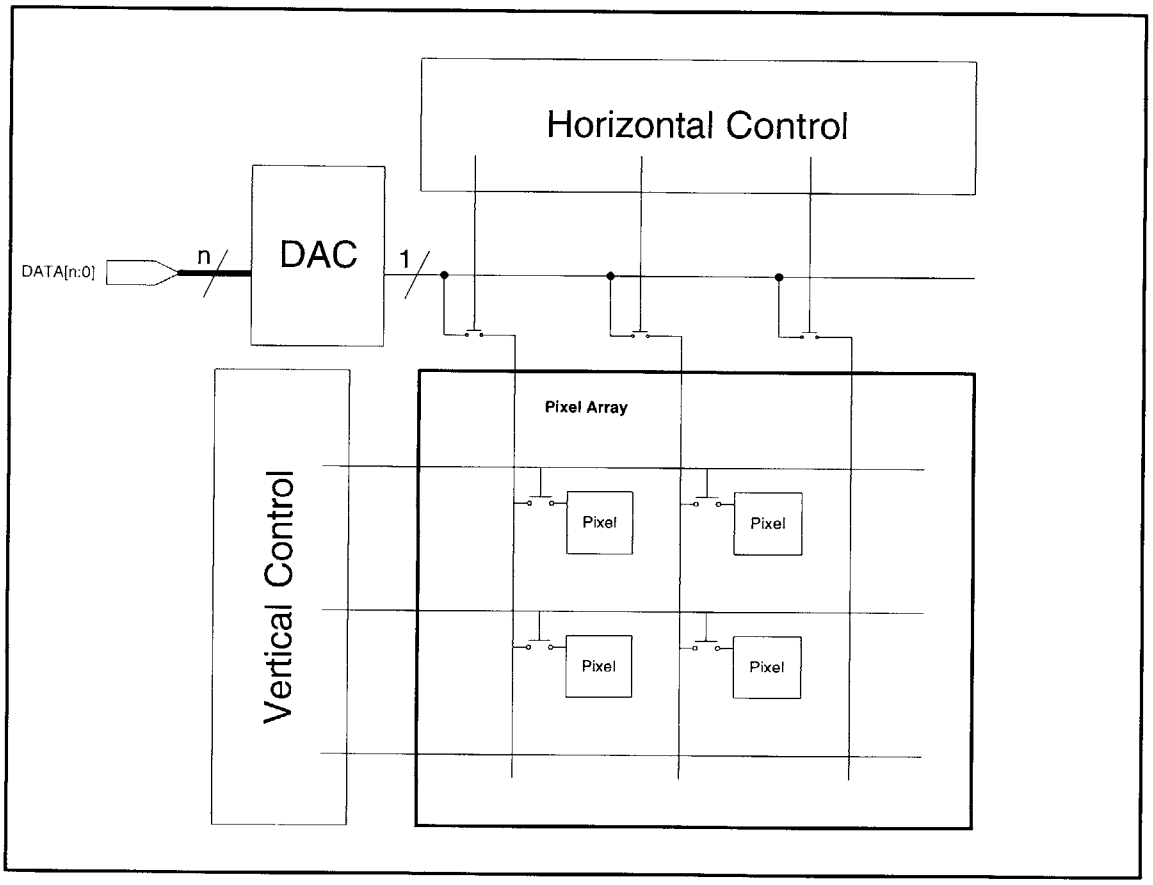

Figure 2-4: Converting Digital Pixel Data using One DAC. This is the most common architecture for LCOS display because it is the easiest way to convert an analog input display to a digital input display. The DAC must operate at the fastest frequency, F_p.

Another architectural extreme is to place a DAC at each pixel (Figure 2-5). Then, the digital distribution system distributes digital data to the columns at F_p, and the

columns distribute digital data to the pixels at F_r. In comparison to the single DAC architecture, the frequency at which data arrives at the pixel is reduced from F_p to F_r. But, having one DAC at each pixel is inefficient from both a space and power perspective. It would be difficult or impossible to fit a DAC within the 25 μm x 25 μm pixel area let alone routing a digital bus through the pixel array to each pixel. Each DAC would have a duty cycle on the order of 10^{-6}, resulting in a gross underuse of resources and unnecessary power dissipation.

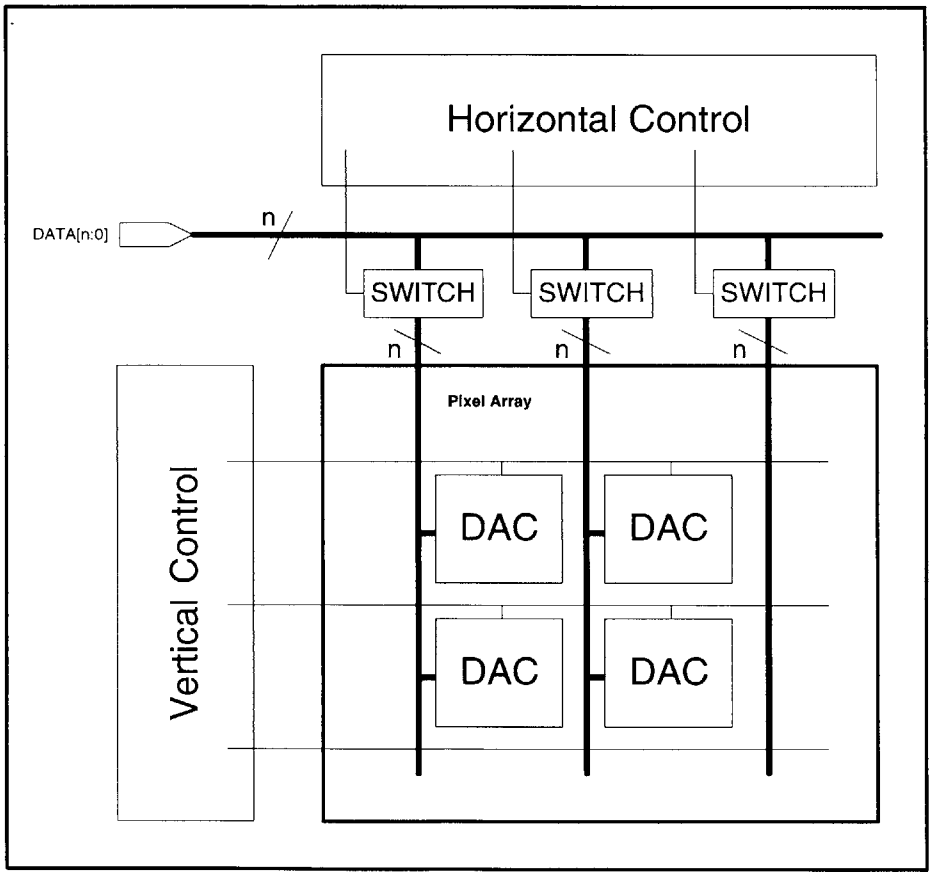

Figure 2-5: Converting Digital Pixel Data using One DAC at each pixel. The number of vertical wires per column is large and may not fit within the narrow pixel pitch. The DACs are underutilized and also may not fit within the pixel pitch.

A logical architectural compromise is to have one DAC for each column (Figure 2-6). The digital path distributes data from the chip input across the horizontal

direction of the chip to each column. The analog path distributes data down the vertical direction of the chip. This data distribution scheme results in each DAC being utilized once every row time. The digital path, where there are many wires, distributes data to the top of each column and therefore does not need to penetrate into the pixel array. The analog path, where there are few wires, is suited to its position inside the pixel array where space is at a premium. The large body of research focusing on low power digital circuits can be applied to the digital distribution path. Recent research in adiabatic charging and energy recovery can be applied to the analog distribution path and the DAC design. Parallelizing the DAC operations at each column allows for the application of low power techniques such as adiabatic charging and energy recovery whose effectiveness improves at lower speeds.

In the preceding discussion, it has been assumed that data is arriving in a raster format because this is the most prevalent form of data delivery. If image data is stored in a local frame buffer, the data can be accessed and displayed in any order (particularly, in raster order). But, if the data is being transmitted from a distance, re-ordering the data from the transmitted order requires the addition of some amount of memory. If data is displayed in the order received, high-power memory components are avoided.

2.2 Low Power Design

The power for an electronic system has four sources:

- **Dynamic power** is the result of charging capacitances in the circuit such as wires and transistor gates.
- **Short circuit power** is a result of resistive paths from power to ground which are formed while circuits are transitioning.
- **Static power** is the result of resistive paths from power to ground when the circuits are not transitioning.
- **Leakage power** is the result of reverse bias leakage between diffusion regions and the substrate.

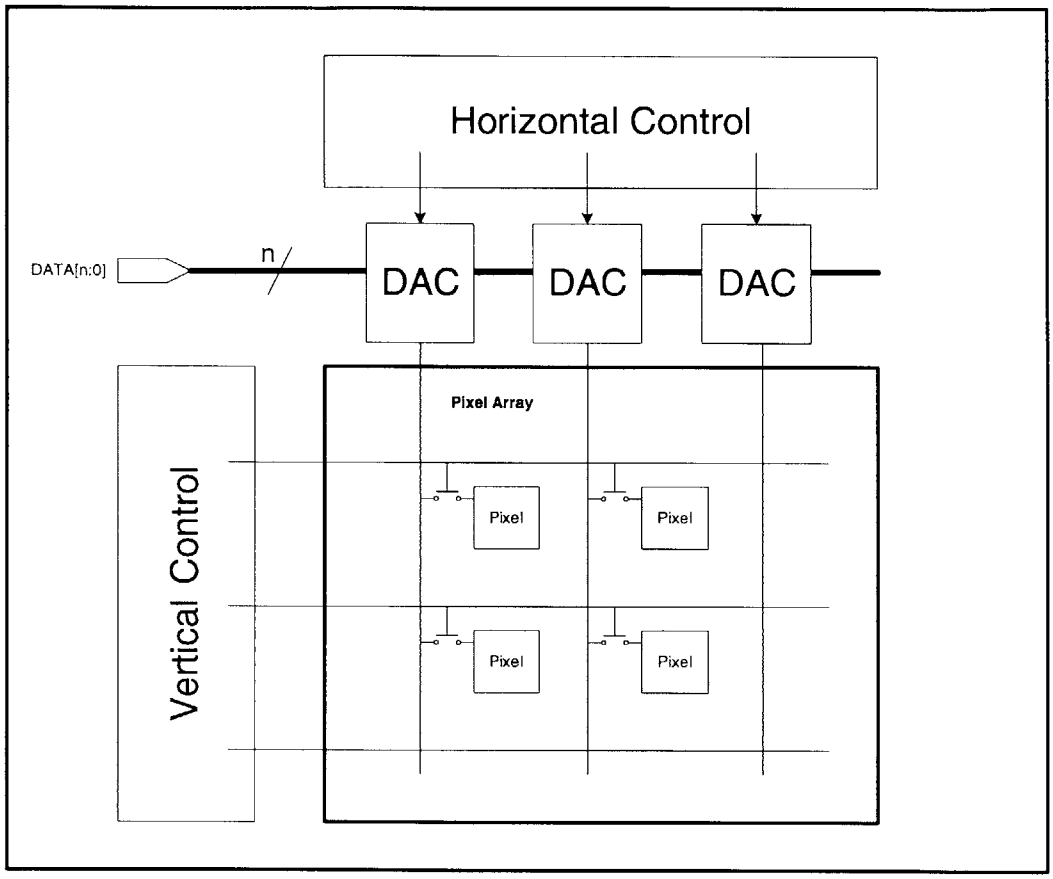

Figure 2-6: Converting Digital Pixel Data using One DAC at each column. This is the architecture used for the ACERDAC. Only one vertical wire per column penetrates the pixel array. Traditional low power techniques are applied to the digital distribution bus.

The total power for a system is the sum of the four components. Since LCOS displays and their surrounding circuitry are fabricated in a standard, relatively high-voltage CMOS process, static power and leakage power are negligible. Therefore, the two major components of power consumption are dynamic and short circuit.

Dynamic power for digital components is governed by the equation

$$P = \alpha_{0\to 1} f C V^2 \tag{2.1}$$

where C is the switched capacitance, V is the voltage, f is the system clock frequency, and $0 \leq \alpha_{0\to 1} \leq 1$ is the fraction of clock periods in which the component switches from a logical zero to one. Low power design encompasses methods of reducing each of these components. The focus of low power design is often to reduce the voltage because of the quadratic, rather than linear, relationship between voltage and power.

Voltage can be reduced by lowering the supply voltage to a component, or by lowering the voltage swing on a part of the component. Lowering the voltage supply of a component reduces the circuit speed, so voltage supply scaling is often combined with parallelism to achieve both acceptable speed and low power. Low voltage swings are often used for highly capacitive wires such as those in global busses. Low swing schemes require special receiver circuits to amplify the low swing signal to full swing for use with the rest of the circuitry.

Dynamic power for analog circuits (circuits which switch between more than two voltages) can be more difficult to analyze than digital circuits. The roles of α and V are less clear than in digital circuits because of the different voltages through which a circuit can transition. As a result, analyzing power for analog circuits requires careful tracking of the specific voltages through which the circuit transitions and the energy required for each of those transitions. If the circuit is switched through several different voltage levels, the power dissipated is the sum of the energy for each transition over a unit time period. Switching a capacitance, C, from a voltage, V_1, to a higher voltage, V_2, requires $CV_2(V_2 - V_1)$ of energy be drained from the supply. Switching a capacitance from a higher voltage to a lower one requires no energy to be

drawn from the supply. The dynamic power of an analog component can be reduced by using some of the same techniques as for digital components such as reducing the capacitance C, the voltage change $V_2 - V_1$, or reducing the number of transitions per second.

Short circuit power for digital CMOS circuits is the result of both a pull-up path and a pull down path being simultaneously on, which occurs when the circuit is transitioning from one state to the other. Short circuit power is dependent on the on-resistance of the pull-up and pull-down paths as well as the input waveforms rise and fall times[10]. Slow rise times on nodes can result in significant (20%) short-circuit dissipation for loaded inverters[10]. Circuits which can never have simultaneously on pull-up and pull-down paths are preferred for low power design because they eliminate short circuit power. However, for circuits that can dissipate short circuit power, sharp transitions should be guaranteed on the input waveforms to reduce the short circuit power component.

2.3 Adiabatic and Energy Recovery Introduction

Adiabatic charging (AC) and energy recovery (ER) are two circuit methods used to reduce the energy dissipated by a system. Typically, these techniques are most effective at reducing dissipation for charging high capacitance loads or loads that require high voltages. Other load types can usually be charged more efficiently by employing voltage-scaling techniques. This section gives an introduction to adiabatic charging and energy recovery beginning with a determination of where energy is dissipated in a simple RC circuit. Energy dissipation is calculated for four different charging sources: a conventional step, a ramp, a sinusoid, and a stair-step ramp.

To evaluate the theoretical efficiency of a waveform shape for charging a capacitive load, the simple RC circuit charged by a voltage source is examined (Figure 2-7). The energy being drawn from the voltage supply,

$$E_s = \int_0^\infty V_s(t) I_s(t) dt. \tag{2.2}$$

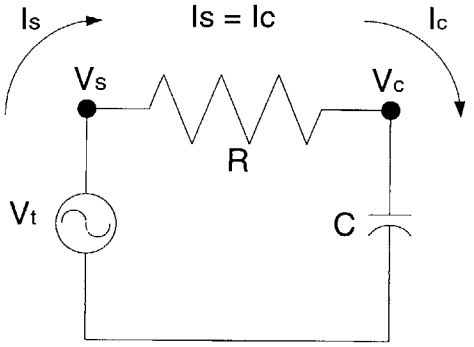

Figure 2-7: A simple RC circuit. Used to analyze the power consumed by different charging schemes.

The energy delivered to the capacitor,

$$E_c = \int_0^\infty V_c(t) I_c(t) dt \tag{2.3}$$

which is always equal to $\frac{1}{2}CV^2$ if the capacitor is charged from GNDto V. (Note that $I_s = I_c$.) Therefore, the energy dissipated in the resistor, E_{diss}, is

$$E_{diss} = E_s - E_c = E_s - \tfrac{1}{2}CV^2. \tag{2.4}$$

also

$$E_{diss} = \int_0^\infty I_s^2(t) R\, dt \tag{2.5}$$

Using Equations 2.4 and 2.5, E_{diss} is calculated for three popular charging waveforms: a step, a linear ramp, and a sine wave (Figure 2-8). The step input shows the typical $\frac{1}{2}CV^2$ dissipation. The linear voltage ramp is the most efficient adiabatic source because it is constant current. When the charging time T approaches infinity, the dissipation approaches zero. The sine wave, adjusted to resonate between 0 and V volts with a charging time of T, has been used in place of a linear ramp [3] [4] [11] because it is simple to generate with a resonating inductor and capacitor circuit. The sine is much more efficient than a step input if the period is sufficiently slow, but only

$\frac{8}{\pi^2}$, or 81% efficient compared to a ramp with the same rise time.

Source	E_{diss}	
V step(t)	$\int_0^\infty (Ve^{-t/RC})^2 R\,dt$	$\frac{1}{2}CV^2$
V ramp(t) : $I_s = \frac{CV}{T}$	$\int_0^T I_s^2(t) R\,dt$	$\frac{RC}{T}CV^2$
$\frac{V}{2}\left(\sin(\frac{\pi}{T}t - \frac{T}{2}) + 1\right)$	$\int_0^T \left(\frac{VT^2\pi C\left(T\sin(\frac{\pi t}{T}) - RC\pi\cos(\frac{\pi t}{T})\right)}{2T^4 + 2\pi^2R^2C^2T^2}\right)^2 R dt$	$\frac{\pi^2}{8}\frac{RC}{T}CV^2$

Figure 2-8: The efficiency of three popular charging waveforms. The step is the most common, but uses the most power. The ramp is the most efficient, but difficult to generate. The sine is easier to generate, more efficient than the step, but less efficient than the ramp.

Another easily-generated charging waveform is a stair-step which is generated by switching the load through N intermediate steps on the way up to the final voltage, V. Using a stair-step input to charge a load is known as step-wise charging. Steps must be spaced at least RC apart in time or the stair-step will look like a step input to the circuit. Assuming the steps are evenly distributed from 0 to V,

$$E_{diss} = \frac{1}{2}NC\left(\frac{V}{N}\right)^2 = \frac{1}{2N}CV^2.$$

As the number of steps approach infinity, the stair-step approaches a ramp, and the dissipation approaches zero just as in the ramp input case. The key concept is that energy is proportional to the voltage squared. Incurring ΔV^2 for N small ΔV transitions is much more efficient than incurring ΔV for one large transition.

Adiabatic charging is achieved when a charging waveform is more efficient than $\frac{1}{2}CV^2$ such as with the ramp or sine waveforms. Energy recovery is achieved when some of the $\frac{1}{2}CV^2$ of energy stored on the charged capacitive load is recovered and reused for later chargings. In order to evaluate the merits of AC and ER, a comparison is made to conventional charging and discharging methods using Equations 2.2

through 2.5.

Conventional charging, also known as the switching method, is used by standard CMOS gates. To charge a load to a voltage V, the load is connected through a resistive switch to a voltage source that is biased at V. To discharge a load to ground, it is connected through a resistive switch to the ground node. These transitions can be modeled as a step inputs (Figure 2-9) because the transition times of the switch input are smaller than the RC time constant of the switch and load.

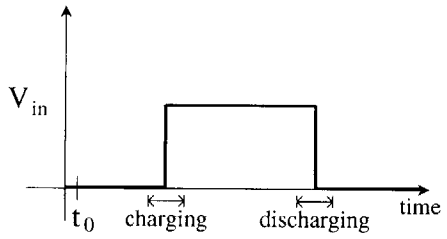

Figure 2-9: Conventional charging and discharging waveform. When a circuit is switching from one voltage supply to another as in standard CMOS circuits, the inputs look like a step.

In this conventional charging example, a load, discharged at t_0, is associated with an energy reservoir of 0 and a voltage source is associated with an energy reservoir of CV^2 (Figure 2-10A).

- **Conventional Charging** (Figure 2-10B) The load is charged with a step input. A large voltage drop across the resistor results in energy dissipation to heat equal to $\frac{1}{2}CV^2$ as calculated in Figure 2-8. In addition $\frac{1}{2}CV^2$ has been transferred to the capacitor.
- **Conventional Discharging** (Figure 2-10C) When discharging the load with a step input, again there is a large voltage drop across the resistor and the $\frac{1}{2}CV^2$ that was stored on the capacitor is dissipated.

Here, the entire CV^2 that began in the voltage supply has been dissipated as heat in the resistor.

A similar procedure is applied to a charge and discharge cycle of adiabatic charg-

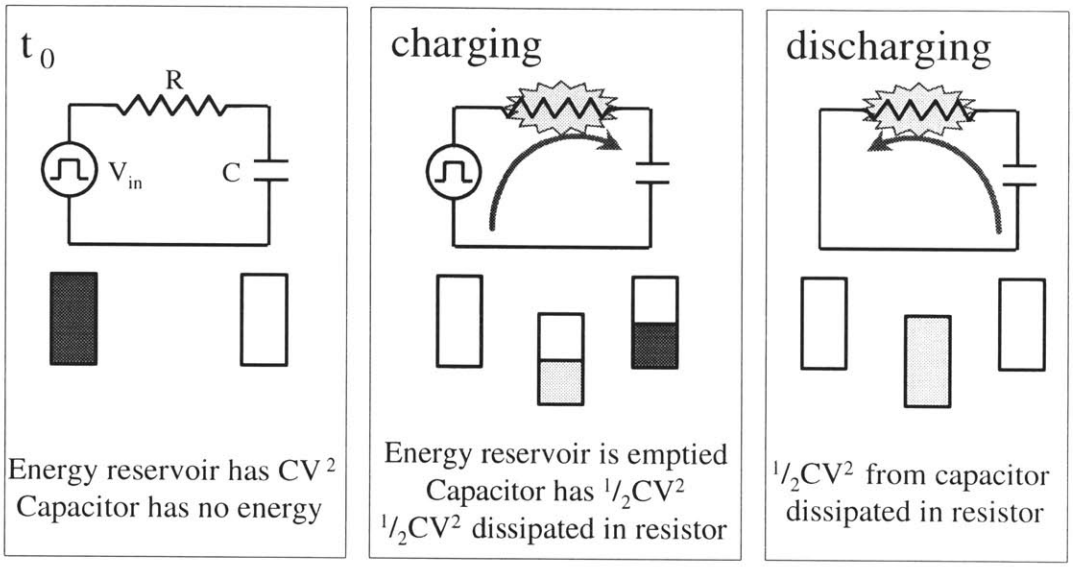

Figure 2-10: Conventional Charging and discharging. Starting at point t_0, an energy reservoir of 0 is associated with the load and an energy reservoir of CV^2 is associated with the supply.

ing. A linear ramp is used for the adiabatic source (Figure 2-11). The ramp transition time, t_c, must be much larger than the RC time constant of the load and switch resistance.

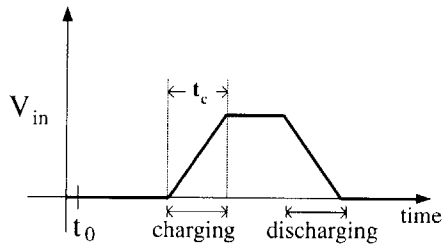

Figure 2-11: Adiabatic charging and discharging waveform. The ramp is the most efficient charging waveform because it is constant current.

As in the previous example, a load, discharged at t_0, is associated with an energy reservoir of 0 and a voltage source is associated with an energy reservoir of CV^2 (Figure 2-12A).

- **Adiabatic Charging** (Figure 2-12B) The load is charged slowly by the ramp. As calculated in Figure 2-8, the dissipation in the resistor is $\frac{RC}{t_c}CV^2$, which is less than the conventional charging case. However, just as in the conventional case, $\frac{1}{2}CV^2$ is transferred to the capacitor. If the ramp transition time is infinitely long, then $\frac{RC}{t_c}$ is zero and only $\frac{1}{2}CV^2$ is drawn from the voltage supply.
- **Adiabatic Discharging** (Figure 2-12C) When discharging the load slowly with a ramp input, again there is a small instantaneous voltage drop across the resistor, so the dissipation in the resistor is $\frac{RC}{t_c}CV^2$. The capacitor began with $\frac{1}{2}CV^2$, and only $\frac{RC}{t_c}CV^2$ is dissipated in the resistor, so $(\frac{1}{2} - \frac{RC}{t_c})CV^2$ is returned to the supply.

A typical voltage supply will shunt any returned energy to ground, dissipating it across some resistance, and rendering the adiabatic discharging no more energy efficient than the conventional case. However, if the supply is a resonant source, it will be able to reclaim the returned energy and use it for subsequent charging. This discharging

method is called energy recovery because the energy transferred to the capacitor is recovered and reused by the supply. If the ramp transition time is infinitely long, then $\frac{RC}{t_c}$ is zero and all of the energy stored on the capacitor is returned to the supply.

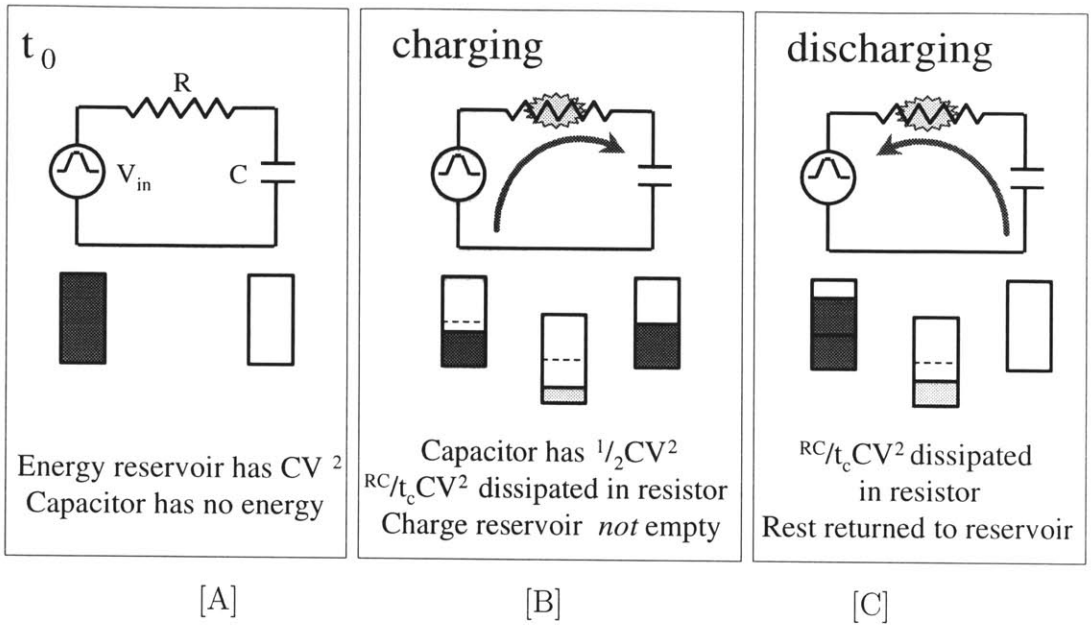

Figure 2-12: Adiabatic charging and discharging. Starting at point t_0, an energy reservoir of 0 is associated with the load, and an energy reservoir of CV^2 is associated with the supply.

Table 2.2 shows the theoretical limits of charging and discharging conventionally and adiabatically. If the load is charged adiabatically with a ramp input and then discharged with a step input, the total energy for the charge/discharge cycle could be half that of the conventional case. Thus, adiabatic charging alone can dissipate only half the power of conventional charging. If the load were charged conventionally with a step input and then discharged adiabatically with a ramp input and energy recovery was performed on the returned energy, the total energy for the charge/discharge cycle could be half that of the conventional case. If the load is charged and discharged adiabatically with ramp inputs and energy recovery is performed, the energy required for the charge/discharge cycle can be zero. In practice, energy savings of 76% for ER over the conventional case have been reported [3].

Charging Method	Discharging Method	Energy per Cycle
Conventional	Conventional	CV^2
Adiabatic	Conventional	$\frac{1}{2}CV^2$
Conventional	Adiabatic	$\frac{1}{2}CV^2$
Adiabatic	Adiabatic	0

Table 2.2: Theoretical Limits for charging and discharging. Purely conventional methods dissipate CV^2, hybrid methods dissipate $\frac{1}{2}CV^2$, and purely adiabatic methods dissipate 0 energy.

Typically, ER has some overhead either in the form of increased circuit complexity or power supply complexity. In addition, the requisite adiabatic charging time is slower than conventional switching, so the system must usually be designed around the adiabatically charged nodes. Thus, ER is often only used when the $\frac{1}{2}CV^2$ of the load is a significant percentage of system energy dissipation. It is often simpler to reduce the voltage, V, or reduce the switched capacitance, C, in order to save power. However when the limits of C and V have been reached (or they are fixed), AC and ER can be powerful tools for reducing the dissipation below $\frac{1}{2}CV^2$.

2.4 Digital to Analog Conversion

One of the main focuses of this work is digital to analog conversion. Digital to analog converter basics are described and several types of DACs are outlined. Then the type of DAC used for the ACERDAC, the sampled ramp DAC, is described.

2.4.1 Digital to Analog Converter Basics

There are eight main components of a DAC transfer curve that are used to rate the static behavior of DACs. Figure 2-13 shows the ideal input/output characteristics of a 3 bit DAC and illustrates some of the quantities below.

- **Resolution** The smallest analog change that can be produced by the DAC.

- **Quantization noise** The uncertainty in digitizing an analog value. Has a value up to $\pm 0.5LSB$.
- **Full scale range** The difference between the maximum and minimum analog voltages produced by the DAC. Denoted S_F.
- **Dynamic range** The ratio of the full scale range to the smallest difference a noiseless DAC can resolve.
- **Signal to noise ratio** Is self-explanatory, and is usually presented as a power ratio in DB.
- **Offset** Quantifies the linear deviation from the ideal transfer curve.
- **Gain** Used with the offset to quantify the linear deviation from the ideal transfer curve.
- **Non-linearity** Can be expressed in two variants, integral or differential. **Integral non-linearity** is the maximum deviation from a straight line drawn from 0 to the full scale range. It is expressed in terms of percent of full scale range or in terms of number of LSBs. **Differential nonlinearity** is the maximum deviation of any of the analog output changes caused by an LSB change from its ideal size of $S_F/2^N$ where S_F is the full scale range and N is the number of bits of precision.

One important dynamic characteristic is the **settling time** which is is the time required for the converter output to respond to a bit change.

There are four basic types of digital to analog converters: current scaling, voltage scaling, charge scaling, sampled ramp and combinations thereof. Each type is defined by the method in which the digital to analog conversion is achieved. Most integrated CMOS DACs use either voltage or charge scaling.

2.4.2 Sampled Ramp DAC Theory of Operation

Sampled ramp DACs are described in more detail than other DAC varieties because they are the DAC type used in this research. The principle behind sampled ramp DACs is to convert time to voltage. As shown in Figure 2-14, an input ramp is sampled at a proper time to achieve the desired voltage.

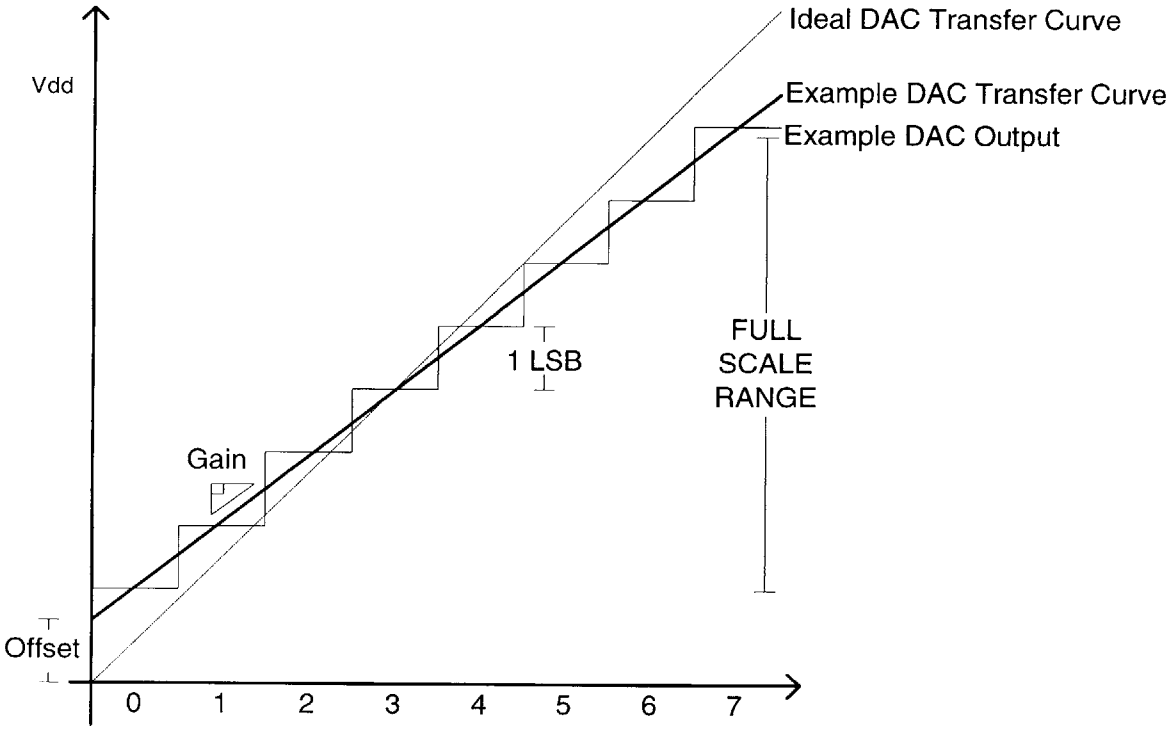

Figure 2-13: The ideal input/output characteristics of a 3 bit DAC. The transfer curve of a non-ideal DAC is also shown. Resolution, Gain, Offset, and Full scale range are illustrated

A notion of time, either local or global, is associated with the ramp so that the DAC can determine when to sample. If the ramp is linear and consistent across temperature and process gradients, a local counter at the DAC can be used along with a global clock to keep track of time. To convert a code, N, the DAC counts N clock cycles and then samples the ramp. If the ramp is not linear or consistent, more complicated schemes can be used to keep time locally, or a global notion of time can be used. In the global scheme, a digital value is broadcast to the DAC which specifies the current value of the ramp. When the broadcast value is equal to the desired value, the DAC samples the ramp. Using a global notion of time requires the DAC to incorporate a minimum of time-keeping logic which simplifies the DAC design and reduces the size.

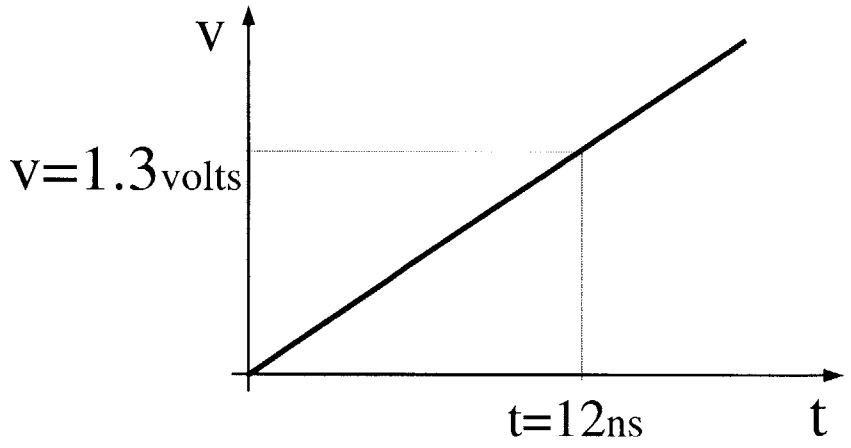

Figure 2-14: Sampled Ramp DAC operation. An input ramp is sampled at a specific time to achieve the desired voltage. A notion of time is kept either locally at the DAC or broadcast globally that allows the DAC to determine when to sample the ramp.

Sampled ramp DACs are used in this research in part because they are particularly suited to display applications. One of these benefits of using sampled ramp DACs in display applications is that sampled ramp DACs are easily parallelized. One ramp source can be distributed to many parallel DACs (Figure 2-15). In addition, the onus of keeping track of sample intervals can be delegated to the ramp driver which further

simplifies the circuitry of each parallel DAC. The major advantage of using a single ramp source is consistency across all parallel components [9]. With the use of a single ramp and a global notion of sampling intervals, the conversion transfer curve will be consistent across all components. Non-linear transfer curves, such as those used to compensate for non-linear display properties, can be implemented and modified by altering the ramp and ramp value synchronization.

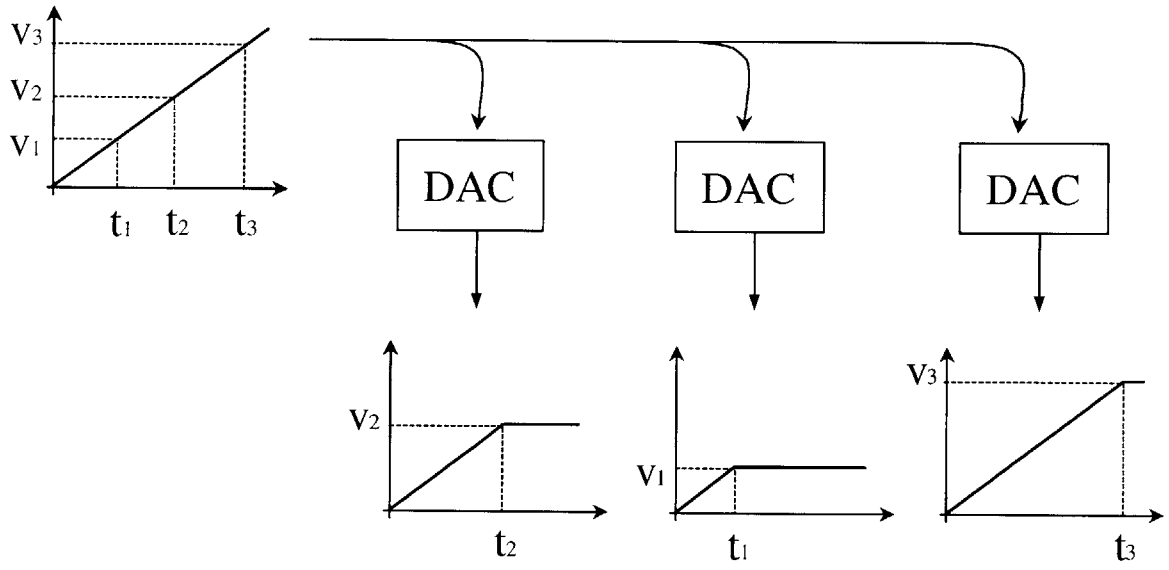

Figure 2-15: Parallel Sampled Ramp DAC. One ramp source drives several parallel DACs. Variations across the DACs are minimized since each uses the same ramp source.

2.4.3 Sampled Ramp DAC with Adiabatic Charging and Energy Recovery

Another reason for using sampled ramp DACs in this research is that they can be combined with adiabatic charging and energy recovery to yield a system that can be low power while maintaining the requisite LC switching voltage of the pixel array. Adiabatic charging is implemented by using the ramp to slowly charge the load.

Energy recovery is performed on the DAC load when the ramp is also used to slowly discharge the load.

In adiabatic mode (AM), the load is attached to the ramp when the ramp is at ground. As the ramp increases, the load is slowly charged by the ramp. When the ramp is at the desired voltage, the load is disconnected from the ramp, and the load holds the proper voltage (Figure 2-16). At the beginning of the next conversion, the ramp and load are reset to GND and the process is repeated.

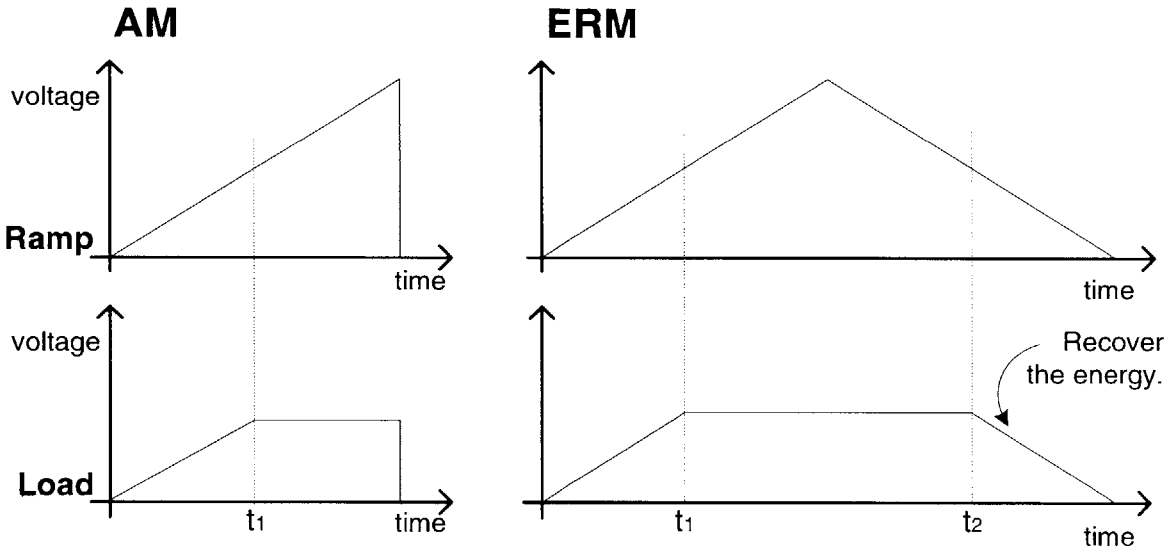

Figure 2-16: Adiabatic Mode and Energy Recovery Mode for a Sampled Ramp DAC. For both modes, the charging cycle is the same. For discharging in adiabatic mode, the ramp and load are reset to GNDat the end of a cycle. In energy recovery mode, a resonant ramp source is used as the ramp and the load is reconnected to the ramp and discharged slowly.

Incorporating energy recovery into this scheme is simple given a ramp source capable of reclaiming energy that is returned to it, such as the Blip circuit [4]. In energy recovery mode (ERM), the load is reconnected to the ramp on the way down when the ramp and the load voltages are equal, so there are no resistive losses as a result of the connection. As the ramp continues to decrease in voltage, the load is discharged adiabatically (Figure 2-16). Any energy that was on the load can be

reclaimed by the ramp source and reused to charge other loads. Adiabatic charging can theoretically reduce the energy of the ramp by half over conventional charging methods. In theory, energy recovery can reduce the ramp energy consumption to zero. However, energy recovery is expected to be 80% more efficient than conventional methods for charging the columns and pixels because of the finite charging time used and inefficiencies of available resonant sources.

2.5 Previous Work

A sampled ramp DAC for driving a TFT display has been developed by IBM [9] (Figure 2-17). Here, the analog sampled ramp is presented to the gate of transistor TM1. Figure 2-18 shows the waveforms for the internal nodes of this circuit where it can be seen that the column wire (data line) is charged instantaneously to match the value stored on the transistor. The load that is charged using the sampled ramp method is the small capacitor CO1. The larger capacitive load, the column wire, is charged instantaneously. Thus, the low power advantages of slow charging are applied only to the small load of sampling transistors and capacitors.

2.6 Conventional TFT D/A's

Conventional flat panel displays are fabricated in a process using thin film transistors (TFTs). These TFTs have poor performance, and the DACs designed for flat panel displays are engineered around this poor performance constraint. Therefore, many of the TFT-based DACs are charge sharing DACs. The charge sharing DAC uses ratioed capacitors to produce an analog voltage (Figure 2-19). This design requires few active transistors, but suffers from the difficulty of requiring carefully matched capacitor values across the display. They also have high power due to high operating voltages and large area. The DAC in Figure 2-19 operates at 15 volts and the capacitor array for each column driver DAC measures more than 21 μm by 1600 μm [6].

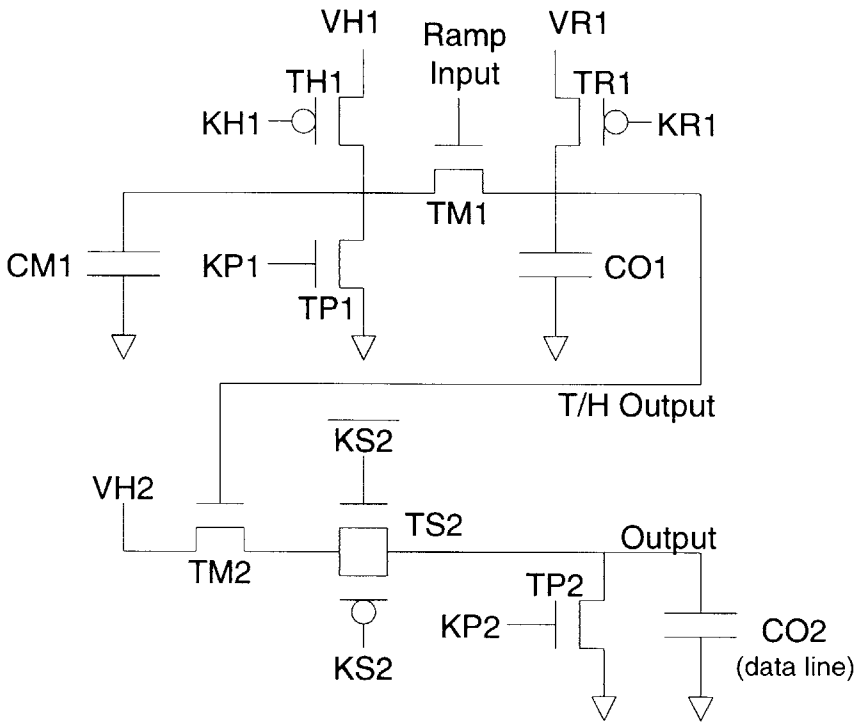

Figure 2-17: A Sampled Ramp DAC for displays. The ramp is sampled onto the gate of transistor, TM1. The column wire is charged through current source, VH2. (Taken from [9].)

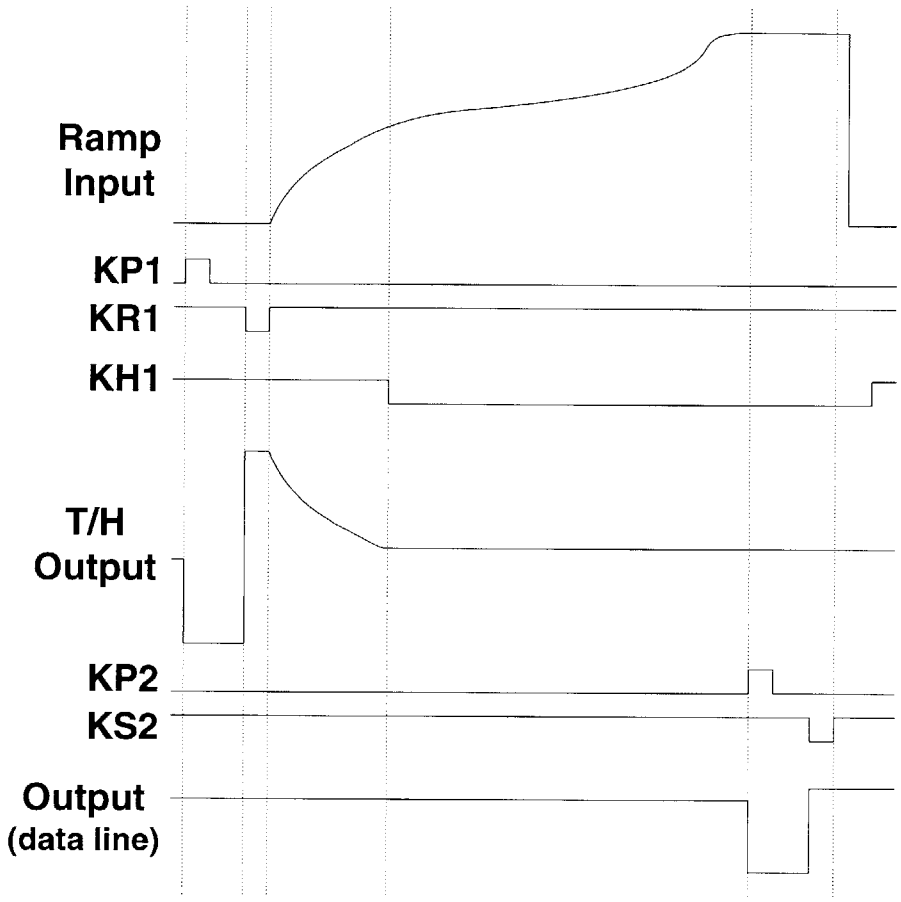

Figure 2-18: Waveforms for Sampled Ramp Above. Column wire is charged instantaneously instead of slowly resulting in additional power dissipation. (Taken from [9].)

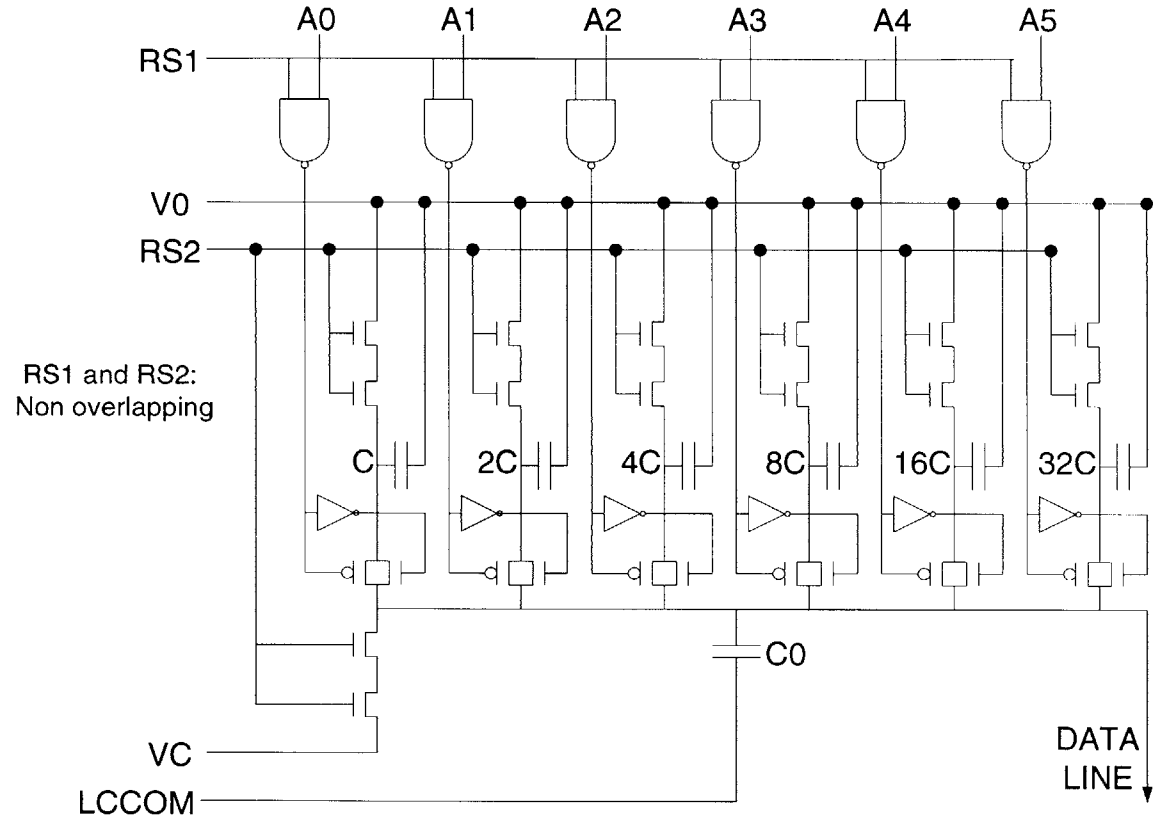

Figure 2-19: Charge Scaling DAC. Typically used in TFT LC displays. (Taken from [6])

Chapter 3

ACERDAC: Overview

The ACERDAC chip uses a combination of the adiabatic sampled ramp DAC and the DAC-per-column architecture. This combination yields both the benefits of the sampled ramp DAC for display applications and the low power benefits of AC and ER. Figure 3-1 shows an ACERDAC display architecture block diagram.

3.1 ACERDAC Operation

The ACERDAC architecture utilizes a digital distribution path to distribute the digital pixel values from the chip input to the column DACs and an analog path to distribute the analog pixel values from the DAC to the pixels.

The digital distribution path consists of a data bus that transports the digital pixel data from the chip input to the column latches. The *capture latches* capture each value in turn. After each column has latched a value, an entire row of pixel data has been loaded and is now ready to be converted to analog values by the DACs.

Pipelining corrects the inefficiencies that arise in this architecture when the conversion circuitry is idle while the broadcast circuitry operates and vice versa. Pipelining utilizes the circuitry more efficiently and allows for conversion to take place slowly over an entire line time. With the addition of a second set of latches, the system captures a row of broadcast data while the conversion of the previously broadcast row takes place. The *preserve latches* store the previously captured data for the con-

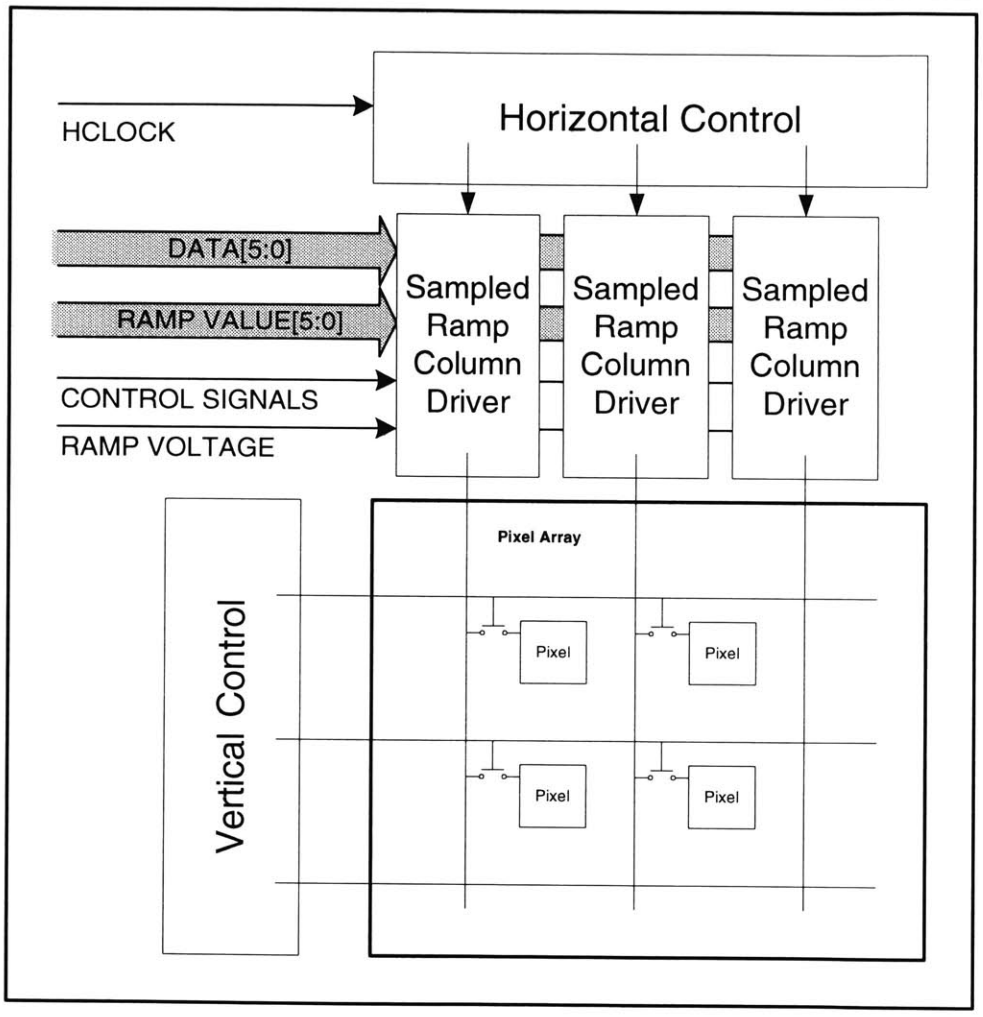

Figure 3-1: Block diagram of an ACERDAC display. There is one sampled ramp DAC column driver for each column of pixels.

version circuitry to operate on. A clock transfers one row of data from the capture registers to the preserver registers once per line period (Figure 3-3).

After a row of data is captured and transfered to the preserve latches, the digital to analog conversion process begins. One row of pixels is connected to their respective column wires and each column wire is connected to the ramp wire via analog switches. As the ramp source rises slowly, each of the columns track the ramp (Figure 3-2). A digital ramp value is broadcast along with, and synchronized to, the analog ramp. A digital comparator at each column compares the globally distributed ramp value to the value stored locally in the preserve latch. When the digital ramp value exceeds the pixel value stored in the preserve latch, the column is disconnected from the ramp wire. The column wire capacitance acts as a storage device of the analog value for the pixel. When the ramp reaches the maximum voltage, all the column wires and one row of pixels hold the desired analog values. The pixels are then disconnected from the column wires leaving the proper voltage on each pixel.

The energy remaining on the column wires is recovered as the ramp voltage returns to GND and the ramp value counts down to zero. Each column wire is reconnected to the ramp when the ramp value equals the stored pixel value. In this fashion, the voltages on the column wires are slowly ramped down to GND. Despite the voltage change on the column wire, the pixel voltage remains at the correct value because it is isolated from the column wire by the analog switch.

When the conversion process is completed for one row, the row enable is incremented, new data waiting in the capture latches is loaded into the preserve latches, and the conversion process repeats.

Figure 3-3 shows the architecture block diagram of the ACERDAC column driver. For each column, there are two sets of latches, a digital comparator, and a column switch connecting the column wire to the global ramp wire. There is a global column enable circuit which controls the digital distribution system, and pixel switches connect a column of pixels to the column wire.

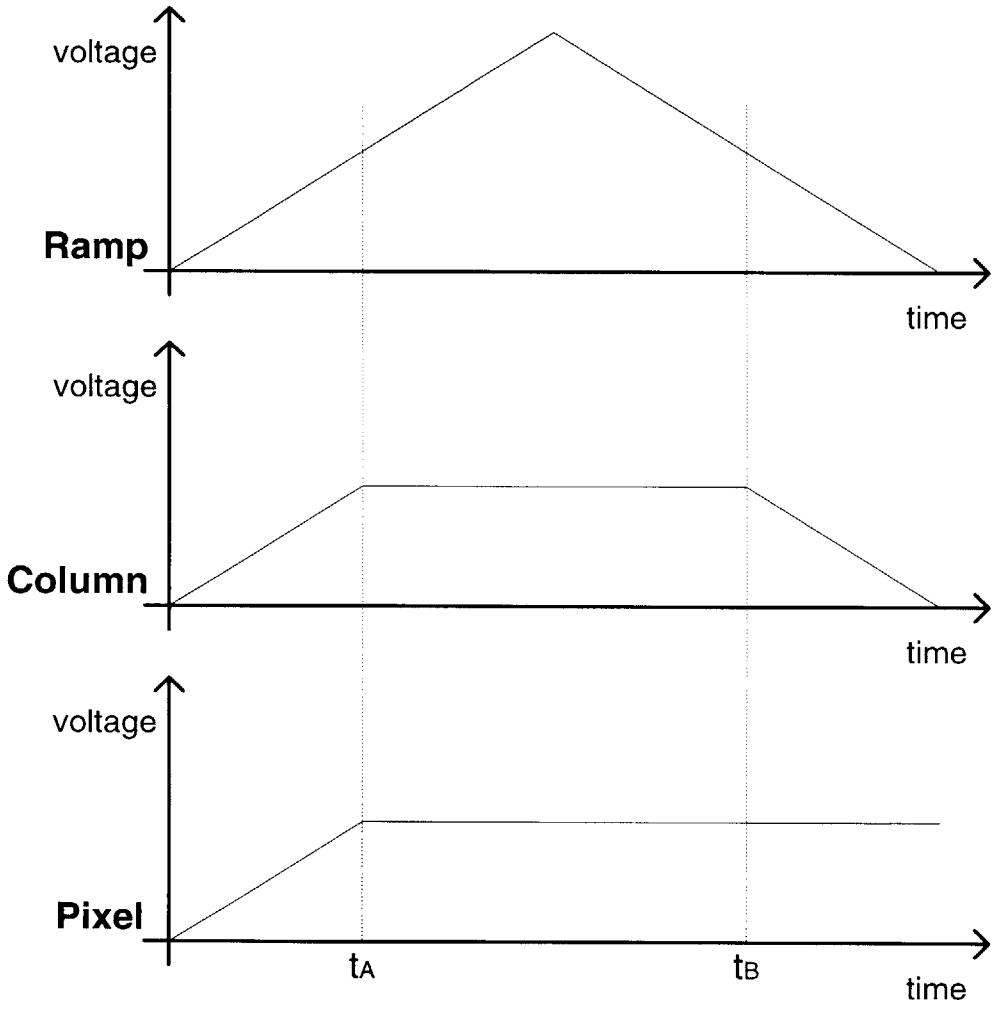

Figure 3-2: Waveforms for the ramp, column and pixel nodes for a sample conversion. The column samples the ramp on the way up and reconnects to the ramp on the way down. The pixel holds the sampled value.

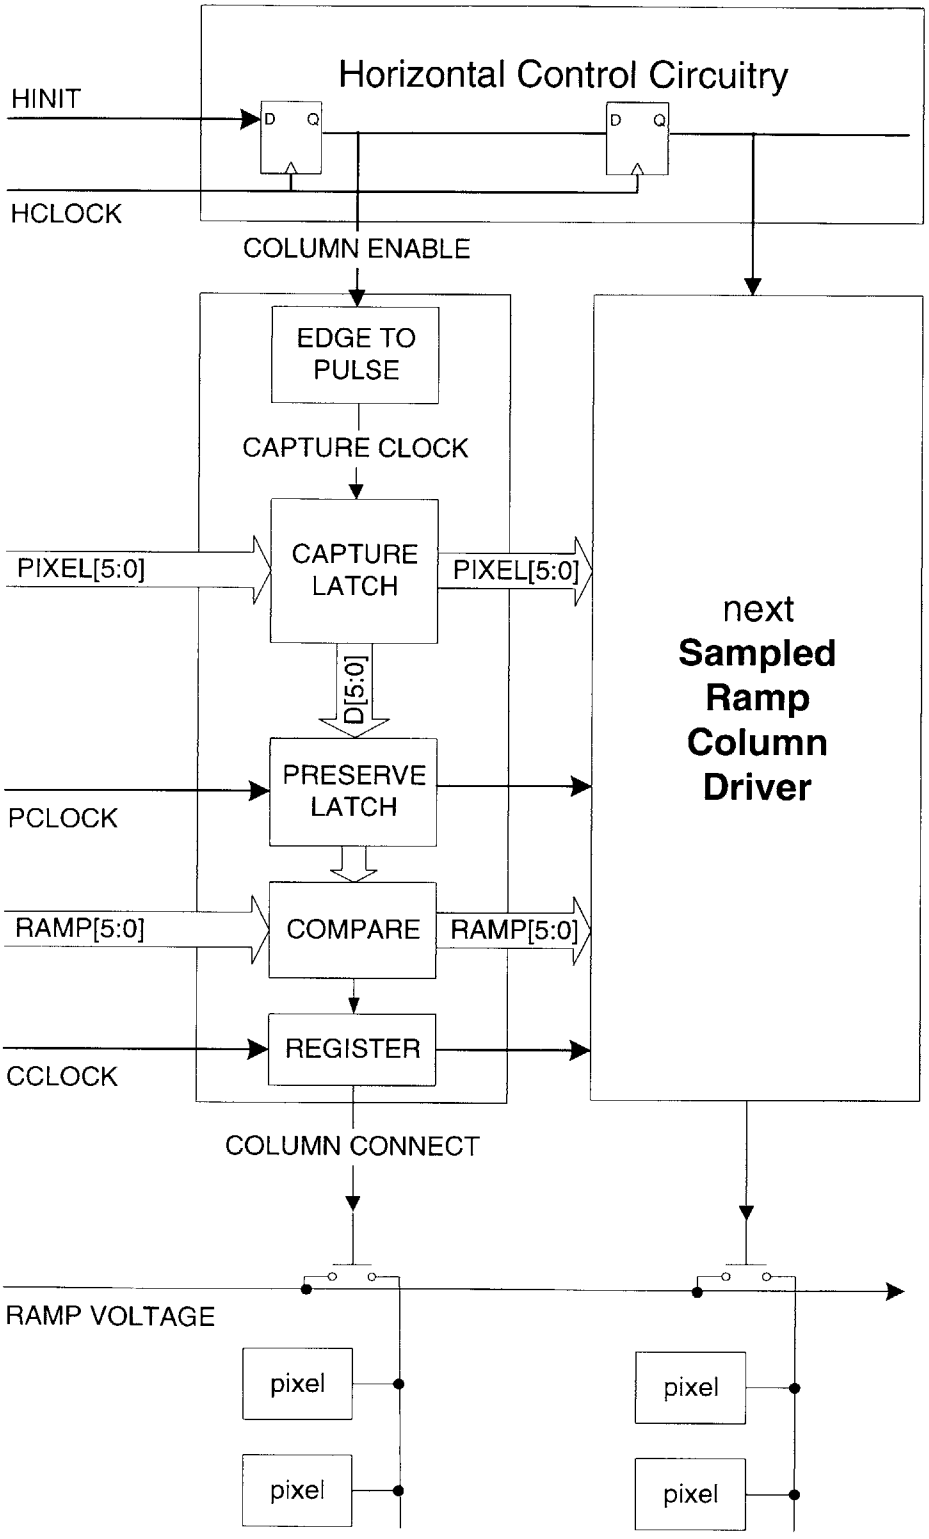

Figure 3-3: Column driver and Sampled ramp DAC circuitry. Two sets of latches, a comparator, an edge to pulse converter, a register, and a switch make for a compact column driver size.

3.2 Performance Advantages

In addition to the aforementioned advantages of using the sampled ramp DAC architecture for display applications, the ACERDAC architecture has additional scalability, low-power, and area advantages. There is a small amount of circuitry at each column (Figure 3-3) which saves silicon area, but more importantly, allows the column driver circuitry to fit within the pixel pitch of the display.

Using the broadcast method of data distribution saves power in the case where images are highly correlated across the rows. For the suite of test images analyzed in §A, horizontal correlation resulted in activity factors at least two times lower than the worst case.

Distributing the digital ramp value globally allows non-linear transfer curves to be implemented by changing the ramp to ramp value synchronization and without modifying the column circuitry. In this way, LC non-linearity correction can easily be applied and modified. Also, some of the LC inversion schemes described in §2.1.2 can be implemented simply by changing the ramp synchronization.

The ACERDAC architecture is easily scalable to larger display resolutions because it utilizes a column-based design. When scaling display resolution, each additional column of pixels is accompanied with an additional column driver. The leaf-cell layout and the layout methodology is not modified.

Aside from the many advantages of the ACERDAC architecture listed above, another advantage of this architecture is low power digital to analog conversion for the display. While AC and ER are touted as the main low power techniques, they are not the dominant cause for power reduction in the system. Pipelining and transition frequency reduction of large wires in the system, initially implemented to increase the efficiency of AC and ER, result in more significant system power reduction.

Although the ACERDAC architecture requires a small amount of circuitry at each column, this amount is still larger than that of the analog display architecture described in §2.1.3. The additional circuitry in the ACERDAC performs the same function as one monolithic DAC would for the analog input display. In order for this

additional circuitry overhead to make sense, it must serve to lower the power of the data conversion and the the ramp wire which is the video wire equivalent.

The analog ramp is the logical video wire replacement in the analog display. The ramp powers a significant portion of the capacitance of the chip including all of the columns and all of the pixels. Examining the power equation, $\alpha f C V^2$ (Equation 2.1), illustrates the difficulty associated with designing an architecture around the ramp wire; C is large, and V is large and fixed which forces α and f to be reduced in order to achieve low power. For the video wire of the analog display, f is high ($f = F_p$) and α is data dependent, but may be small due to image correlation. Nonetheless, large f, C, and V translate into significant power consumption.

By operating several DACs in parallel, it is ensured that the ramp need only transition once every line time (transition is loosely defined as charging from GND to VDD and discharging back to GND). In effect, the frequency has been reduced from F_p to F_r – a reduction factor of over a hundred (depending on the display resolution). The column wires still transition the same number of times (once per line time), but the time over which these wires are charged has increased from $1/F_p$ to $1/F_r$ resulting in lower power.

The last power equation component to be analyzed is α. As previously mentioned, and as will be formally discussed in §A, images tend to be highly correlated in both the horizontal and vertical dimensions. In the ACERDAC architecture, the columns are reset to zero between lines which may eliminate the vertical correlation advantages. However, this decorrelation may not have a significant effect if some LC inversion schemes are being used. In the pixel and row inversion schemes, all vertical correlation is already eliminated by the fact that neighboring vertical pixels are drawn with alternating polarities. In the frame inversion scheme, vertically-neighboring pixels are drawn with the same polarity so correlations are preserved. As a result, a conventional display exhibits a smaller vertical activity factor. However, the use of AC minimizes the penalty of having a larger activity factor in the ACERDAC and thus the column wire dissipation is less than the conventional display even when operating in frame inversion mode.

Chapter 4

Implementation

Several constraints governed the implementation of the ACERDAC architecture. The two primary goals of the implementation are to verify the proposed advantages of this architecture and to accurately predict the performance of this architecture for different display resolutions and bit precisions. While the main performance consideration is low power, ensuring the basic operation of the chip is essential. Therefore, only the key architectural features and several low-risk optimizations were implemented. The implementation brings most crucial signals off-chip so that maximal flexibility is achieved. (e.g., different drive schemes can be evaluated, and on-chip components can be isolated for independent study.)

Considerable layout constraints drive the architecture implementation. In particular, the pixel pitch defines the column circuitry width. The pixel pitch for this implementation is 25 μm, but designs as low as 10 μm may be required for future displays. The narrow layout constraints require that few wires travel the length of the column circuitry. Thus, each block of circuitry must rely on minimal information propagation to and from its neighbors.

This ACERDAC implementation realizes a 160 x 120 display with 6 bits of grayscale precision (Die Photo in Figure 4-1). It was fabricated in a 3-metal, 0.8 μm Hewlett-Packard process that is both 5 volt compliant and relatively inexpensive. The resolution was chosen because it is the smallest of the "standard" display resolutions and therefore the most inexpensive to fabricate.

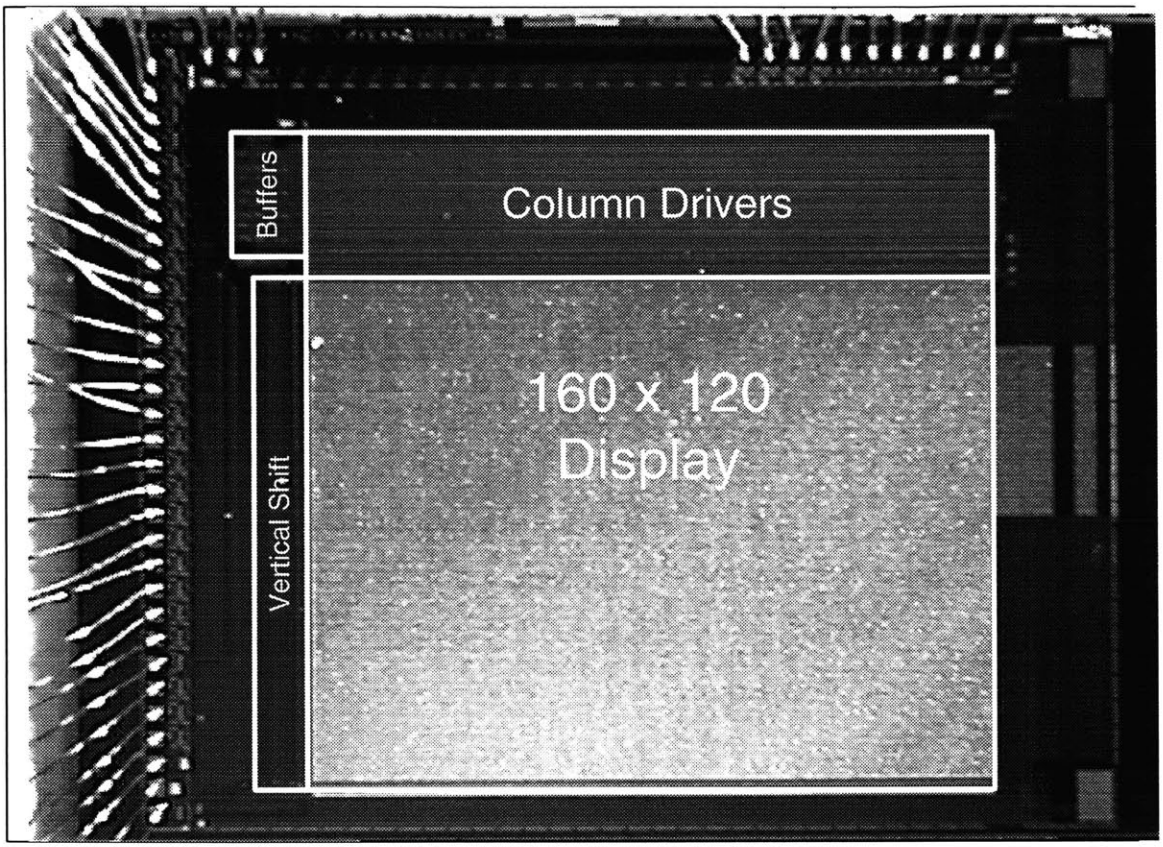

Figure 4-1: Die Photo of 160 by 120 pixel ACERDAC. The active display area is marked. The Die measures 5.5 mm by 4.5 mm.

4.1 Wire Analysis

The charging and discharging of high capacitance wires accounts for the majority of the power consumption of an integrated display system. Hence, overall system power is reduced by careful analysis and optimization of each power-intensive wire. A key characteristic that determines wire power is transition frequency. Assuming C and V are constant, wires that transition more frequently will consume more power. The wires that transition at frame rates (F_f) contribute such a small amount to the total system power consumption that no effort is spent in optimizing them. Much of the engineering effort should be spent on the wires that transition at pixel rates (F_p) because they transition most frequently.

Wire power comes not only from the dynamic power, but also from short circuit power in circuitry driven by the wire. If the circuit on the receiving end of the wire will dissipate short circuit power when the wire is at an intermediate voltage, sharp transitions must be guaranteed for that wire.

These sources of power lead to a simple, two-criteria wire classification system. Each digital wire in the system is classified according to transition frequency and potential for short circuit power consumption. A classification of the different wire types and several low power methods that are applicable to each type is given in Table 4.1.

Frequency	Short Circuit	Type	Low Power Techniques
fast	yes	A	Supply Scaling
fast	no	B	Low Swing, Energy Recovery
slow	yes	C	Guarantee Quick Transitions
slow	no	D	Not Necessary

Table 4.1: The four digital wire types and possible low power methods. Type A wires consume the most power; Type D consume the least.

Type A wires consume the most power since they switch frequently and can cause short circuit current when left at intermediate voltages. A type A wire example

is a one that transitions at F_p and connects to an inverter input. Short circuit power is reduced by ensuring that signals run rail to rail. For these wires, reducing the supply voltage, and thus the rail to rail voltage swing, reduces overall power consumption. This reduction may be difficult because lowering the voltage supply lowers the circuitry speed. If A is required to be fast, then the circuitry receiving it must have similar constraints. To reduce short circuit current, the transitions on Type A wires should be sharp. This minimizes the time spent at an intermediate voltage which minimizes short circuit power. Other low power techniques such as reducing the switched capacitance, or coding the data for minimum transition activity can be applied. However, the best way to engineer a Type A wire for low power operation is to convert it to a Type B, C, or D wire.

Type B wires transition frequently, but can be left at intermediate voltages without dissipating short circuit current. A type B wire example is one that transitions at F_p, but connects to the gate of a pass transistor. The voltage swing on Type B wires can usually be lowered without any modifications to the receiving circuitry because they already tolerate intermediate voltages. Sometimes, only minor modifications are needed. Since these wires can tolerate intermediate voltages, they are also prime candidates for adiabatic charging and energy recovery.

Type C wires transition infrequently, but can dissipate short circuit current. A type C wire example is one that transitions at F_r and connects to an inverter input. Often, type C wires are ignored in power analysis since their dynamic power is low, but the short circuit power of these wire can dominate if they are engineered poorly. Guaranteeing sharp transitions on the wire will alleviate short circuit current. Then, type C wires usually do not contribute significantly to the system power dissipation.

Type D wires consume the least power as they transition infrequently and do not dissipate short circuit power. A type D wire example is one that transitions at F_r and connects to the gate of a pass transistor. Both conventional low power techniques (e.g., reducing switched capacitance or reducing the voltage swing/supply voltage) and AC/ER techniques can be applied. However, Type D wires are not where the design effort should be spent as they contribute only a small percentage of the total

system power consumption.

Type A and C wires will cause short circuit current dissipation when at an intermediate voltage. Short circuit current for these wires is reduced by guaranteeing sharp transitions on the wires. One method of accelerating transitions on a wire is to decrease its RC delay. Resistance decreases with width while area capacitance increases with width, so for very wide wires where the area capacitance dominates, the RC delay will not scale with wire width. However, in narrow wires, fringe capacitance contributes significantly to the total capacitance. Therefore, by slightly increasing the width of a narrow wire, the RC delay can be reduced. Figure 4-2 shows the RC delay of a typical wire versus width. The RC delay of the wire can be significantly improved by increasing the width up to widths of about 4 μm. For widths greater than 4 μm, the area capacitance begins to dominate and increasing the width does not benefit the RC delay. One must be careful when widening wires because power is proportional to capacitance. Hence, wire-widening reduces the short circuit power dissipation, but increases the dynamic power dissipation. The two constraints must be carefully weighed.

Each digital wire in the ACERDAC system is classified using this Type A, B, C and D wire model and evaluated to reduce power. The system wires and their type are identified in Table 4.2. Each wire and any low power method used in this implementation are described in detail along with its surrounding circuitry in the following sections.

Wire	Type
HCLOCK	A
DATA	B
PCLOCK	D
RVALUE	B
CCLOCK	A

Table 4.2: The ACERDAC system wires. Each major wire in the system is classified as Type A, B, C, or D. Type A wires consume the most power; type D wires consume the least.

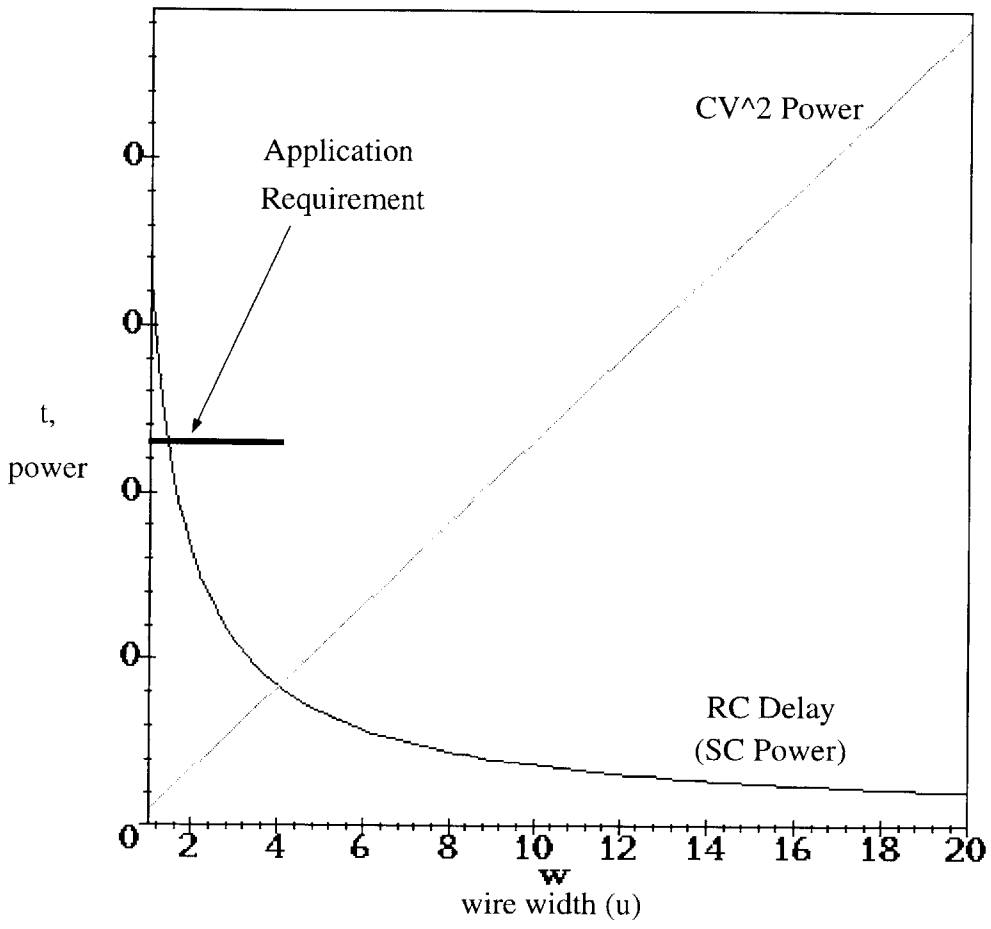

Figure 4-2: RC delay as a function of wire width. Because of fringe capacitance dominance, increasing the wire width for narrow wires can decrease the RC delay and thus decrease the short circuit current.

4.2 Floor Planning Implications

Displays have a very regular structure in which pixel cells are arrayed horizontally and vertically to achieve the desired resolution. Control circuitry associated with each column or row is aligned with the first column or row and arrayed horizontally with the pixels. The chip is described in a hierarchy where large components are divided into several smaller components recursively until each component consists of a simple circuit block. In this way, the complexity of the chip is reduced the complexity of a small number of circuit elements which are arrayed to achieve the proper display resolution and functionality.

The chip core is made up of three components: the horizontal control logic, the vertical control logic and the pixel array. The vertical control is at the left of the pixel array and the horizontal control is on top of the pixel array (Figure 4-3). The pixel array is an array of pixel cells with the requisite number of rows and columns. The vertical control logic is an array of identical row drivers (one for each row) along with some global wire buffers. The horizontal control logic includes an array of column drivers (one for each column), and some global wire buffers. Each column driver includes column enable circuitry, an ACERDAC column driver, and a column switch. Each ACERDAC column driver includes a six-bit capture latch, a six-bit preserve latch, a six-bit digital comparator, and a register (Figure 4-4).

One bit of receive latch, one bit of preserve latch, and one bit of digital comparator are combined into a single block, the receive/compare block (RCB), which is stacked to achieve the desired bit precision for the column driver (Figure 4-4). This design can accommodate future pixel widths as narrow as 10 μm by stacking the two latches vertically rather than horizontally within the bitslice (Figure 4-5). Each receive/compare block measures 25 μm by 76 μm; the assembled column driver measures 25 μm by 900 μm. If any layout change was required in the RCB, the change could be made manually to the block itself, and the rest of the blocks can be arrayed relative to the modified block.

Clustering circuitry in the RCB minimizes the number of wires that need to travel

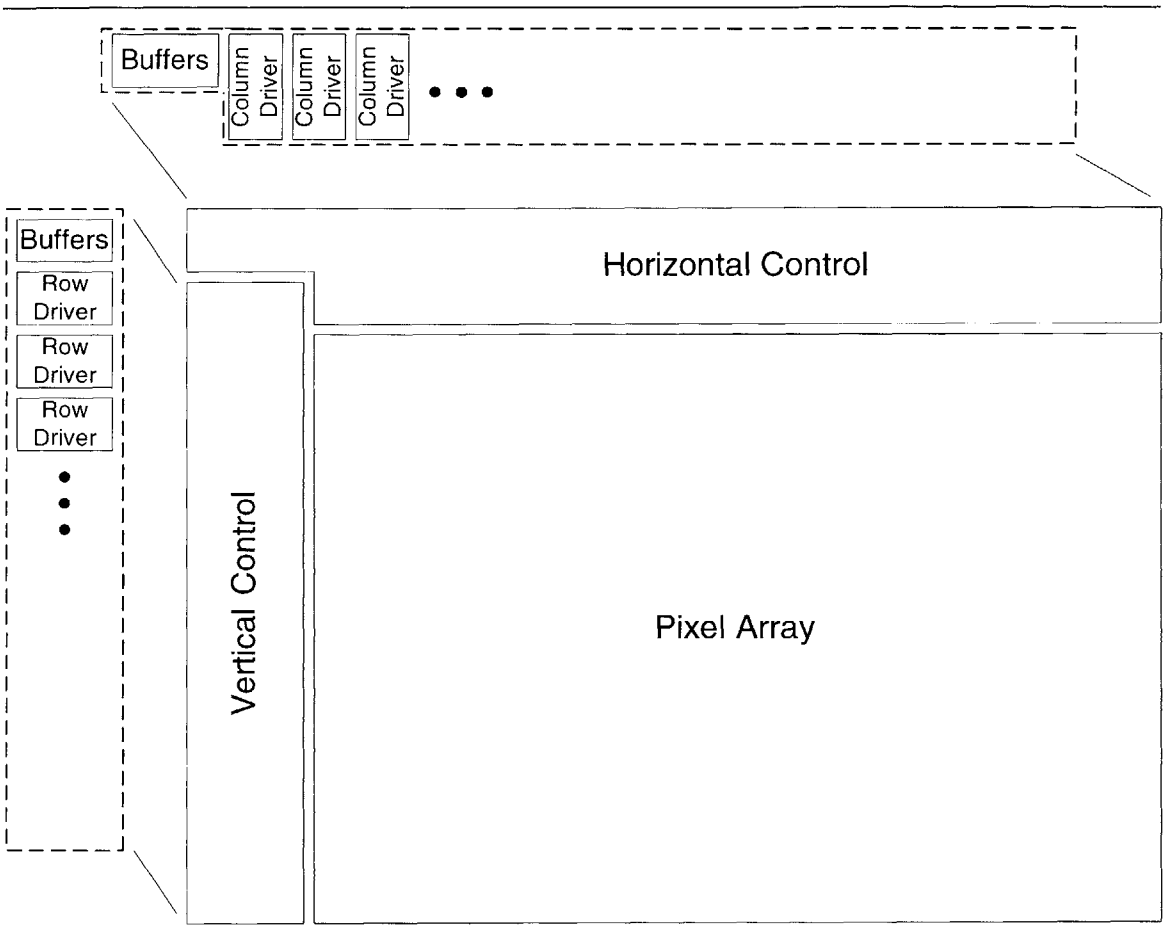

Figure 4-3: Chip Core Floor Planning. The chip core consists of three elements: pixel array, vertical control logic, and horizontal control logic

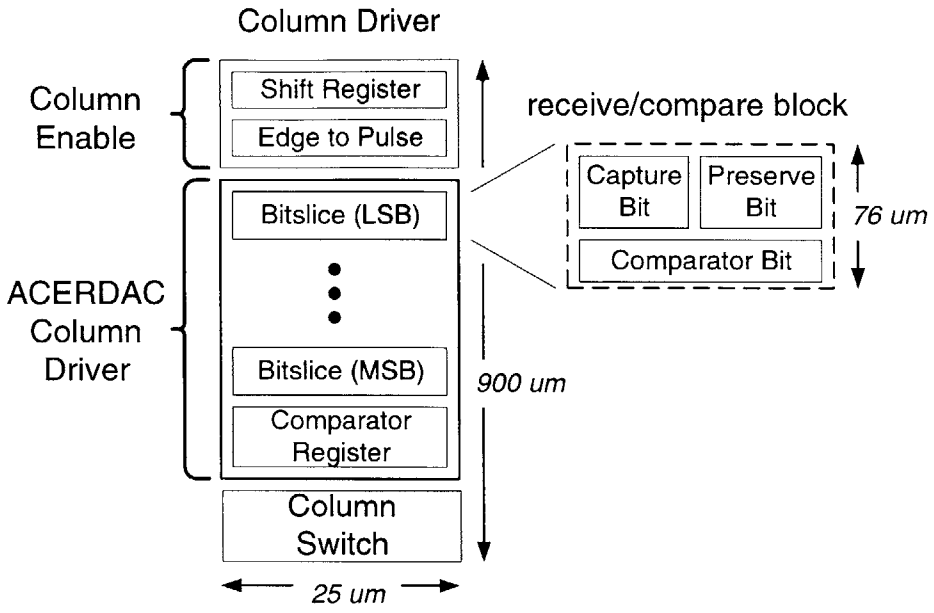

Figure 4-4: Column Logic Floor Planning. Each column consists of column enable circuitry, an ACERDAC column driver, and a column switch. The ACERDAC column driver includes a six-bit capture latch, a six-bit preserve latch, a six-bit digital comparator, and a register.

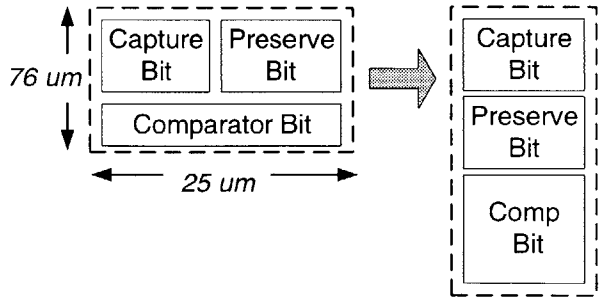

Figure 4-5: Future receive/compare block. Accommodates narrower pixel pitches by stacking the latches vertically rather than horizontally.

long distances within the column driver, allowing the circuitry to fit within the pixel pitch. Each bitslice requires two intracolumn inputs: the column clock signal from the column enable circuitry, and the carry out of the comparator in the previous bitslice. Thus, only two wires travel vertically between bitslices (Figure 4-8). The 25 μm pitch allows only three vertical wires to travel between RCBs: clock, carry out, and VDD.

4.2.1 CAD Support for Hierarchical Layout

To facilitate this hierarchical layout procedure and to reduce human error, the DISplay COmpiler (DISCO), developed by The MicroDisplay Corporation, is used which accepts a chip description and automatically lays out the pixels and surrounding circuitry to ensure proper alignment. The tasks done manually are the design and layout of each circuit block so that it can be abutted directly to its neighboring blocks by DISCO.

Designing the circuitry with the power of DISCO in mind can greatly reduce the manual layout requirements, can greatly facilitate layout changes, and can simplify verification. The use of DISCO promotes the use of hierarchy to describe the chip structure and the design of simple circuit blocks which can be arrayed to achieve full functionality.

4.3 Digital Data Distribution

Conventional display drivers use a shift register distribution system which dissipates power even when the incoming data is the same as the current data. Since images are highly correlated, the shift register method results in unnecessary power dissipation. The broadcast method of data distribution does not dissipate power when incoming data is the same as the current data and is therefore chosen for the ACERDAC. The ACERDAC data distribution system is a 6 bit bus that runs horizontally across the length of the display. The wires are driven by low-swing buffers in order to help to reduce the bus power consumption. Each column has a receive circuit which amplifies

data from the bus to full swing and stores it for use in the digital to analog converter.

4.3.1 Receiver/Latch

The receiver/latch is designed to have a compact layout and low power in keeping with the goals of the ACERDAC implementation. A traditional clocked-inverter latch design is inappropriate in this application for two reasons. First, this type of latch requires at least 13 transistors to implement resulting in a large physical area. Second, this type of latch dissipates power every time it is clocked regardless of whether the data is changing or not. Also, the clocked-inverter latch dissipates short circuit power when an intermediate voltage appears on the input. A more power and space efficient latch is a constant-read SRAM latch. The SRAM cell requires just 6 transistors to implement, and only consumes power when it is switching state. Two NFET pass gates couple the dual rail bus wires to the cross coupled inverters of the SRAM latch (Figure 4-6).

Normally, SRAM cells are difficult to design because there is a small design space in which the SRAM cell is both readable and writable. A constant-read SRAM can be optimized for the write to be both reliable and take little power. A constant-read SRAM is one where the SRAM cell internal state is buffered and constantly available for read. There is no read-cycle or read-enable signal as with standard arrays of SRAM cells.

Different write styles were employed for the capture register and the pipeline register to simplify interactions with their inputs. The capture register receives its data from a large global differential low swing bus when two pass-NFETs pass the low swing signal from the data bus to the SRAM internal nodes. The low swing signals are level-restored by the cross-coupled inverters of the SRAM. The source/drain load presented to the global digital distribution bus by the pass-NFETS is lower than a gate capacitance load. Therefore, the bus power is lower because the data bus capacitive load is smaller than if gates were attached to the bus. The number of global bus wires can be reduced further in future implementations by utilizing a single-ended write implementation, which will reduce the capacitance by a factor of two.

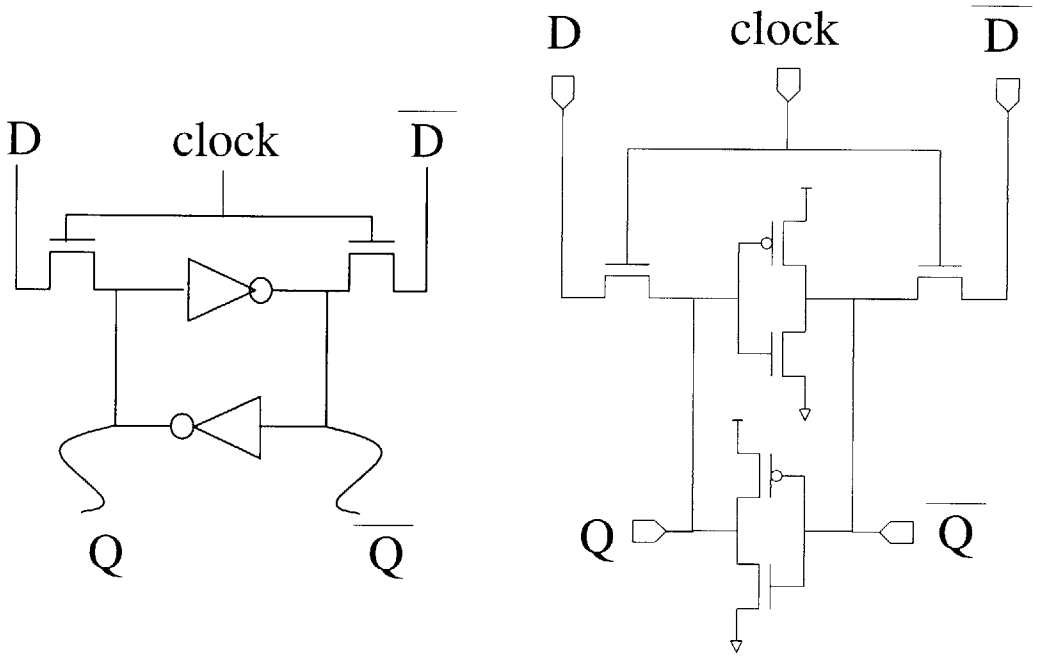

Figure 4-6: Capture Latch Schematic. SRAM style latches use only 6 transistors for a compact layout. Source/drain bus connections reduce the data bus capacitance.

Unlike the capture latch, the preserve latch is physically close to the circuit supplying its inputs. The capture latch outputs are the preserve latch inputs. Since the capture latch output nodes could be corrupted if they were accidentally driven, the preserve latch much be carefully designed not to disturb its inputs. Therefore, the data inputs to the preserve latch are presented to gates of transistors (Figure 4-7). When the SRAM cell is enabled, the inputs to the SRAM cell act to pull down one side of the cell to ground. This sets the SRAM cell to the correct value. The capacitive load of the gate input to the preserve latch does not greatly increase the power of their driving circuitry, as each capture latch switches once per line time.

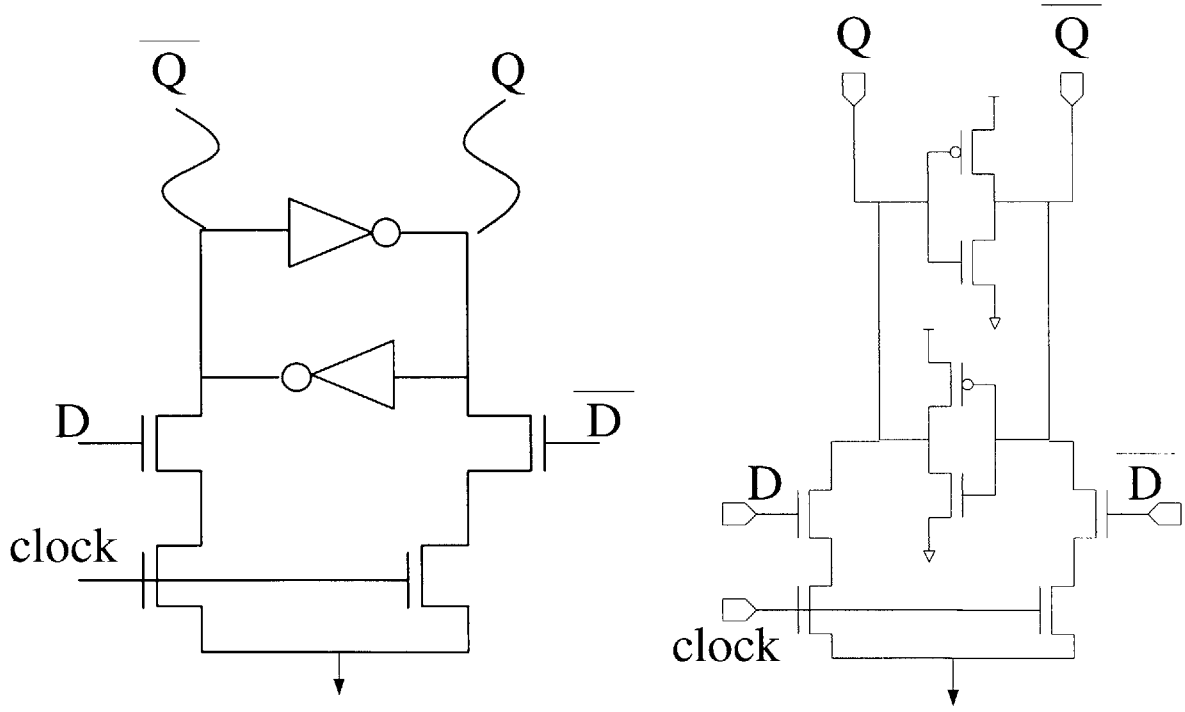

Figure 4-7: Preserve Latch Schematic. Data inputs are presented to gates so that the data values are not disturbed by the preserve latch circuitry.

4.3.2 Wires

Figure 4-8 shows how the capture latch and the preserve latch fit together in the ACERDAC system and identifies the wires which connect to this circuitry. The PCLOCK transitions once per line time to transfer data from the capture latches to the preserve latches. It is a Type D wire since it connects the enable gate on the preserve latch. This enable gate acts to pull down one side of an SRAM cell in order to set the state. The enable gate need not transition full swing in order to flip the SRAM cell.

The digital DATA bus is composed of Type B wires. The digital data transitions at pixel rates, but the SRAM style latch receiving the data will not dissipate short circuit current (other than the short circuit current while the cell is being flipped) while the SRAM clock is disabled. Since the digital data represents an image, it is often highly correlated, so the transition frequency, α, is generally low. The switched capacitance is lowered by connecting the SRAM latches through a source/drain instead of a gate which has higher capacitance. Low swing drivers are also used to reduce the bus power. The data is distributed dual rail to allow the correct receiver operation.

4.3.3 Low Swing Buffering

Large buffers take the digital data from the chip input pads and regenerate them to be driven across the chip. To reduce the bus power consumption, low voltage swings are used. A buffer with an NFET in the pull-up path and a PFET in the pull-down path (Figure 4-8) limits the swing on the bus from V_{tp} to $\mathrm{V_{DD}} - V_{tn}$.

4.3.4 Column Enable

A horizontal shift register is used to sequentially enable the column capture latches as described in §2.1.3. A pixel clock, HCLOCK, synchronized to the data advances the enable bit along the shift register. An edge to pulse converter (Figure 4-9) is used which takes the shift register output and produces a short pulse enable signal. If the SRAM internal nodes remained connected to the bus for the entire HCLOCK

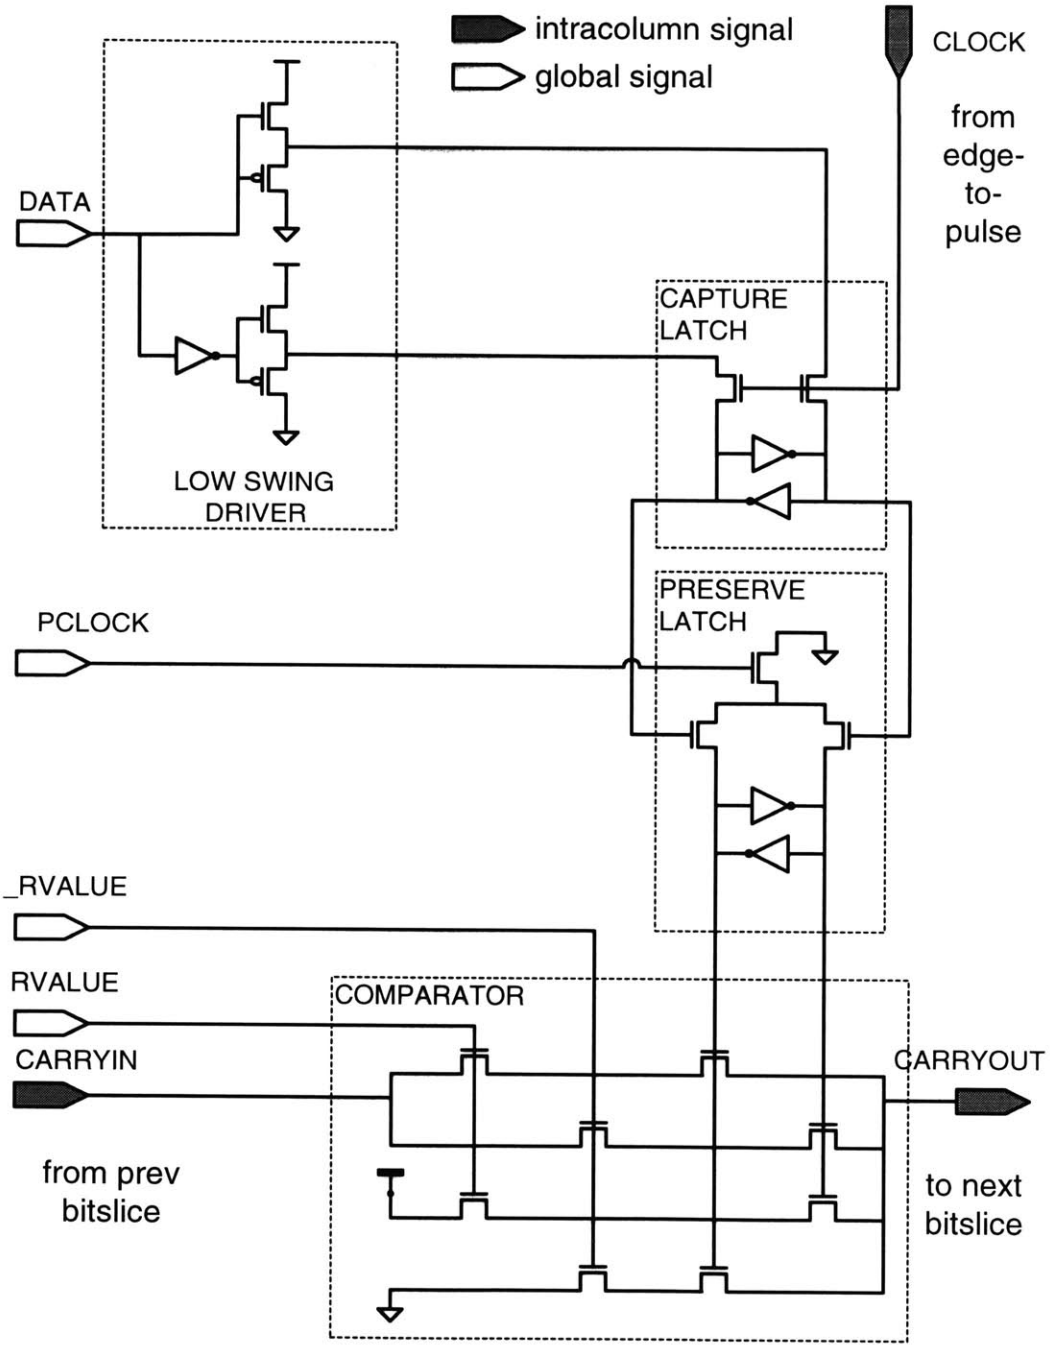

Figure 4-8: Data Receive circuitry. Capture latches connect to the global data bus, and preserve latches connect to the digital comparator. Each is implemented using constant-read SRAM latches, but different write styles are used.

period, they would begin to power the bus, attempting to bring the bus to full swing, and causing unnecessary power dissipation. The edge to pulse converter shortens the enable pulse and allows the SRAM internal nodes to be connected to the bus only long enough for the SRAM to switch state.

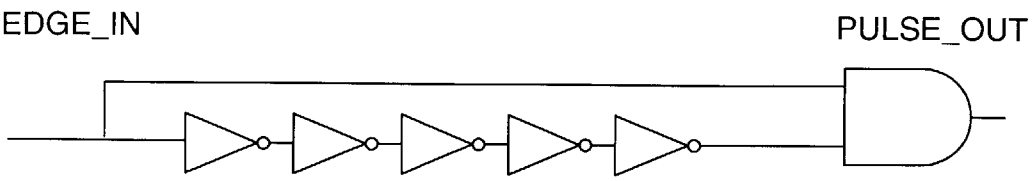

Figure 4-9: Edge to pulse converter. Reduces the time that the capture latches are connected to the data bus. This serves to limit the latch power consumption.

The HCLOCK wire is Type A because it transitions at pixel rates, and the receiving shift registers will dissipate short circuit power when HCLOCK is at an intermediate voltage. In the future, the voltage supply to the shift registers can be lowered significantly without affecting performance since the shift registers are already faster than necessary.

4.4 Digital Comparators

The digital comparator compares a locally stored digital pixel data with a globally distributed digital ramp value. The digital pixel data remains constant in the preserve latch while the ramp value counts up and then down. The digital comparator uses NFET only pass gate logic for low power operation. Each bit is compared separately and one ripple carry bit is propagated between the bit slices.

Each bit of the comparator (Figure 4-10) examines the digital data input (input A) to the digital ramp value input (input B). If they are equal, the comparator defers to the next least significant bit by passing the carry from that neighboring bit.

Inputs A and B to the comparator are differential. If both A and $\overline{A}$, or both B and $\overline{B}$ are above V_T, there can be short circuit current dissipation in the comparator.

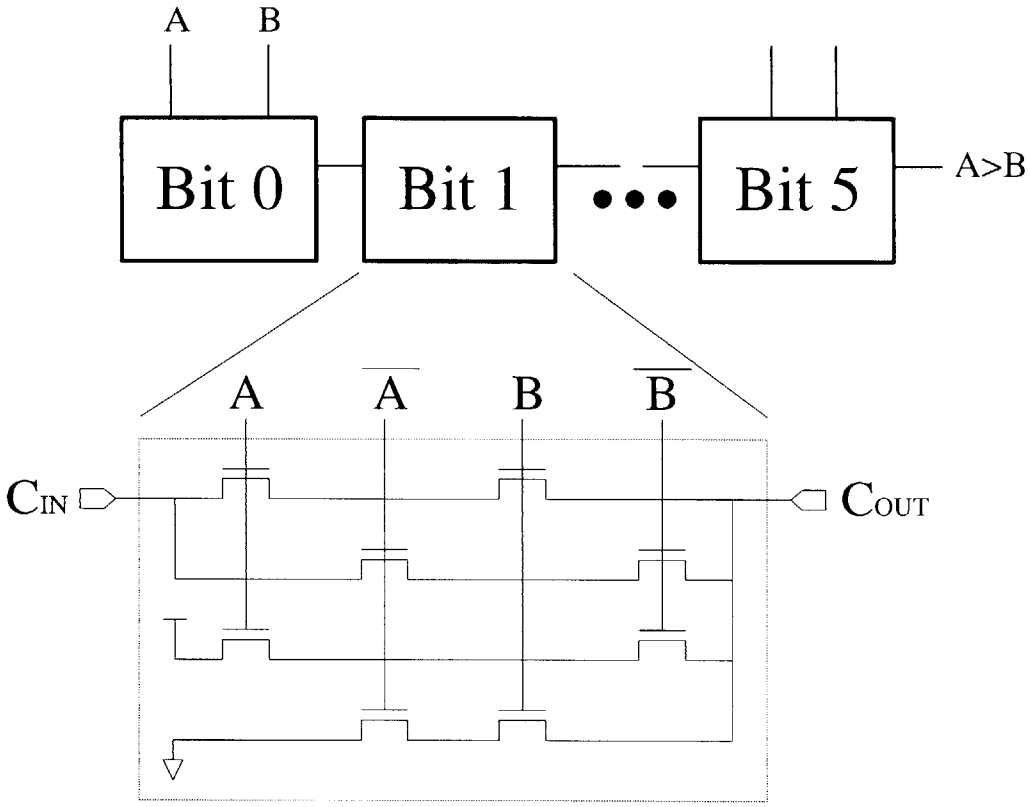

Figure 4-10: Digital Comparator Schematic. NFET only pass transistor logic yield a low power and small layout. The bitsliced design allows the comparator to fit within the 25 μm pixel pitch.

Figure 4-11 shows a short circuit path in the comparator when B is held at ground, $\overline{B}$ is held at VDD, and both A and $\overline{A}$ are at intermediate voltages greater than V_T. A similar path exists when A and $\overline{A}$ are held at the rails while B and $\overline{B}$ are at intermediate voltages. As long as one wire of each differential pair is held firmly below V_T, there can be no short circuit dissipation even if the other wire is not at VDD.

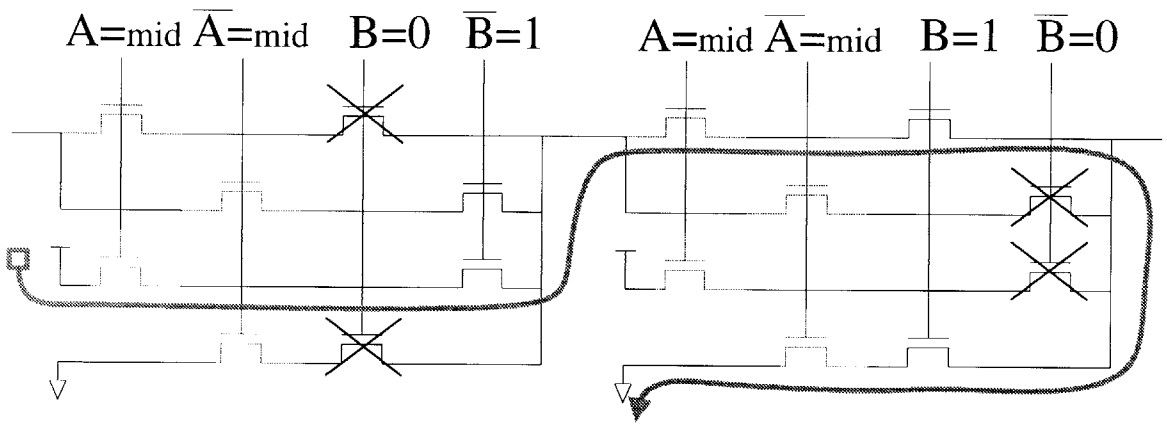

Figure 4-11: Short Circuit In Comparator Circuit. Can occur when two inputs are not held firmly at ground. As long as one wire from each bit pair is held below V_T, no short circuit current is dissipated.

Another possible short circuit current path is at the output of the comparator. The comparator output does not achieve full voltage swings due to the NFET only design and could cause short circuit current dissipation in subsequent circuitry. To restore the voltage level to full swing and to prevent glitches from causing the column and ramp to connect spuriously, the comparator output is registered using a standard clocked inverter register.

The CCLOCK signal controls this register. It transitions as frequently as the RVALUE bus and dissipates short circuit current when left at an intermediate voltage. Hence, CCLOCK is a Type A wire. Large buffers were used to drive the wire to ensure sharp transitions.

4.5 Ramp Distribution

The analog ramp and the digital ramp value are distributed globally as specified by the ACERDAC architecture. The synchronization circuitry for the two signals is implemented off-chip to allow for the greatest flexibility in testing different schemes. This synchronizing circuitry is ultimately responsible for the accuracy of the pixel voltages because it determines the timing of the ramp sampling.

4.5.1 Analog Ramp

The analog ramp is drawn directly from an off-chip source through an unbuffered analog input pad. The ramp wire resistance is a key design parameter if a resonant driver is part of the system, as these circuits often require a low series resistance to sustain resonance. Therefore, the analog ramp wire width was designed to be 20 μm to reduce the wire resistance. If a resonant driver is not used, the ramp wire can be made much narrower.

4.5.2 Digital Ramp Value

The digital ramp value is buffered on chip by large buffers whose final stage is 100 times the size of a minimum inverter. The ramp value is driven differentially across the chip to the digital comparators at each column, and is synchronized to the comparator register clock, CCLOCK.

The RVALUE bus wires are made up of Type B wires. Depending on the display resolution, the DAC bit accuracy, and whether energy recovery is being used, the RVALUE bus may transition more or less frequently than the DATA bus. Since the bus connects to the NFET pass gates of the comparator, no short circuit power will be dissipated as long as one wire of each bit pair is held below V_T as illustrated in figure 4-11. In addition, the values on the RVALUE bus either increment or decrement by one, therefore, the average transition frequency, α, is approximately $\frac{1}{3}$ for a six bit bus. For future designs, Gray coding can be used to limit α to *exactly* one bit pair per transition ($\alpha = \frac{1}{6}$).

4.6 Adiabatic Mode Versus Energy Recovery Mode

In the ACERDAC design, after the ramp reaches its maximum voltage, there are two options for the next phase. In adiabatic mode (AM), the ramp is simply reset to the minimum voltage. In energy-recovery (ERM) mode, the video line is ramped down slowly and the columns reconnect to the ramp when the column and video voltages are equal (Figure 2-16). This requires that the digital value stored in the preserve register remain until after energy recovery is complete. As shown in §2.4.3, little power is dissipated since the potential across the switch is zero when the load is reconnected. If an energy-recovering source is driving the ramp, the column energy will be returned to the source and reused. However, ERM doubles the RVALUE frequency as it must count up and down. Also, an energy recovering source can have significant overhead which increases the total system power.

Simulation and analysis shows that energy recovery is not advantageous for a 160x120 display because the video and column wire power is only a small percentage of the system power. Nevertheless, the scalability of this architecture enables its use at higher LC voltages and display resolutions where energy recovery becomes useful.

4.7 Ramp and Ramp Value Synchronization

The accuracy of the digital to analog conversion heavily depends on the synchronization between the ramp and ramp value. In this implementation, the synchronization circuitry is implemented off chip and the ramp value is broadcast globally as specified by the ACERDAC architecture. This configuration allows more flexibility with the ramp source, requiring only that it pass through all the desired analog voltages, but not that it be linear, monotonic, or consistent across time, process or temperature (§2.4.2). The ramp shape only affects the conversion energy efficiency, as a linear voltage ramp is theoretically the most energy-efficient source with which to charge a linear capacitor load. Arbitrary DAC transfer curves can be achieved by altering the ramp count and ramp synchronization, allowing LC transfer curve correction to be

implemented by the ramp generation circuitry.

An example conversion is shown in Figure 4-12 for a sample value of 4. The broadcast ramp value compensates for the ramp nonlinearity. In addition, any delay in the comparator can be pre-compensated for in the ramp value which allows the result of the comparator to be known exactly when the ramp should be sampled.

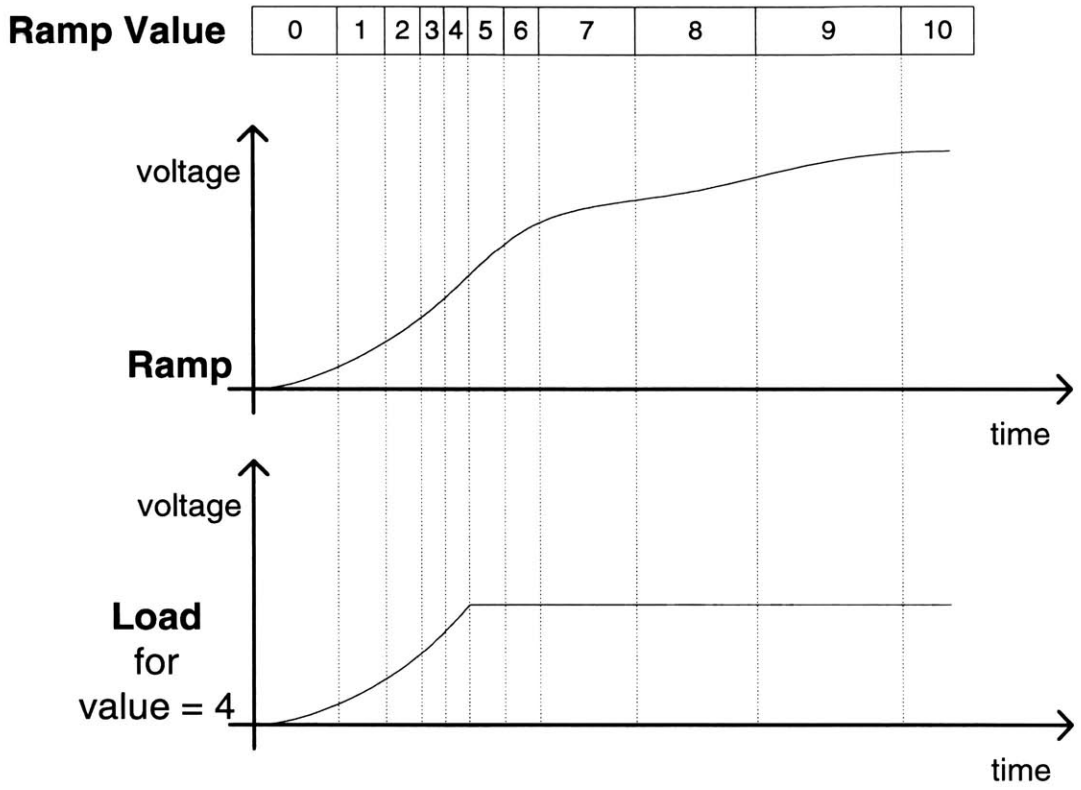

Figure 4-12: An example conversion for a sample value of 4. The digital ramp value notes when to disconnect from the ramp. Non-linearity in the ramp is compensated for in the ramp value timing.

Chapter 5

Analysis and Testing

The goal of this work is to design a low power digital display driver architecture, fabricate it, and use it to predict the power of other display resolutions and configurations. Therefore, the goals in testing the ACERDAC chip is to verify the correct operation, measure the power, pinpoint implementation areas where the power can be reduced, and use the implementation to predict the power for future implementations with different display resolutions and configurations. As a result, testing the ACERDAC chip consisted of two phases: functional testing to verify the correct operation of the chip, and performance testing. Performance is broken into two further categories: DAC accuracy and power consumption. System power consumption measurements allow the ACERDAC to be compared to other display implementations. Component-level power measurements allow power estimates for other display configurations to be made and helps to pinpoint areas where the ACERDAC can be improved. A theoretical power consumption model is made and compared to the measurements. Then, the model is used to predict the power consumption of future display systems.

5.1 Theoretical Power Consumption Model

In tandem with the fabricated ACERDAC chip, a theoretical model of power consumption in the ACERDAC architecture was built. The model combines parasitic capacitance and resistance estimations with simulations and calculations to predict

the power consumption of any ACERDAC implementation. Varying display drive schemes and circuit techniques within the model allows tradeoffs to be evaluated for different display resolutions, grayscale resolutions, and liquid-crystal inversion schemes. The validity of this model can be verified by comparison with the physical implementation.

The model calculates the parasitic capacitances and resistances of the wires in the ACERDAC implementation and uses these parasitics to estimate the system's power consumption using $\alpha f C V^2$ for each digital wire. The large buffers in the system were simulated to determine the energy dissipated in one transition. These results were added to the software model and multiplied by the frequency of transition for each buffer.

Parasitic resistances and capacitances were computed using parameters provided by the foundry. Sheet resistances, area and fringe capacitances, and transistor gate, junction, and sidewall capacitances were used along with approximate geometries to model the circuit parasitics. The resistance and capacitance of each wire were combined with the frequency and transition activity factor to estimate power.

All of the large buffers on chip were constructed similarly. This allowed a few simulations to yield accurate power estimations for all of them. The energy per transition given by the simulation was multiplied by the frequency of transition for each buffer. The sum of wire parasitic power and buffer power yielded a power estimate for the chip.

5.2 Test Setup

Testing the ACERDAC chip consisted of two phases, functional testing and performance testing. Performance is broken into two further categories: DAC accuracy and power consumption. The test setup included circuitry to drive data and control signals onto the ACERDAC chip, monitor the chip operation, measure the DAC accuracy and power consumption, and generate the synchronized ramp and ramp value since this function was not included on-chip. Some testing circuitry was included

on-chip which worked with the external circuitry to perform both functional testing and performance measurements.

A block diagram of the test setup is shown in Figure 5-1. The field programmable gate array (FPGA), from the Altera MAX7000S series, was used to drive signals onto the chip. A counter contained within the FPGA along with a DAC provided the synchronized ramp value and stair-step ramp.

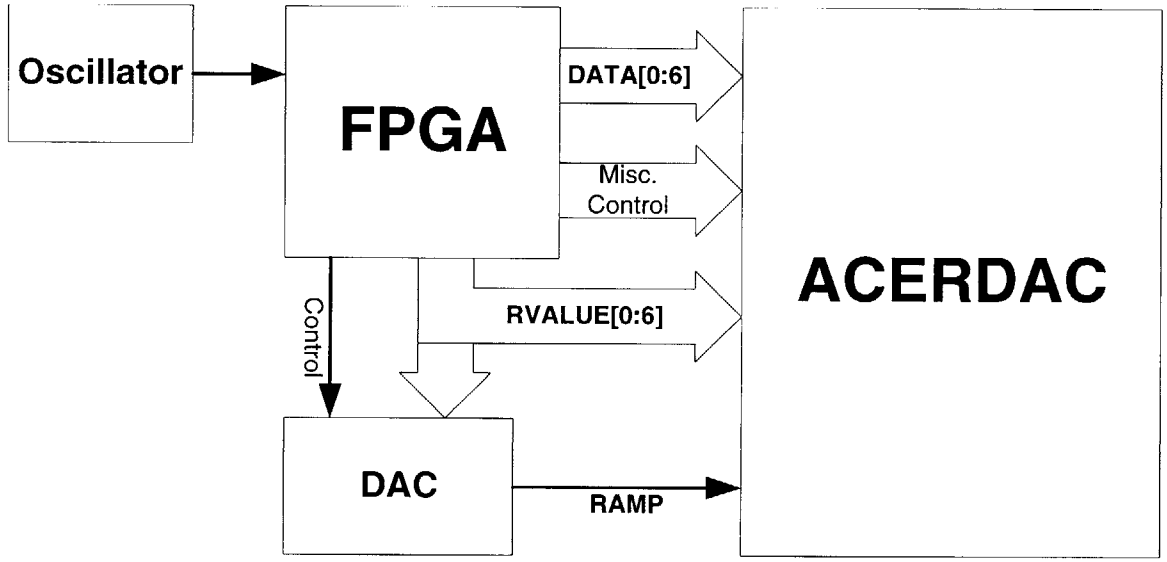

Figure 5-1: Drive Circuitry Schematic. An FPGA drives signals onto the ACERDAC chip. A discrete DAC provides the ramp.

A Keithly source meter was used to supply the digital $V_{DD} = 5V$ to the ACER-DAC chip. The analog circuitry was driven by a separate supply. Since the analog ramp connects directly from off chip to the column wires and the pixels, no energy is drawn from the chip V_{DD} to supply this node. Therefore, the power for the analog ramp is not included in the power measurements of the chip. The power consumed by the ramp source was not measured, but was estimated using the theoretical power consumption model.

5.2.1 On-chip Testing Circuitry

For testing purposes, a dummy column was included in the pixel array. Several critical signals in the dummy column are brought to output pads so that they may be observed for test purposes. Some of the signals are used to verify functionality of the ACERDAC chip, while others are used to measure the DAC accuracy.

The dummy column shares an enable signal with the last column, latching in the same piece of data as the last column. The dummy column has one pixel which shares an enable with the first row (5-2). However, the dummy pixel is not visible when viewing an image on the pixel array because it is physically separated from the array.

Several digital signals used for functional testing are buffered and connected to pads. The set of digital output signals was chosen to give the most information about the operation of the chip while using the fewest pads. Outputs were chosen at various points along the data path so that in the case of failure, the error could be pinpointed.

The first signal, HSH, is the horizontal shift register output. HSH is the first signal tested on every chip because it is the simplest way to verify basic operation of the chip, which will rule out any catastrophic errors such as power/ground shorts. Similar to the HSH signal, the vertical shift register output, VSH is brought to a pad. Proper vertical shift register operation can be verified this way.

The next signal in the data path, QC, is the capture latch low order bit for the last column. Proper operation on this signal allows verification of correct timing on the HCLOCK wire and DATA bus, and correct operation of the capture latches and the edge to pulse converter. A signal analogous to QC, QP is the preserve latch low order bit for the last column. QP allows the correct operation of PCLOCK and the preserve latches to be verified.

The following test signal, COL_EN, is the comparator register output of the dummy column. Monitoring COL_EN allows the correct operation of the digital comparator and comparator latch to be verified.

The analog output signals are used for testing the digital to analog conversion

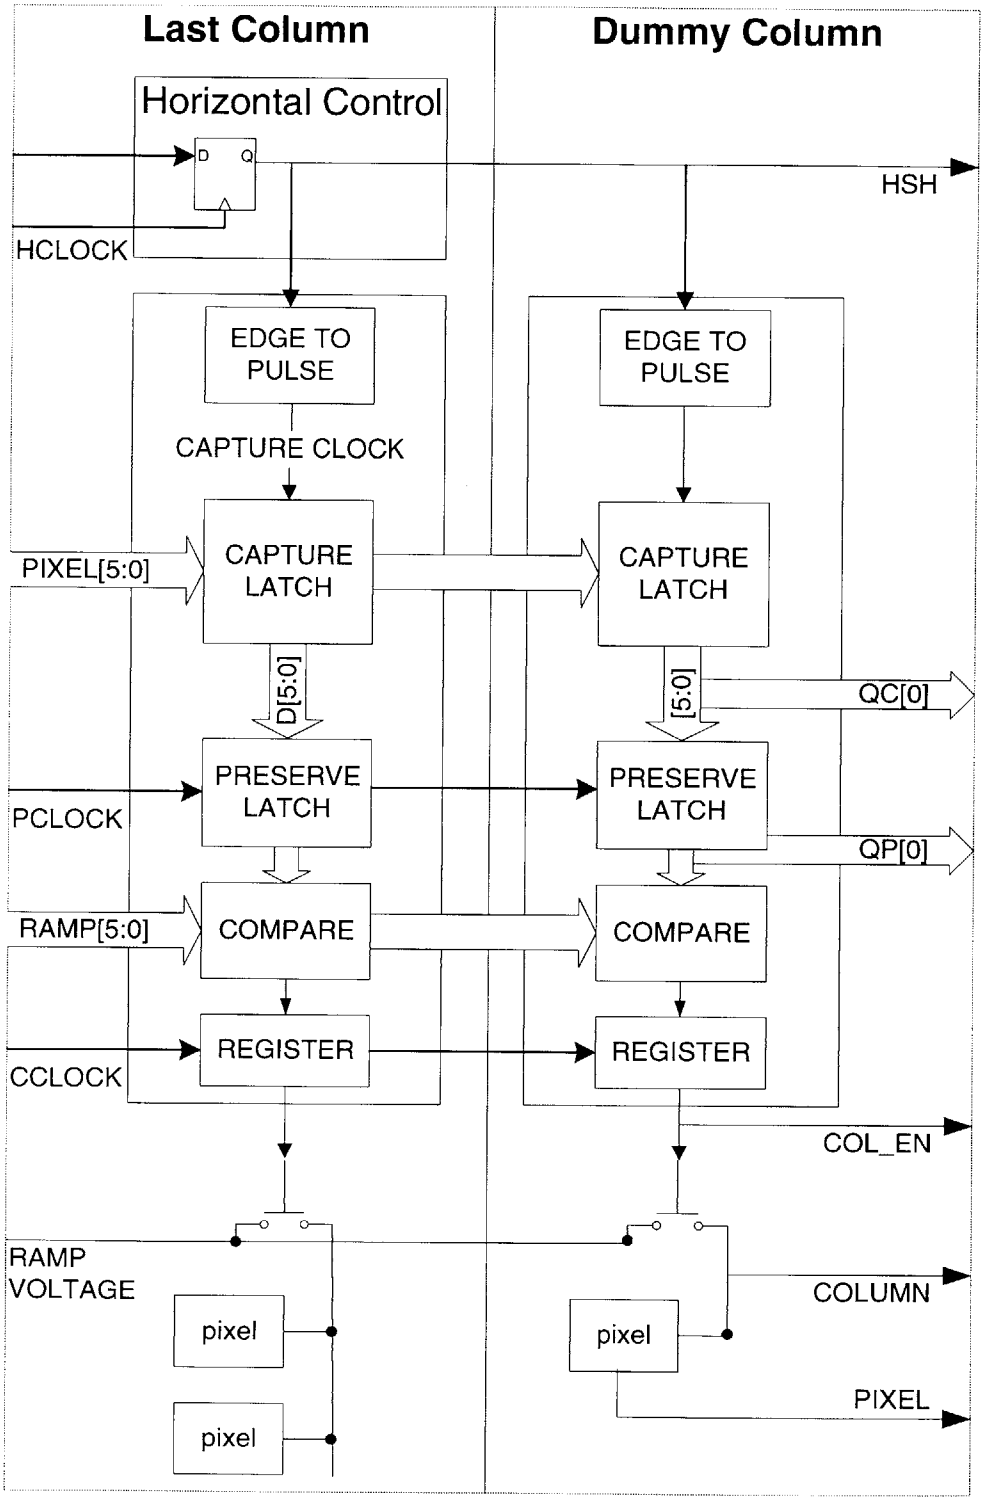

Figure 5-2: Dummy Column Circuitry. Used to monitor the operation of the ACER-DAC chip. A dummy pixel, used to measure DAC accuracy, shares an enable with the first row.

accuracy. Because they require analog buffers, they are not as simple to output as the digital signals. In the interest of design time, simple source-follower analog taps were used; the analog signal of interest is attached to a gate of a transistor and the voltage on the gate is mirrored in the output pad (Figure 5-3). This simple design does not accurately mirror the voltage presented to it, and the output deviation from the input varies with temperature and fabrication parameters. Therefore, a means of calibration is necessary. The ACERDAC chip includes an additional source-follower analog tap, the calibration tap, which has its input connected directly to a pad. This allows the exact transfer curve of all the analog taps to be obtained. A transfer curve derived by simulation (Figure 5-4) shows the general shape of the curves. The inverse of the measured transfer curve can be used to determine the exact voltage presented to the input of an analog tap given its output.

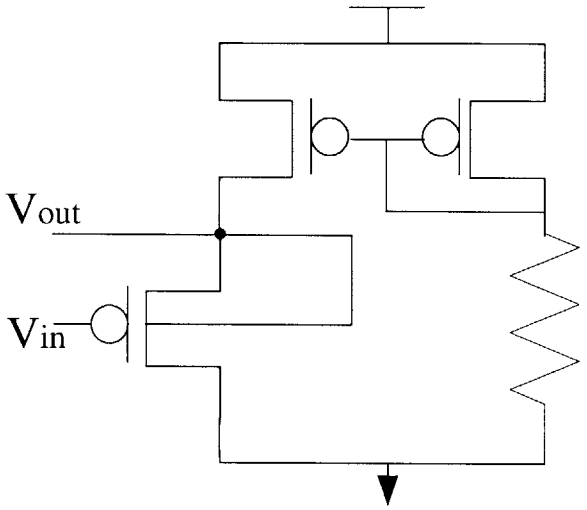

Figure 5-3: Analog Source Follower. Used to buffer pixel and column voltages for testing. The desired voltage is connected to the transistor gate and mirrored to an output pad.

The dummy column wire voltage and the dummy pixel voltage are brought off-chip for inspection. The sample and hold behavior of the sampled ramp DAC is first verified here and the row enable signal is verified by when the column and the pixel

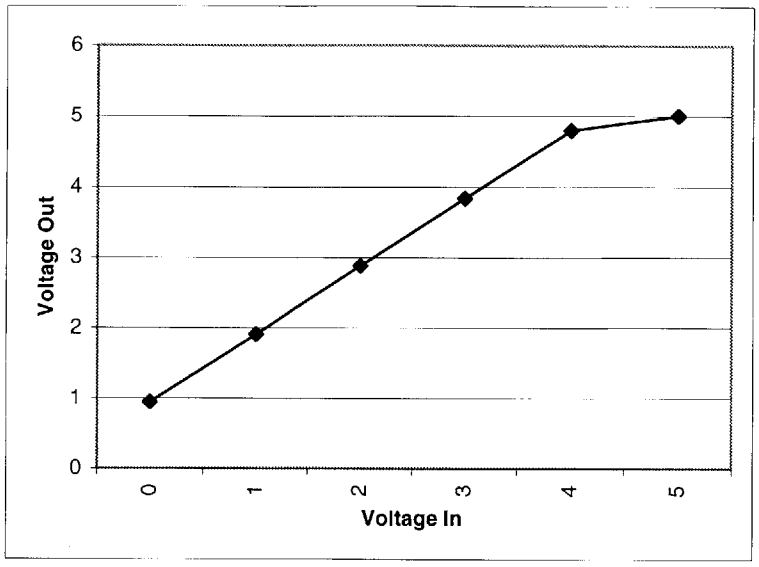

Figure 5-4: Simulated Analog Source Follower Transfer Curve. The nonlinearity will distort the readings from the chip. This curve will be used to correct measurements of analog signals on the chip.

are connected or disconnected at the right time. But, most importantly, the voltage on the pixel is the final DAC result, so the DAC accuracy is measured here.

5.3 Test Types

Tests were designed to isolate the various system components and measure their power individually. Since one common power supply was used for the entire chip, asserting some control signals while holding others constant allowed one subsystem to be exercised while activity in other subsystems is prevented. After subsystems are tested separately, the entire display power is measured under different input patterns and the DAC accuracy is measured. A discussion of the results follows at the end of this chapter. In particular, a comparison to the theoretical power dissipation model predictions is made to verify the accuracy of the model.

5.3.1 Individual Component Power Measurements

Test 1: Digital Bus In this test, the power required to swing the data bus is measured. This is separate from the power to run the horizontal shift registers and to latch the data in to the capture latches. This test includes only the power of the data bus input pads, the data bus low swing buffers, and the data bus wires. To perform this test, HCLOCK, PCLOCK, and all other signals are deasserted while the data bus bits are toggled at a frequency of 4MHz. Current draw from the VDD supply is measured under varying number of bits toggling.

Test 1 Results A certain amount of quiescent current draw is expected when no signals on the chip are toggling. This component of the system power is due to charge leakage and imperfections in the transistors that may result in a small amount of short circuit power. This quiescent power is measured in this test when no signals are toggling.

Each digital DATA bus bit should consume identical power since each is a replicated copy of the others. Only process variations or different capacitance on the input pads and bonding wires should cause different current draw for different bits. Therefore, a plot of number of bits toggling versus power consumption is expected to be a line. The slope of this line equals the additional power required to add another bit of precision to this bus. Figure 5-5 shows a plot of the power versus number of bits toggling.

Test 2: Digital Data Distribution System The Data distribution system is designed to consume little power in the common case of correlated images, specifically, the capture latches consume little power in the case of vertically correlated images, such as vertical stripes. In this test, a black and white vertical stripe pattern is displayed. The data captured by each capture latch remains constant throughout the test due to the vertical correlation, but the values on the DATA bus toggle at full activity as in Test 1 to create the stripes. The power required to swing the digital bus and latch data into the capture registers is measured. This includes the power to

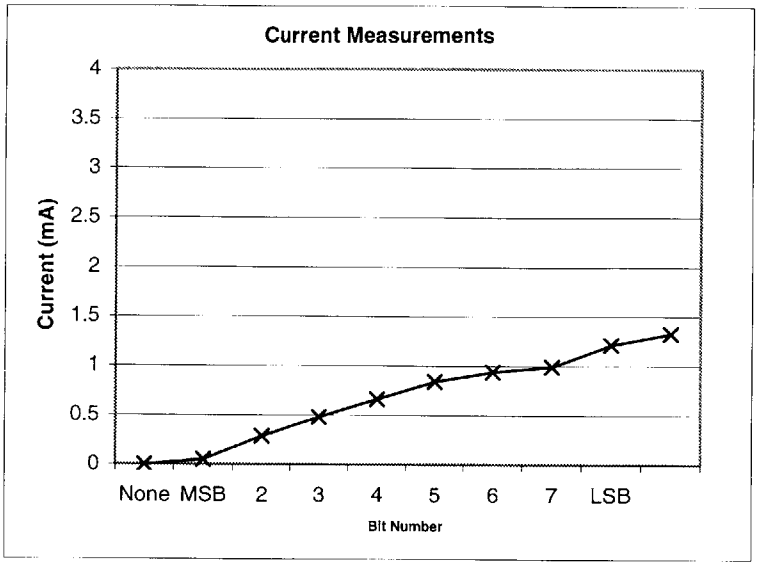

Figure 5-5: DATA bus power versus number of DATA bits toggling. Measurements show fidelity to the expected linear shape.

run the horizontal shift registers and the edge to pulse converters. To perform this test, all signals including PCLOCK are deasserted. HCLOCK is pulsed, and one 1 is injected into the horizontal shift register using HINIT. The data bus is toggled as in Test 1. Current draw from the VDD supply is measured under varying number of bits toggling at a frequency for HCLOCK and the DATA bus at 4MHz. Since PCLOCK is never asserted, the preserve latches never toggle. And since the RVALUE bus is never toggled, the comparators and subsequent circuitry are dormant. Therefore, it is certain that this test only measures the power of the digital distribution system.

Test 2 Results When no DATA bits are toggling, the current draw of the chip is entirely due to activity in the horizontal shift register. When DATA bits are toggling, current draw is due to the shift register, the data bus, and the capture latches. The DATA bus power is subtracted using the results of Test 1 allowing the power for the horizontal shift register and the capture latches to be calculated (Figure 5-6). As expected, the power required for the capture latches and shift registers (measurements

taken with bits toggling) is relatively similar to the power required for the shift registers alone (measurements taken with no bits toggling) in the case of vertically correlated images.

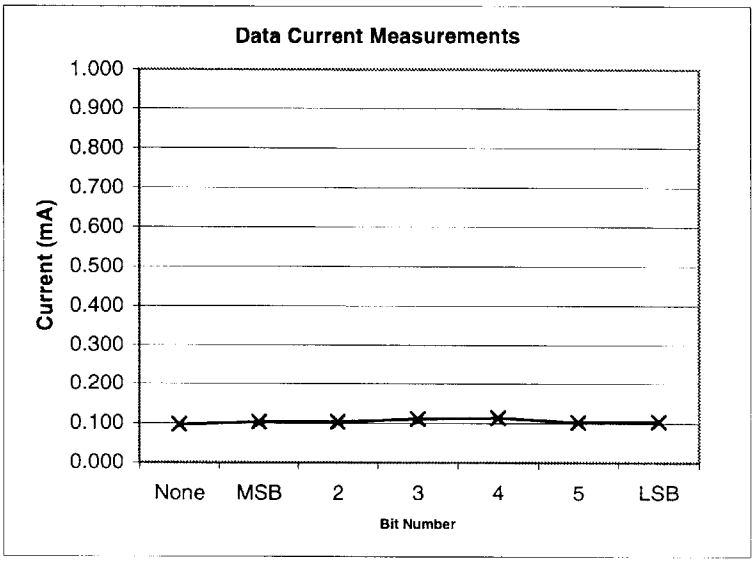

Figure 5-6: Data distribution system power versus number of DATA bits toggling. With Data bus power from Test 1 subtracted out. The curve is fairly constant as expected for this vertically correlated image.

Test 3: Digital Ramp Value Distribution and Comparator The power consumed by the digital ramp value distribution and the comparator are impossible to separate given the control signals available on this ACERDAC chip. Even for digital ramp value inputs which do not alter the final comparator output, the comparator internal nodes can toggle substantially. The comparator output is prevented from propagating to subsequent circuitry by deasserting CCLOCK which controls the register on the comparator output. Here, the analog RAMP is held at ground, and all digital signals including HCLOCK, DATA bus, and PCLOCK are deasserted. The digital RVALUE bus is toggled at a frequency of 4MHz. Current draw from the power supply is measured under varying number of bits toggling.

Test 3 Results Like the DATA bus wires, the addition of a toggling RVALUE wire should cause the power to increase linearly. However, the power in the comparator is difficult to predict due to uncertainty of the short circuit current component. Also, the addition of another toggling RVALUE bit will not necessarily increase the comparator power linearly because the internal nodes of the comparator can change even when the comparator final result is not changing. Therefore, a linear relationship between the number of RAMP VALUE bits toggling and the measured power is not expected. Isolating these components would lead to a better understanding of the power consumption in each of these elements and subsequently to an understanding of where to improve the circuit for next time. However, the available control signals do not make this possible. Figure 5-7 shows a plot of the power versus number of bits toggling.

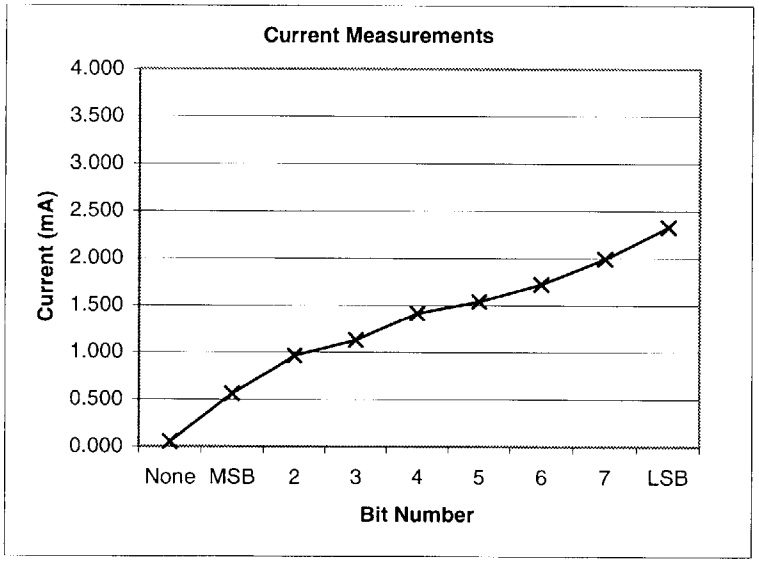

Figure 5-7: RVALUE distribution power and comparator power versus number of RVALUE bits toggling. This curve is not linear due to the activity of the comparator whose internal nodes may change even when the output does not.

5.3.2 Display Power Measurements

The following two test were selected to determine the best and worst case power of the ACERDAC display. Image characteristics greatly affect the power consumption of the ACERDAC chip because the digital data bus activity factor, and therefore, the DATA bus power is directly proportional to the degree of horizontal correlation in the display. Lowest power consumption is achieved when a uniform constant image is being displayed, and worst case power consumption is achieved when a pixel-by-pixel black/white checkerboard is being displayed.

Test 4: Display Power for a Constant Uniform Image The digital data bus is held constant while the rest of the display operates normally. The analog ramp continues to ramp up and down and the columns connect and disconnect as per normal display operation. The current draw from the ACERDAC chip power supply is measured.

Test 5: Display Power for a Constant Uniform Image The digital data bus is toggled at F_p and adjacent rows are initialized with alternating starting values (either full white or full black) to guarantee a pixel-by-pixel black/white checkerboard on the display. The current draw from the chip power supply is measured.

Tests 4 and 5 Results The power consumption of the chip when displaying a constant uniform image is 5.9 mW. The theoretical model predicts 4.91 mW for this configuration which translates into an error of 17 %. The power consumption of the chip when displaying a pixel-by-pixel black/white checkerboard is 7.25 mW. The theoretical model predicts 6.68 mW for this configuration which translates into an error of 11 %. As shown in §A, typical images are highly correlated. Therefore, the typical power consumption of the display showing real images will be closer to the best case power consumption than the checkerboard.

5.3.3 DAC Accuracy Measurements

Test 6: DAC accuracy The sampled ramp DAC output is the voltage stored on the pixel electrode. The accuracy of this DAC is therefore measured by measuring the pixel voltage for every digital input. The digital DATA bus is incremented through all the digital inputs with one new data value appearing each line time. The pixel voltage is measured once per line time. Since the pixel electrode voltage passes through an analog tap before it is presented to the output pad, the measurements are corrected for the analog tap transfer curve. The analog tap transfer curve is measured by increasing the voltage to the calibration tap's input and measuring the output.

Test 6 Results The plot of the measured analog tap input voltage versus output voltage is shown in figure 5-8. A comparison of this curve to the simulated transfer curve shown in figure 5-4 shows that they are similar. A plot of the DAC's digital input versus measured output is shown in figure 5-9. The transfer curve of the analog tap is corrected for in the DAC output with the results shown in figure 5-10.

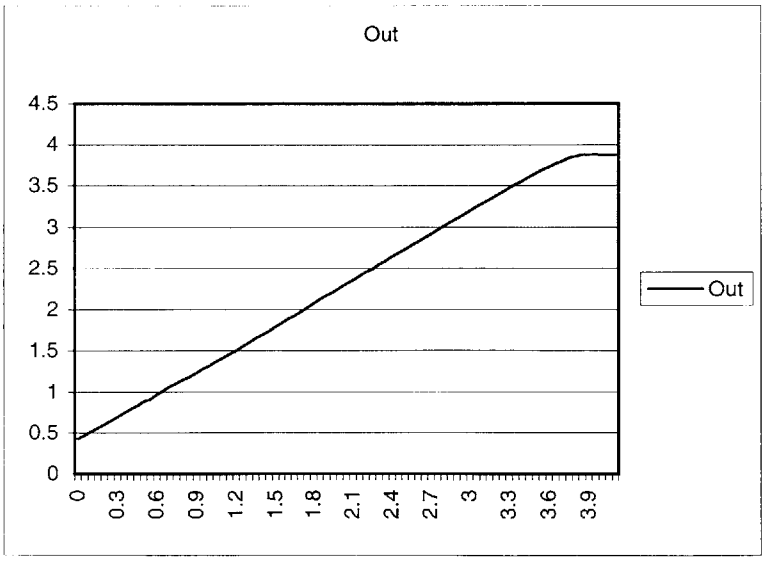

Figure 5-8: Measured Analog tap transfer curve. It matches the simulated results well. This allows corrections to be applied to analog measurements from the chip

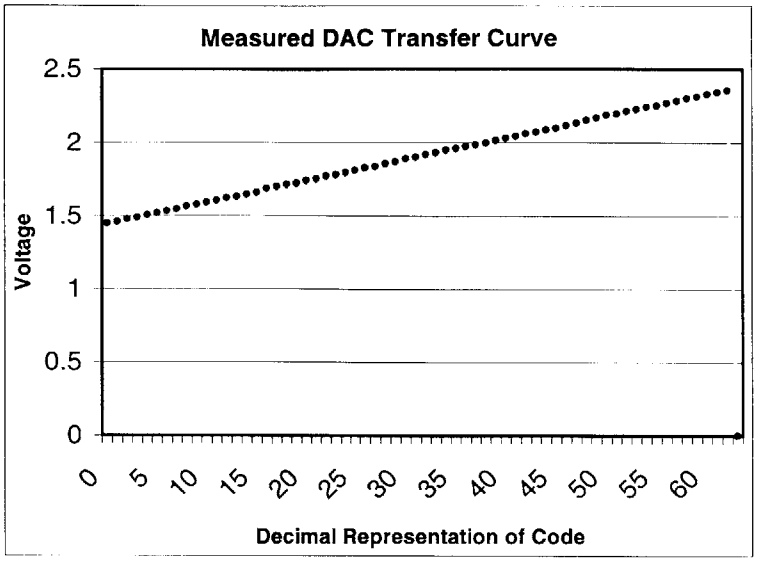

Figure 5-9: Measured DAC output versus digital input value. Values conform to 6-bit accuracy

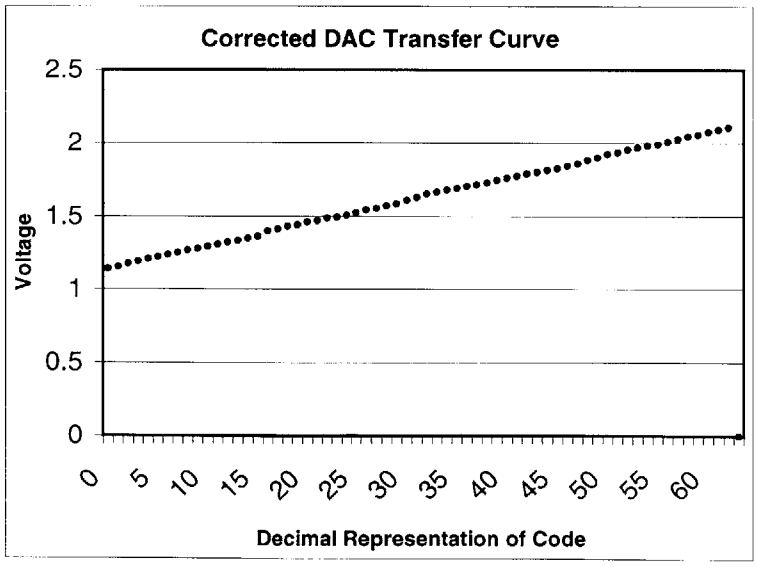

Figure 5-10: DAC output after correction for analog tap distortion. Measurements are corrected by application of the inverse of the transfer curve obtained by analog tap measurements. Corrected values also conform to 6-bit accuracy.

5.4 Discussion

There are two aspects of these tests that warrant discussion: interpretations and consequences of the test results and the testing process itself.

A consequence of comparing the measured results to the theoretical power dissipation model was the conclusion that the buffers driving the global wires were grossly oversized. In fact, the power required to switch the internal buffer node capacitance was greater than the power required to switch the global wire capacitance itself.

The measurements also highlighted that spurious transitions on the data and ramp value wires can significantly raise the system power consumption because these wires have large capacitances which are exaggerated by the oversized buffers. Simple drive schemes coded into the FPGA left the DATA and RVALUE bus toggling unnecessarily. More complicated drive schemes that gated the busses so that they did not toggle spuriously greatly reduced the system power consumption.

Examinations of the DAC accuracy readings showed that the measurements taken from the dummy column and pixel were inaccurate since the dummy column included only one pixel. Therefore, the dummy column wire had much less capacitance than an actual column wire. A different experiment must be designed to measure the DAC accuracy of an actual column and pixel.

The tests results show adequate fidelity to the theoretical model which allows the model's predictions to be trusted for other display configurations. The software model predicted within 17 % for the display power under varying image characteristics.

Interpretation of the DAC accuracy is limited since the accuracy mostly depends on the synchronization of the analog RAMP and the digital RVALUE. In the test setup, a discrete DAC is being used to generate the analog ramp. Therefore, the DAC transfer curve is not indicative of the quality of the system engineering since the discrete DAC, designed by a Analog Devices, is performing the difficult analog conversion. The accuracy of the measurements form the ACERDAC dummy pixel implies that the ACERDAC circuitry is engineered well enough to not disturb the discrete DAC's output by more than an LSB. DAC accuracy measurements will be

more important when a custom low power ramp and ramp value synchronization circuit is built.

The testing process itself was greatly facilitated by the large number of control signals brought off-chip. As a result, testing different drive schemes and testing different ramp and ramp value synchronization circuits was possible. This ability is crucial to the development of more efficient ramp and ramp value synchronization circuits which will be incorporated on-chip in the next ACERDAC implementation.

There were a few signals not brought off chip, that should have been. Testing would have been greatly simplified if each circuit subsystem had a separate power supply. Then, it would be more certain exactly where the power was being consumed. For instance, it was difficult to separate out the digital ramp value bus power from the comparator power. Most importantly, it would have been extremely useful to have the input and output pads use their own separate supply as the power dissipated in the input pads is dependent on the quality of the signals being supplied to the chip. The 3 volt swing signals being supplied from the FPGA had slow rise times that were causing extraneous short circuit current to be dissipated in the input pads.

Measuring the DAC output from the analog tap to the pixel was difficult because the analog taps share a common ground with the digital circuitry. So, every time the HCLOCK transitioned, there was noise injected into the analog tap. In order to obtain uncorrupted measurements, the voltage on the analog tap was measured during periods of digital system inactivity.

Chapter 6

Results and Future Work

The measurements described in the previous section prove that the ACERDAC functions correctly but also point out areas for improvement in the ACERDAC implementation. Possible improvements are presented and a second version of the ACERDAC implementation is proposed. Along with a proposal for version 2 of ACERDAC, future research directions are outlined.

This second version of ACERDAC can be compared with existing systems to evaluate the advantages of the ACERDAC architecture as a low power digital input display driver. In this chapter, a comparison to an existing system is made. The reference system is the combination of an analog input display with a monolithic DAC at the input.

6.1 ACERDAC Version 2

The primary goal of this first ACERDAC implementation (ACERDAC1) was to produce a working prototype to prove the advantages of the architectural design. Now that a working prototype exists, a second implementation of the ACERDAC architecture (ACERDAC2) is proposed which incorporates lessons learned from this implementation.

There were several mistakes made in ACERDAC1 which are corrected in ACERDAC2. Some features were left out of ACERDAC1 because they were deemed

too risky for the first implementation. Now that a better understanding of the ACERDAC architecture is available, these previously risky features are implemented in ACERDAC2.

Some problems with ACERDAC1 are discussed and possible solutions are outlined. Also, improvements discovered in the testing and analysis of ACERDAC1 are discussed. One key component missing from ACERDAC1 was an on chip ramp and ramp value generator. Several options for this circuit are presented. Finally, the ACERDAC2 implementation is proposed and analyzed. The theoretical power dissipation model validated by the ACERDAC1 measurements is used to predict the power consumption of ACERDAC2 with confidence.

6.1.1 Buffer and wire sizing Issues

As has been shown, the global wire buffers were unnecessarily large and dissipated short circuit power. Several improvements are made to reduce the power consumption of these buffers:

- The buffer fanout factor is reduced to decrease the number and size of each stage.
- On the final buffer stage, separate pull-up and pull-down controls are provided which reduces the short circuit current.
- Some of the buffers (especially the RVALUE wires) are converted to low swing buffers.

Some of the wires in ACERDAC1 were oversized and are thinned for ACERDAC2. This decreases the power of each of these wires significantly. The improvement is due in part because thinner wires are spread further apart without increasing the chip size thus reducing the coupling capacitance between wires.

6.1.2 Decoupling Digital Logic Voltage from LC Voltage

The requirement of a high voltage for the LC does not need to translate into a high voltage for the ACERDAC digital logic. A central theme to this thesis has been the use of conventional low power techniques such as low voltage in tandem with

unconventional low power techniques such as adiabatic charging and energy recovery. The main reason for using AC and ER was the high LC voltage requirement; the digital circuitry can be run at lower voltages.

Research in low power electronics has shown great success in lowering the voltage supply on digital logic. The place in the ACERDAC circuitry where the digital logic and analog logic meet is in the column switch. As long as the column switch control signal runs full swing, the digital logic preceding the column switch can be run at a lower voltage. To amplify the low swing column switch control signals, SRAM style receiver circuits are used. In fact, the very latch design used for the capture latch is duplicated here. Experiments in lowering the power of different components is facilitated by separating out the power supplies for each major component of the ACERDAC2 chip.

6.1.3 Variable Load Issues

The variability of the capacitance attached to the analog ramp (C_{ar}) can be quite substantial and is a result of the column wires connecting and disconnecting from the analog ramp at different times. The capacitance attached to the analog ramp ranges from the capacitance of the analog ramp wire itself (C_{rw}) to the capacitance of one column wire (C_{cw}) times the number of columns (N_c) plus C_{rw}.

$$C_{rw} \leq C_{ar} \leq N_c * C_{cw} + C_{rw}$$

It is difficult to operate in the presence of this data-dependent variability because it may degrade the digital to analog conversion accuracy. Some sources can operate accurately in the presence of this variability, but these types of sources usually consume large amounts of power and are therefore unacceptable for this application. There are several methods used to decrease this variability so that more power efficient sources can be used to drive this load.

One method for decreasing the capacitive variability is to decrease the capacitance of the column wire and the pixel. This is often difficult to do without affecting the

display and DAC performance. Charge leakage from a pixel with a small capacitance results in a higher voltage error, therefore a higher capacitance pixel is more desirable. Decreasing the column wire capacitance is difficult because it is heavily constrained by the display geometry. The column wire must run the entire height of the display, and it must have a source/drain connection for every pixel in the column. The column wire width must be increased in order to reduce the resistance and meet the RC charging time requirements of the display. Wider wires results in larger column wire capacitance which leads to greater ramp wire capacitance variability. Several standard capacitance reducing layout tricks can be used for the column wire such as reducing the source/drain area of the pixel switches, and reducing the coupling capacitance of the ramp wire by physically separating it from other wires in the layout. However, these layout tricks will not eliminate the large column wire capacitance and therefore will not eliminate the variability problems for the analog ramp.

Yet another method of reducing this variability is to charge only a subset of the columns at a time. As the number of simultaneously charged columns is decreased, the capacitive variability of the ramp wire is decreased. Since there is a fixed amount of time to convert one row of data, dividing the conversion into N groups necessitates that each group be converted N times as fast. Thus DAC accuracy is traded off for higher power consumption. However, in the case of pixel inversion drive mode, half the pixels in a row are drawn with one polarity relative to the coverglass, and the other half are drawn with the opposite polarity. Therefore, the column wires are automatically split into two groups when using the optimal image quality inversion scheme – pixel inversion.

There are two situations under which charging N times as slowly will not take proportionally more power. The first condition is when using a blip circuit which resonates within a range of frequencies related to the resistance, inductance and capacitance of the circuit. The frequency of operation is difficult to control and is often faster than a line time. Therefore, the frequency may need not be increased to accommodate N groups of conversions per line time. The other condition which will not use proportionally more power is when using a stair-step ramp. Each step in

the ramp need only be an RC time constant apart in time. Transitioning any slower does not save power. So the reduction in conversion time as a result of having several groups of conversions may not translate into proportionally more power consumption.

Several methods can be used to divide the columns into two or more groups. The k groups can be distinguished by the $\log k$ high order bits; this decreases the bit precision by $\log k$ or increases the circuitry overhead by $\log k$ bits. Or, the groups can have separate enable signals. An appropriate place for these enable signals is at the de-glitch register on the column enable signals. However as it is implemented, the amount of information needed to distinguish between k groups is $\log k$, and therefore, $\log k$ bits of additional information will need to be distributed globally in order to enable the groups at separate times.

The reason the capacitive variability needs to be reduced is to increase the DAC accuracy, or decrease the analog ramp driver complexity. DAC accuracy is determined on the sampling phase of the ramp. During the charge-recovery phase of the ramp, ramp accuracy is less important because the accuracy during this phase only determines the energy recovery efficiency and has no effect on the optical performance of the display. Therefore, even if one line of data needs to be converted in two groups for DAC accuracy purposes, they can be recovered in one group and the circuit will operate correctly, but less efficiently.

The stair-step ramp alleviates the variability problems because of the timing of columns connecting and disconnecting from the ramp. Columns disconnect from the ramp only after the rising edge of the CCLOCK. As long as this clock happens well before the ramp begins transitioning, there is a fixed capacitance during each transition. As long as the ramp source can accurately drive the highest capacitance, but is not expecting a specific capacitance for every transition, the circuit will function correctly.

The sampled ramp DAC implementation described in §2.5 does not connect the column wires directly to the ramp source, instead the gates of sampling transistors are connected to the ramp which presents a constant load to the ramp source. This implementation can not benefit from the power saved by using the ramp to slowly

charge the column wires. But, this implementation does not suffer from the load variability issues of the ACERDAC.

There are several methods of dealing with the ramp wire capacitance variability. The accuracy of different ramp generation schemes is affected differently by the variability. Therefore, the scheme of dealing with the variability will be tailored to the generation scheme chosen.

6.1.4 Ramp Source Possibilities

Conventional DAC with digital counter The ramp for ACERDAC1 is implemented using a counter and a discrete DAC. The counter is fed to the DAC's digital input and also to the RVALUE input of the chip. The DAC analog output is connected to the analog RAMP input of the ACERDAC. Therefore, the DAC drives a very large load which is an inefficient use of a high-precision, high-power analog device.

Conventional DAC with digital counter, current source, and analog comparator So that the DAC does not have to drive the large pixel array, a current source is used to charge the load slowly (Figure 6-1). A DAC is used to pre-compute the next voltage to which the ramp should rise. While the ramp is being charged by the current source, the DAC output is compared with the ramp value. When the two voltages are equal, the ramp is at the proper voltage, and the current source is turned off. The code presented to the DAC input is used as the ramp value just as in the previous example.

If the conventional DAC did not operate to full rail voltage, this system could still work. The comparator could scale the ramp voltage into the range of the DAC before making the comparison. This allows the high-powered discrete DAC to operate at lower voltages and thus consume less power.

Blip and ADC In order to perform energy recovery, a resonant supply is used. The Blip circuit uses inductors and the load capacitance to act as a resonant supply [4]

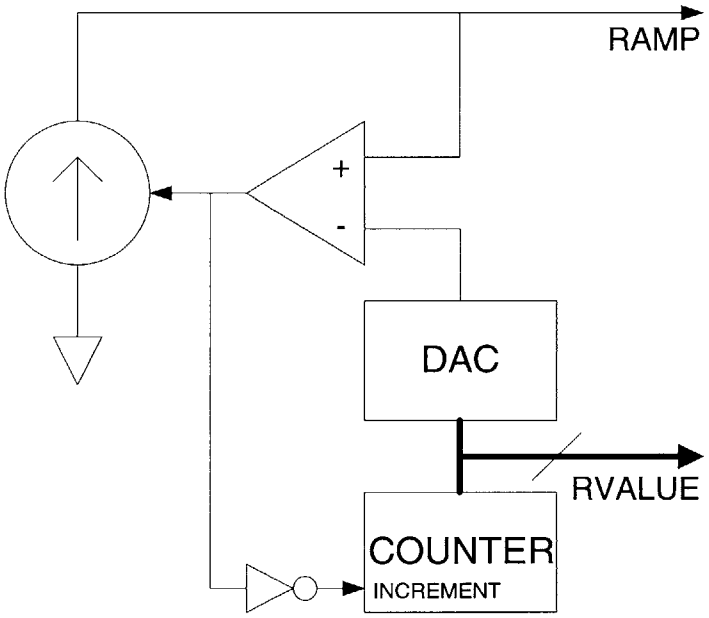

Figure 6-1: Ramp Source Possibility. Consisting of a conventional DAC with a digital counter, current source and analog comparator. It uses a current source to charge the load instead of using the DAC.

(Figure 6-2). The Blip circuit frequency of oscillation is highly uncontrollable when driving the variable capacitive load of the ramp because different capacitances will cause a different frequency of oscillation, and some combinations of capacitance and inductance might cause the circuit to stop resonating completely. An ADC is used to compute the ramp value from the blip circuit's voltage. Therefore the ramp value is slightly delayed relative to the ramp. This may cause offset error in the resulting DAC transfer curve.

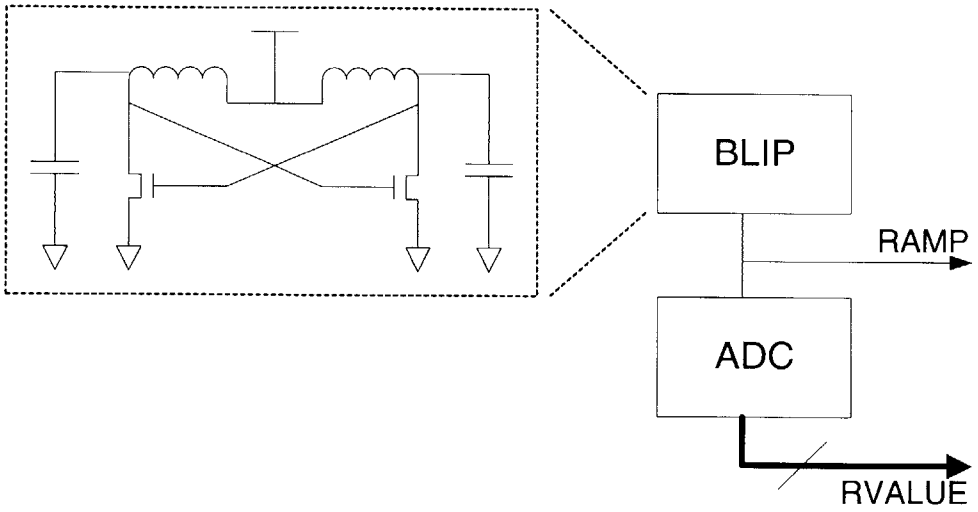

Figure 6-2: Ramp Source Possibility. Consisting of a Blip circuit and an ADC. The Blip circuit resonates and can perform energy recovery on the load. (Blip circuit taken from [4].)

Exploiting Regularity for Low Power The DACs and ADCs described above can exploit the regularity of the inputs to save power. The DACs can trade off long latency for low power because they can anticipate the next code to be converted (the next value is the previous value plus one). The ADCs can also anticipate codes and, as a result, save power.

Charge Transfer DAC A charge transfer DAC dissipates very little power because it requires fewer amplifiers resulting in less stand-by current draw[7]. However, the

charge transfer DAC has a long latency (Figure 6-3). This latency is due to the large RC time constant required to equalize charge between two large capacitors which is performed once for each bit of DAC precision. The long latency of the charge transfer DAC is not an issue if the desired code to be converted is known well in advance of its use.

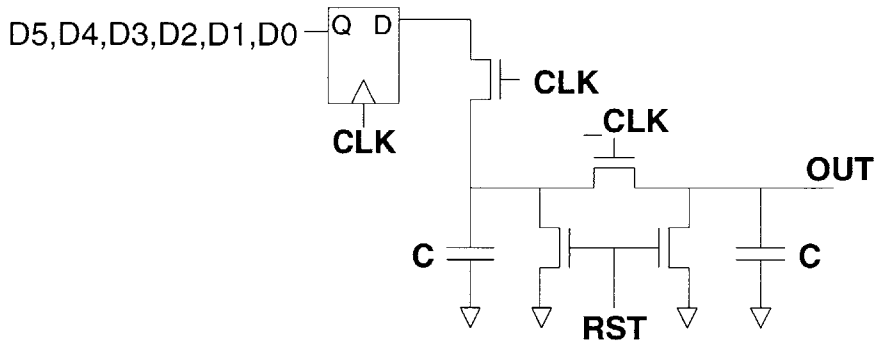

Figure 6-3: Schematic of a Charge Transfer DAC. An analog value is produced by shuffling charge between two capacitors, once for each bit of DAC precision. The long latency is not an issue if the desired code is known well in advance of its use.

6.2 Technology Directions

The ACERDAC architecture scales to larger resolutions. Additional column drivers for each new column presents few scalability problems. Increasing buffer sizes for the now larger global wires, and increasing the comparator speed is all that is required.

In addition to moving to larger resolutions, the ACERDAC architecture allows LCOS display quality to be improved by using LC inversion schemes which would consume too much power if implemented on conventional displays.

6.2.1 Larger Resolutions

The parallel aspect of the ACERDAC architecture allows it to scale to larger resolutions without drastic implementation modifications. An increase of a factor of K in

each display dimension requires K times the number of column drivers each running K times faster. Whereas a monolithic DAC requires a K^2 increase in speed. The ACERDAC linear speedup requires a faster ramp resulting in more power consumption, but does not require a drastic column driver redesign. However, redesigning a monolithic DAC for a K^2 speedup probably requires a new architecture.

For a display resolution increase of K in each direction, the pixel rate increases by a factor of K^2 and the line rate increases by K. The number of ramp values per line stays the same, so the ramp value rate increases by a factor of K. Figure 6-4 shows how the data rate increases relative to the ramp value rate for different display resolutions operating in energy recovery mode. For a 640x480 display and an 8 bit accuracy, the data rate is approximately equal to the ramp value rate. As display resolution increases, the data rate becomes significantly larger than the ramp value. For larger resolutions, the percentage power overhead of doubling the ramp value rate to accommodate energy recovery mode is less.

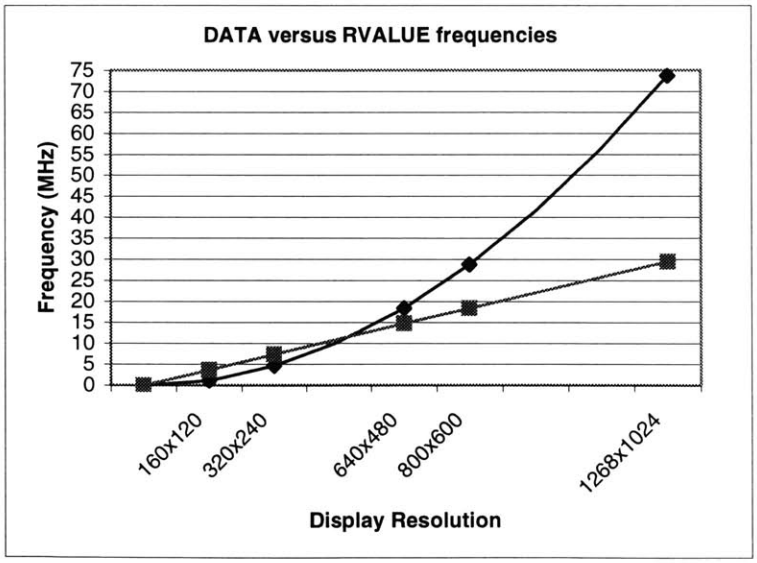

Figure 6-4: Growth rates of data rate and ramp value rate for different resolutions. The data given is for 8-bit accuracy and energy recovery mode. The data rate grows quadratically while the ramp value rate grows linearly.

The pixel array for larger resolutions is more capacitive. Specifically, the capacitive load for each column driver is increased linearly in K. Therefore, the pixel array power increases relative to the digital circuitry power. In the 160 x 120 display, the pixel array power was such a small percentage of the system power, that energy recovery was not economical. But, in larger resolution displays, the pixel array power becomes a greater percentage of the system power and energy recovery becomes a viable option for reducing the system power. Therefore, the percentage overhead of energy recovery decreases, and the recoverable energy increases making ER more advantageous the larger the display.

6.2.2 Revised Implementation Performance

This section explores the potential performance of the ACERDAC architecture for a 640 by 480 resolution given the incorporation of a few low power methods which are straight forward to implement. The ACERDAC1 implementation is scaled to a 640 by 480 resolution using the theoretical model. A few additional low power techniques were assumed. First, the digital logic voltage was reduced to 2 volts. Second, the low swing data drivers were reduced to 1 volt swings. For 6 bits of accuracy in the DAC operating in adiabatic mode, this scheme consumes 5.72 mW. Table 6.1 shows a breakdown of the power for the digital components of this revised display implementation. The analog ramp and column wires consume 0.6 mW.

	Wires		Buffers		Total
System	(mW)	%	(mW)	%	(mW)
Pixel Data	2.26	39.6	1.3	22.8	3.56
Ramp Value	0.39	6.9	0.13	2.3	0.52
Comparator Clock	0.21	3.6	0.09	1.5	0.30
HCLOCK	0.47	8.1	0.86	15.0	1.33

Table 6.1: Power breakdown for a 640 x 480 6-bit display. Total power is 5.71 mW. Data broadcast consumes the most power (62.4%).

6.3 Comparison to Conventional

This section compares the performance of driving an analog input LCOS display with a video rate on-chip monolithic DAC to the ACERDAC with one of the ramp and ramp value generators described above. A typical resolution of 640 by 480 is used in the calculations.

A conventional display is assumed to consist of a subset of the circuitry used in ACERDAC including: the horizontal shift register, the pixel array, and the ramp wire which is used as the video wire. All of the digital control circuitry and latches of the ACERDAC are removed. It is assumed that, like the ACERDAC implementation described above, the conventional display uses the lower 2V supply for digital circuitry. An on-chip video rate DAC converts the digital input data to analog values to be put on the video wire. A typical low power video rate monolithic DAC running at 1.2 MHz, and 5 volts, consumes 21 mW[1][8].

The ACERDAC1 does not have a ramp and ramp value generator on-chip, but for purposes of comparison, it is assumed to be included on-chip. A discrete DAC, counter, current source, and an analog comparator are used to generate the ramp and synchronize it to the ramp value as was described in §6.1.4. The important thing to note is that the discrete DAC need not run with a full swing. Also, the discrete DAC in this system is required to run at ramp value rates which are more than half as slow as data rates on a 640 x 480 display for six bits of DAC precision in AC mode. A typical low power DAC running at 5 volts consumes 20 mW at the ramp value rate of 500 KHz[2][8]. Thus, the ACERDAC display including the digital circuitry, analog ramp, and ramp generator consumes $5.72mW + 0.6mW + 20mW = 26.32mW$.

The conventional display consumes 8.37 mW at 60 Hz. A breakdown of the power for this display is shown in table 6.2. The total for the conventional display plus monolithic DAC is $8.37mW + 21mW = 29.37mW$.

[1]The DAC draws 4 mA in addition to the current draw necessary to drive the output load which is calculated to be 0.23 mA. At 5 volts, 4.23 mA results in a dissipation of 21 mW.

[2]Again, the DAC draws 4 mA in addition to the current draw necessary to drive the output load which is calculated to be 0.014 mA. At 5 volts, 4.014 mA results in a dissipation of 20 mW.

	Wires		Buffers		Total
System	(mW)	%	(mW)	%	(mW)
HCLOCK	0.47	5.6	0.86	10.3	1.33
Video Wire	6.47	77.3	0.0	0.0	6.47
Column Wires	0.57	6.8	0.0	0.0	0.57

Table 6.2: Power Breakdown for a conventional 640 x 480 Display. Total power is 8.37 mW. Video wire power dominates (77.3%).

If a lower power video rate DAC was found for use with the conventional display, the same DAC could be run at less than half the speed and used for the ramp and ramp value generator for the ACERDAC display. Therefore, with better low power DAC technology, the ACERDAC remains lower power than the conventional display.

Chapter 7

Conclusion

This work developed a low-power high-resolution digital input display architecture for LCOS displays that operated at the 5 volts required by the liquid-crystal material. This architecture will improve the usability of portable devices by allowing them to effectively display the increasing bandwidth of information available to them. The sampled ramp DAC method was used and combined with the parallel DAC-per-column architecture and adiabatic charging and energy recovery to achieve both low power and a compact design that integrated well into the narrow display pixel pitch. The display can operate in both adiabatic and energy recovery modes which allows the lowest power mode for a particular display configuration to be evaluated and used. The ACERDAC architecture incorporates previously-known advantages of using sampled ramp DACs for displays in addition to the scalability, low-power, and area advantages discovered by this work.

The ACERDAC architecture replaces the fast analog wires of the analog input LCOS display with slower digital wires and takes advantage of image correlation which is overlooked in conventional display drivers. The architecture also allows the use of optimal LC drive schemes which were previously too power-intensive to be implemented. Low power circuit techniques were employed to accommodate the high voltage requirement of the LC material. Thus, the rapid progress of the low-power electronics community is leveraged instead of the slower development of low voltage LC material science.

A working implementation of the ACERDAC architecture allows the advantages of the architecture to be verified. The 160 x 120 pixel display consumes 7.5 mW (worst case) with a 4 MHz pixel clock. A theoretical power consumption model, validated by the ACERDAC implementation, allows the power consumption of other display resolutions, grayscale resolutions, and drive methods to be evaluated. Using the results from the first ACERDAC implementation, a second version is proposed which is lower power for a 640 x 480 resolution than a conventional analog input LCOS display combined with a monolithic DAC. Even with improved low power DAC technology, the ACERDAC remains lower power than the conventional display.

This work extends the field of low power electronics by including environments of high voltage and high capacitance in the domain of low power digital to analog conversion. This work also shows that combining conventional low power circuit techniques with adiabatic charging and energy recovery can be practical.

Appendix A

Image Correlation and Effect on Power

Many of the circuit techniques in the ACERDAC architecture are efficient in the case of correlated images. This section shows a high degree of correlation to be evident in a set of test images by computing the activity factor, α, for these images. The worst case activity factor is analyzed for both analog and digital displays. Then, the activity factor is calculated assuming pixels drawn randomly from a uniform distribution. The activity factor assuming a random uniform distribution of pixels is at least a factor of two lower than the worst case, and activity factor for each test image is lower than that of the random uniform distribution.

A.1 How activity factor is computed

A conventional drive scheme is being analyzed therefore no energy recovery is being performed. As a result, once energy has left the power supply, it must eventually be dissipated in the display system. Therefore, a simple way of calculating how much energy is dissipated in the display system is to calculate how much energy is drawn from the power supply. Using this method, only the transitions which involve charging the column wire from a lower voltage to a higher one will draw energy from the power supply. The fraction of time these power-consuming transitions occur is α, and the

average energy per transition in this system is αCV^2.

In a digital system, the only transition from a lower voltage to a higher one is the $0 \rightarrow 1$ transition. Therefore, the activity factor is the fraction of transitions that are $0 \rightarrow 1$ transitions, which is why the activity factor for a digital system is often denoted $\alpha_{0\rightarrow 1}$.

In an analog system, there are several different ways of transitioning between a lower and a higher voltage. Conventionally charging a capacitance, C, from a voltage, V_1, to a voltage, V_2, requires $CV_2(V_2 - V_1)$ of energy be drained from the supply. In keeping with the convention that α is multiplied by CV^2 to obtain the average transition energy, α for an analog system is

$$\alpha = \text{avg}_{(V_2 > V_1)}\left(CV_2(V_2 - V_1)/\text{VDD}^2\right). \tag{A.1}$$

A.2 Worst Case Analysis

In a digital system, the worst case activity factor is one half since a circuit can transition from $0 \rightarrow 1$ a maximum of half the time.

In an analog system, the worst case activity factor is also one half. This is a bit more complicated to show, but is true because the energy for a $0 \rightarrow \text{VDD}$ and then a $\text{VDD} \rightarrow 0$ transition requires $C\text{VDD}^2$ of energy which is more energy than any other possible set of two sequential transitions. To see this more clearly, examine the energy for two sequential transitions, $X\text{VDD} \rightarrow Y\text{VDD} \rightarrow Z\text{VDD}$ for $0 \leq X, Y, Z \leq 1$, which is

$$C\text{VDD}^2(Y(Y - X) + Z(Z - Y)). \tag{A.2}$$

Equation A.2 is maximized when $X = 0$ and $Z = 1$ and when $X \leq Y \leq Z$ because both quantities in the sum are non-negative. Using this information, the energy can be rewritten as

$$C\text{VDD}^2(Y(Y - 0) + 1(1 - Y)) = C\text{VDD}^2(Y^2 + Y - 1)) \tag{A.3}$$

which is less than $C\text{VDD}^2$ for $0 < Y < 1$. Therefore, the worst case energy is achieved when the circuit transitions between 0 and VDD and the worst case activity factor is one half.

A.3 Random Uniform Distribution Analysis

In a digital system, with pixel values drawn from a uniform distribution, all transitions are equally likely. Therefore, the activity factor, $\alpha_{0\to1}$, is one quarter because the $0 \to 1$ transition happens a quarter of the time.

As before, the analog system requires a more detailed analysis. Define the probability that a pixel will have value $x * \text{VDD}$ for $0 \leq x \leq 1$ to be $p(x)$. Since $p(x)$ is a uniform distribution, $p(x) = 0$ for $x < 0$ or $x > 1$ and $p(x) = 1$ for $0 \leq x \leq 1$. Define Δ to be the difference between two randomly selected pixel values x and y. If $|x - y| = \Delta$, x and y can be represented as k and $k + \Delta$. Then,

$$P\{X - Y = \Delta\} = \int_{k=\Delta}^{\infty} P\{X = k, Y = k - \Delta\} \tag{A.4}$$

and assuming x and y are independent,

$$P\{X - Y = \Delta\} = \int_{k=\Delta}^{\infty} P\{X = k\}P\{Y = k - \Delta\}. \tag{A.5}$$

Then, the activity factor of this system using Equation A.1,

$$\alpha = \int_{\Delta>0} \int_{k=\Delta}^{\infty} p(k)p(k - \Delta)(k)(\Delta)dk\, d\Delta. \tag{A.6}$$

For $p(x)$, the uniform distribution, $p(k)p(k - \Delta)$ is nonzero only when $-1 < \Delta < 1$ and $0 < k < 1$ because $p(k)$ is nonzero only between 0 and 1. Since the first integral calls for $\Delta > 0$, the integral is evaluated over the range $0 \leq \Delta \leq 1$. Therefore, Equation A.6 reduces to:

$$\alpha = \int_{\Delta=0}^{1} \int_{k=\Delta}^{1} p(k)p(k - \Delta)(k)(\Delta)dk\, d\Delta = \frac{1}{2} \int_{\Delta=0}^{1} \Delta(1 - \Delta^2)\, d\Delta = \frac{1}{8}. \tag{A.7}$$

Therefore, the activity factor for an analog system is 1/8, which is half that of a digital system when values are drawn from a random uniform distribution.

A.4 Typical Images Analysis

For a set of test images, the digital and analog activity factor is calculated. In the digital domain, the activity factor is calculated for horizontally adjacent pixels which is the order in which they would be driven onto the data bus. For the analog domain, activity factor is calculated for both horizontally and vertically adjacent pixels which is the order in which they would be driven onto the ramp wire and column wires respectively. Digital activity factor is calculated by computing the fraction of $0 \rightarrow 1$ transitions. Analog activity factor is computed using Equation A.1.

Table A.1 shows the horizontal and vertical analog activity factors of each of five test images. The table also shows activity factors for an image with pixel values drawn from a uniform distribution which agrees with $\frac{1}{8}$ from the formal analysis above. The last row of the table lists the worst case activity factor which was computed above.

	Analog α	
Image Name	Horizontal	Vertical
Baboon	0.01	0.02
Barbara	0.01	0.01
Goldhill	0.005	0.006
LAX	0.01	0.01
Woman	0.006	0.004
Uniform Dist.	0.126	0.125
Worst Case	0.5	0.5

Table A.1: Analog activity factors for each of five test images and one computer generated image.

Table A.2 shows the horizontal digital activity factors of each of five test images assuming pixel values with 6 bits of precision. The activity factor for each bus bit is displayed as well as the average activity factor across all bits of the bus. It is

expected that the MSB will have a lower activity factor than the LSB because digital values whose analog voltages are similar will have similar MSB's, but will differ in the LSB's. The table also shows activity factors for an image with pixel values drawn from a uniform distribution which agrees with $\frac{1}{4}$ from the formal analysis above. The last row of the table lists the worst case activity factor which was computed above.

	Horizontal Digital α						
Image Name	MSB					LSB	Average
Baboon	0.07	0.10	0.16	0.22	0.24	0.25	0.17
Barbara	0.05	0.08	0.13	0.17	0.21	0.24	0.15
Goldhill	0.02	0.05	0.10	0.15	0.21	0.24	0.13
LAX	0.05	0.12	0.15	0.22	0.25	0.25	0.17
Woman	0.02	0.05	0.08	0.13	0.19	0.24	0.12
Uniform Dist.	0.25	0.25	0.25	0.24	0.25	0.25	0.25
Worst Case	0.5	0.5	0.5	0.5	0.5	0.5	0.5

Table A.2: Digital activity factors for each bit of a 6 bit data bus for five test images and one computer generated image.

Each of the test images have both digital and analog activity factors lower than the uniform distribution which is over half as low as the worst case.

The five test images are displayed below.

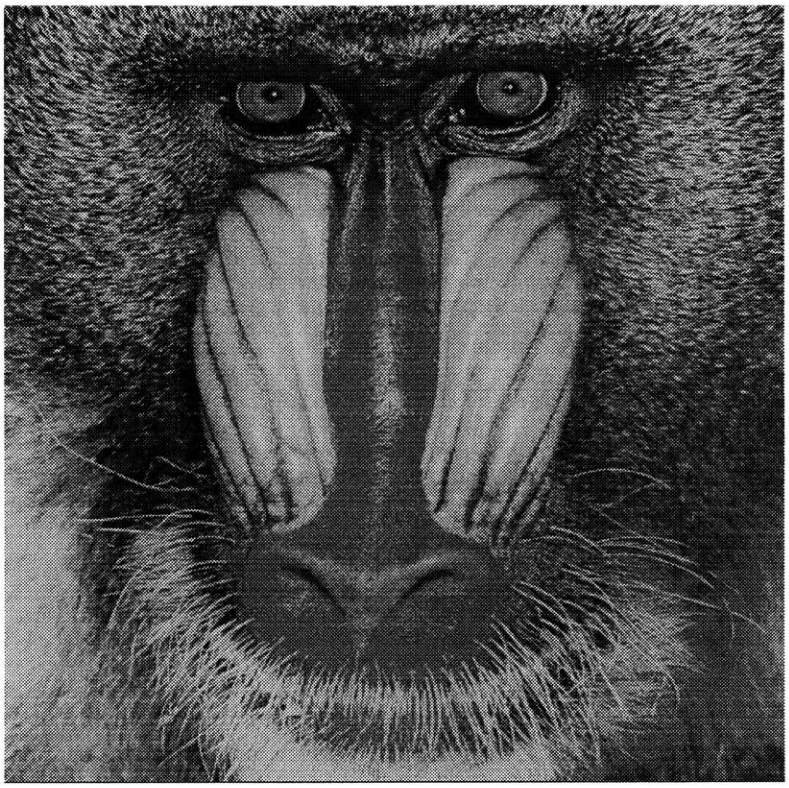

Figure A-1: Test Image 1: Baboon.

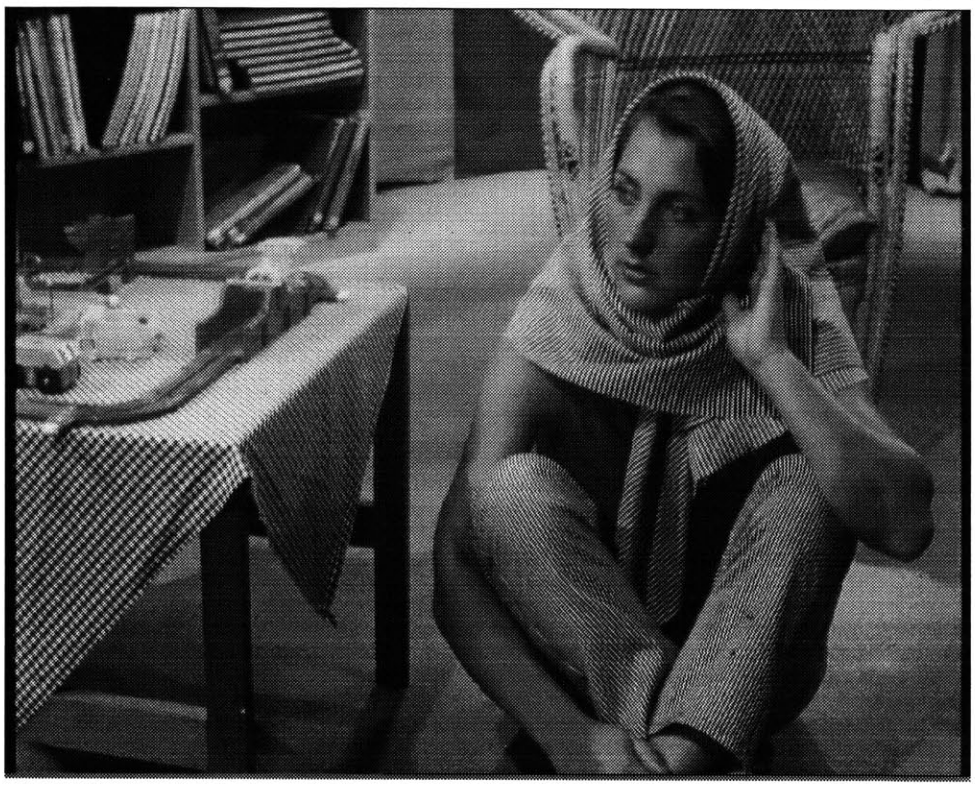

Figure A-2: Test Image 2: Barbara.

Figure A-3: Test Image 3: Goldhill.

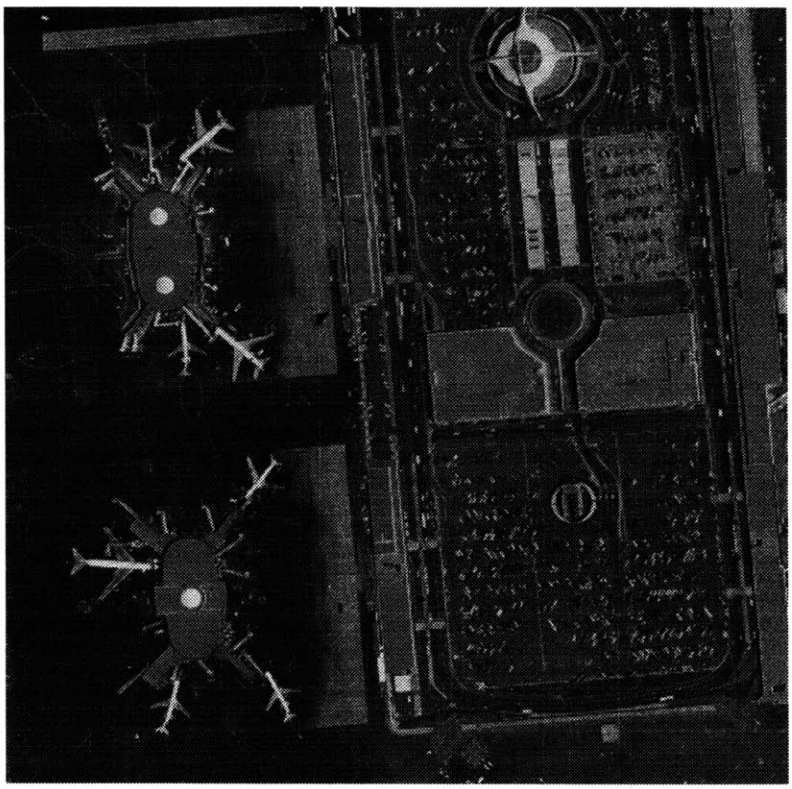

Figure A-4: Test Image 4: LAX.

Figure A-5: Test Image 5: Woman.

Bibliography

[1] M.J. Ammer, et al., *A 160 by 120 Pixel Liquid-Crystal-on-Silicon Microdisplay with an Adiabatic DAC.* To be published in the proceedings of the Interntational Solid State Circuits Conference. Feb. 1999.

[2] A. Chandrakasan and R. Brodersen, *Low Power Digital CMOS Design.* Kluwer Academic Press, 1996.

[3] W. Athas, J. Koller, and L. Svensson, *An Energy-Efficient CMOS Line Driver Using Adiabatic Switching.* In the Proceedings of the Fourth Great Lakes Symposium on VLSI Design, IEEE Press, pp. 159-164, Mar. 1994.

[4] W. Athas, et al., *AC-1: A Clock-Powered Microprocessor.* In International Symposium on Low Power Electronics and Design. pp. 18-20, 1997.

[5] M. Bolotski and P. Alvelda, *Low-Power Miniaturized Information Display Systems.* In International Symposium on Low Power Electronics and Design, 1998.

[6] Y. Matsueda, et al., *Low Temperature Poly-Si TFT-LCD with Integrated 6-bit Digital Data Drivers.* 1996 SID Technical Digest. pp. 21-24.

[7] J. McCreary and P. Gray, *All-MOS Charge Redistribution Analog-to-Digital Conversion Techniques – Part II.* In Journal of Solid State Circuits, vol. SC-10, pp. 379-385, Dec. 1975.

[8] D. Mercer and L. singer, *12-b 125 MSPS CMOS D/A Designed for Spectral Performance.* In International Symposium on Low Power Electronics and Design, 1996.

[9] E. Schlig and J. Sanford, *New Circuits for AMLCD Data Line Drivers.* International Display Research Conference, Monterey, CA, Oct 10-13, 1994, pages 386-389.

[10] N. Weste and K. Eshraghian,*Principles of CMOS VLSI Design* Second Edition. Addison-Wesley Press. 1993.

[11] S. Younis and T. Knight, *Asymptotically Zero Energy Split-Level Charge Recovery Logic.* International Workshop on Low Power Design, 1994. pp. 177-182