NOISE REDUCTION TECHNIQUES FOR CMOS SRAMS

APPROVED BY SUPERVISORY COMMITTEE:

Elizabeth E. Chain, Supervising Professor

Karl Renner, Supervising Professor

Blake E. Cherrington

NOISE REDUCTION TECHNIQUES FOR CMOS SRAMS

by

DORAN DAVID, BSEE

THESIS

Presented to the Faculty of

The University of Texas at Dallas

in Partial Fulfillment

of the Requirements

for the Degree of

MASTER OF SCIENCE

IN

ELECTRICAL ENGINEERING

THE UNIVERSITY OF TEXAS AT DALLAS

May, 1989

Acknowledgments

First and foremost, I wish to express my deepest appreciation to my advisors, Elizabeth Chain and Karl Renner, for their guidance, enthusiasm, and encouragement during the research and writing of this thesis. I wish to acknowledge and express appreciation to Ron Taylor and George Giles for their ideas and support of this research in noise reduction.

Appreciation is also expressed to Harrell Meador for his assistance in facilitating a lab set-up for the 8Kx9 SRAM.

Special thanks to Ari Reubin and Cal Knickerbocker for their assistance in the organization of this thesis.

Finally, I would like to thank my parents, Mr. and Mrs. Matana David, whose encouragement, patience, and understanding helped to make this work possible.

April, 1989

NOISE REDUCTION TECHNIQUES FOR CMOS SRAMS

Doran David, MSEE

The University of Texas at Dallas, 1989

Supervising Professors : Elizabeth Chain and Karl Renner

A 35-ns 8Kx9 CMOS SRAM is presented for analysis and design under the influence of power line noise, resulting from high-current output devices. Device architecture, functional operation, and performance are analyzed with respect to the noise problem. Beginning at the output driver stage, noise is reduced, and outputs are more quickly stabilized by the insertion of small-valued polycide resistances. These are placed in series with the inherent, parasitic package inductance. The resistors operate in conjunction with sensor devices, which selectively increase output driver impedance in accordance with voltage oscillations produced by the parasitic inductors. This circuit modification provides up to 50% reduction in power line noise at a cost of 2 ns speed loss for a 35-ns part. This method is enhanced by other circuit changes, which have been shown to correct functional test failures due to noise. These circuit enhancements are currently used on Texas Instruments' 0.9 μm 64K military SRAMs. Speed vs. power trade-offs are discussed for each design change. Design simulations and actual lab data illustrate the effects of noise before and after the circuit modifications are implemented.

Contents

List of Figures

Chapter 1

Introduction

This effort provides design solutions to the problem of power line noise for devices with transistor-transistor logic inputs. This is done using Texas Instruments' 64K military Static Random Access Memories, operating at 30–50 MHz in worst case noise conditions. These devices are fabricated using 0.9 μm silicon gate Complementary Metal Oxide Semiconductor technology with a 6-transistor memory cell. Static Random Access Memory configurations of 8Kx9, 16Kx4, and 8Kx8 are used to illustrate design changes for noise reduction.

There are five chapters :

- Chapter 1 defines noise margins for transistor-transistor logic (TTL) inputs; introduces package inductance; briefly describes Static Random Access Memory (SRAM) operation; and also defines some of the terminology and symbols used throughout the text.

- Chapter 2 provides a detailed discussion of the 8Kx9 SRAM architecture and operation. This includes the memory structure, the data I/O path, and a timing set which illustrates the AC/DC electrical specs.

- Chapter 3 defines where the noise problem originates, and shows three methods which can be used to reduce the noise magnitude.

- Chapter 4 continues with full-circuit simulations of 64K SRAMs, showing further circuit enhancements to add more noise immunity. Bench data is also shown to illustrate device performance with and without noise reduction.

- Chapter 5 concludes with a cumulative summary of the results of the circuit enhancements presented in chapters 3 and 4.

Chapter summaries are provided to highlight the main topics and results of each chapter.

1.1 Bipolar vs. CMOS Technology

Over the past 20 years, TTL has become the most widely used form of digital integrated circuit (IC), especially in the areas of small-scale and medium-scale integration (SSI and MSI). The popularity of this bipolar logic family, which includes standard TTL, low-power Schottky (Series 54/74LS and 54/74ALS), Advanced Schottky (Series 54/74AS), the Fast Series (54/74F), and emitter-coupled logic (ECL), led to a decline in interest in the Metal-Oxide-Semiconductor (MOS) transistor. However, due to the increasing demand for high-density parts with low power consumption, MOS technology became the basis for the large-scale integrated (LSI) digital memory and microprocessor circuits.

The system designer's choice of logic family is influenced by many factors. Among these are speed, power, cost, and compatibility with other parts of the system. As a result, many systems mix logic families, e.g., TTL and Complementary MOS (CMOS). In order to maintain compatibility with the already established TTL family, the IC designer faces the problem of TTL input noise margins for CMOS devices.

1.2 Noise Margins and TTL Inputs

Noise margin is a measure of the maximum allowable input noise voltage which will not affect the output behavior of a device. This specification is typically defined in terms of two parameters – the low noise margin, NM_L, and the high noise margin, NM_H. TTL input and output voltage specs determine the noise margins. The TTL specs currently used for Texas Instruments' (TI's) CMOS SRAMs are :

- VIL = 0.8V is the maximum input low voltage
- VIH = 2.2V is the minimum input high voltage
- VOL = 0.4V is the maximum output low voltage
- VOH = 2.4V is the minimum output high voltage

The noise margins are typically defined as the differences between the input and output voltage specs [5, p. 22] :

$$NM_L = VIL - VOL \tag{1.1}$$

$$NM_H = VOH - VIH \tag{1.2}$$

These equations give the spec noise margins for the case of one TTL gate driving another TTL gate. For the TTL voltage levels defined above, the spec noise margins are $NM_L = 0.4V$ and $NM_H = 0.2V$. However, for this effort, the major focus is on noise margins for TTL inputs with respect to power line noise, i.e., ground noise.

Meeting the TTL output voltage spec is generally not a problem. The major noise problem occurs when 1 to 2 volts of power line noise is fed back to TTL inputs that are designed to trip at 1.5V, the center of the VIL/VIH spec. This results in functional test failures. Active low device inputs, normally set to 0.8V, would have a low noise margin (NM_L) of 0.7V before they trip the input to the opposite state. Similarly, active high device inputs, normally set to 2.2V, would have a high noise margin (NM_H) of 0.7V

before tripping the input. For example, a 2.2V input level would be interpreted as 1.2V by a TTL input when ground carries +1.0V noise. Since 1.2V is below the 1.5V trip level, this input is mistaken for a low voltage level rather than a high voltage level. Similarly, a 0.8V input level would be interpreted as 1.8V when VSS carries −1.0V noise. This input would be mistaken for a high level rather than a low level, since 1.8V exceeds the 1.5V trip level.

1.3 Package Inductance

Due to the leadframe, bond wires, external routing wires, and other current path connections between the device and externally applied signals, a parasitic inductance results within these paths. The inductance creates an unwanted oscillation in the power lines. These voltage oscillations are fed back to the TTL inputs through the power lines and result in the noise problems discussed in the previous section. Section 3.2 addresses this problem in detail, and also shows how the power line oscillations create a constant noise feedback path to the TTL inputs.

1.4 SRAM Operation

Figure 1.1 shows a simplified block diagram of the address decode and data path for an 8Kx9 SRAM. Seven row addresses select one of 128 rows; six column addresses select one of 64 9-bit words. The ninth bit is generally a parity bit for improved system reliability. Thus, the SRAM is organized as 8,192 9-bit words.

Figure 1.1 shows a common I/O configuration, in which the nine DQ terminals serve as data inputs during a write cycle, as well as data outputs during a read cycle. Read or write mode is selected by the write enable input, W_. There is also an output enable terminal, G_, which sets the outputs to the high-impedance state, simplifying data bus design. The DQ terminals provide direct TTL compatibility with a fanout of 20 Series 54LS or 54ALS TTL gates, 16 Series 54AS TTL gates, or 13

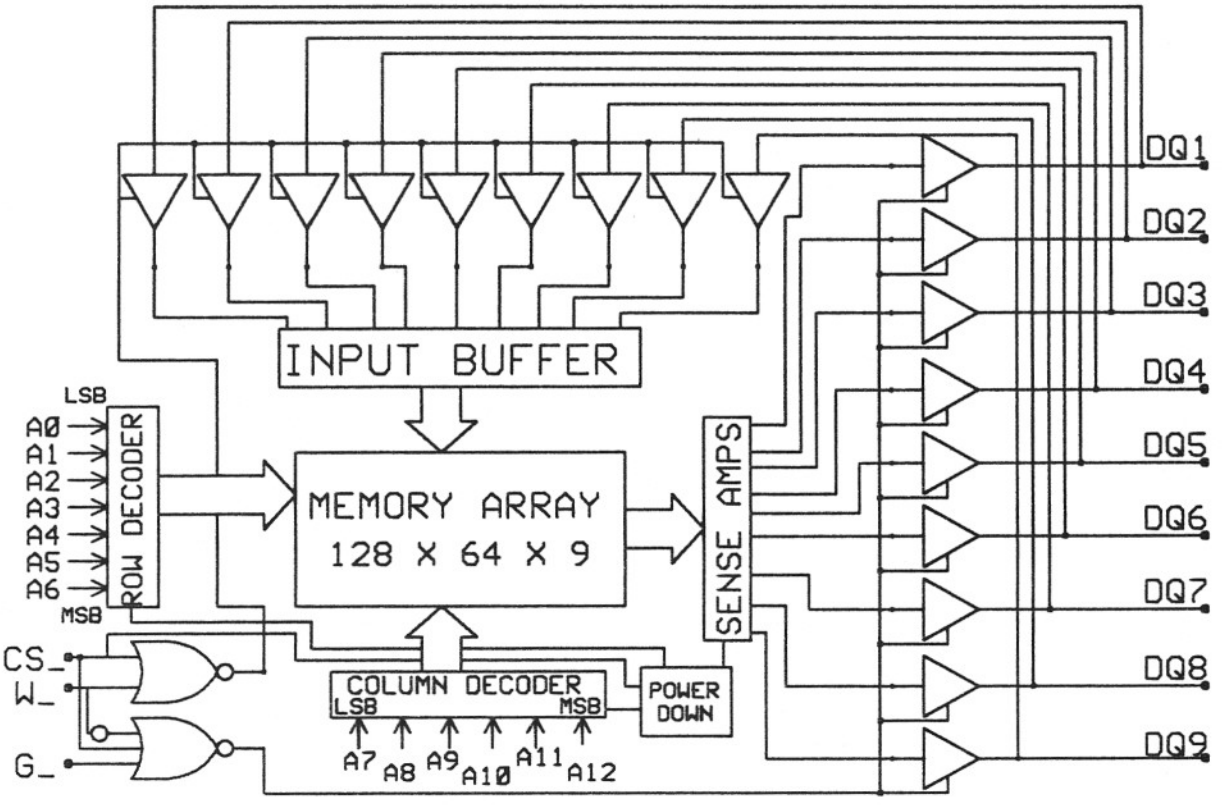

Figure 1.1: 8Kx9 SRAM Functional Block Diagram

Series 54F TTL gates. Because noise magnitude increases with the number of outputs switching simultaneously, the 8Kx9 SRAM configuration was chosen for most of the noise simulations in Chapter 3.

The device also includes a chip select/power down input, CS_. This input is used to disable the entire device, while retaining memory. This is also referred to as standby mode. Access time is classified by the maximum read cycle access from address or chip select, and is currently 25–45 ns.

Power dissipation is classified according to the active and standby modes of operation, as follows :

- Active Power.....660 mW MAX
- Standby Power....55 mW MAX (TTL Inputs)

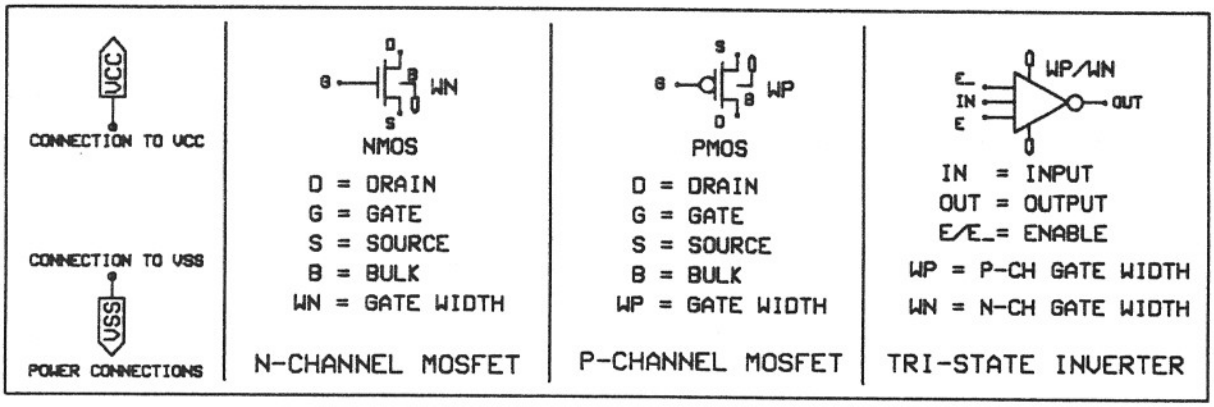

Figure 1.2: Symbol Definition

- Standby Power....5.5 mW MAX (CMOS Inputs)

Active power refers to the operating power consumption of the fastest (25 ns) part. Standby power refers to the power consumption during standby mode, and is at a minimum when CMOS input levels are used. By using a CMOS input level instead of a TTL level (e.g., 5.0V instead of 2.2V) on the chip select input, standby power decreases by a factor of 10.

1.5 Symbol Definition and Nomenclature

The symbols used in the schematics of this text include basic CMOS logic gates, with minor modifications, as defined in Figure 1.2. Standard CMOS logic gates (i.e., inverter, NAND, NOR) have been omitted from Figure 1.2.

1.5.1 Voltage Levels

Device power and ground will be VCC and VSS, respectively. Typical values are VCC=5.0V and VSS=0.0V. VCC and VSS also refer to internal device signal levels. The VCC voltage level will be referred to as a logic 1 level, and the VSS voltage level will be referred to as a logic 0 level.

For device inputs, VIL refers to the input low voltage level, and VIH refers to the input high voltage level. For device outputs, VOL refers to the output low voltage level,

and VOH refers to the output high voltage level. Voltage levels with VCC/VSS logic swings are referred to as CMOS levels. TTL voltage levels refer to the input/output voltage levels defined in Section 1.2. There is also a high-impedance state (or tri-state level) for the outputs, when they are disabled. This level is determined by the Thevenin equivalence of the tester load (see Section 3.1).

1.5.2 Devices and Naming Convention

After the power symbols in Figure 1.2 are the n-channel and p-channel transistor symbols. All n-channel devices have their substrates (bulk) tied to VSS, and all p-channel devices have their substrates tied to VCC. WN and WP are the n-channel and p-channel transistor gate widths, specified in microns. All transistor gate lengths are currently 0.9 μm for TI's 64K/72K military SRAMs. The gate lengths have been omitted from all figures.

Signal names that end in an underscore are the logical complement of that signal without the underscore; e.g., signals OE1 and OE1_ in Figure 3.1 are logical complements of each other. All logic gates will be referred to by their output node; e.g., the 2-input NAND gate in Figure 3.1 will be referred to by its output at node D. All n-channel and p-channel mosfets will have unique names, such as M1 and M2 in the same figure. Figure 3.1 also shows the gate widths next to the transistor symbols.

Figure 1.2 also shows a tri-state inverter. When the enable signals E and E_ are set to VCC and VSS, respectively, the output is the complement of the input, as in a standard inverter. When the enable signals are set to the opposite state, the output is in high-impedance, regardless of the input. This device is used to perform a signal latching function in Section 3.4.

1.5.3 Process Corner

Process corner refers to transistor process and environment conditions which affect device performance. Various process corners are used to characterize appropriate parameters for device performance. For example, a slow process corner is used to determine worst case access time, while a fast process corner is used to show worst case noise levels.

Fast process refers to process conditions which yield the fastest part (e.g., +3 sigma, or 3 standard deviations in process distribution towards the fastest part). Similarly, a slow process may represent a −3 sigma part. Environmental conditions include temperature and power supply voltage. High and low temperature will be 150C and −55C, respectively (military spec). Power supply (VCC) will range from 4.5V to 5.5V. Process and environment conditions taken together can be used to define the following three process corners :

- Fast process, −55C, 5.5V supply refers to a fast process corner
- Nominal process, 25C, 5.0V supply refers to a nominal process corner
- Slow process, 150C, 4.5V supply refers to a slow process corner

1.6 Chapter Summary

This chapter defines the TTL input specs and TTL noise margins, introduces package inductance noise and basic SRAM operation, and also defines some terminology to be used in the following chapters. Section 1.2 indicates that ± 0.7V of ground noise can trip a TTL input.

Chapter 2

Device Structure and Operation

2.1 Device Architecture

The basic SRAM architecture is shown in Figure 2.1. There are four essential functions :

1. Data I/O path
2. Memory matrix and column decoding
3. Row and column addressing
4. Control buffer circuitry

The data I/O path and control buffer circuitry will be the major points of focus for noise reduction, and are discussed in the following sections.

The 8Kx9 SRAM is configured into 16 blocks of memory, each containing 4,608 bits (or 512 nine-bit words). Row and column addressing select one of the 16 blocks at any time, leaving the other 15 blocks powered down. This is done to decrease power consumption during a read or write operation. Each block matrix contains 64 rows and 72 columns. The 72 columns are subdivided into nine groups of eight columns each. A selected *word line* refers to a 72-bit row within a block, in the active logic 1 state. Each column within a block is also referred to as a bit-line pair, BLT and BLC, for bit line true and bit line complement, respectively.

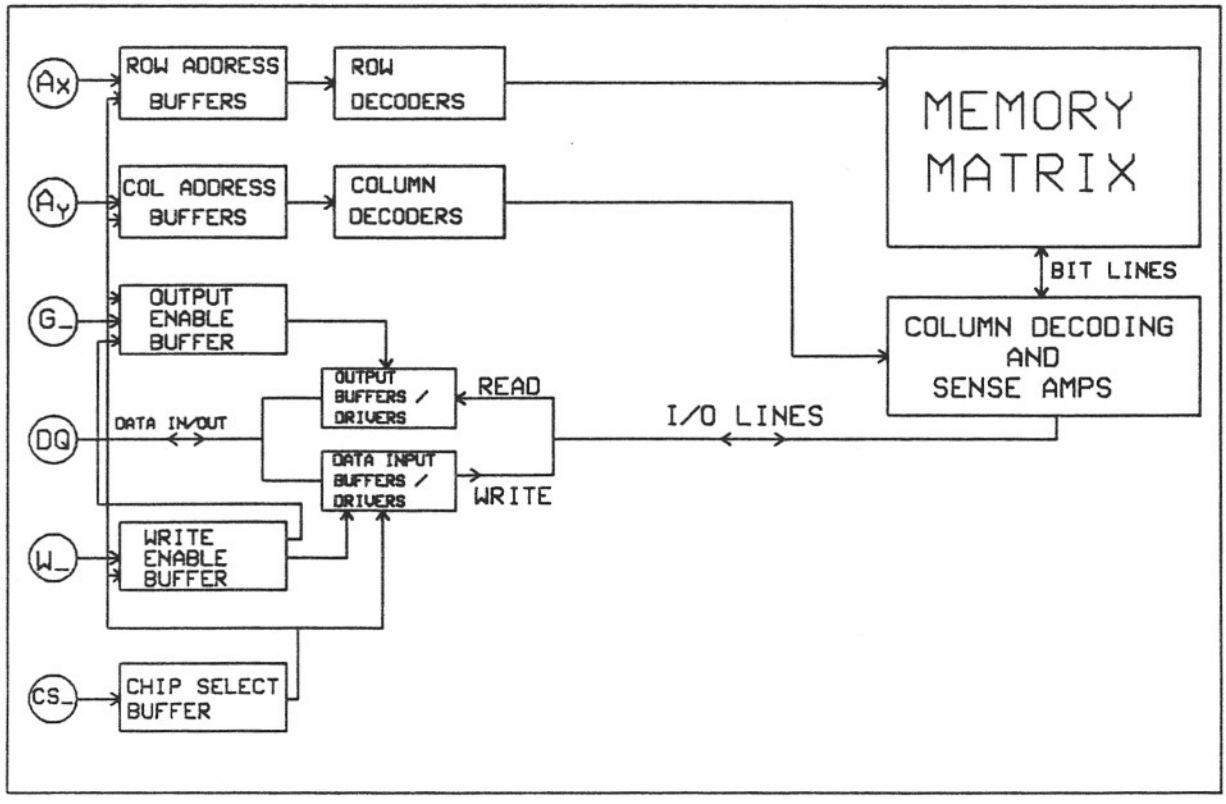

Figure 2.1: SRAM Architecture

During a read or write cycle, one word line is accessed, along with one of eight columns from every group of eight within the selected block. This gives a total of nine selected bits for each word to be read from or written to the memory matrix.

The I/O lines shown in this figure carry write data from the data input drivers to the memory matrix. These lines are also used during a read cycle to carry data from the sense amps to the output buffer/driver circuits. The control buffer section consists of the chip select, write enable, and output enable buffer, and is discussed in Section 2.3.

2.2 Input Buffer Functions

The input buffer circuits serve three main functions. Due to the TTL compatibility requirement of the device, TTL to CMOS voltage level conversion must be done for

all external device inputs. The TTL spec currently used for 64K/72K SRAMs (See Section 1.2) calls for a maximum input low level (VIL) of 0.8V, and a minimum input high level (VIH) of 2.2V. Since CMOS logic levels of VCC/VSS (typically 5.0V/0.0V) are used for all major internal signals as well as supply lines to the device, a voltage level conversion is required for all inputs.

The other functions of the input buffer circuits are to amplify signal strength and generate critical timing paths required by the read and write functions. These may include delaying the start of write until row and column are properly accessed, or, on the other hand, speeding up the end of write before the next row or column access. These conditions are required to pass the address setup, $t_{su}(A)$, and the address hold, $t_h(A)$ timing specs defined in Section 2.6.1.

2.3 Control Buffers

The control buffers set the operational state of the device. When the device is in standby mode, with CS_ at the VIH level, the outputs are set to the high-impedance state. All row/column addressing is disabled, as is the write command. In this mode, the device retains memory in the entire matrix and consumes under 55 mW of static power. All other functions are performed with the device enabled (i.e., CS_ at VIL).

The output enable feature (G_ device input) is used for bus control, and only affects the data output terminals. When G_ is set to VIH, the output terminals are set to the high-impedance state. When G_ is set to VIL, the outputs (see Figs. 1.1 and 2.1) read a 9-bit word from the matrix. The read is only performed when write command W_ is at the VIH level. When W_ is set to VIL (write mode), the outputs are set to the high-impedance state, regardless of the state of output enable G_. This allows the DQ terminals to write a desired 9-bit word to the matrix.

VSS noise fed back to either W_, G_, or CS_ during a read cycle drives the outputs to the high-impedance state. This creates oscillations between the read state

and the high-impedance state and generates more noise that can be fed back to the inputs through the power lines. This feedback oscillation will be shown in the next chapter.

2.4 Data I/O

The data I/O path is governed by the write enable buffer, the data input buffer, and the output enable buffer.

During a read cycle, G_ and CS_ are set to VIL, W_ is set to VIH, and the nine DQ outputs read a word from the nine memory cells accessed by the row/column address pins. As data leaves the memory matrix, sense amps (see Figure 2.1) restore CMOS logic levels to the data from the bit lines to the I/O lines. The I/O lines drive the data output buffers/drivers.

During a write cycle, CS_ and W_ are set to VIL, and the output drivers are cut off. Write data enters at these terminals, and is sent to the data input buffers via the write command. The information is driven to the I/O lines, and then directly to the matrix cell locations specified by the row and column addressing.

2.5 Memory Cell, Bit Line Pull-Up, and Equilibration

2.5.1 Memory Cell and Bit Line Pull-up

The memory array consists of 6-transistor cells, as shown in Figure 2.2. The figure also shows a bit line pull-up circuit, which is used to precharge the bit lines.

The memory cell is accessed by setting the word line (WL) to the logic 1 level [7, p. 350]. This enables access mosfets MN2 and MN3, which transfer the true and complement data from the cell to bit lines BLT and BLC, respectively. The other four transistors form an inverter-type latch, which stores one bit of data. Transistor sizings (i.e., gate widths) are chosen so that data can be written to the cell only when

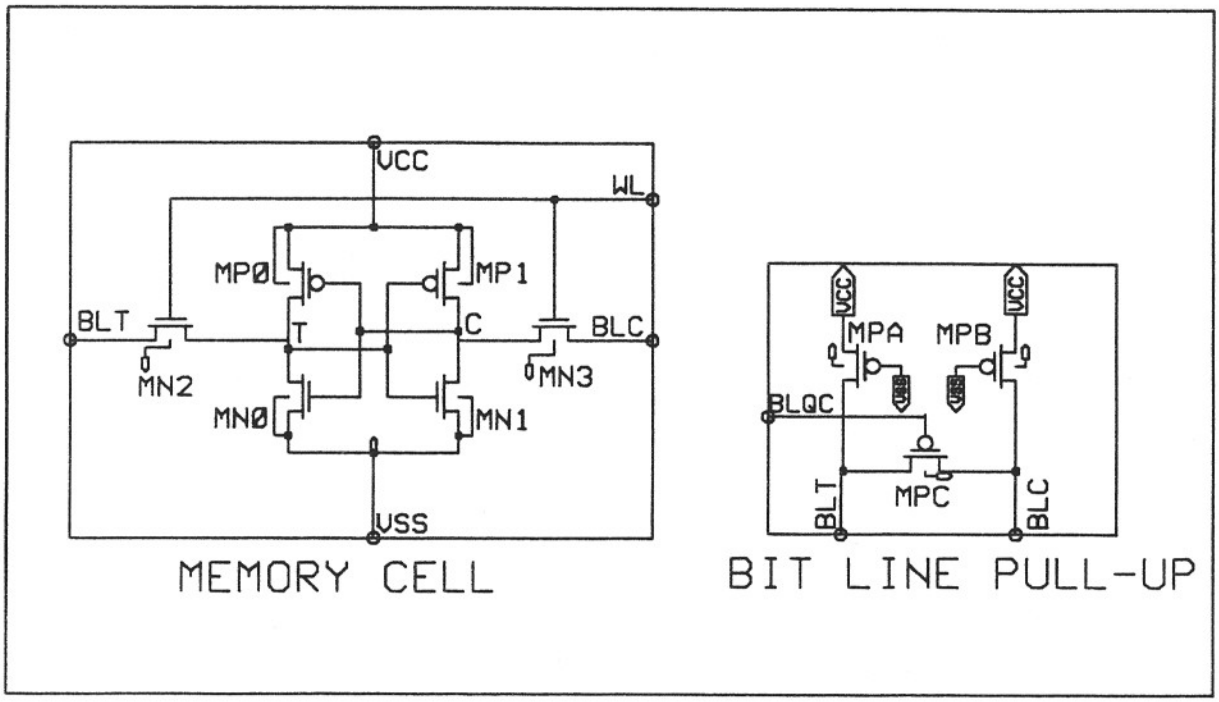

Figure 2.2: Memory Cell and Bit Line Pull-Up Circuit

one side (true or complement) is pulled down below a threshold, say 20% of VCC [7, p. 353]. This allows the cell to remain in its stable memory state when the word line is active and both bit lines are in the logic 1 state. This condition occurs for the 63 unused cells along a 72-bit active word line, when the other nine cells are being read or written.

The bit line pull-up circuit is placed on top of each column (or bit line pair) within each block of memory. Pulling the bit lines high helps to prevent writing undesired cells, as discussed above. The bit line pull ups, MPA and MPB, are sized so as not to overpower the series latch transistors MN0 and MN2 (or MN1 and MN3) during a read. This allows the true and complement bit line pairs to develop large enough voltage differentials for the sense amps during a read cycle.

Another purpose for the bit line pull-ups is to speed up access by preventing large logic swings between the true and complement bit line pairs. During a read cycle, with VCC set to 5.0V, typical bit line low and high levels would be 3.5V and 5.0V,

respectively. Since the bit line read levels are not CMOS levels, sense amps are required to restore these levels back to CMOS levels, as shown in Figure 2.1. The sense amps drive the I/O lines to the corresponding output buffers. During a write cycle, the sense amps are disabled, and the I/O line write data is sent to the bit lines. The write is dominated by the bit line with the logic 0 voltage level.

In order for the write data to overcome the latch, the data input drivers (see Figure 2.1) must be sized large enough to drive the I/O lines and bit lines (columns), and invert the memory cell. The logic 0 write data has to overcome the logic 1 state of the corresponding bit line pull-up device, MPA or MPB in Figure 2.2, and the corresponding memory cell pull-up, MP0 or MP1 in the same figure. This is done by using large transistor widths (e.g., 50–100 μm) for the data input driver pull-downs, which drive the I/O lines during a write cycle.

2.5.2 Write Recovery and Bit Line Equilibration

The bit line pull-up circuit of Figure 2.2 also shows a signal, BLQC, fed to a p-channel device between the true and complement bit lines. This *equilibration* signal helps prevent write cycle failures, and also speeds up access time.

BLQC serves the purpose of *bit line equilibration*, which pulls the bit lines together when it is at logic 0. This action helps the bit lines quickly recover from their large voltage differential after a write cycle. If the bit lines do not recover (or pull up) fast enough after the write cycle, they may have enough potential to destroy (write) the next address location selected by the following read cycle. By quickly restoring the bit lines to the logic 1 (non-destructive) state, the data in the next cell location can be safely read.

Equalization of bit lines after an address change is also a common practice for speeding up access time. This is also referred to as Address Transition Detection (ATD), where the array is equalized and the RAM is shut down after an address

change [4, p.208]. Bit line equilibration prevents large logic swings on the bit lines after each address change. This resets the bit lines to a common logic state, ready to respond more rapidly to the next data change. In order for ATD to speed up access, the equilibration signals must occur faster than a word line or column access.

2.6 Electrical and Timing Specs

2.6.1 Spec Definition

The 8Kx9 SRAM is designed to meet military specs at various process and test conditions. Power consumption is classified by the active and standby power, as defined in Section 1.4. All input/output voltage levels meet TTL requirements, as described in Section 1.2. Device speed is classified by the read cycle address access time, $ta(A)$. TI is currently offering 25, 35, and 45 ns parts. Other timing parameters include (for 35ns parts) :

$t_a(CS)$ Access time from chip select low; current spec is 35 ns. This is the read access time after the falling edge of CS_.

$t_{dis}(CS)$ Output disable time from chip select high; current spec is 15 ns. This is the time delay for the outputs to shut off after the rising edge of CS_.

$t_a(OE)$ Access time from output enable low; current spec is 20 ns. This is the read access time after the falling edge of G_.

$t_{dis}(OE)$ Output disable time from output enable high; current spec is 15 ns. This is the time delay for the outputs to shut off after the rising edge of G_.

$t_w(W)$ Write enable pulse duration; current spec is 20 ns. This is the time required for W_ to remain low in order to perform a write cycle.

$t_{su}(A)$ Address setup time to write start; current spec is 0 ns. This means that address may change at the same time as the falling edge of W_.

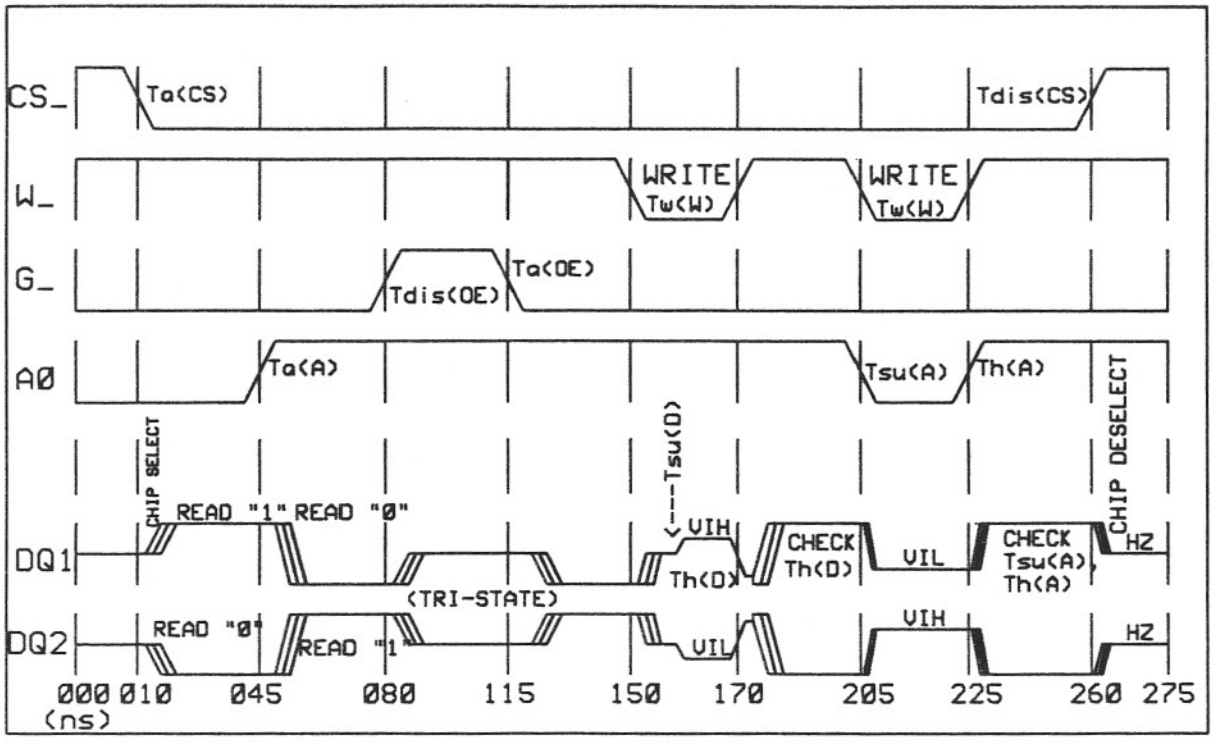

Figure 2.3: SRAM Timing Diagram

$t_h(A)$ Address hold time from write end; current spec is 0 ns. This means that address may change at same time as rising edge of W_.

$t_{su}(D)$ Data setup time to write end; current spec is 15 ns. This means that write data must be valid at the DQ terminals 15 ns prior to rising edge of W_.

$t_h(D)$ Data hold time from write end; current spec is 0 ns. This means that input data may be changed at same time as rising edge of W_.

2.6.2 Graphical Illustration of Timing Specs

The above timing parameters are shown in the timing diagram, Figure 2.3. The diagram shows how chip select powers up the device at the start of the timing set, and also shows the device powered down at the end of the timing set. This first write cycle illustrates data setup and hold; the second write cycle illustrates address setup and hold. The read cycles are used for access time measurements, and also for verification

of the write cycles. Two outputs, DQ1 and DQ2, which read and write complementary data, are shown. This is done to show both low-to-high and high-to-low transitions at each cycle. The shaded regions on these signals represent output transitions, where data changes to the opposite state.

All device inputs use TTL levels. These include control inputs CS_, W_, G_, and address A0. DQ1 and DQ2 also use the TTL VIL/VIH input levels for the write data in each write cycle. During a read cycle, the DQ output levels meet the TTL VOL/VOH output spec; i.e., low level below 0.4V, and high level above 2.4V.

Chip Select and Address Access

The timing set begins with the device disabled by chip select (CS_). This sets DQ1 and DQ2 to the high-impedance state. At 10 ns, CS_ enables the device into the first read cycle. Chip select access ($t_a(CS)$) is measured from 10 ns to the time the outputs reach their correct voltage levels. In this case, DQ1 must reach 2.4V (VOH), and DQ2 must reach 0.4V (VOL).

At 45 ns, A0 changes state and a row access is performed. This selects a different word line within the current memory block. The timing set shows the outputs reading opposite data in the new location. Address access time ($ta(A)$) is measured from 45 ns to the proper DQ1 or DQ2 output level. In this example, A0 is assumed to be a row address. However, A0 may also represent a column address, or even multiple addresses, switching.

Output Enable Access

The output disable time ($t_{dis}(OE)$) can be measured in the next cycle, where G_ switches high at 80 ns. This delay is measured from 80 ns to the time the data output drivers are cut off. At 115 ns, G_ releases the outputs. The G_ access time ($t_a(OE)$) is measured from 115 ns to the time the outputs reach their previous read state levels.

Data Setup and Hold

The first write cycle is performed at 150 ns. The data in the current matrix locations, addressed by A0 at VIH, is written to the opposite state by DQ1 and DQ2.

The data setup spec, $t_{su}(D)$, is shown with the write data valid at this time spec before the rising edge of W_. Zero data hold ($t_h(D) = 0$ ns) is also shown with DQ1 and DQ2 changing at the end of the write cycle. Negative data hold time would require the data to change before the end of write. This condition may rewrite the cells back to their original state, failing the write cycle. The zero data hold spec requires the write cycle to pass if data changes at the same time as the end of write (i.e., rising edge of W_).

The written data is verified by the following read cycle at 170 ns. A data hold failure occurs if DQ1 and DQ2 do not read the same data that was written.

Address Setup and Hold

The next write cycle at 205 ns checks address setup and hold rather than data setup and hold. DQ1 writes a logic 0, and DQ2 writes a logic 1 to the location addressed by A0 at VIL.

A0 switching from high to low at the start of write illustrates zero address setup time. Similarly, zero address hold time is illustrated by A0 switching from low to high at the end of the write cycle. In both cases, the data in the read location, addressed by A0 at VIH, is in danger of being written in addition to the desired write location, addressed by A0 at VIL.

The outcome is seen in the following read cycle at 225 ns. Zero address setup and hold times are passed if the data from the previous read cycle (addressed by A0 at VIH) is read again.

If the last read cycle fails, the failure is classified as either an address setup ($t_{su}(A)$), or an address hold ($t_h(A)$) failure. This can be determined by checking the cell contents

of the read location at the start and end of the write cycle. If this cell is written after 205 ns, the failure mode is *address setup to write start.* If this cell is written after 225 ns, the failure mode is *address hold from write end.* In order to guarantee a spec of 0 ns for both parameters, the part is designed around a tighter spec, such as -2 ns. In this case, A0 would switch high to low at 207 ns, and then switch low to high at 223 ns.

Power Down

Finally, the device is powered down by CS_ at 260 ns. The time it takes to shut off the output drivers is the output disable time from CS_ high, $t_{dis}(CS)$. This is typically under 10 ns for a 35-ns part.

2.7 Chapter Summary

The operation of the 8Kx9 SRAM has been presented. The read and write operations were analyzed with respect to the memory matrix, the data I/O path, and the timing specs. The next chapter takes a close look at noise in the output stage of the 8Kx9 SRAM.

Chapter 3

Noise Reduction in Output Stage

3.1 Source of Noise and Inductive Isolation

The 8Kx9 output terminals are connected to a pair of high-current CMOS transistors. These are shown as M1 and M2 in Figure 3.1, and make up the output driver section for each output pad. These drivers are, in effect, current amplifiers for the low current outputs of the memory array. The output buffer/driver circuit is basically a tri-state pad, controlled by the OE1/OE1_ enable signals [7, p. 229].

The output current spec requires that the drivers source 4 ma and sink 8 ma of current. This spec, combined with the VOL/VOH requirement for the output voltage levels, determines a Thevenin equivalent for the output driver I–V characteristics as follows :

$$\frac{2.4V - VTHEV}{RTHEV} = 4\ ma \tag{3.1}$$

$$\frac{VTHEV - 0.4V}{RTHEV} = 8\ ma \tag{3.2}$$

These equations yield VTHEV=1.73V and RTHEV=167Ω. These values make up the tester load circuit shown in Figure 3.2. This figure illustrates all power line connections between the device, the package and the tester, and also shows the package inductance responsible for the power line noise. The 30 pf capacitor is the timing spec load for speed classification. In order to meet the current requirements, the output drivers are sized from 500 to 1000 μm gate width, allowing them to sink or source static current

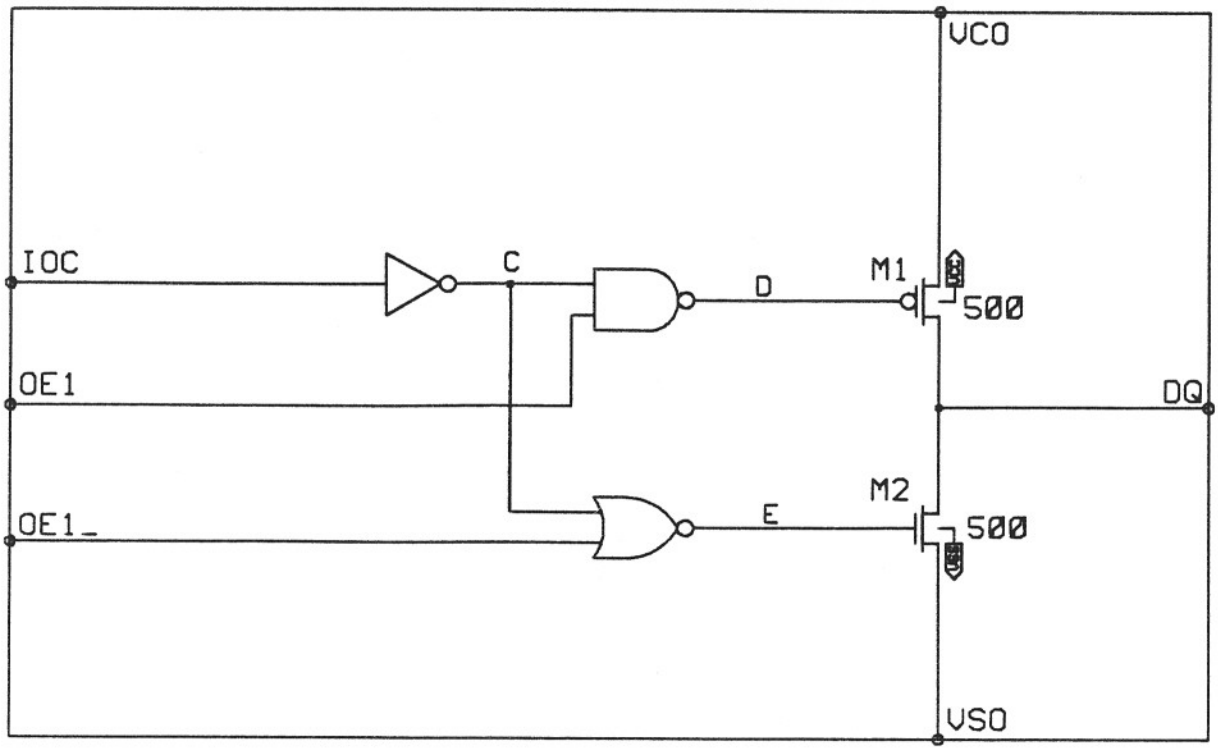

Figure 3.1: Output Buffer and Driver Circuit

levels from a few milliamps to tens of milliamps. During output transitions, current gains up to 100 are common for these drivers.

Due to the leadframe, bond wires, tester leads, and other current path connections between the device and externally applied signals, a parasitic inductance results. The inductance slows down the effective speed of the drivers. This inductance creates an unwanted oscillation in the output of the drivers and the driver supply wiring. These oscillations are fed back to the inputs through the power lines. A noise feedback path results, which induces even more oscillations in the outputs. This is discussed in the next section, where the power line noise is fed back to the G_ input pin.

A common method used to help isolate the inductance noise from the VCC/VSS power lines is the use of a separate supply connection for the output drivers. These new output driver supply nodes are designated by VCO and VSO in Figures 3.1 and 3.2. As shown in Figure 3.2, separate bond wires and leadframe paths are taken from the

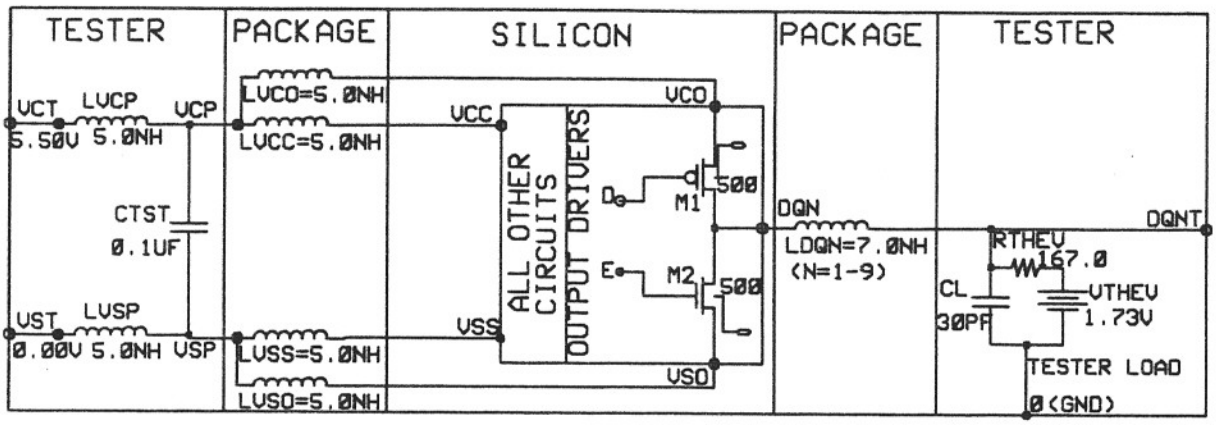

Figure 3.2: Package Inductance Circuit

package pin supplies, VCP and VSP. These are lumped into the equivalent inductances LVCO and LVSO. Similarly, LVCC and LVSS are the bond wire and leadframe inductances for device power and ground. Thus, the VCO/VSO output driver supplies are inductively isolated from the VCC/VSS supplies used by the rest of the device.

All supply inductances have a spec of 5.0 nh. Each of the nine DQ output inductances has a spec of 7.0 nh. These values are used in simulation to give worst case noise levels, comparable to the levels actually measured during production testing of the device without noise reduction. Other elements in Figure 3.2 include a 0.1 μf tester bypass capacitor across the package VCP and VSP power pins, and the tester supply power line inductances, LVCP and LVSP.

3.2 Read Path Simulation

3.2.1 Output Enable Function

Figure 3.3 shows a simplified output enable buffer. This circuit is used to control the output state, depending on the CS_, W_, and G_ device inputs.

The first stage is a TTL NOR gate, which is specially sized to handle the TTL to CMOS voltage level conversion for the G_ device input. CS1_ is a global chip select signal, and is the output of the chip select buffer (see Fig. 2.1). This signal has the

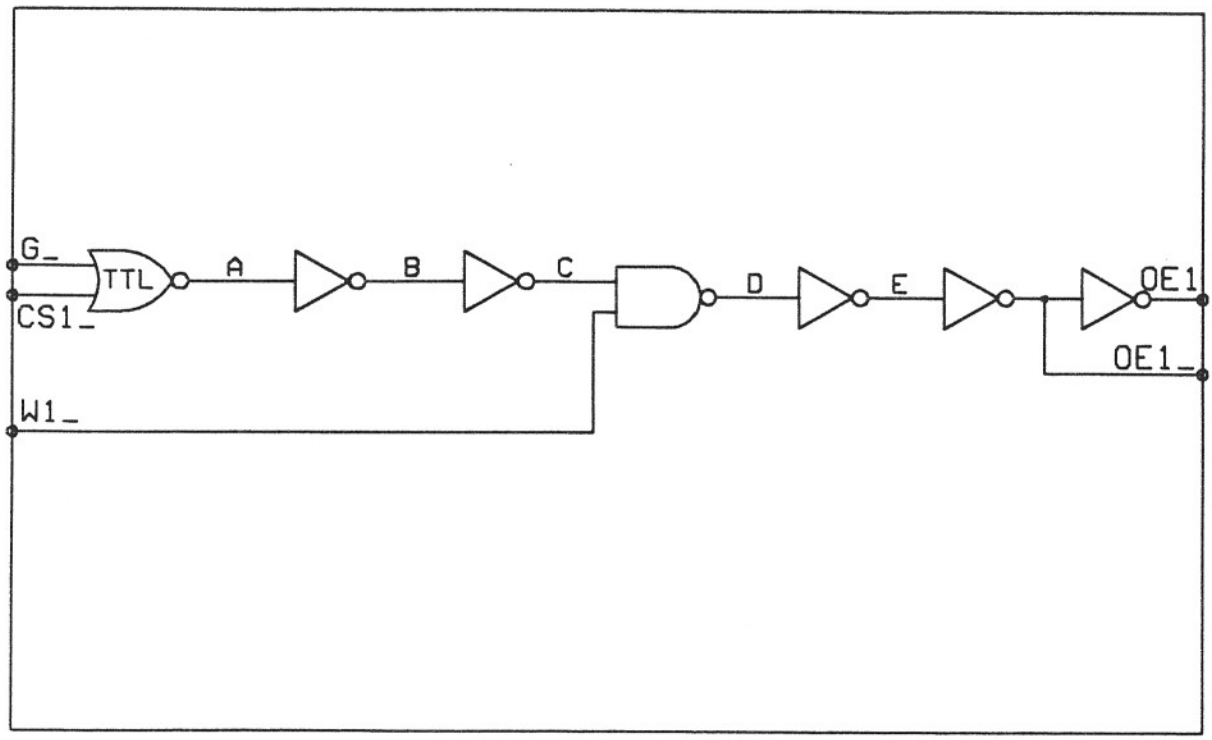

Figure 3.3: Output Enable Buffer Circuit

same logic as the CS_ input (active low), but includes the TTL to CMOS voltage level conversion with enough current drive to power down all circuits in the device. For the purpose of output buffer simulation, CS1_ is set to VSS in the circuit of Figure 3.3. Similarly, W1_ is an output from the write enable buffer, and is at logic 1 during a read. This signal is set to VCC for the simulation.

The outputs of Figure 3.3 are the output enable signals, OE1 and its complement, OE1_. These signals control the state of the output buffer of Figure 3.1. The output is enabled as long as OE1 is at logic 1, and OE1_ is at logic 0.

The outputs are disabled (OE1 at logic 0) when any one of the following three conditions occurs :

- G_ input is set to VIH (output disable)
- W_ input is set to VIL (write mode)
- CS_ input is set to VIH (power down/standby mode)

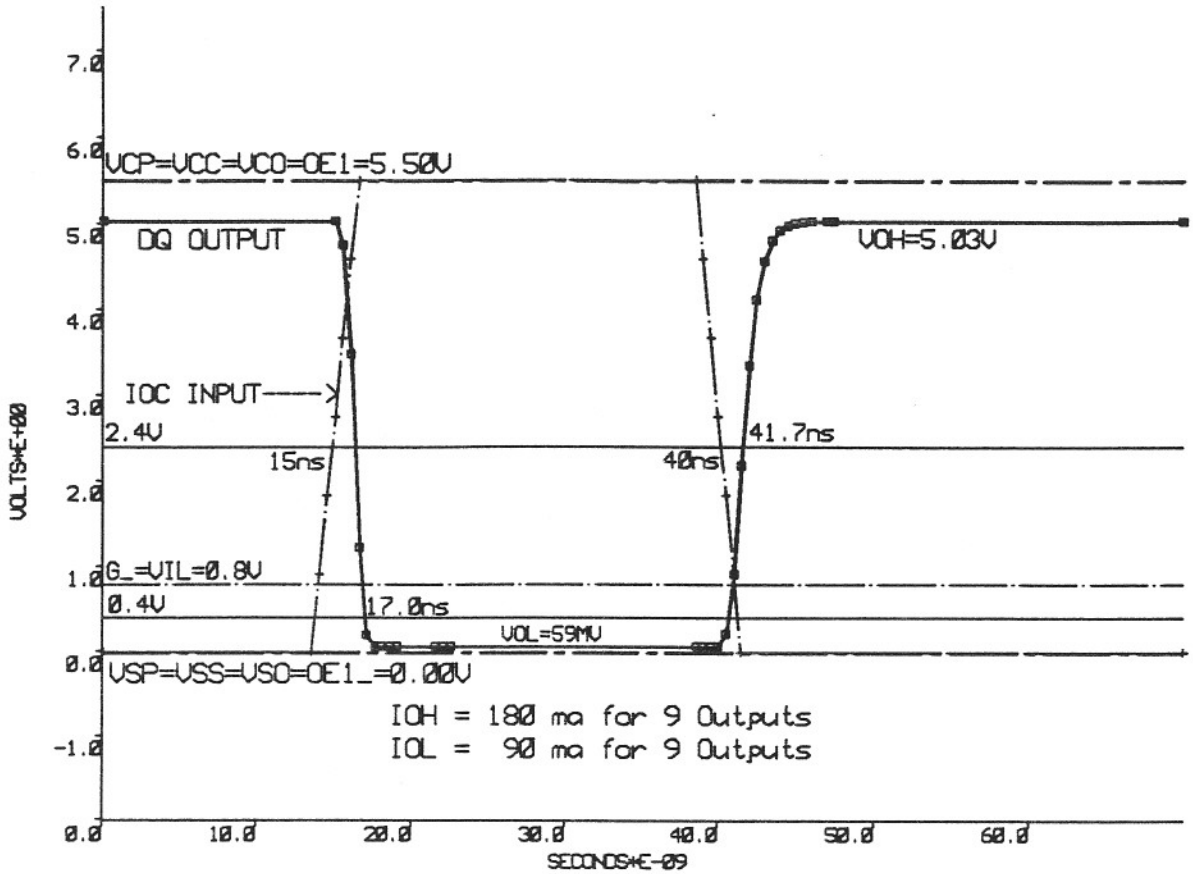

Figure 3.4: Output Buffer Fast Simulation, Ideal Case

When the OE1/OE1_ signals are in the disable state, node D in Figure 3.1 is at logic 1, while node E is at logic 0. This condition cuts off the output drivers, setting the DQ outputs to the high-impedance state.

When OE1 and OE1_ are in the enabled state, the DQ1-DQ9 outputs become the logical complement of the corresponding IOC line (IO1C-IO9C), and the read cycle is complete. The IOC lines are the outputs of the sense amps, and are shown in the architecture diagram of Figure 2.1.

3.2.2 Ideal Simulation

A simulation for the output buffer of Figure 3.1 and the output enable buffer of Figure 3.3 is shown in Figure 3.4. This figure illustrates the ideal behavior, without

package inductance.

The output enable buffer is set to the active state with G_ at 0.8V, CS1_ at VSS, and W1_ at VCC. This combination sets signals OE1 to VCC and OE1_ to VSS, which enable the drivers in the output buffer. In order to simulate a worst case noise test, the fast process corner is used (ref. Sec. 1.5.3). In this process corner, VCC is set to 5.5V, slightly higher than the nominal VCC of 5.0V. Fast process transistor models are used, and ambient temperature is set to -55C, which is the minimum temperature for military operation. These conditions will give the fastest part with the most noise.

When all seven parasitic inductors shown in Figure 3.2 are removed (shorted), the ideal device behavior can be simulated. This is shown in Figure 3.4. Note that the three supply points, package supply VCP, device supply VCC, and output driver supply VCO all remain at the constant voltage level of 5.5V. Similarly, VSP, VSS, and VSO all remain at the 0.0V ground level.

The simulation runs for 70 ns, with the IOC input switching at 15 and 40 ns. The DQ output switches to the correct state, with noise-free transitions. By using a SPICE multiplication factor (i.e., M=9), the single DQ output waveform simulates the effect of all nine outputs switching together. This gives nine times the current drive of a single output. The tester load Thevenin resistance is divided by nine to maintain the same output voltage levels seen by one output.

Verification of Output Current

The simulation gives a source current of 180 ma and a sink current of 90 ma for nine outputs with 500 μm width drivers. This corresponds to 20 ma source current, and 10 ma sink current for a single output. These numbers can also be verified by a hand calculation, using the linear resistance equations for the 500 μm n-channel and p-channel drivers as follows :

The channel resistance of each device in the linear region is given by [7, pp. 40–42] :

$$R_C = \frac{1}{\beta|V_{GS} - VT|} \tag{3.3}$$

where

$$\beta = \frac{\mu\epsilon}{t_{ox}}\frac{W}{L} \tag{3.4}$$

and the factor

$$KP = \frac{\mu\epsilon}{t_{ox}} \tag{3.5}$$

can be extracted from the device models file along with the n- and p-channel VT values. These values from the fast models file are :

- VTN=0.86V and KPN=$47.7\mu a/V^2$ for the n-channel device
- VTP=-0.67V and KPP=$10.2\mu a/V^2$ for the p-channel device

Using $W = 500\ \mu m$, $L = 1\ \mu m$, and $|V_{GS}| = 5.5V$ for the linear region, the channel resistances become :

- $R_{CP} = 40.6\Omega$ p-channel resistance
- $R_{CN} = 9.0\Omega$ n-channel resistance

The source current is given by

$$IOH = \frac{VCO - VTHEV}{RTHEV + R_{CP}}, \tag{3.6}$$

and the sink current is given by

$$IOL = \frac{VTHEV}{RTHEV + R_{CN}}. \tag{3.7}$$

Using $VTHEV = 1.73V$, $RTHEV = 167\Omega$, $VCO = 5.5V$, and the channel resistances calculated above, the calculated values for IOH and IOL using Equations 3.6 and 3.7 are 18.2 ma and 9.8 ma, respectively. These values are within 9% of the simulated values of 20 ma and 10 ma.

These values differ from the IOH/IOL spec currents in Equations 3.1 and 3.2. The reason is that the spec currents assume the worst case VOH/VOL levels of 2.4V/0.4V, which are usually surpassed due to the CMOS stage. Thus, IOH=180 ma, IOL=90 ma (for nine outputs), and output levels are VOH=5.03V and VOL=59 mv, as shown in Figure 3.4. The output enable signals, OE1 and OE1_, remain at their active levels of 5.5V and 0.0V, respectively.

The output low-to-high delay time is the time it takes the output to reach the VOH level of 2.4V, after the falling edge of the IOC input. In the ideal case, this is from 40.0 to 41.7 ns, or 1.7 ns time delay. This is not a spec delay time. It is only the time it takes the IOC transition to propagate through the three logic stages of Figure 3.1.

3.2.3 Package Inductance Simulation

Figure 3.5 shows the same simulation after all seven parasitic inductances have been added. VCC shows noise spikes as high as five volts, while VSS shows up to 2.8 volts of noise.

An example of output noise oscillation is also shown in this figure. After the second IOC input transition, the output goes into oscillation around 3.7V, as it tries to reach its steady state value of 5.03V, achieved before the first transition. The output remains in constant oscillation even 30 ns after the second transition. Note that the output low-to-high delay time has increased from 1.7 ns in the ideal case to 5.2 ns after the package inductance was added.

The part is not likely to fail DC levels after the second transition, because of the low VOH spec of 2.4V. However, the VSS noise induced by the VSO supply inductance (see Fig. 3.2) causes the output to cross the 0.4V VOL line 10 ns after the first crosspoint at 18.0 ns. This is a definite failure due to ground noise. Because the source current is greater than the sink current, the VCC noise has higher amplitude than the VSS noise (static values at 20 ma IOH and 10 ma IOL per output).

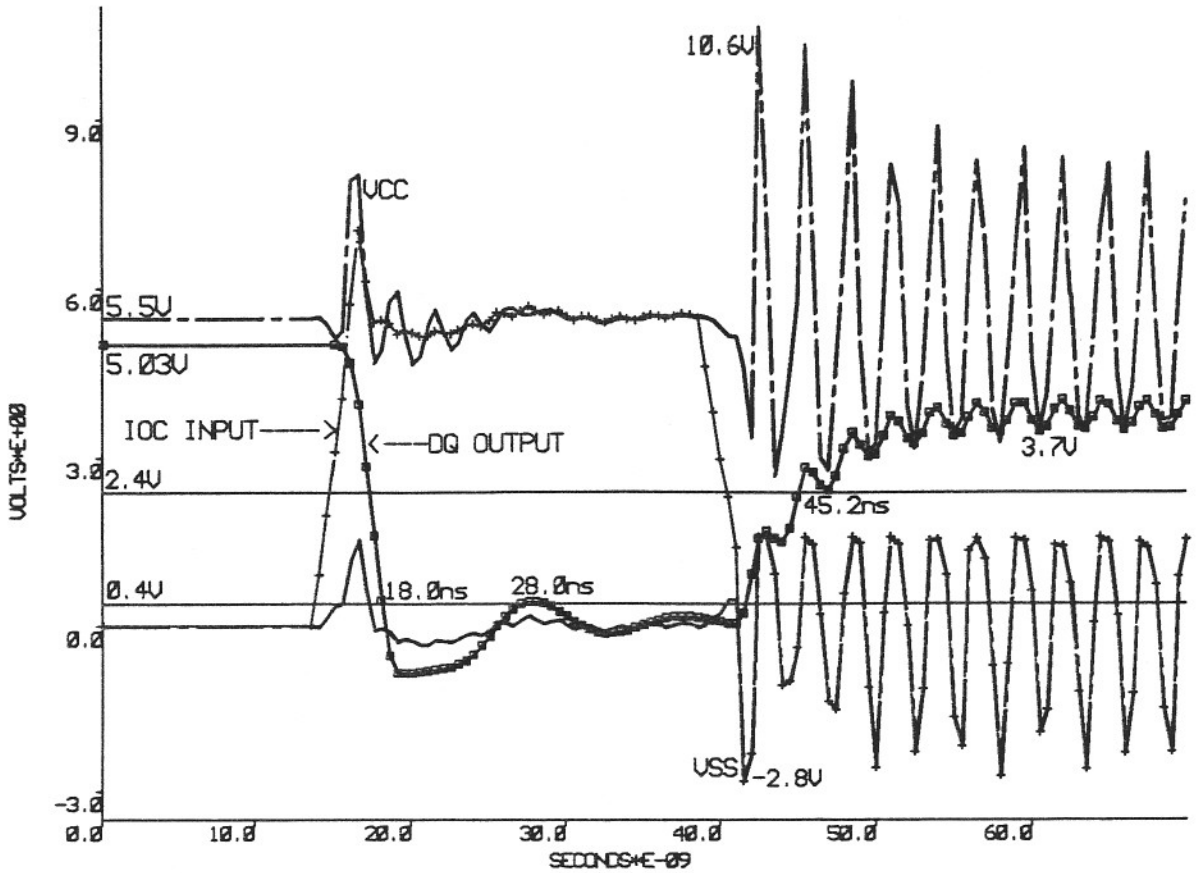

Figure 3.5: Output Buffer Fast Simulation with Package Inductance

Figure 3.6 shows an expanded view of the low-to-high output transition in Figure 3.5. The cause of the oscillation is seen by the addition of the OE1 signal. Noise from the VCC/VSS power lines has prevented OE1 from maintaining its enable logic 1 level of 5.5V. As a result, the DQ output gets repeatedly thrown into the high-impedance state, as it tries to read a logic 1 level. This causes oscillations and more noise.

3.2.4 Bypass Capacitor

This section looks at the placement of the 0.1 μf bypass capacitor (CTST) of Figure 3.2, and also the effect of the tester lead lengths.

The bypass capacitor is used to help filter out power line noise right at the package

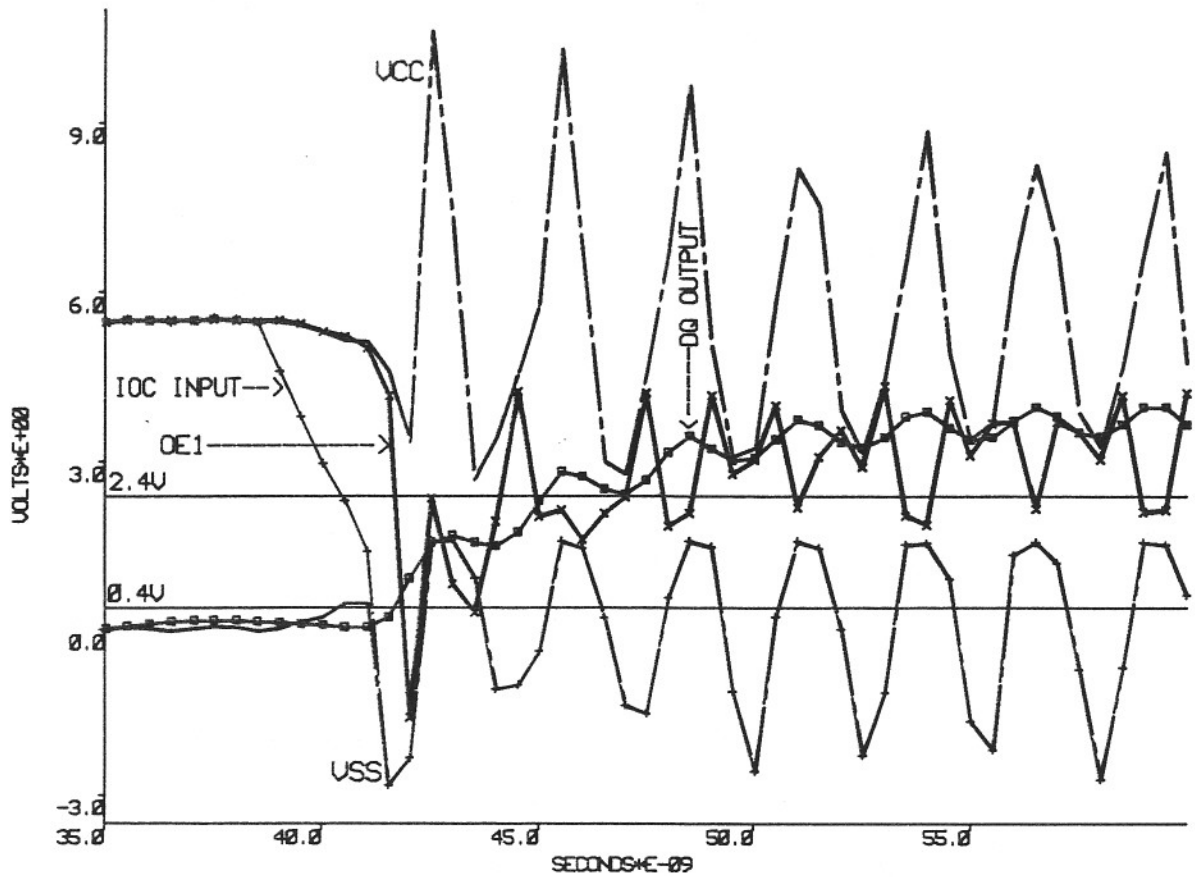

Figure 3.6: Output Buffer Simulation Showing Noise Oscillation

power supplies, VCP and VSP, during production testing. This capacitor is usually soldered to the performance board, which connects all tester supplies to the device package pins. Because of the power line inductance effect of tester ground lead LVSP (see Fig. 3.2), VSP is not a good place to bypass noise. Figure 3.7 shows another simulation, similar to Figure 3.5, but with capacitor CTST between nodes VCP and VST, instead of nodes VCP and VSP. In addition to this change, another 0.1 μf capacitor was placed between nodes VSP and VST.

The resulting simulation shows a significant improvement over the reference simulation of Figure 3.5. The output quickly reaches its steady state value of 5.0V, with no serious oscillations along the way. The improvement is due to the fact that the noise is now bypassed to tester node VST, which is a more stable ground than package

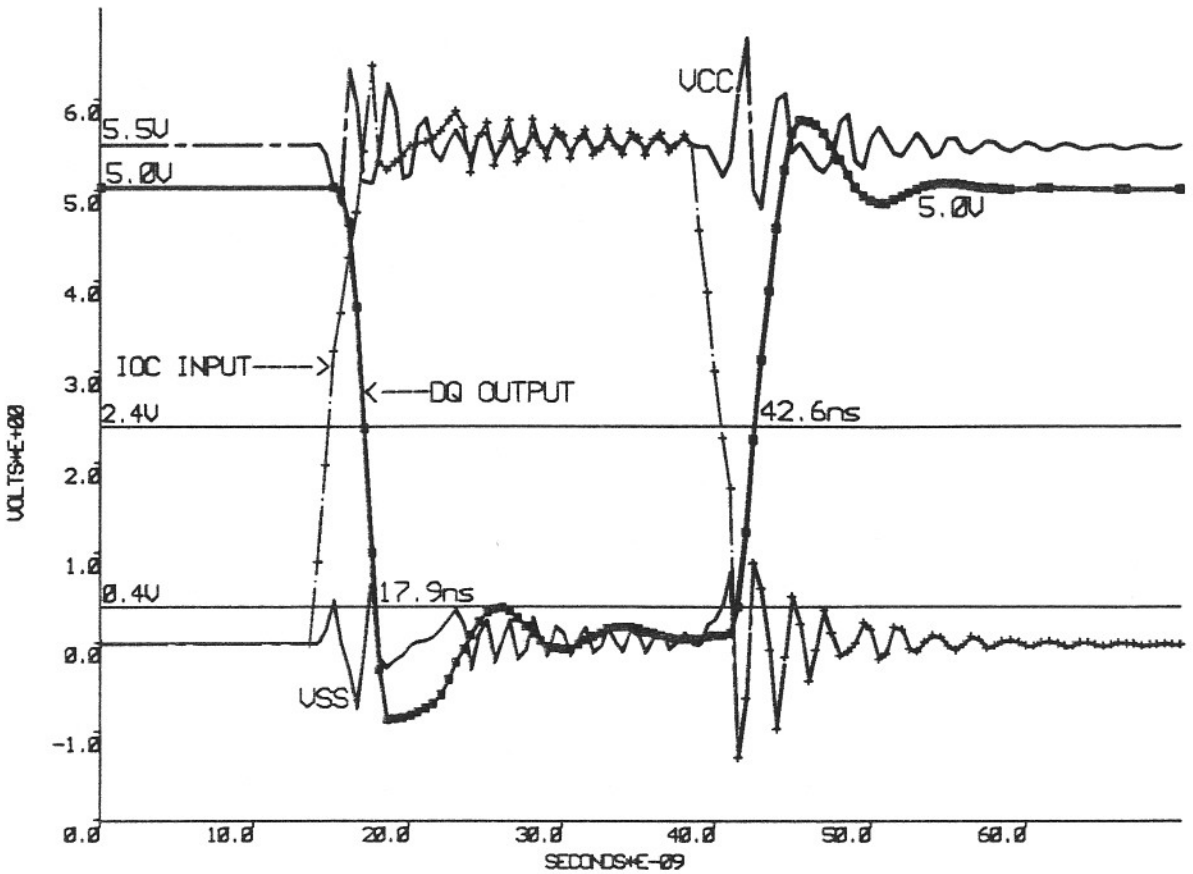

Figure 3.7: Noise Simulation with Added Bypass Capacitor

node VSP. Similar results can be obtained by removing the tester power line inductors, LVCP and LVSP. This assumes that the tester lead lengths can be neglected. In order to design around worst case conditions, these two inductances will be kept in the simulation, making Figure 3.5 the reference noise simulation to be worked on.

3.3 Power Line Resistors

One method currently used to reduce the noise of Figure 3.5 involves the insertion of small-valued polycide resistors in the power lines. These are shown as RVCO, RVCC, and RVSO in Figure 3.8, which is a modification of previous Figure 3.2. These resistors serve the purpose of dampening the noise oscillations on the power lines [1, p. 3]. By using small resistor values, there is minimal effect on output levels and access time.

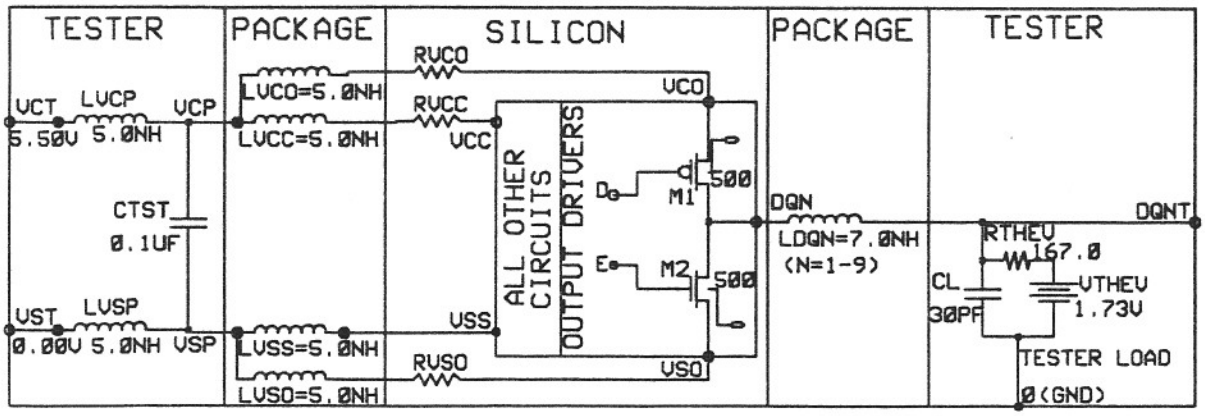

Figure 3.8: Package Inductance Circuit Showing Power Line Resistors

The upper limits for the resistor values come from the output voltage levels seen in the slow process corner simulation of Figure 3.9 (without resistors). This figure shows the same simulation of Figure 3.4, but uses slower conditions such as lower VCC supply, high temperature (125C military spec), and slow process models. In this process corner, the output voltage levels come closer to the VOH/VOL specs of 2.4V/0.4V. This is mainly due to the lower VCC supply of 4.5V.

For this process corner, the output characteristics are :

- VOH = 3.71V
- VOL = 180 mv
- IOH = 107 ma (for 9 outputs)
- IOL = 83 ma (for 9 outputs)

The output voltage levels determine the maximum resistor values which can be used on the power lines before the outputs fail the VOH/VOL specs.

The output high level (VOH) is determined by the VCO supply value and the p-channel driver channel resistance (see Fig. 3.8). Since VOH is 3.71V for this process corner, RVCO may have a voltage drop up to $3.71V - 2.4V = 1.31V$. Similarly, the

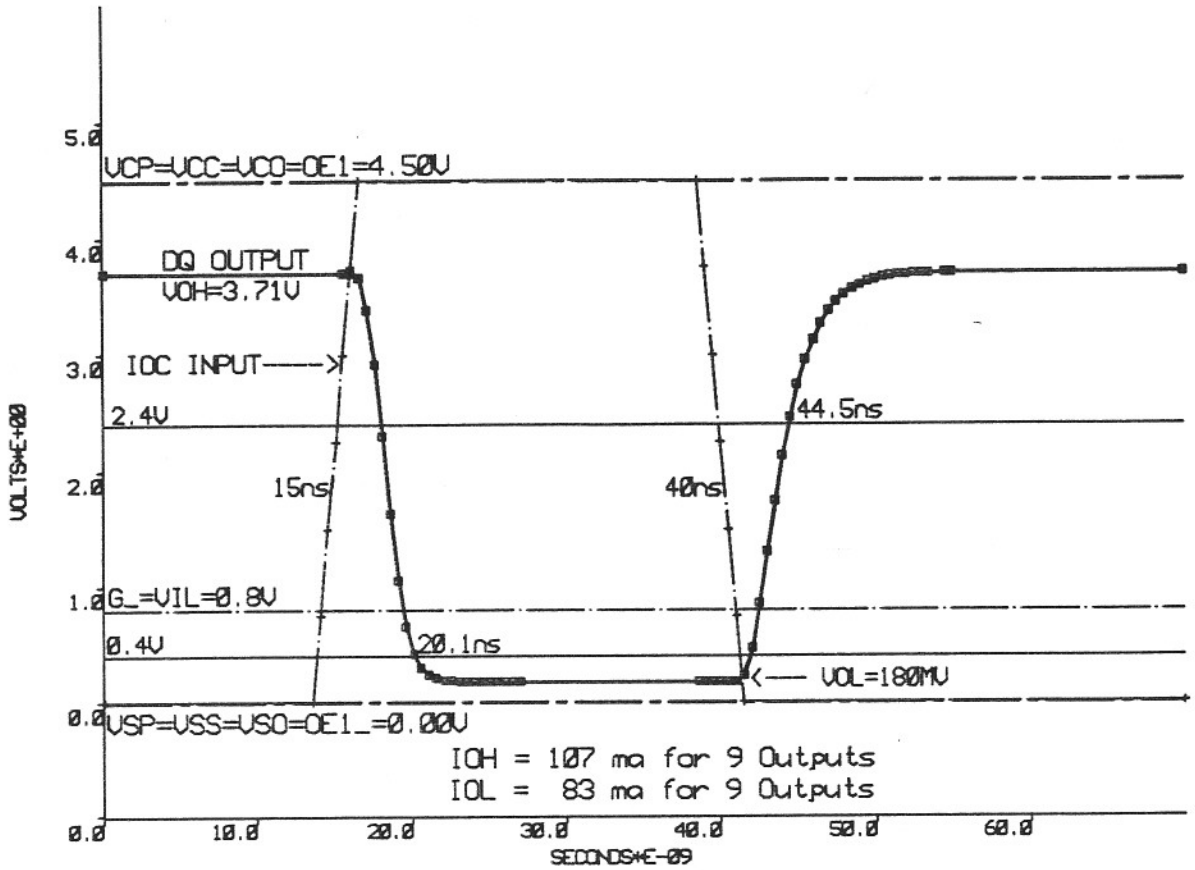

Figure 3.9: Output Buffer Slow Simulation, Ideal Case

output low level may be raised by an RVSO voltage drop of no more than 400 mv − 180 mv = 220 mv.

These maximum resistor voltage drops, taken with the VOH/VOL specs and the tester load Thevenin equivalence, determine the maximum resistor values as follows :

$$\frac{2.4V - 1.73V}{167/9\Omega} = \frac{1.31V}{RVCO} \tag{3.8}$$

$$\frac{1.73V - 0.4V}{167/9\Omega} = \frac{220\ mv}{RVSO} \tag{3.9}$$

The Thevenin resistance is divided by 9 since RVCO and RVSO must source/sink current through nine identical output driver stages. Solving these equations gives the maximum values for the polycide resistors :

- RVCO = 36.3Ω = 18.0 squares max

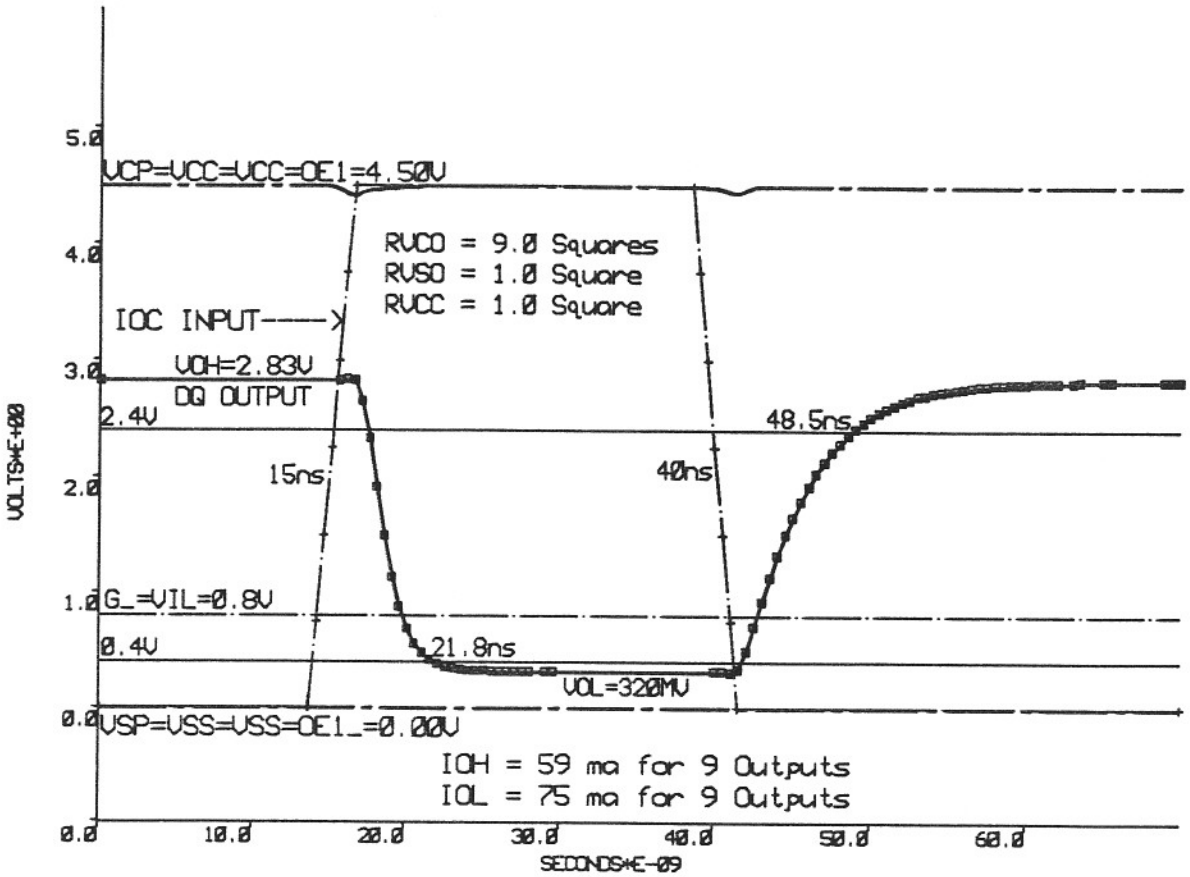

Figure 3.10: Output Buffer Ideal Slow Simulation with Power Line Resistors

- RVSO = 3.1Ω = 1.5 squares max

A polycide sheet resistance of 2.0Ω/square for the slow process models gives the resistance in terms of squares. These are the upper limits for RVCO and RVSO.

Figure 3.10 is a rerun of the slow corner simulation of Figure 3.9 with added power line resistances of :

- RVCO = 9 squares
- RVSO = 1 square
- RVCC = 1 square

Since the maximum resistor values would give output levels equivalent to the spec

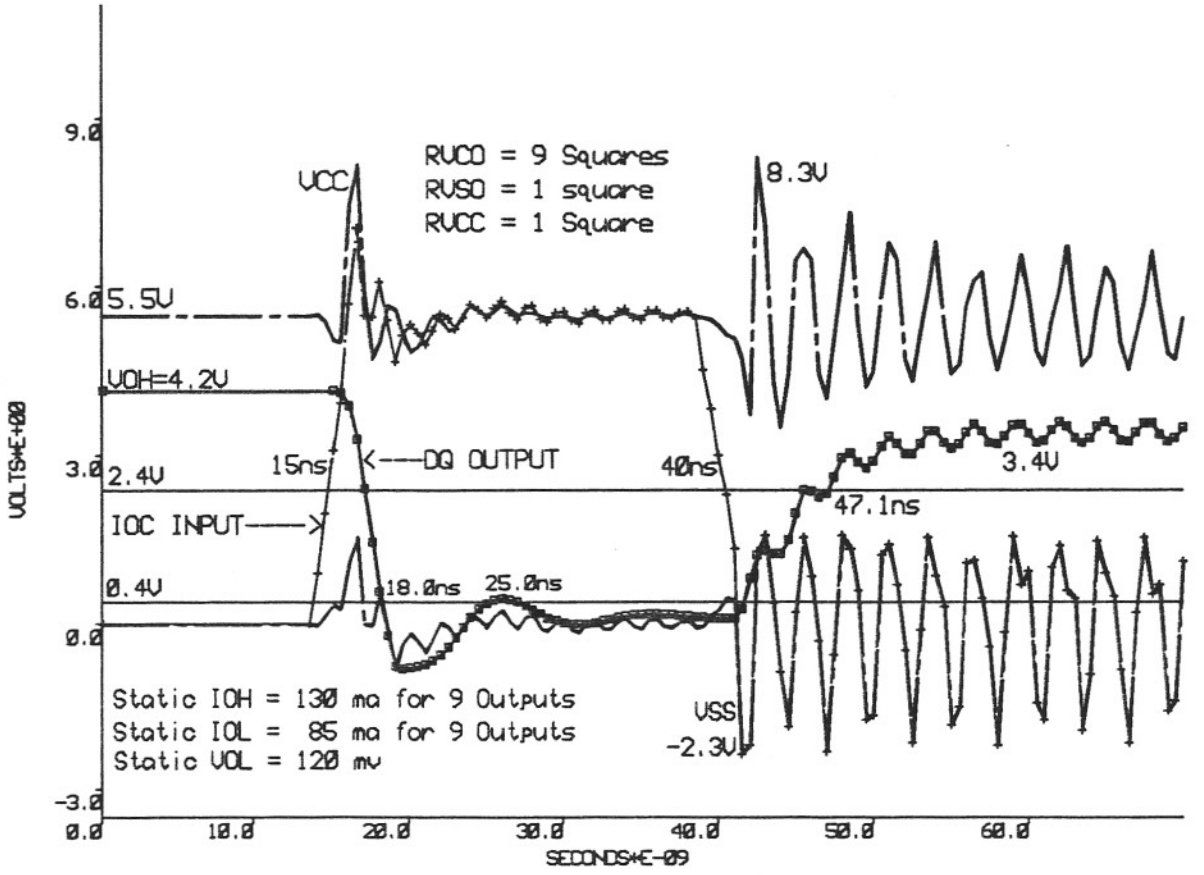

Figure 3.11: Output Buffer Fast Simulation with Power Line Resistors

VOH/VOL levels (2.4/0.4V), the simulation uses one half of these values for RVCO and RVSO.

RVCC is taken as 1 square. A resistor is not used for the VSS power line, since this line supplies the ground level to all TTL input buffers. Raising the ground level with a resistor would worsen the noise margin for active high TTL inputs.

Comparing Figure 3.10 to Figure 3.9, the AC/DC characteristics are :

- VOH dropped from 3.71V to 2.83V
- VOL increased from 180 mv to 320 mv
- IOH decreased from 107 ma to 59 ma for 9 outputs
- IOL decreased from 83 ma to 75 ma for 9 outputs

- Output high-to-low delay increased from 5.1 ns to 6.8 ns
- Output low-to-high delay increased from 4.5 ns to 8.5 ns

Figure 3.10 shows worst case speed and output voltage level performance due to the addition of the three resistors. Package inductance is not included in order to focus on the DC characteristics. The inductance noise is usually not a problem in this corner due to the decrease in speed and switching current.

All four DC characteristics, VOH, VOL, IOH, and IOL are still within spec. However, the high-to-low delay increased by 1.7 ns, and the low-to-high delay increased by 4.0 ns due to the 9.0 square RVCO resistor. This speed trade-off is acceptable for the slow process corner. When the delay time is compared for nominal process conditions, the 4.0 ns increase in delay due to RVCO drops to a 2.2 ns increase. This delay increase affects only the output buffer circuit, which is the only one connected to the VCO/VSO supplies that use the 9.0 square resistor.

Now that the slow corner shows satisfactory performance, the power line resistors are used in the fast corner, where most of the noise occurs. Figure 3.11 is a rerun of the fast corner package inductance simulation of Figure 3.5 with the three added power line resistors. Comparing these two figures, along with Figure 3.4 for DC levels, the changes are :

- VOH dropped from 5.0V to 4.2V
- VOL increased from 59 mv to 120 mv
- IOH decreased from 180 ma to 130 ma for 9 outputs
- IOL decreased from 90 ma to 85 ma for 9 outputs
- Output high-to-low delay did not change
- Output low-to-high delay increased from 5.2 ns to 7.1 ns

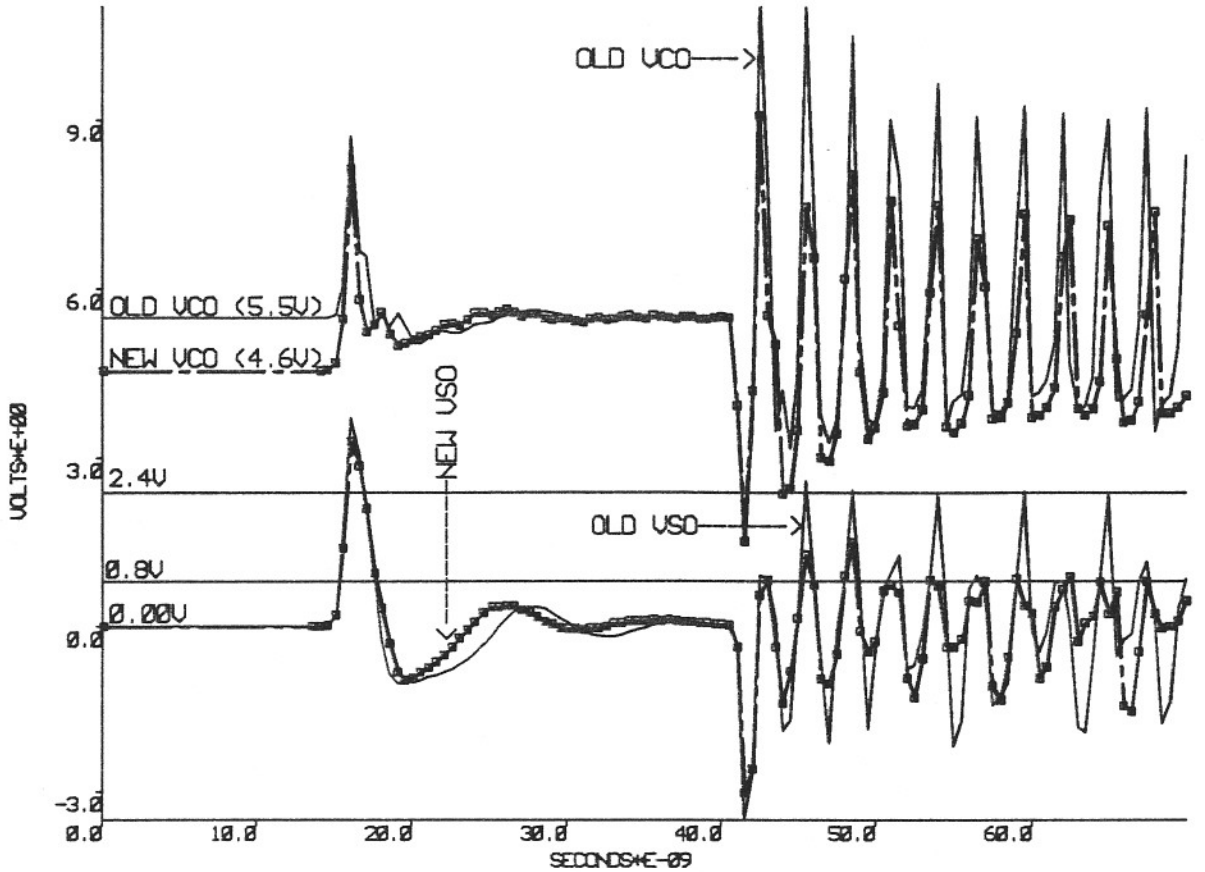

Figure 3.12: Output Driver Supply Noise Comparison After Resistors Were Added

The major improvement here is the 50% decrease in magnitude of the VCC noise, showing the effect of the 9 square VCO resistor. By reducing the VCO output supply noise, the VCC noise is also reduced with minimal performance degradation. There is essentially no improvement on the ground noise due to the smaller size of RVSO.

Figure 3.12 compares the output driver supplies VCO and VSO (not shown in previous figures) before and after the resistors are added. The new VCO/VSO signals, from the simulation with resistors, are the bolder lines (with the boxes). The lighter lines are the old VCO/VSO signals, corresponding to the simulation of Figure 3.5 without resistors. For these supply lines, there is approximately 2 volts noise reduction on the VCO supply, and a 1.5-volt noise reduction on the VSO supply. Note that the old VSO supply jumps all the way up to the 2.4V VOH level during the low-to-high

output transition region. This is where the larger output source current starts driving the tester load of Figure 3.8.

3.4 Output Enable Latch Function

Due to the restriction on the size of resistor RVSO discussed in the last section, there is no significant improvement in the ground noise, as shown in Figure 3.11. This section looks at the cause of the oscillation and shows a circuit modification which eliminates the oscillation.

3.4.1 Cause of Oscillation

Referring to Figure 3.3, the G_ input must be kept at the VIL level of 0.8V in order to enable the outputs. CS1_, the global chip select signal, is the second input to the TTL NOR gate. This signal has a level equal to VSS, including the VSS noise, since it is a driver output from the chip select buffer. However, the G_ input does not include the ground noise, since it is a package pin input, hard-wired to a separate tester (or system) power supply. This input is kept at a VIL level of 0.8V to enable the output.

The problem occurs when VSS carries up to −2V noise to the source connection of the lower n-channel transistor within the TTL NOR gate. The TTL gate recognizes 0.8V as VIL with respect to a VSS level of 0.0V, and not -2.0V. In the latter case the G_ input is seen as 2.8V rather than 0.8V. This causes the entire path generating the OE1 and OE1_ signals to trip to the opposite state. This creates the oscillation problem seen in Figure 3.11 for the low-to-high output transition.

Figure 3.13 shows the same simulation, but includes an internal signal from the output enable buffer of Figure 3.3. The power lines and the input have been removed to focus on the new signal. The signal is node B from the schematic of Figure 3.3. Since both inputs to the TTL NOR gate are low, node B should be at logic 0, since it appears two stages later. However, this signal oscillates between VCC and VSS, due

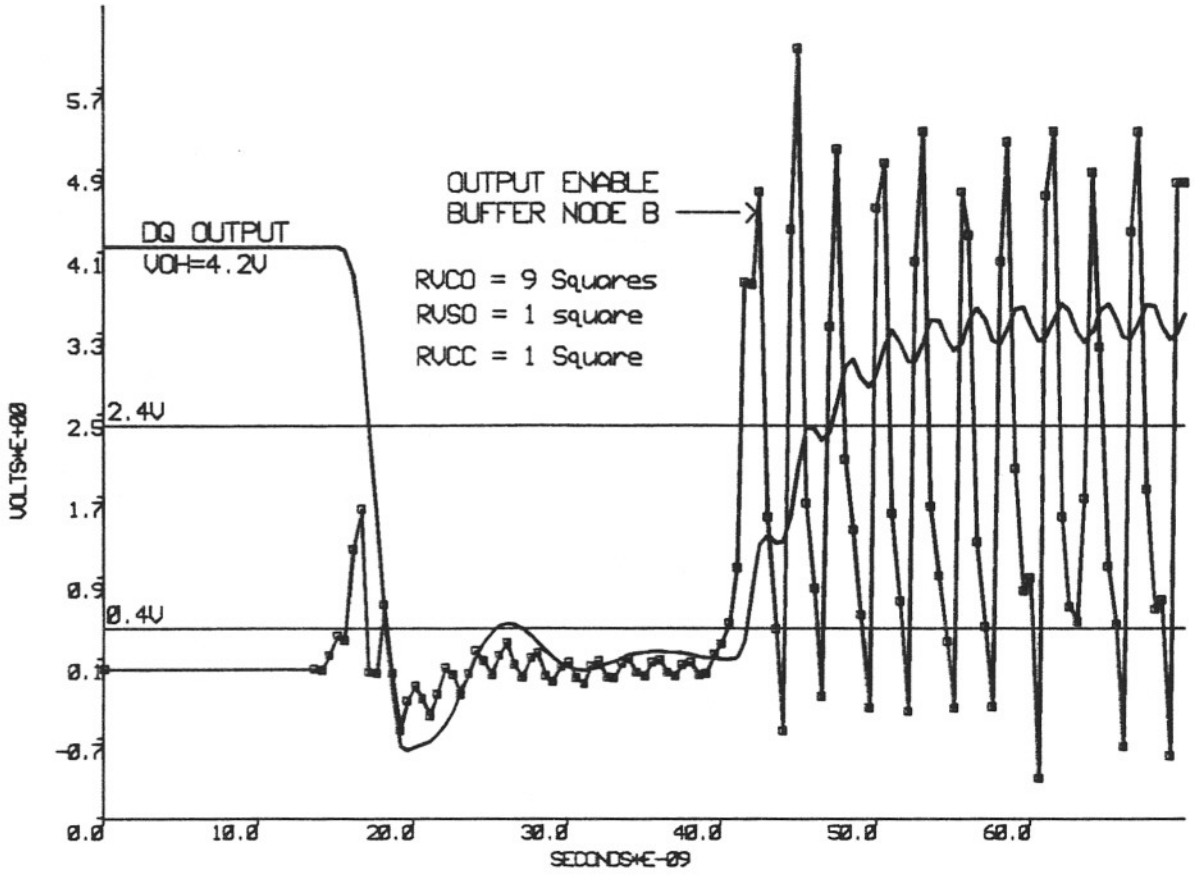

Figure 3.13: Noise Simulation Showing Output Enable Internal Node B

to the ground noise tripping the G_ input.

3.4.2 Latch Function

One method of preventing this oscillation from reaching the output buffer stage involves the use of a latch signal, which makes use of the tri-state inverter shown in Figure 1.2. A modification to the output enable buffer is shown in Figure 3.14. This circuit is a modified version of the one shown in Figure 3.3, and includes the latch circuitry. This consists of two tri-state inverters at node E, along with a latch generator which generates latch signals LAT and LAT_. These are wired to the enable inputs of the tri-state inverters.

During normal operation the first tri-state inverter, with input node D, is enabled

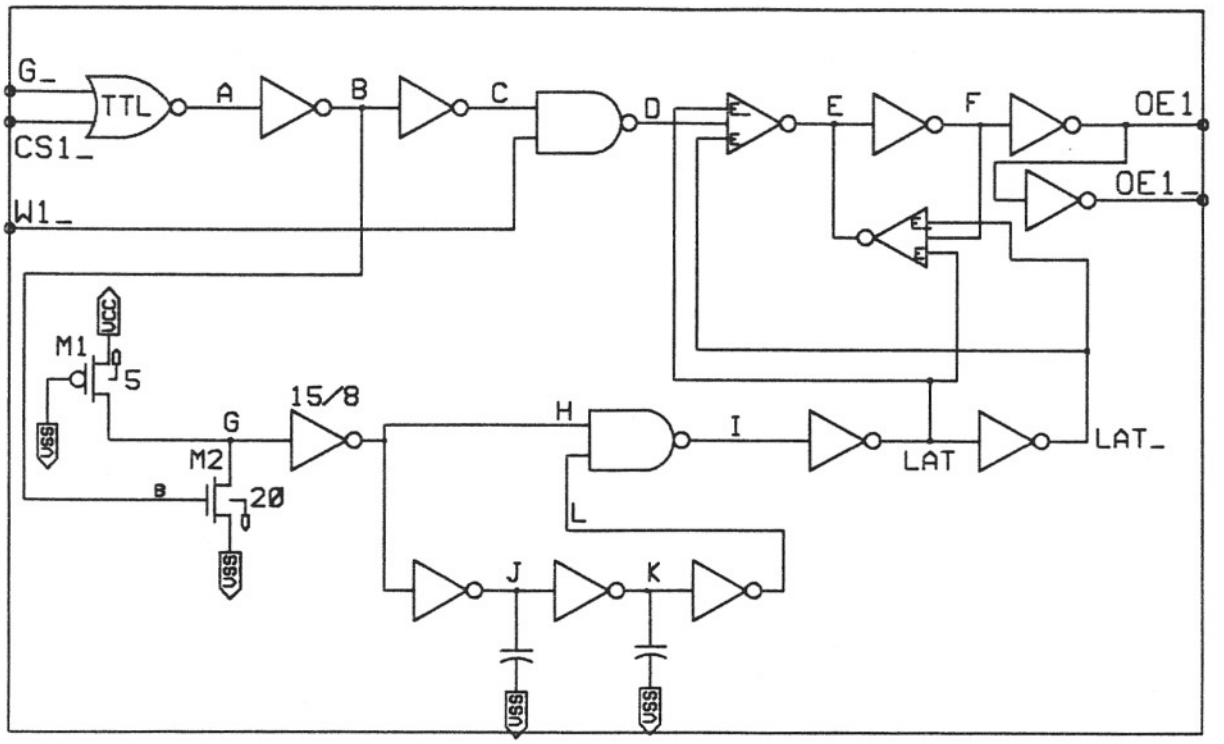

Figure 3.14: Output Enable Buffer with Added Latch Function

by logic 0 from the LAT signal, which connects to the the "E_" input, and logic 1 from the LAT_ signal, which connects to the "E" input. This inverts the signal at node D, as does a standard inverter. The lower tri-state inverter, with input at node F, is disabled since the LAT and LAT_ signals are swapped at the enable inputs. After node E, the signal propagates through the last three stages, giving the correct output enable logic at OE1 and OE1_.

Since node B is stable at either logic 1 or logic 0 in the absence of noise, the logic level at node H becomes the same as at node B. If node B is low, transistor M2 is cut off, and pull-up transistor M1 sets node G to logic 1. Node G is then inverted, making node H low. If node B is high, the 20 μm pull-down transistor (M2) overpowers the 5 μm pull-up transistor (M1). This sets node G low and node H high.

Since Node L is taken three inversions after node H, the NAND gate output at node I is always high, regardless of the states of the G_, CS1_, and W1_ inputs. This

sets the LAT and LAT_ signals to logic 0 and logic 1, respectively, indicating that no latching is to be performed by the tri-state inverters in the absence of noise.

When the outputs are enabled, the correct state for node B is logic 0 (G_ at VIL). In the absence of ground noise, this node retains its low state, setting node H low and node L high. When ground noise is present, node B oscillates, causing the M2 pull-down transistor to momentarily reverse the logic levels at nodes G and H. Thus, node H makes a low-to-high transition during the oscillation. Due to the inverter/capacitor delay string at nodes J and K, node L does not quickly respond with a high-to-low transition. This gives a short time period where both NAND gate inputs, H and L, are at logic 1. This creates a one-shot pulse at node I, having a pulse duration determined by the inverter/capacitor delay string.

The pulse is triggered each time VSS develops negative-going noise spikes, which trip the G_ input. This pulse is then inverted to the proper logic for the LAT/LAT_ signals, setting LAT to logic 1 and LAT_ to logic 0 during the latch pulse duration. As a result, the LAT/LAT_ signals latch the current state at nodes E and F, before the noise propagates to these nodes.

When this occurs, the first tri-state inverter cuts off the noise up to node D, while the lower one is enabled. This forms a latch (or feedback loop) between nodes E and F. The result is that the OE1 and OE1_ outputs will not oscillate when nodes A through D oscillate from ground noise. Simulation results are shown in Figure 3.15, which still includes the power line resistors of Figure 3.11.

The new simulation shows that only two latch pulses are required to stop the oscillation. The first VSS noise spike, after the 40 ns input transition, is the one responsible for starting the output oscillation. This triggers the first latch pulse, which latches the OE1 & OE1_ signals in the active state. This prevents the output from oscillating between its logic 1 destination and the high-impedance state. Since the feedback oscillation path has been eliminated, the second VSS noise spike is now only due to the

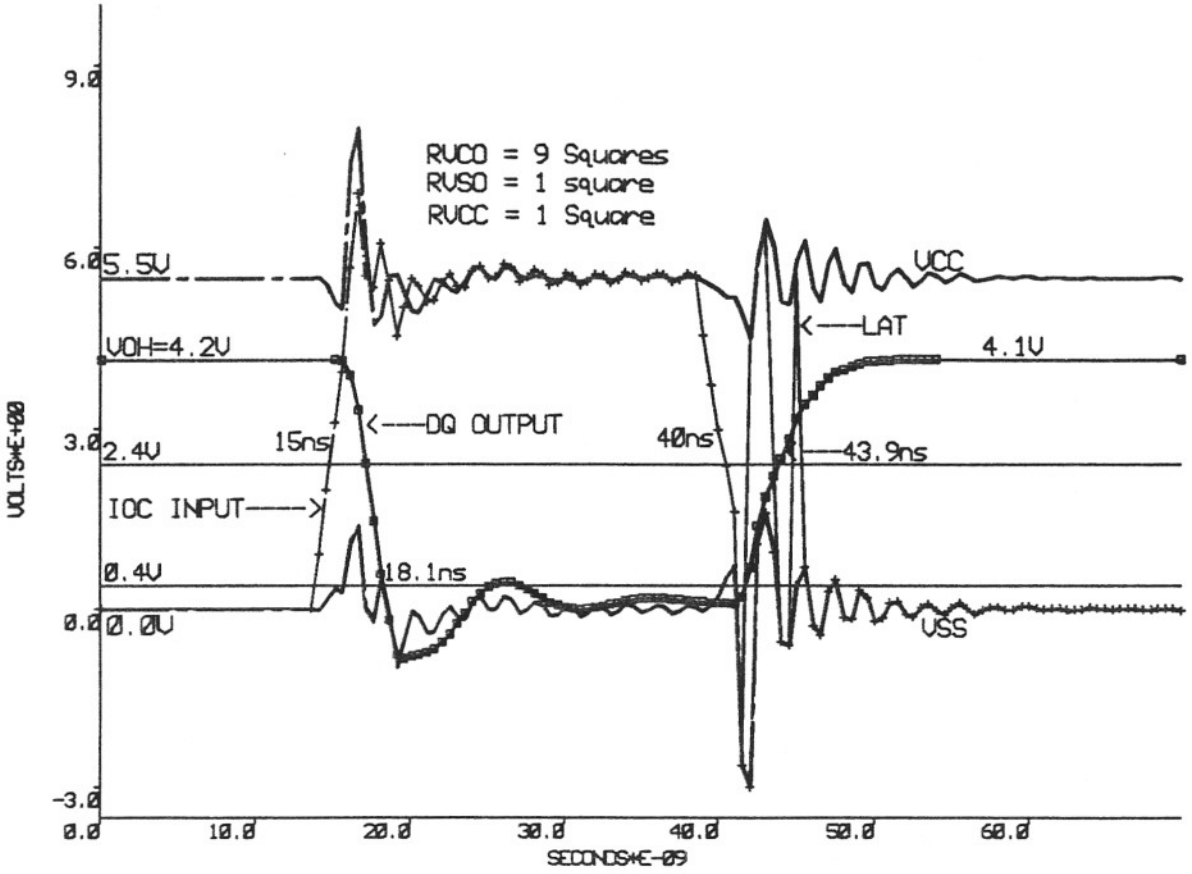

Figure 3.15: Noise Simulation with Output Enable Latch Function

low-to-high output transition, no longer influenced by the output oscillation. Thus, the VSS noise rapidly decays, allowing the output to make a smoother transition to the logic 1 level. Figure 3.15 also shows that the new low-to-high delay is 3.9 ns. Comparing this to Figure 3.11 shows that 3.2 ns delay time, which was lost to inductance noise and the RVCO resistor, has been recovered.

3.5 Noise Sensor Devices

There are times when cutting off the output actually helps reduce noise. In the previous section, the output enable signals were latched to their current state as soon as ground noise was detected in the G_ input path. By quickly latching these signals, they were prevented from disabling the outputs and starting a noise feedback loop to the G_

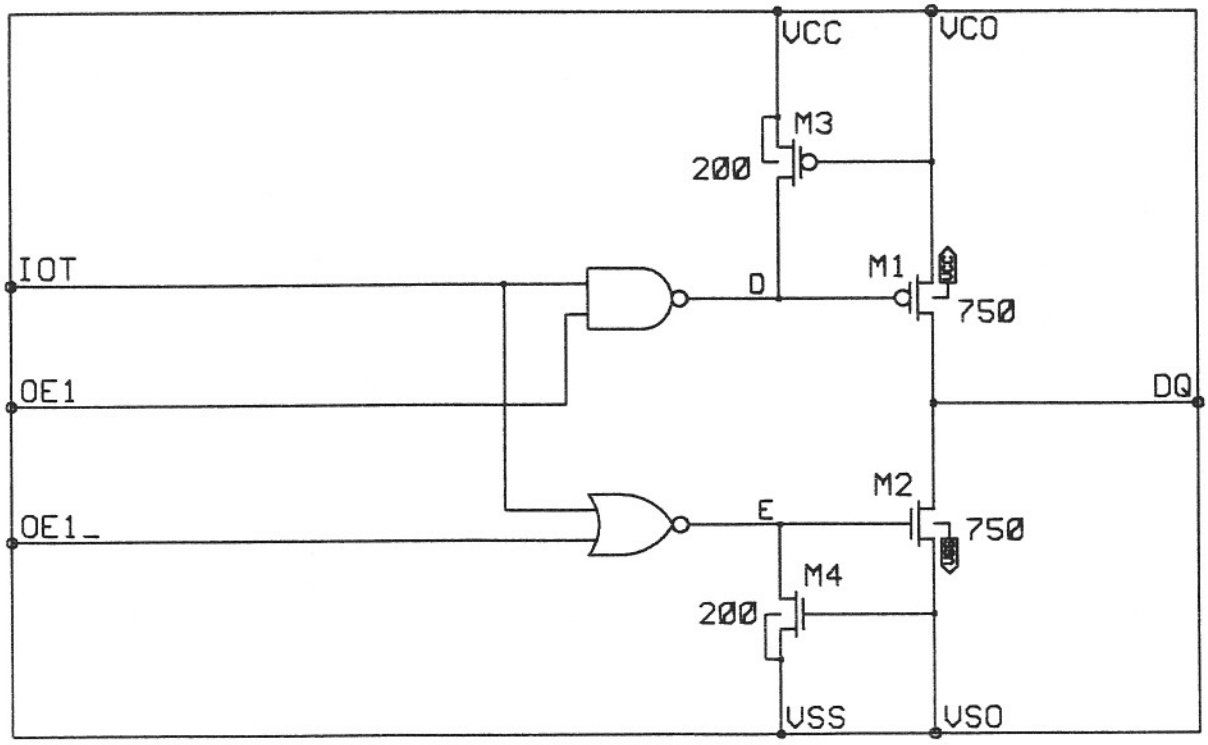

Figure 3.16: Output Buffer with Noise Sensor Devices

input. This section shows a method to reduce the output driver switching current by increasing the output impedance when noise appears on the VCO/VSO output driver supply lines. This method reduces the power line noise magnitude and is illustrated using a full-circuit simulation of a 16Kx4 SRAM, rather than the previous 2-circuit simulation of the 8Kx9 SRAM read path.

3.5.1 Circuit Changes and Sensor Device Operation

Figure 3.16 shows a modification of the output buffer circuit of Figure 3.1. The new circuit shows a transistor added to each gate (nodes D and E) of the output drivers. These 200 μm devices (M3 and M4) serve as noise sensors, which increase output impedance when the VCO and VSO supply lines begin to oscillate [1, p. 3].

Transistor M3 is used to detect noise when the VCO supply drops below VCC. Ac-

cording to Equation 3.10, this occurs when the output makes a low-to-high transition :

$$VCO = VCP - LVCO\frac{d(ILVCO)}{dt} \tag{3.10}$$

ILVCO is the magnitude of the current through inductor LVCO (see Fig. 3.8). During the low-to-high output transition, ILVCO increases, making $\frac{d(ILVCO)}{dt}$ positive. This results in a decrease in the output driver supply, VCO, during the low-to-high output transition. As VCO drops, the gate-to-source voltage drop on M3 increases in magnitude, in accordance with the noise level on VCO. When VCO drops beyond the p-channel threshold, transistor M3 acts as a pull-up on node D and decreases the current drive of transistor M1.

Similarly, transistor M4 is used to detect noise when the VSO supply rises above VSS. Equation 3.11 indicates that this occurs during the high-to-low output transition, in which the current magnitude through inductor LVSO (ILVSO) is increasing :

$$VSO = VSP + LVSO\frac{d(ILVSO)}{dt} \tag{3.11}$$

As the VSO supply rises beyond an n-channel threshold above VSS, transistor M4 reduces the current drive of transistor M2.

The effect of increasing output impedance during noisy transitions reduces the switching current of the output drivers. This reduces the noise magnitude on the VCO/VSO driver supply lines, as well as the magnitude of the noise fed back to the main device power lines, VCC and VSS.

3.5.2 Simulation of Sensor Devices

The cost of implementation of the sensor devices is the effect of the added capacitance to the gates of transistors M1 and M2. For the circuit of Figure 3.16, this results in the following reduction in speed and switching current after transistors M3 and M4 are added :

- Low-to-high delay increases from 8.7 to 9.4 ns (8%) in slow process corner

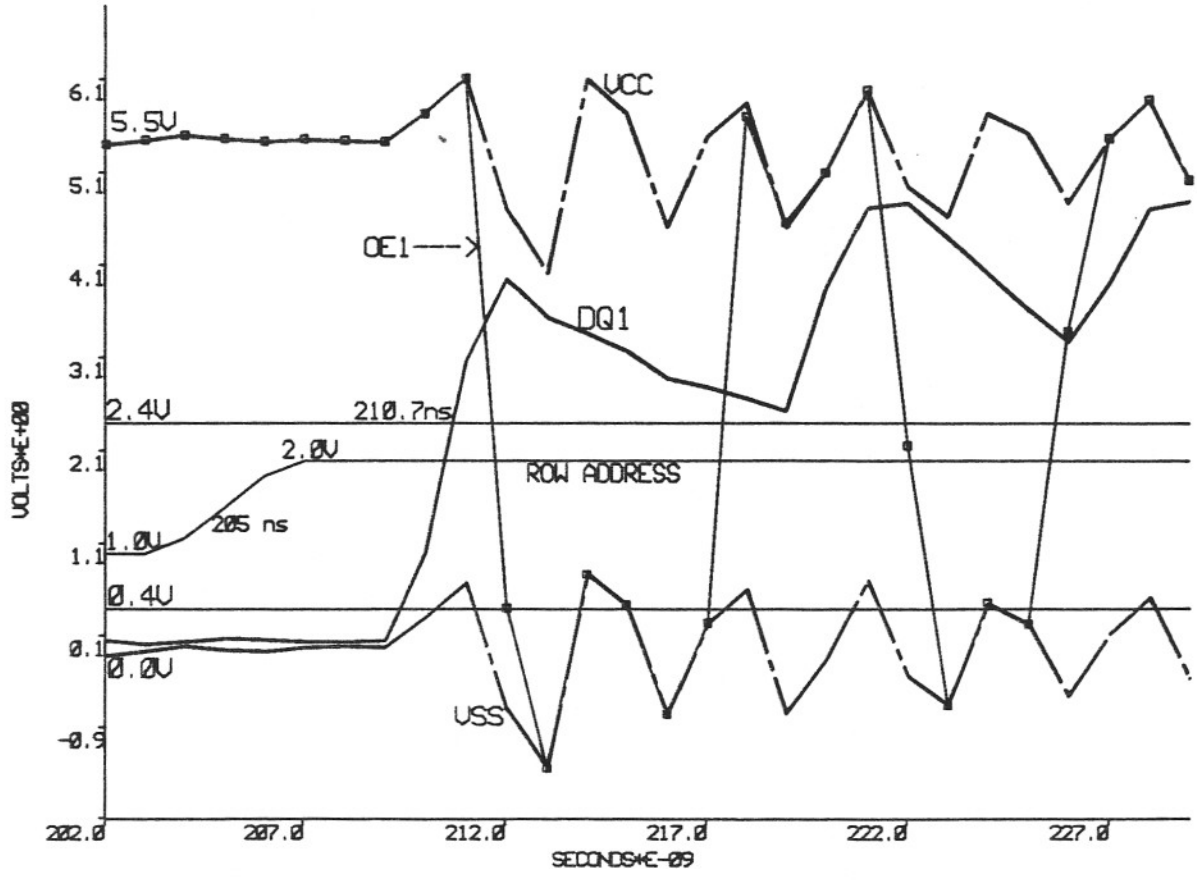

Figure 3.17: 16Kx4 Simulation without Noise Sensor Devices

- VCO source current decreases from 189 to 144 ma (24%) when 9 outputs switch low to high in fast process corner
- VSO sink current decreases from 384 to 317 ma (17%) when 9 outputs switch high to low in fast process corner

For the above results, the slow process corner was used to determine the worst case speed loss. The fast process corner was used to determine the decrease in switching currents. Transistor gate widths used were 750 μm for the output drivers, and 200 μm for the noise sensors, as shown in Figure 3.16.

Figure 3.17 shows a full-circuit inductance simulation of a 16Kx4 SRAM without the noise sensors. The expanded region shows a row access, which starts at 205 ns. TTL input levels are tightened to VIL=1.0V and VIH=2.0V. The circuit includes

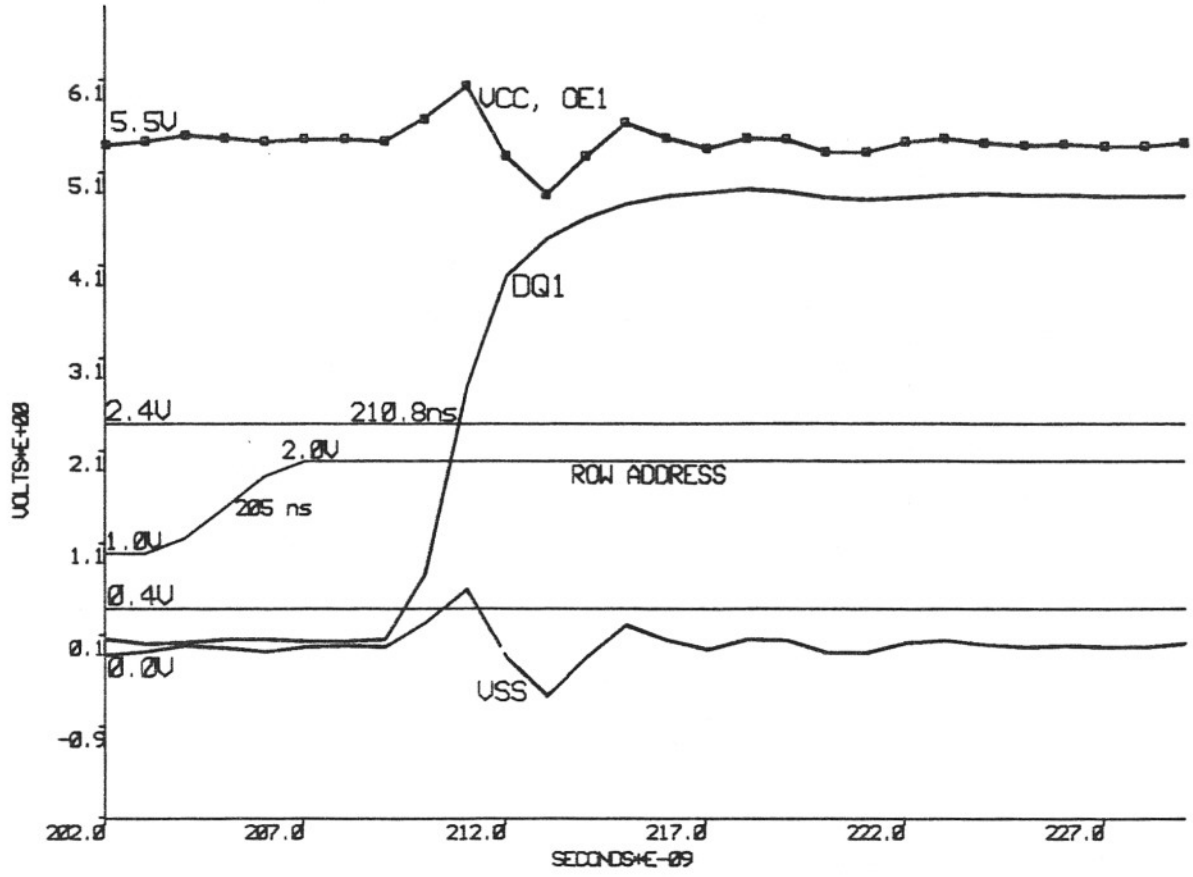

Figure 3.18: 16Kx4 Simulation with Noise Sensor Devices

the power line resistors of Section 3.3, but does not include the output enable latch discussed in the last section.

This device does not have the G_ output enable option, but still uses an output enable circuit similar to that of Figure 3.3, with the G_ input tied to VSS. The figure shows the same problem seen in Section 3.4, where power line noise fed back to control buffer TTL inputs trips the output enable signals, OE1 and OE1_. This causes the output to oscillate during the low-to-high transition, even without the G_ device input. In this case, the noise feeds back through the chip select (CS_) and write enable (W_) TTL inputs.

Figure 3.18 shows the same simulation after the noise sensors are added to the output buffer. During the low-to-high output transition, the noise sensor devices limit

the switching current, and reduce the power line noise levels that feed back to the control inputs. The results are :

- The output makes a clean low-to-high transition, without indication of being cut off by the noise sensors
- Output enable OE1 is identical to VCC, free of feedback noise
- Power line noise is reduced for the VCC and VSS supplies

The noise sensors have successfully reduced the switching current, giving the desired output behavior.

3.6 Chapter Summary

The circuit modifications presented in this chapter apply mainly to noise in the output path. Power line noise generated by the output drivers and package inductance has been shown to feed back to the output enable circuit through the power lines. Due to the ground noise margin of $\pm$ 0.7V for TTL inputs, VSS noise tripped the G_ input and created output oscillations.

Power line resistors were used to reduce the VCC noise by 50%. Due to the size restriction of the output driver ground supply resistor, a latch function was used to stabilize the output enable signals in the presence of ground noise. The latch circuit prevented the output noise oscillation during the low-to-high output transition and reduced the amount of ground noise feeding back to the G_ input pin.

Noise sensor devices were also shown to be helpful in reducing the noise magnitude. These devices were placed between the VCC/VSS power lines and the gates of the output drivers. By gating these devices with the VCO/VSO output driver supplies, output impedance was increased in accordance with the noise voltage levels on the VCO/VSO supply lines. This reduced the switching current, which, in turn, reduced the noise magnitude.

As circuit complexity increases, it may not be desirable to rely on one-shots (i.e., latch pulses) to fix noise problems. Metal runs and the placement of control circuits add uncertainty to the timing of the latch circuits, which may be required for the CS_ and W_ control inputs as well as for the G_ input. The next chapter illustrates additional methods which can be used to add more noise immunity to TTL inputs.

Chapter 4

Noise Reduction for Control Circuits

4.1 Control Circuit Operation

The main functions of SRAM control circuits are to set the operational state of the device, and to set all read and write cycle timing parameters in accordance with the device specification. Figure 4.1 illustrates the major connections between the three control circuits – chip select, write enable and output enable. These circuits generate critical timing paths required to pass the timing specs of Section 2.6. Changes in these timing delays can be used to reduce some of the output noise.

The chip select buffer powers up all circuits in the device when the CS_ input makes a high-to-low transition. When CS_ makes a low-to-high transition, all circuits are disabled and the outputs are set to the high-impedance state. The timing specs associated with the CS_ input are the chip select access time, $t_a(CS)$, and the chip select disable time, $t_{dis}(CS)$. The time period during chip select access usually contains noise, since all circuits power up simultaneously.

The write enable buffer controls the start and end of the write cycle. The W1 and W1_ signals in Figure 4.1 are used to enable the data input drivers to write the data at the DQ terminals. This data is driven to the matrix I/O lines during the write cycle. At the end of the write cycle, the data input drivers are disabled. Then, the

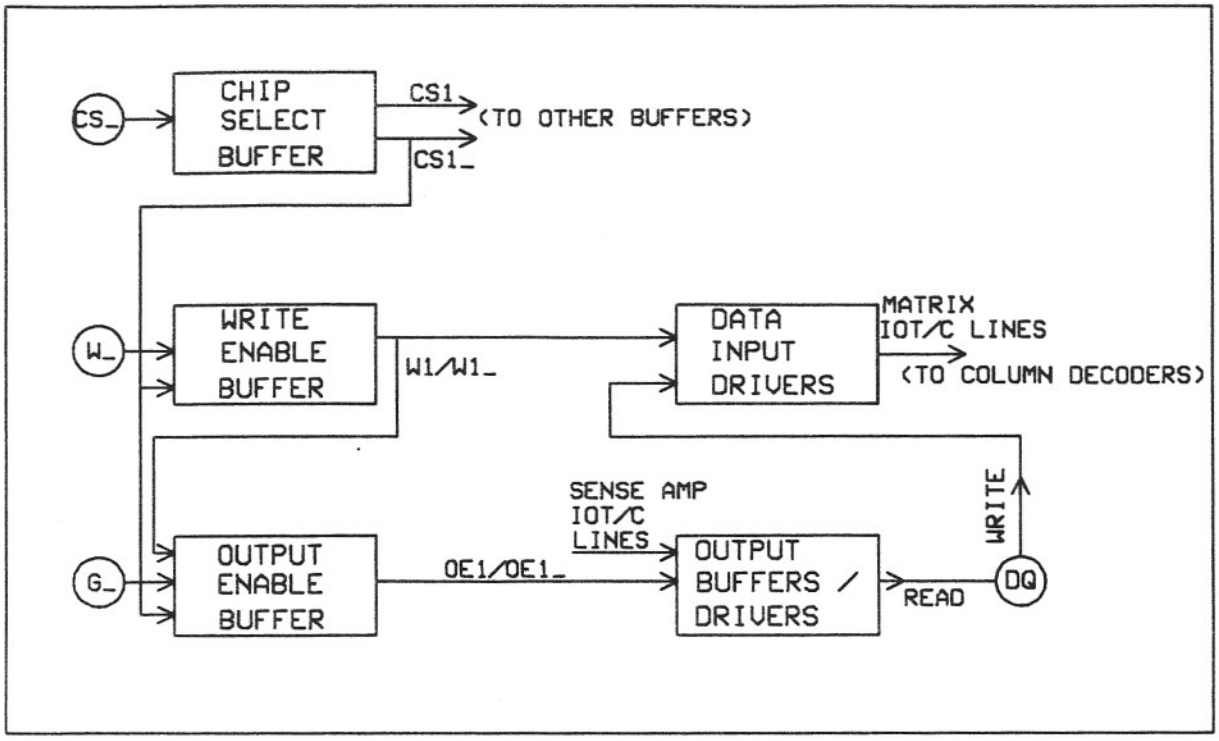

Figure 4.1: SRAM Control Circuit

output enable signals, OE1 and OE1_, turn on the output buffers, which read the sense amp I/O line data. This read cycle after a write cycle is also a critical noise region, especially when the outputs are quickly enabled to read data that is opposite to the data which was written.

The output enable buffer combines the chip select function, the write enable function, and the G_ input to control when the outputs are to be turned on, and when they are to be shut off. Any of these three functions can disable the outputs when they are in the deselect or write state (ref. Sec. 2.3). To enable the outputs, all three functions must be in the active read state, i.e., CS_ at VIL, W_ at VIH (read mode), and output enable G_ at VIL. The outputs begin to oscillate during read mode when VSS noise trips any one of the three control inputs.

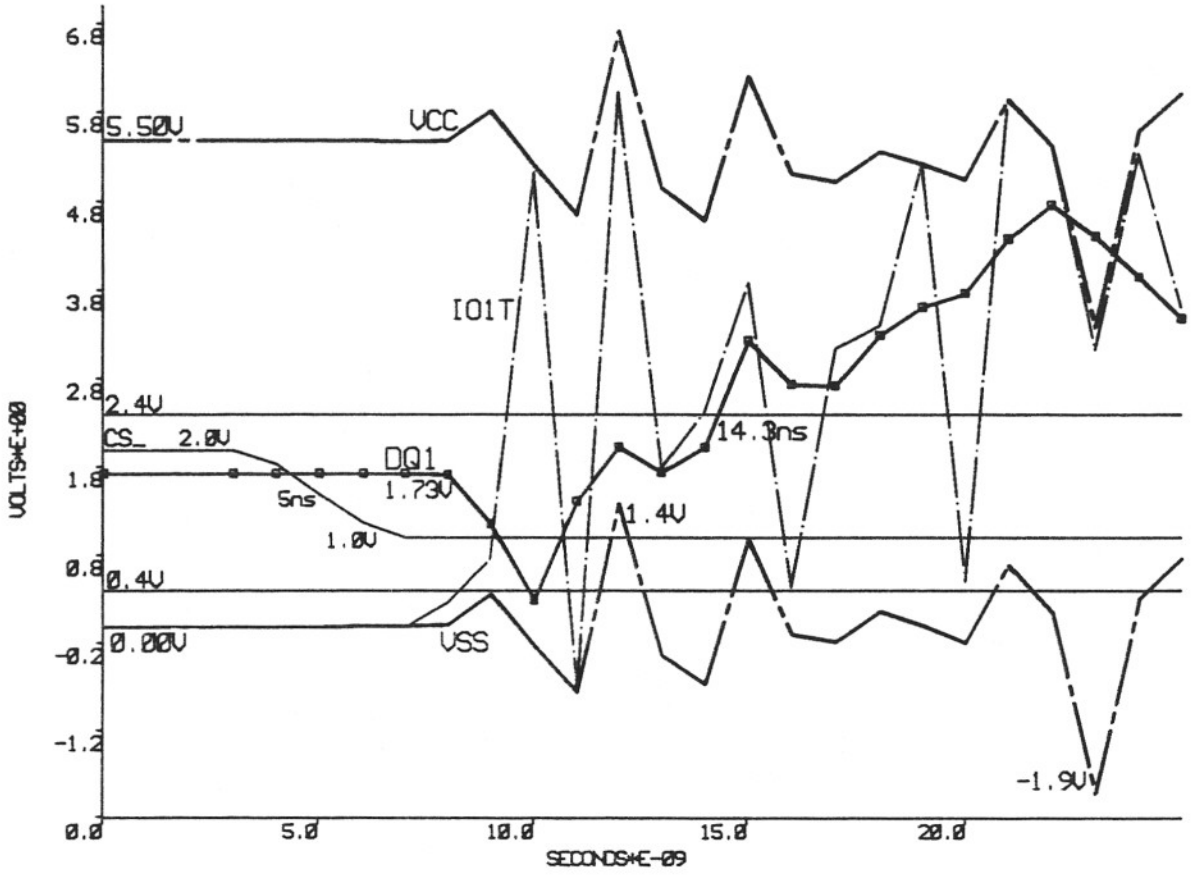

Figure 4.2: 16Kx4 Chip Select Access with Fast Output Enable

4.2 Effect of Output Enable Delay on Chip Select Access

It is desirable to have a long turn-on time for the outputs during chip select access. During this power-up period, there is a slight delay before the I/O line data from the sense amps is available to the output buffers. The I/O line data may also oscillate before reaching the correct read state. If the outputs are enabled during this time, the I/O line oscillations will be amplified by the output drivers.

4.2.1 Fast Output Enable

Figure 4.2 shows the chip select access from a full-circuit simulation of TI's new 20-ns 16Kx4 SRAM design at the start of noise reduction. The simulation is run in the

fast process corner to show worst case noise. It includes the power line resistors of Section 3.3, as well as the VCO/VSO power line noise sensors of Section 3.5. This particular device contains four common I/O terminals, similar to the 8Kx9, but does not have the output enable (G_) input.

The output oscillation of Section 3.4, where the ground noise was tripping the G_ input, can still cause the outputs to oscillate even though this part does not have the G_ input. In this case, the power line noise feeds back to the chip select and write paths, which trip the major internal control signals, CS1/CS1_ and W1/W1_ of Figure 4.1. As a result the OE1/OE1_ output enable signals also start to oscillate and trip the output, as seen in Section 3.4.

The oscillation of Figure 4.2 starts as the DQ1 output gets released from its high-impedance state before IO1T has stabilized to the logic 1 level. The IO1T signal for this device is the I/O line input to the first of four output buffer/driver circuits (ref. Fig. 3.16). The DQ1 output starts at the high-impedance level of 1.73V, determined by the tester load circuit. Then, at 5 ns, chip select CS_ powers up the device and generates a noise glitch on IO1T. Since the output is enabled before the noise glitch on IO1T, the noise gets amplified by the output drivers. This results in noise oscillations on the I/O lines, as well as the chip select and write paths discussed above.

4.2.2 Delayed Output Enable

A latch function similar to the output enable latch of Section 3.4 would not be feasible in this case. In order for a latch signal to stop the oscillation in Figure 4.2, noise would have to be detected in the chip select and write enable paths, as well as in the output enable path. This may not be possible due to the location of each control buffer, which may lead to long metal runs and unpredictable latch timing.

As an alternative to the output enable latch of Section 3.4, the problem can be avoided by delaying the output turn-on time just enough to get past the first noise

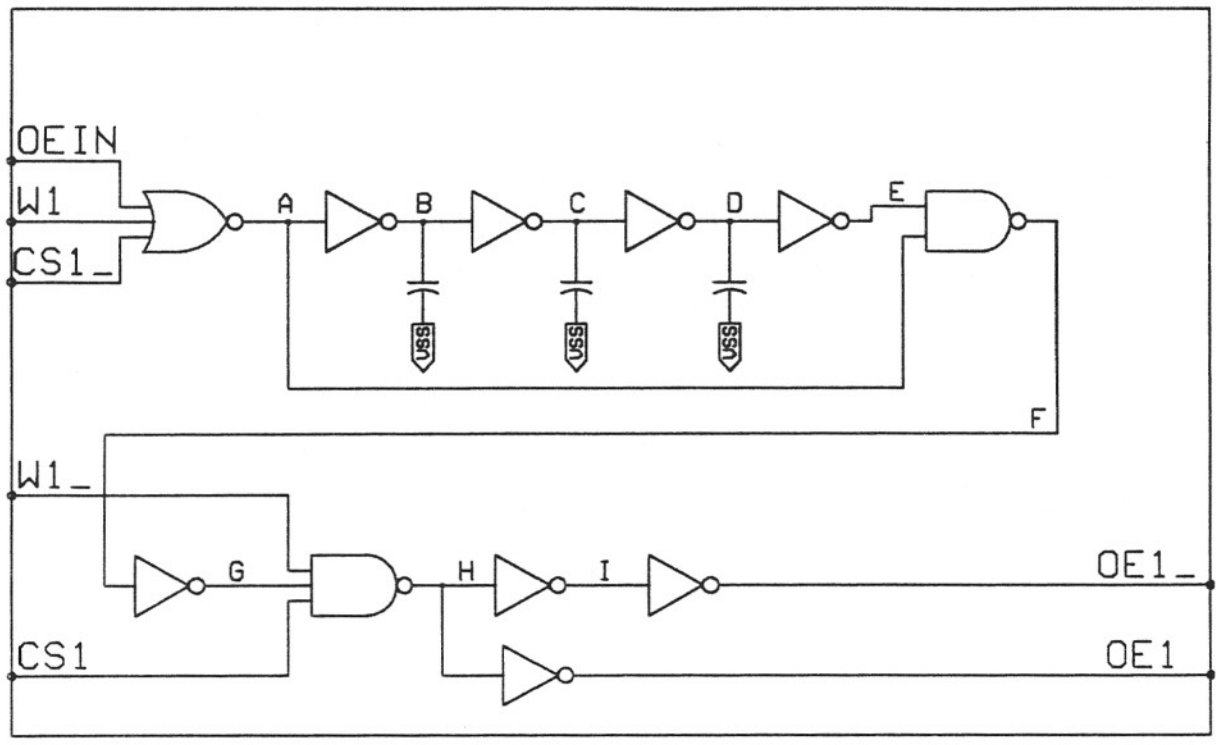

Figure 4.3: Output Enable Circuit with Conditional Delay

glitch, which is usually where most of the oscillation begins. Figure 4.3 shows a method which can be used to increase the enable time for the OE1/OE1_ signals without changing the disable time.

OEIN is an active low signal taken one stage after the TTL input NOR gate in the G_ input path. OEIN is tied to VSS for this device option, which does not have the G_ input. W1 and W1_ are internal signals from the write enable circuit. CS1 and CS1_ are internal signals from the chip select circuit (see Fig. 4.1).

The output disable path is taken when any one of the three NOR gate inputs makes a low-to-high transition. This delay path includes the 3-input NOR gate, the NAND gate at node F, and the gate delays through nodes G, H, and I. This sets OE1 to logic 0, and OE1_ to logic 1.

The output enable path is taken when all three NOR gate inputs are in the logic 0 state. This delay path is the same as the disable delay path, except for the addition of

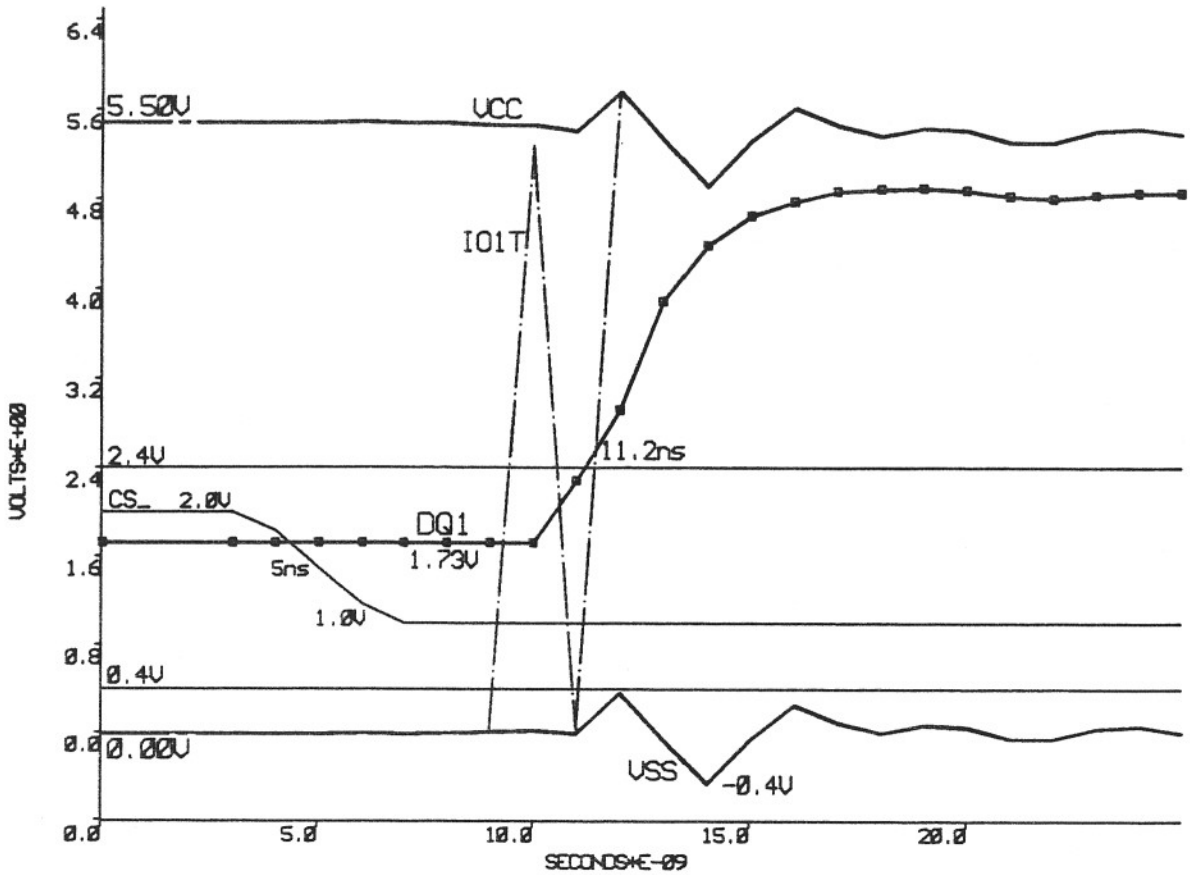

Figure 4.4: 16Kx4 Chip Select Access with Delayed Output Enable

the inverter/capacitor delay string. This behavior results from the fact that a logic 1 at node A must propagate through the inverter/capacitor delay before the 2-input NAND gate trips. For the disable path, a logic 0 at node A trips the 2-input NAND gate regardless of the length of the inverter/capacitor delay string.

Figure 4.4 shows the same simulation of Figure 4.2, but includes the new output enable circuit with the enable path delay string. By holding the output in high-impedance just past the first noise glitch on IO1T, the output drivers are prevented from reading (or amplifying) the noise glitch. As a result, feedback noise oscillation is reduced and the IO1T and DQ1 signals achieve their steady state voltage levels after the first noise glitch. For this process corner, the chip select access time decreases from 9.3 ns to 6.2 ns, as seen in Figures 4.2 and 4.4.

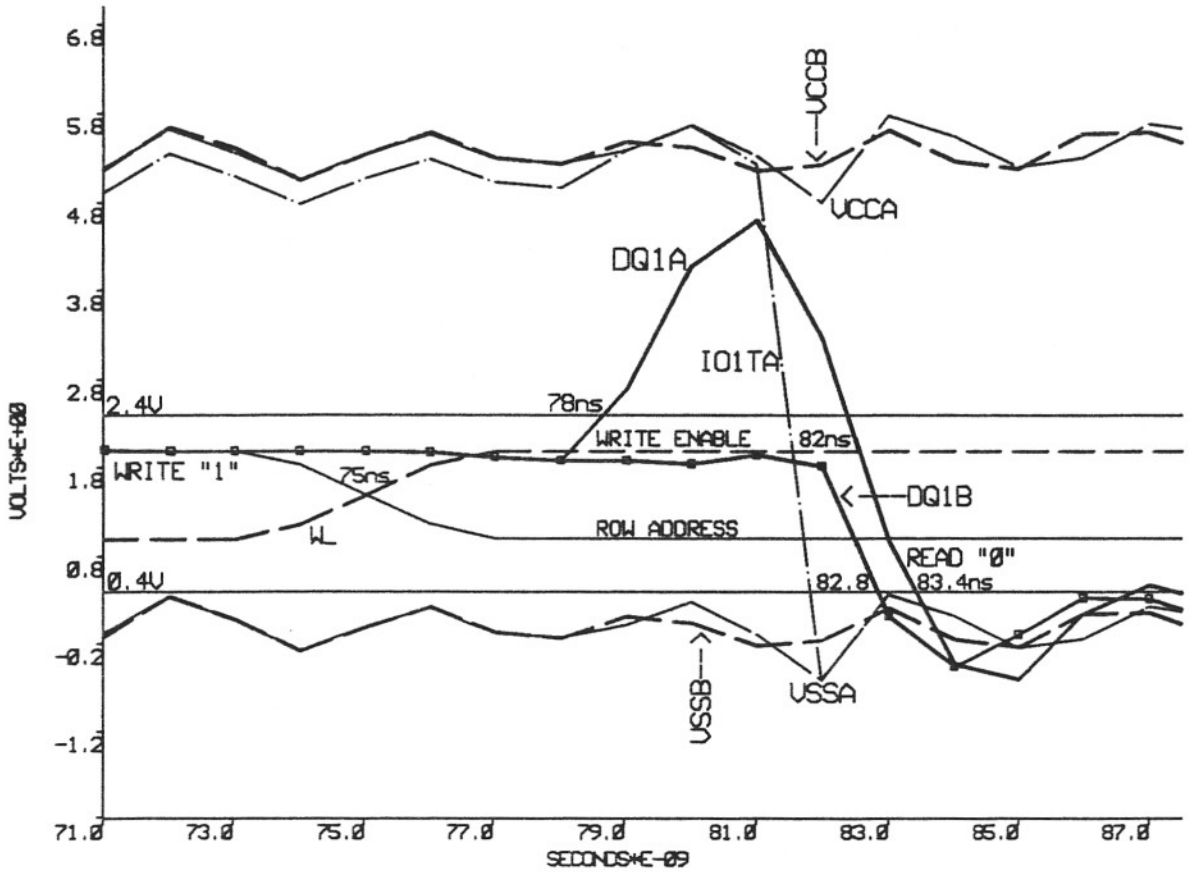

Figure 4.5: 16Kx4 Read After Write, Before and After Output Enable Delay

4.3 Effect of Output Enable Delay on Read After Write

A read cycle after a write cycle is also a critical noise region, as discussed in Section 4.1. Figure 4.5 shows DQ1 writing a logic 1 to a specified row address, and then reading the complement level from another row address. Signals ending in "A" are from the simulation without the delayed output enable. Signals ending in "B" are from the simulation with the delayed output enable.

Before the write cycle ends at 75 ns, DQ1A and DQ1B are set at 2.0V, which is the 2.2V VIH spec guard-banded by 0.2 volts. This level is used to write a logic 1 to the address location specified by row address at 2.0V. At the end of the write cycle, the DQ1 terminal becomes an output and attempts to read a logic 0 from the address

location specified by row address at 1.0V. Signal IO1TA shows the logic 0 data to be read by the output buffer.

In the absence of the output enable delay circuit DQ1A gets released at 78 ns, before IO1TA is valid. This causes DQ1A to rise to a logic 1. After IO1TA reaches its correct logic 0 state, DQ1A reverses direction and reaches the VOL level at 83.4 ns. This behavior on DQ1A causes a slight power line oscillation on VCCA and VSSA.

The same figure also shows what happens after the inverter/capacitor delay string is added to the output enable buffer. The new DQ1 output, DQ1B, is held in high-impedance for an additional 4 ns. At 82 ns, after the I/O line data is valid, DQ1B is released. This allows DQ1B to make one transition from high-impedance to the VOL state, rather than the two transitions seen on DQ1A. By eliminating the first transition for DQ1B, there is a 0.6-ns reduction in the logic 0 access time. During the output transition, the VCCB/VSSB power lines show approximately 0.5 volts reduction in noise magnitude, in comparison to the VCCA/VSSA power lines.

4.4 Delay Constraint

The delay circuit, which was used to slow down the output turn-on time in the last two sections, is limited by the access time in the slow process corner. The results of Sections 4.2 and 4.3 indicate a decrease in access time for chip select access and read after write access, respectively. This improvement in speed results from the reduction in output oscillation in the fast process corner, where most of the noise occurs. However, this is not necessarily true for the slow process corner, where output noise is not a problem. A slower part may show longer output access times, resulting from the inverter/capacitor delay string.

This increase in access time can result in read access failures, since access times are determined by the slow process corner. Due to this constraint, the inverter/capacitor delay string should be sized in the slow process corner, so as not to fail slow corner

access time, before it is used in the fast process corner. Although this does not guarantee enough delay in the fast corner, it determines the maximum delay that may be added. A rerun of the fast corner with this delay determines whether the change is worth the effort.

4.5 Output Edge Control

Spurlin and Stein [6, pp. 33-39] used a processing change to grade the output turn-on for TI's AC11244 (8-output Advanced CMOS buffer/line driver). The output transistors were split into many small subtransistors by removing portions of polysilicon gate segments, forming a serpentine pattern. The resistance of the polysilicon and the capacitance of each gate segment were used to form a distributed R–C network, which slows down the turn-on of each succeeding gate segment. By driving the gate from one end, a graded turn-on results in the series of subtransistors. This splits the total switching current into a series of smaller currents distributed over time and reduces the effective di/dt.

This method is valid for non-silicided polysilicon gate material, which contains enough distributed resistance and capacitance to provide an R–C delay that is sufficient to grade the turn-on. However, a silicided polysilicon gate material is used in TI's CMOS process to reduce the parasitic interconnection resistances. The R–C time delay for this process is too short to effectively grade the turn-on.

4.6 Bench Data

This section shows the results of a lab set-up, which illustrates device performance of TI's 8Kx9 SRAM before and after the use of noise reduction. The noise reduction methods include the circuit changes presented in the preceding sections and in Chapter 3, except for the output enable latch of Sec. 3.4. After writing all ones to one address location and all zeros to another address location, a row access is used to switch all nine outputs from logic 0 to logic 1. This low-to-high output transition contains the most noise since VCC and VSS drop in voltage at the start of the low-to-high transition (see Fig. 3.5, p. 28). When VSS drops below −0.7V (ref. Sec. 1.2), it trips all active low TTL inputs, e.g., CS_ and G_. For the figures of this section, "new 8Kx9" refers to an 8Kx9 SRAM with noise reduction. "Old 8Kx9" refers to an 8Kx9 SRAM, processed two years earlier, without noise reduction.

4.6.1 Output Enable

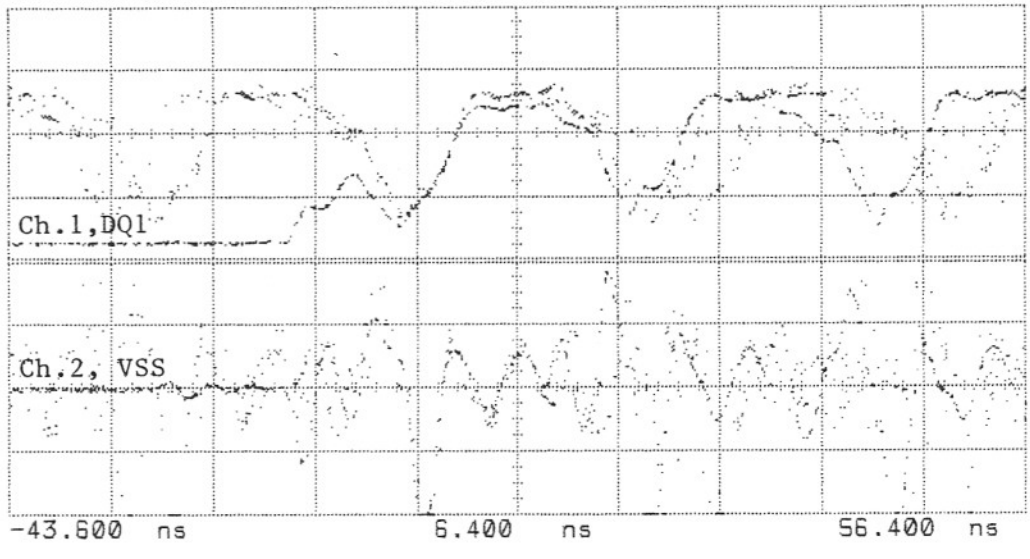

Figure 4.6: Old 8Kx9 with G_ =−1.0V and VCC=3.9V

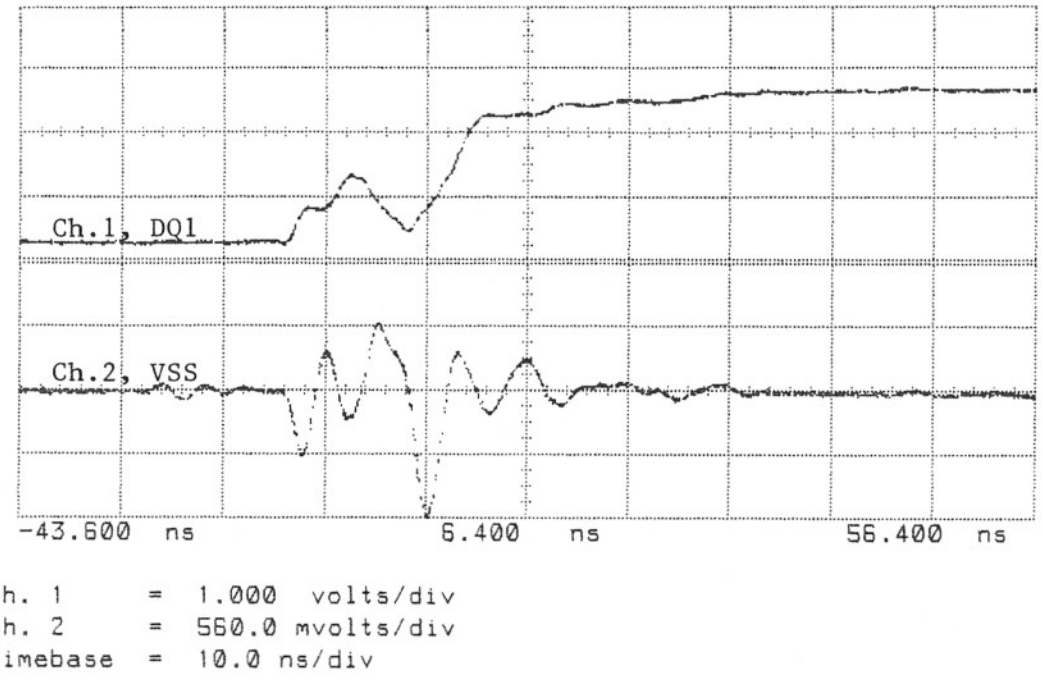

Figure 4.7: Old 8Kx9 with G_ =−1.0V and VCC=3.7V

Figures 4.6 and 4.7 show the effect of the VCC power supply level on the noise margin for the G_ input pin of an old 8Kx9 SRAM. Chip select CS_ is tied to ground for this lab set-up. Figure 4.6 shows oscilloscope waveforms of DQ1 and VSS for an old 8Kx9 SRAM. The figure shows output oscillations with G_ set to −1.0V and VCC set to 3.9V. VSS noise ranges from −1.1V to +1.0V. When VCC is lowered to 3.7V, the oscillation stops, as shown in Figure 4.7. This indicates that a VCC supply higher than 3.7V speeds up the part to the point where ground noise trips output enable G_ when G_ is set to a VIL level of −1.0V.

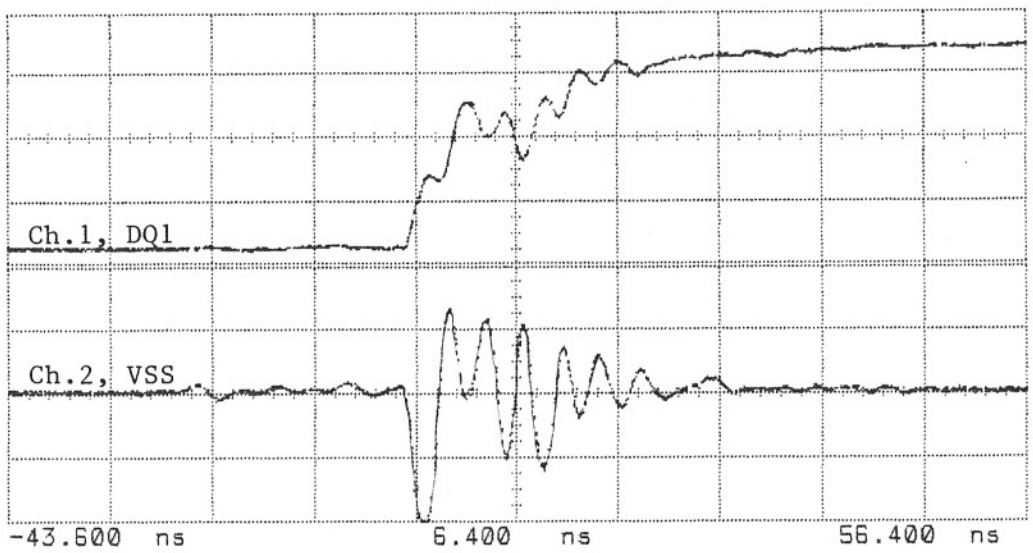

Figure 4.8: New 8Kx9 with $G_{\_}$ =+1.0V and VCC=5.5V

Figure 4.8 shows the improvement in the $G_{\_}$ noise margin for an 8Kx9 SRAM with noise reduction. VCC is set to the 5.5V level used in the fast process corner SPICE simulations. With a VIL level of +1.1V on the $G_{\_}$ input pin, the part shows good noise immunity to VSS, which ranges from −1.1V to +0.6V during the low-to-high output transition region.

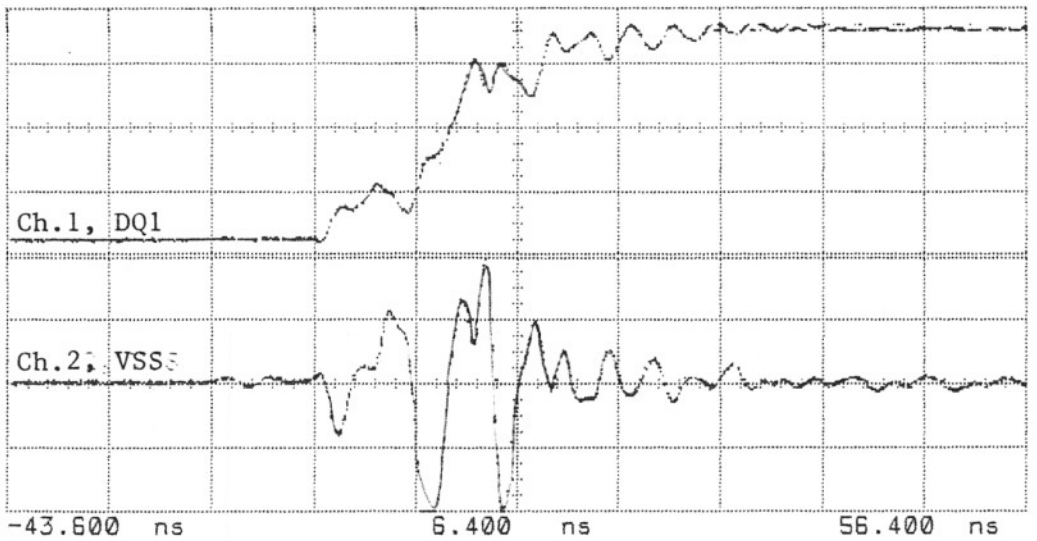

Figure 4.9: Old 8Kx9 with CS_ =−0.7V and VCC=5.5V

4.6.2 Chip Select

The effect of ground noise on chip select CS_ is shown in Figure 4.9. With VSS ranging from −1.1V to +1.1V, the VIL level of the CS_ input had to be set to −0.7V in order for the output to switch without oscillating.

Figure 4.10 shows the same row access of Figure 4.9 after the old 8Kx9 SRAM was swapped with a new part, which includes noise reduction. For approximately the same VSS noise magnitude, the new part passed with a VIL level of +1.1V at the CS_ input.

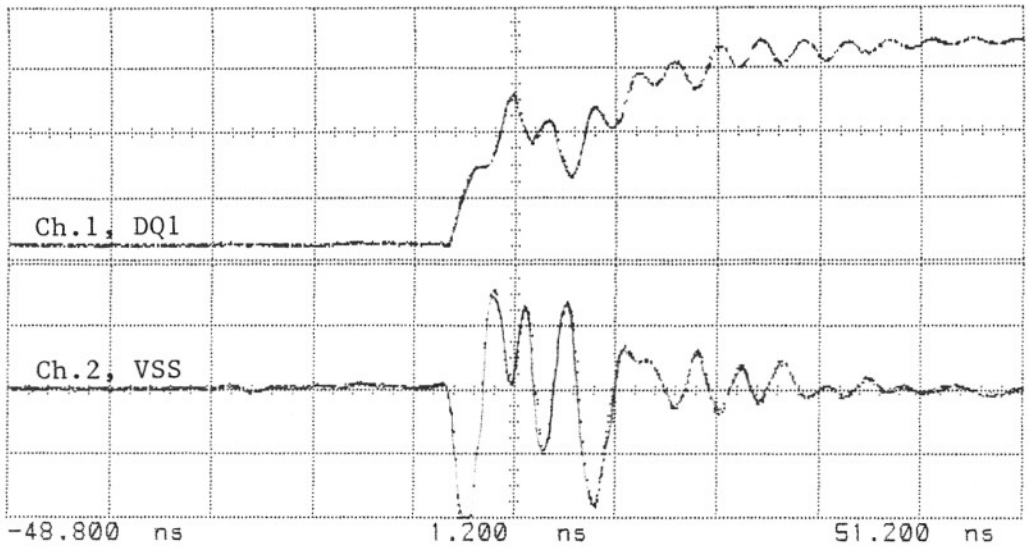

Figure 4.10: New 8Kx9 with CS_ =+1.1V and VCC=5.5V

4.6.3 Row Address

During the 8Kx9 lab set-up, the problem with noise margins was only seen on the active low control inputs. All addresses passed the the 0.8V/2.2V TTL input spec, regardless of whether the part had noise reduction. Figure 4.11 shows row address A0 and output DQ1 during the low-to-high output transition of an 8Kx9 SRAM with noise reduction. The VIL level for CS_ was set to +1.1V; G_ was tied to VSS; input levels for row address A0 were VIL=+1.0V and VIH=+2.0V. Figure 4.11 shows a 30 ns address access for DQ1 to reach the 2.4V VOH level.

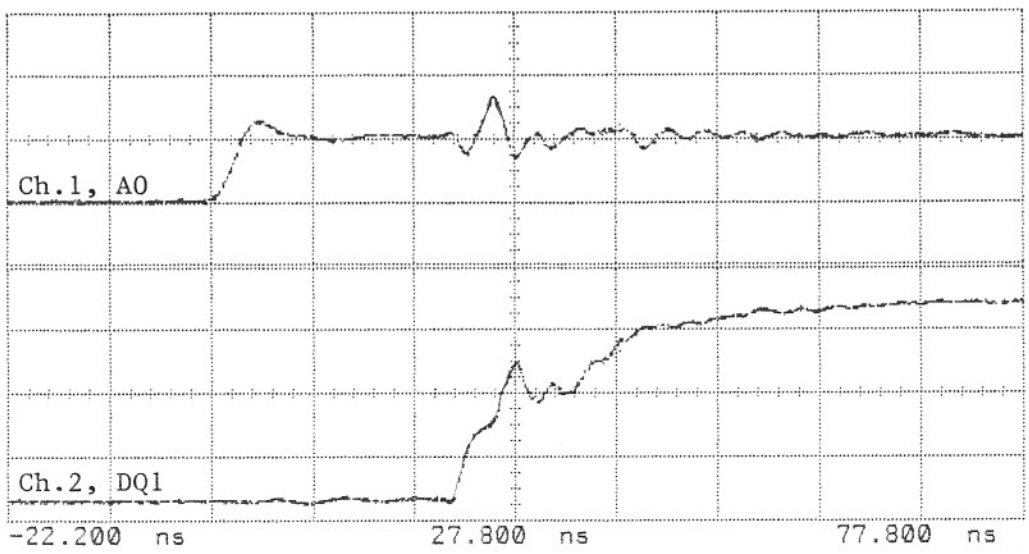

Figure 4.11: New 8Kx9 with CS_ =+1.1V and A0=1.0/2.0V

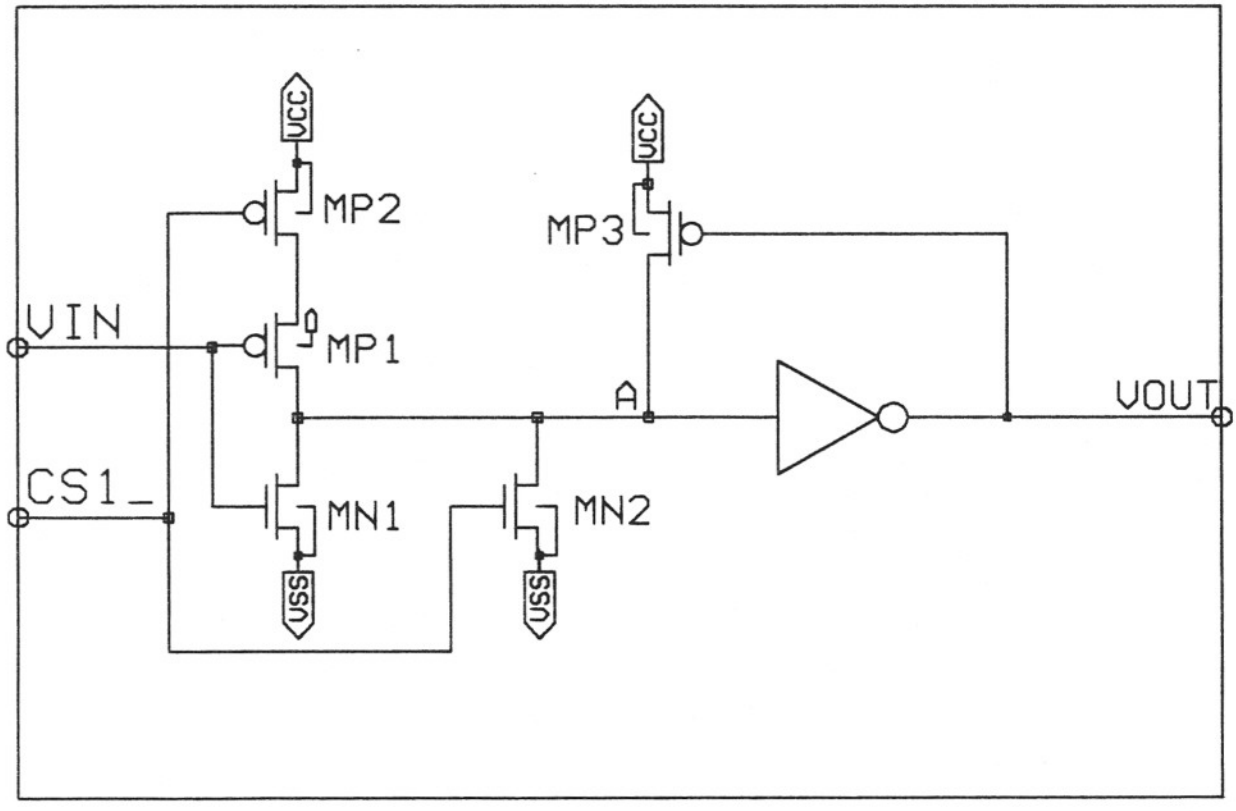

Figure 4.12: Input Buffer with Hysteresis

4.7 Hysteresis

During the latter stages of noise reduction, it may be necessary to add more noise immunity to the TTL input buffers. This section shows the results of a study on the use of hysteresis to gain additional noise margin for TTL inputs [3, p. 151].

4.7.1 Hysteresis Circuit Operation

Figure 4.12 shows a non-inverting TTL input buffer with hysteresis. VIN is the external TTL input and may be any address pin, I/O pin, or control pin input. CS1_ is the internal active low chip select signal. VOUT is the buffered output, with the same logic as VIN.

The first stage of Figure 4.12 consists of a conventional NOR gate used as an input buffer. When the SRAM is deselected, CS1_ is at logic 1 so that transistor MN2

conducts and transistor MP2 is cut off. This holds node A at logic 0 regardless of the state of the input, VIN. In this condition, all input buffers are deselected with VOUT at logic 1. The condition of interest is when CS1_ is at logic 0. In this case, MP2 conducts and MN2 is cut off so that node A responds to signal VIN. The NOR gate is comprised of transistors MP1, MP2, MN1, and MN2 and is essentially a clocked inverter which is enabled when CS1_ is active [2, col. 4,9-22].

When VIN is at logic 0, node A is at logic 1 and the inverter sets the output to logic 0. This output level causes transistor MP3 to conduct and enforces the logic 1 level at node A, which is significant as VIN approaches the trip level of the MP1/MN1 input to the NOR gate. The trip level of the NOR gate is determined primarily by the transistor gate widths (or gains) of MP1, MP2, and MN1. Transistor MP2 is sized significantly larger than MP1. This is done so that the ratio of the sizes of transistors MP1 and MN1 is the most significant determinant of the trip level and the effect of transistor MP2 is relatively minor. For active operation, transistor MN1 is sized significantly larger than MP1 to achieve a TTL trip level. For example, when transistor MN1 is sized four times larger than transistor MP1, the MP1/MN1 inverter pair trip at approximately 1.5V, the center of the VIL/VIH input spec [2, col. 5,1-22].

Transistor MP3 serves as a feedback device, which affects the trip level as signal VIN makes a low-to-high transition. When transistor MP3 is conductive, the trip level of inverter pair MP1/MN1 is increased because of the additional current required to flow through transistor MN1 as it becomes conductive. As signal VIN increases from its VIL (or logic 0) level, transistor MN1 starts to conduct when VIN exceeds the threshold voltage of MN1. This causes node A to drop in voltage towards VSS. Transistor MP3 being conductive retards the rate at which transistor MN1 can reduce the voltage at node A. Thus, the presence of transistor MP3 has the effect of increasing the trip level for the case of VIN switching from logic 0 to logic 1 (or VIL to VIH).

The effect of transistor MP3 on the trip level of the MP1/MN1 inverter pair is

determined by current magnitude between VCC and node A. This current magnitude is in turn related to the gate width (gain) of transistor MP3. The gate width of transistor MP3 can be selected to provide the desired change in trip level for the case of signal VIN making a low-to-high transition [2, col. 5,24-45].

When signal VIN makes a high-to-low transition, signal VOUT sets the gate of transistor MP3 to logic 1 at the beginning of the transition on VIN. During this input transition, the change in trip level is less than the change during the low-to-high transition because transistor MP3 is not conductive when VOUT begins at logic 1. As a result, the trip level is higher for the case of VIN switching low-to-high than it is for the case of VIN switching high-to-low. This difference in the trip level, dependent upon the direction of the input signal, is known as hysteresis. The hysteresis provides an added margin of safety in preventing the voltage on node A from changing direction for a very small change in input voltage due to ground noise, which moves the input voltage potential closer to the TTL trip level [2, col. 5,46-61].

4.7.2 Hysteresis Simulation

The effect of this circuit is shown in Figure 4.13. This figure shows the before/after hysteresis results for a chip select access from an 8Kx8 SRAM noise simulation. This simulation includes power line resistors, output driver noise sensors, and a delayed output enable. The device control inputs include CS_, W_, and G_. The signals that end with an "A" are from the simulation without hysteresis. Signals ending in a "B" are from the simulation with hysteresis.

The improvement in this simulation is seen in the OE1A and OE1B signals. Without hysteresis, ground noise tripped the control inputs and caused output enable OE1A to trip during the low-to-high output transition. This caused output DQ1A to oscillate during the transition. After adding hysteresis to all TTL input buffers, output enable OE1B stabilized, preventing output DQ1B from oscillating.

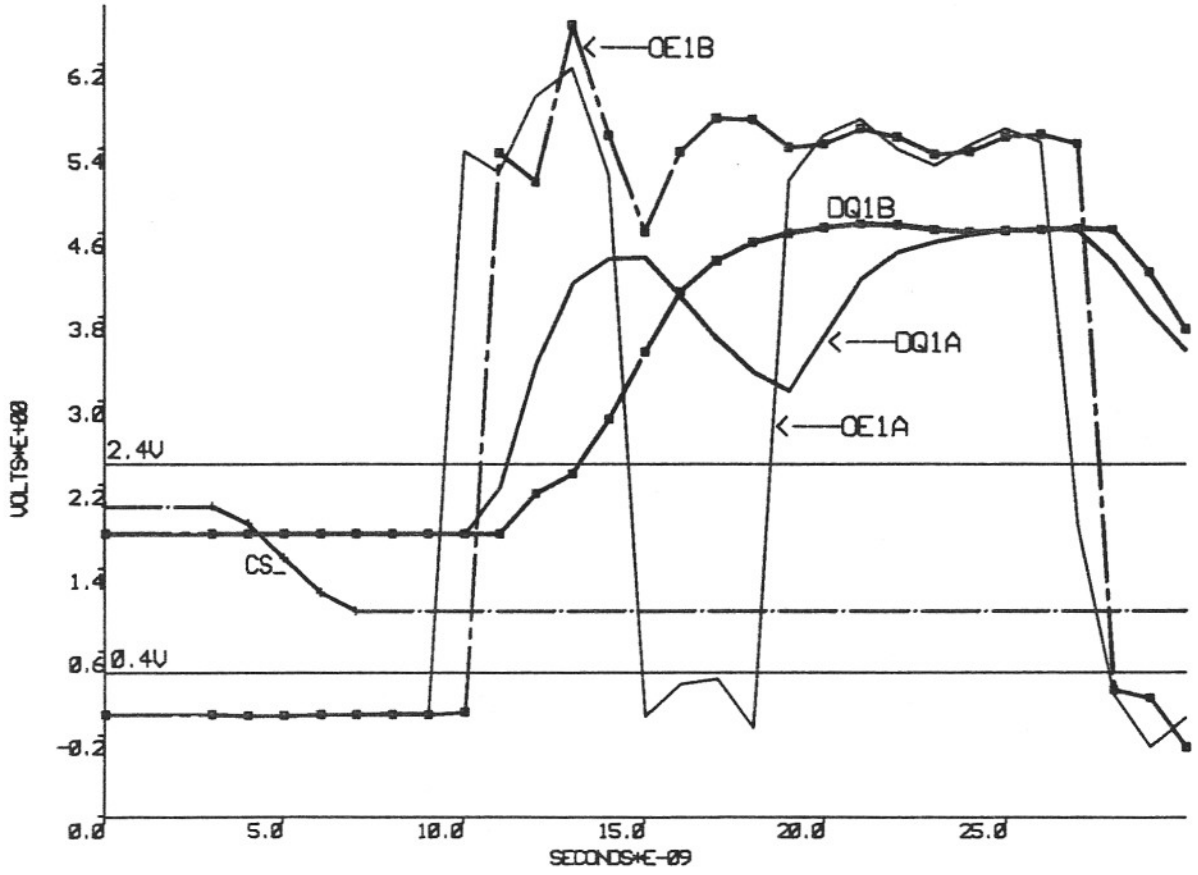

Figure 4.13: 8Kx8 Noise Simulation Before/After Hysteresis

4.8 Chapter Summary

This chapter focused on the feedback noise oscillations in the low-to-high output transition, introduced in Chapter 3, and presented alternatives to the output enable latch of Section 3.4.

The control circuit was analyzed with respect to ground noise tripping the three control inputs, CS_, W_, and G_. A 16Kx4 full-circuit simulation was used to show output behavior after chip select access, and also during a read after a write cycle. The control buffers were affected by ground noise even without the G_ input. The ground noise tripped internal signals in the chip select and write paths. As a result, output enable signals OE1 and OE1_ started to oscillate, tripping the DQ outputs to the high-impedance state. By adding a time delay to the output enable path, the

output drivers were held in high-impedance before being released to read (amplify) the oscillations. This reduced the power line oscillations, which, in turn, reduced the amount of ground noise fed back to the TTL control inputs.

Results from a lab set-up of the 8Kx9 SRAM were also used to illustrate the improved TTL input immunity to ground noise. A low-to-high transition on all nine outputs was used to illustrate noise immunity for active low control inputs CS_ and G_. Hysteresis was also discussed for additional TTL noise margin.

Chapter 5

Conclusion

This final chapter summarizes the results of the noise reduction methods of Chapters 3 and 4. These methods apply to any logic family with high-current outputs, e.g., CMOS, TTL, BICMOS.

5.1 Power Line Resistors and Noise Sensors

The power line resistors of Section 3.3 reduced the VCC noise magnitude by 50% for the case of nine outputs switching low to high. For resistor values, Section 3.3 used half of the maximum allowable values determined by the slow process corner conditions. The values used were :

- RVCC = 9 squares (18Ω max)
- RVCO = 1 square (2Ω max)
- RVSO = 1 square (2Ω max)

Speed degradation for these resistor values was :

- 4.0 ns loss for slow process conditions (9% for 45 ns part)
- 2.2 ns loss for nominal process conditions (6% for 35 ns part)
- 1.9 ns loss for fast process conditions (8% for 25 ns part)

The noise sensor devices of Section 3.5 worked in conjunction with the power line resistors. During output transitions, the VCO/VSO output driver supply voltage levels were determined by the inductive voltage on the power lines, as well as the power line resistor values. The power line resistors enhanced the contribution of the noise sensors, which were used to increase output driver impedance when the inductive voltage exceeded the threshold voltage of the noise sensors.

For worst case speed degradation from the use of the noise sensors, Section 3.5.2 reported a 0.7 ns increase in delay (8.7 ns to 9.4 ns) for slow process conditions. However, this was only for the output stage delay. For a 45 ns part, this delay increase amounts to 2%. Thus, the noise sensor devices show virtually no performance degradation.

The improvement in performance was shown in Figures 3.17 and 3.18, where a 6-ns row access was shown from a fast process 16Kx4 SRAM simulation. During the low-to-high output transition, the noise sensors reduced the output switching current, thereby reducing power line noise levels so as not to trip control buffer inputs set to TTL levels of 1.0/2.0V. As a result, output enable OE1 stabilized at logic 1, allowing the output to make a noise-free transition.

5.2 Output Enable Latch

Section 3.4 demonstrated the use of a latch circuit to eliminate the noise feedback path between the output drivers and the G_ input. This was done by detecting noise at an early stage in the G_ input path and latching the output enable signals before they get tripped by the noise. The results were shown in Figure 3.15, where two latch pulses stopped the lengthy oscillation seen in the previous figures. The use of the latch function redeemed 3.2 ns of access time, which was lost due to inductance noise and the use of power line resistors. Since process variations and control buffer circuit placements can lead to uncontrollable timing for the latch pulse, the latch circuit was not kept in the succeeding simulations.

5.3 Output Enable Delay

Section 4.2 presented a method which stretches out the high-impedance state of the outputs, as they are enabled into a read cycle. This allowed additional time for internal signals, such as the sense amp data entering the output buffer, to stabilize before being amplified by the output drivers. The improvement in performance was shown in Figures 4.2 and 4.4. In addition to eliminating the output oscillation, access time decreased by 3.1 ns for the chip select access.

The amount of delay which can be added to the output enable path is limited by the output enable access in slow process corner conditions. The output enable delay must be adjusted so as not to fail chip select access, read access after a write cycle, and G_ access in the slow process corner.

5.4 Bench Results

Section 4.6 confirmed the use of the power line resistors, the noise sensors, and the output enable delay to improve TTL noise margins for the 8Kx9 SRAM. A lab set-up was used to compare input noise margins for an 8Kx9 device without noise reduction to an 8Kx9 device with noise reduction. The low-to-high output transition, with all outputs switching simultaneously, was used to determine the worst case input levels for active low control inputs for 30 ns parts at 5.5V power supply. The results indicated that control inputs which had to be set to a VIL of -1.0V for a part without noise reduction may be set to $+1.1$V on a part which includes noise reduction.

Bibliography

[1] U.S. Patent Pending, Application # 126,539, *Noise Reduction for Output Drivers* Filed Nov. 30, 1987.

[2] U.S. Patent # 4,807,198, *Memory Input Buffer with Hysteresis and DC Margin* Feb. 21, 1989.

[3] Bernard C. Cole, "25 ns SRAMs Use Only 100 mW"
Electronics, July 24, 1986, Vol. 59, No. 27.

[4] Stephen T. Flannagan et al., "Two 64K CMOS SRAMs with 13 ns Access Time"
ISSCC Digest of Technical Papers, Feb., 1986.

[5] David A. Hodges and Horace G. Jackson, *Analysis and Design of Digital Integrated Circuits*
McGraw-Hill Book Company,New York, 1983.

[6] Craig Spurlin and Dale Stein, "EPIC Advanced CMOS Logic Output Edge Control"
Texas Instruments Technical Journal, April 1989, Vol. 6, No. 2.

[7] Neil H. E. Weste and K. Eshraghian, *Principles of CMOS VLSI Design.*
Addison-Wesley,Reading,Massachusetts, October 1985.

VITA

Doran David was born in Ramat-Gan, Israel, on May 4, 1962, the son of Matana David and Nagia David. In 1968, he moved to Chicago. After finishing his work at Maine Township High School North, Des Plaines, Illinois, in 1980, he entered Purdue University at West Lafayette, Indiana. He majored in Electrical Engineering, and continued his work at The University of Illinois, Urbana, Illinois, in 1982. He specialized in IC Fabrication, and received the degree of Bachelor of Science from The University of Illinois in May 1984. He was employed by Texas Instruments as Product Engineer in June 1984, and then as SRAM Design Engineer in July 1986. In September 1986, he entered the Graduate School of The University of Texas at Dallas.