Improving Toggle Rate in a Rail-to-Rail Comparator Output Stage

by

John D. Morris

Submitted to the Department of Electrical Engineering and Computer Science

in Partial Fulfillment of the Requirements for the Degrees of

Bachelor of Science in Electrical Engineering and Computer Science

and

Master of Engineering in Electrical Engineering and Computer Science

at the

Massachusetts Institute of Technology

May 23, 2001

Author _____
Department of Electrical Engineering and Computer Science
May 23, 2001

Certified by _____
Professor Hae-Seung Lee
Thesis Supervisor

Certified by _____
Brian D. Hamilton
VI-A Company Thesis Supervisor

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

Improving Toggle Rate in a Rail-to-Rail Comparator Output Stage

by

John D. Morris

Submitted to the Department of Electrical Engineering and Computer Science

in Partial Fulfillment of the Requirements for the Degrees of

Bachelor of Science in Electrical Engineering and Computer Science

and

Master of Engineering in Electrical Engineering and Computer Science

at the

Massachusetts Institute of Technology

May 23, 2001

ABSTRACT

The demand for high-speed components has driven an increase in the speed of analog comparators, a building block for many analog circuits. This paper describes the modification of one of Linear Technology's low-cost, high-speed comparators to increase the output toggle rate beyond the one hundred megahertz range. The essential modifications to the output stage mainly relate to overcoming quasi-saturation of the output devices by extracting unwanted, stored base charge. Additional features were added into a dual comparator package with a tiny footprint to increase consumer interest and to diversify it from other comparators in Linear Technology's line. A final circuit design and physical silicon layout were designed using computer design tools, and the IC was fabricated and tested. The first silicon was tested extensively and worked successfully with only minor undesired discrepancies that were deemed acceptable. The LT1715 design was successful since the design itself accomplished all the desired specifications and the part is now available for sale.

Thesis Supervisor: Hae-Seung Lee
Title: Professor, Electrical Engineering

Table of Contents

1 Introduction

1.1 Who Cares About Comparators?

"Comparators may be the most underrated and underutilized monolithic linear component," Jim Williams proclaims in Linear Technology Corporation Application Note 13. He then continues on to show numerous examples of oscillators, voltage to frequency converters, sample and holds, pulse stretchers, and more, all using comparators. The comparator is as much a building block in analog circuitry as the op-amp. Because it is a building block, the speed of the comparator must be increased to follow the trend of today's technology.

1.2 Demand for High Speed

As technology has advanced, people have demanded more data, and have demanded it faster. In 1985, Linear Technology released a new comparator, the LT1016, boasting a very fast 10ns propagation delay.[1] The comparator had lower propagation delay than the TTL or LS logic of its day, and was also capable of toggling in the tens of megahertz range, comparable to the speed of digital logic at the time.[2] Analog circuitry presented little barrier to increasing data rates.

Today, however, the situation has drastically changed. Microprocessor logic has eclipsed the 1GHz range.[3] Technologies for data transfer have reached bandwidths of 400Mbps in a single serial cable, a remarkable advancement from the data rates of the mid-1980s.[4] While digital circuits have been keeping up with advancements, if not driving them, analog circuitry has begun to fall behind. Today a demand has developed for much faster

[1] The LT1016 datasheet can be viewed on Linear Technology's website at: http://www.linear.com/pdf/1016fb.pdf .
[2] See On Semiconductor's website at http://www.onsemi.com for datasheets of old LS-logic IC's.
[3] Both AMD and Intel offer processors for the general public exceeding 1GHz
[4] A description of Firewire technlogy can be found at http://www.apple.com/firewire/ .

comparators that can transmit or translate data at the high speeds present in a mixed-signal system. However, very few comparators were designed to operate at high frequencies and those that do are often difficult to use or very demanding for power. This thesis will explore the process of redesigning a comparator to work in modern high-speed applications where existing comparators cannot.

1.3 Organization

This thesis explains the design of a new comparator IC from conception through verification, in a manner closely correlated to the chronological order that the design progressed, and with a primary focus on the toggle rate limits of the output stage:

1. **Existing Parts and Their Limitations (Chapter 2):** Linear Technology carries a wide range of high-speed comparators; the LT1719, LT1720 and LT1721 series are among the simplest and fastest with just 4.5ns of propagation delay and less than 5mA quiescent current per comparator. Customers have requested a similar part with a higher toggle rate than currently available, but retaining the features of the existing LT1720 series. A new part was conceived that would fit into the same line and fulfill customers' needs.

2. **Understanding Toggle Rate Limitations (Chapter 3):** The toggle rate of the existing comparators is limited primarily by the near saturation of its rail-to-rail output stage. The actual near-saturation operation of the output emitter followers is termed quasi-saturation, and the AC effect of the same condition is termed AC soft saturation. Understanding this operating region is essential to increasing toggle rate.

3. **Design Solutions (Chapter 4):** While several initial designs failed for various reasons, a final compromise was reached and an effective circuit devised. This chapter studies the failed circuits briefly and the final circuit design in detail with analysis across process variation and temperature.

4. **Other Circuit Design Issues (Chapter 5):** To maximize the speed of the comparator, slew currents were increased in the output stage. Trims and hysteresis were improved to increase DC performance. These modifications would allow the comparator to outperform its predecessors for specifications beyond toggle speed.

5. **Testing Methods and Procedures (Chapter 6):** AC, DC and debugging test fixtures were designed to test bare silicon wafers, as well as packaged parts. Attention was paid that measurements would be accurate representations of what the parts were doing, and not artifacts or irregularities caused by equipment.

6. **Toggle Rate Testing (Chapter 7):** The specifications for toggle rate were pieced together from those of mainstream logic families. Parts were only tested by hand; a digitally incremented phase locked loop circuit was designed for production testing of the comparator's maximum toggle rate but was not completed.

7. **Characterization and Correlation (Chapter 8):** On arrival, initial silicon was tested on a probe station to ensure proper fundamental operation. Soon after, wafers were sorted by performance to DC specifications. Packaged parts were sorted based on the same tests as well as AC tests. The results of testing and characterization are tabulated and evaluated.

8. **Error Sources and Discussion (Chapter 9):** Although the part was successfully released, it does have flaws. The hysteresis did not match perfectly between channels on the same die, the two comparators interacted at high speed, and the temperature stability of the DC trip points was poor. Each of these problems is examined and a solution proposed, so that in future revisions all defects can be eliminated.

9. **Comparisons and Conclusion (Chapter 10):** The LT1715 successfully improved on the other members of its family and was a success from the standpoint that it met all its design criteria and completely eliminated the constraint of quasi-saturation. Additionally, by meeting or surpassing the specifications of its competitors, it is guaranteed success in its market.

2 Existing Parts and Their Limitations

2.1 Linear Technology's High Speed Comparators

Linear Technology's comparator line is anchored by the popular LT1016, one of the first modern high-speed comparators. The LT1016 is a single or dual supply comparator with TTL inverting and non-inverting outputs, no current supply spiking in the linear region, and output latch capability. For years it was the industry standard, but it did lack one very useful feature: ground was outside the single supply common mode range. The LT1116 evolved from the LT1016, with the primary change being the ability to sense ground.[5] This feature came at only a minute cost to performance, but the upper limit to the common mode range was decreased by the same amount that the lower limit was brought toward ground.

Some years later, process technology improved and it became apparent that a better speed-power product was attainable for a new comparator design. The LT1394 is designed on a newer bipolar process but is pin compatible with the original LT1016.[6] Not only does the LT1394 have much lower quiescent current and faster propagation delay, but it also regains a volt of common mode at the top of the range while maintaining ground sense capability.

After introducing this modern high-speed comparator, Linear Technology decided to expand its comparator product tree to fit other markets with the addition of two new comparator families. One family of comparators would emphasize features and flexibility. The other family would consist of simple, low power, low cost parts. A rail-to-rail output stage was deemed necessary for the new comparators in order for them to reach valid logic levels on the low supply modern faster and more efficient digital systems run on.

[5] This datasheet is shown on Linear Technology's website at http://www.linear.com/pdf/lt1116.pdf .
[6] This datasheet is also shown online, at http://www.linear.com/pdf/1394f.pdf .

The high-end line of products was first released with the LT1714.[7] It is the top-of-the-line comparator Linear has to offer, a part with rail-to-rail inputs and rail-to-rail complimentary outputs, individual latches for both of the comparators in a dual, a linear output stage without supply current peaking in the linear region, and operation on single or dual supplies ranging from 2.7V to 12.6V.

The low-cost series of comparators started with the LT1720, another high-speed dual comparator.[8] This series is spared the rail-to-rail input stage, latch, and linear output stage, but retains the rail-to-rail output stage. Its simplicity increases its speed, bringing propagation delay down to only 4.5ns. At the same time, well-controlled internal hysteresis simplifies usage. In addition to a quad-comparator version, the LT1721, an innovative single comparator adaptation was also designed. The LT1719 is a single of the LT1720, but it adds a low power shutdown mode as well as separate and independent input and output supplies that allow the comparator to level shift between analog and digital supplies.[9]

The concept for the new comparator would be a dual, like the LT1720, using the same topology as the other comparators in the low-cost series, but with a substantially increased toggle rate and the separate supplies of the LT1719. The part would not have shutdown or a latch, and would fit both comparators in a tiny MS10 package.

2.2 LT1720 Simplified Circuit Description

The block diagram of the LT1720 is shown in Figure 2-1. The LT1720 has a pair of differential inputs, a single output, as well as single supply terminals V_{CC} and ground. Each

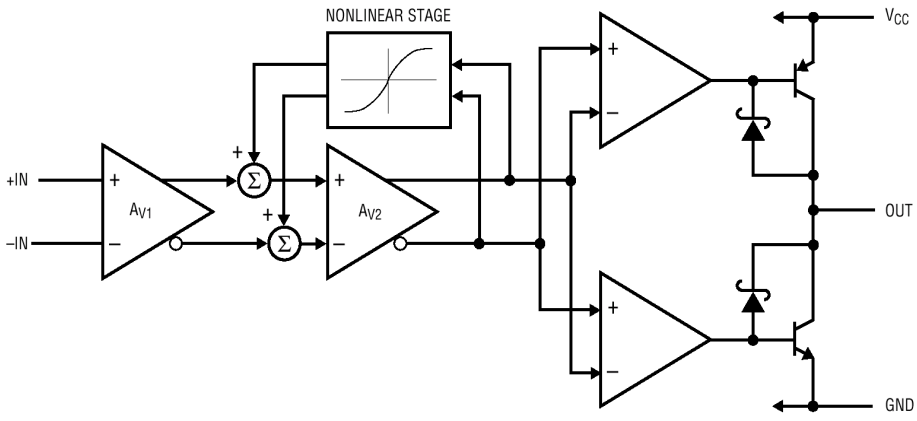

Figure 2-1: LT1720 Block Diagram

[7] This datasheet is online at Linear Technology's website: http://www.linear.com/pdf/1714i.pdf
[8] See this datasheet online at: http://www.linear.com/pdf/17201f.pdf .
[9] This datasheet is found at: http://www.linear.com/pdf/1719f.pdf .

of the two comparators is independent inside the package, sharing only the power and ground pins. The comparator design can be divided into the input stage, the gain stage, and the complementary output stage.

2.2.1 Input Stage

A simplified schematic of the input stage can be seen in Figure 2-2. The input uses a PNP differential pair tied together by Schottky diodes at the emitters. The differences between the input stage in the LT1720 and a simple differential pair with a current source are for improvements in common mode range and input voltage range.

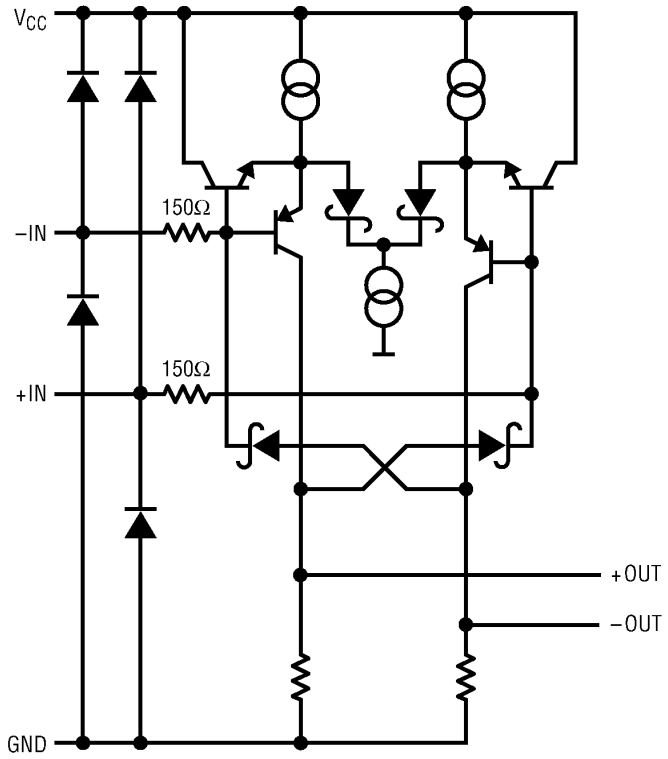

Figure 2-2: Input Stage Topology

First, the pair of Schottky diodes at the emitters of the input devices protects the input devices from reverse base emitter breakdown when the differential input voltage is large. If a differential pair is tied with diodes as shown, and the current source in the cathodes of the diodes is less than the current in each of the input devices, then the input devices will always be forward biased, and will always have some collector current. The side of the PNP mirror that is higher will have the diode at its emitter

conducting some fraction of the input device bias current, while the other diode in the pair will be reverse biased and the lower input device conducting its full bias current. This is beneficial for protection of the input devices. When one input falls well below the other, the reverse V_{BE} breakdown voltage could easily be reached. However, the Schottky diode will carry the entire reverse drop to protect the transistor in this case.

To improve the operation of the differential pair at the top of the common mode range, NPN's are added in parallel to the PNP input devices. Additionally, the current source at the cathodes of the diodes is now larger than the bias currents of each input device. As the two PNP input devices approach the top rail, they will eventually saturate their bias current sources. When the two inputs are more than two diode drops apart, the parallel NPN will forward bias on the higher input. At that point, the diode connected to the opposing input device's emitter will turn off in turn. With the NPN input device on, the input can now go almost to the rail without saturating the current source since the emitter of the input device is now a PNP drop as well as an NPN drop further from the rail than it would have been before, leaving an NPN V_{BE} of headroom. Additionally, once the current source does begin to saturate, the NPN can supply current to the Schottky diode at it's emitter all the way to a V_{BE} above the rail, at which point the operating range will have been thoroughly exceeded and the ESD devices turned on.

The other pair of Schottky diodes, which connect the input devices' bases to the opposing devices' gain resistors, prevent phase reversal at the bottom of the common mode range. Without them, if one input were to fall below the bottom rail while the other input stayed within the common mode range, the low output device would begin to saturate. As the device saturates, the voltage on the gain resistor would decrease until reaching the negative rail, potentially a lower voltage than the opposing device. While this situation cannot be prevented, the input can be protected from inversion. By tying the opposing gain resistor through a Schottky diode to the input, the opposing "output" of the input stage will be pulled low with the falling input. The difference of the PNP V_{BE} and the Schottky diode drop is greater than the voltage on the gain resistor, so the opposing collector will be pulled below the negative rail before the falling input transistor saturates and the voltage on its collector falls toward the rail.

A final detail of the input stage is the input resistors of 150 Ω after the ESD diodes. They provide additional protection from ESD strikes into the input pins.

2.2.2 Gain Stage

The gain stage can be seen in simplified form in Figure 2-3. The stage is primarily composed of an NPN diff pair. A pair of emitter followers buffers the output from the input stage, before the signal is level shifted by Schottky diodes to create enough headroom for the diff pair current source to avoid saturation. At the diff pair's collectors, the signal is taken again through another pair of emitter followers before entering the output stage.

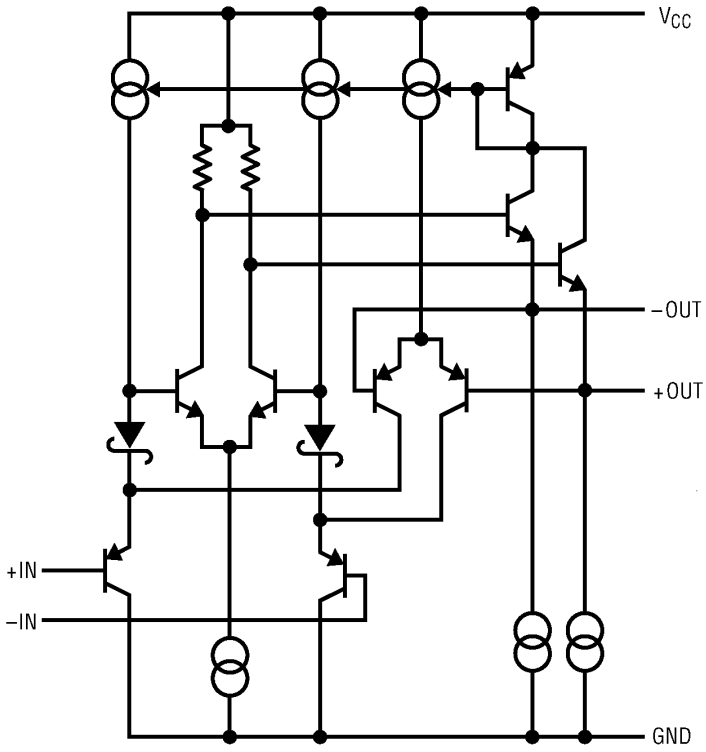

Figure 2-3: Gain Stage Topology

Hysteresis is implemented by switching current onto the emitter followers at the input of the gain stage to change their forward bias voltage, creating some minimal and well-controlled hysteresis. The diff pair responsible is switched by the output of the gain stage and the currents are summed in at the Schottky diode level shifters' cathodes, into the PNP followers emitters. The current for the hysteresis is mirrored

by the same source as the emitter follower legs, so the ratio of hysteresis is unaffected by process for a given temperature.

2.2.3 Output Stage

Figure 2-4 shows a simplified output stage schematic. At the input to the output stage, the differential signal is split into two independent drive signals for the complementary output devices. This is accomplished by switching the differential signal with two complementary differential pairs, one pair referred to each rail. The complementary signal is then level shifted one diode drop further from the rail through an emitter follower, and then buffered through another emitter follower for current gain to the output device. The output device is a common emitter referenced to the rail with a Baker Schottky clamp to its collector to prevent saturation. Reverse biased ESD diodes at the output protect the circuit from ESD strikes in the output line.

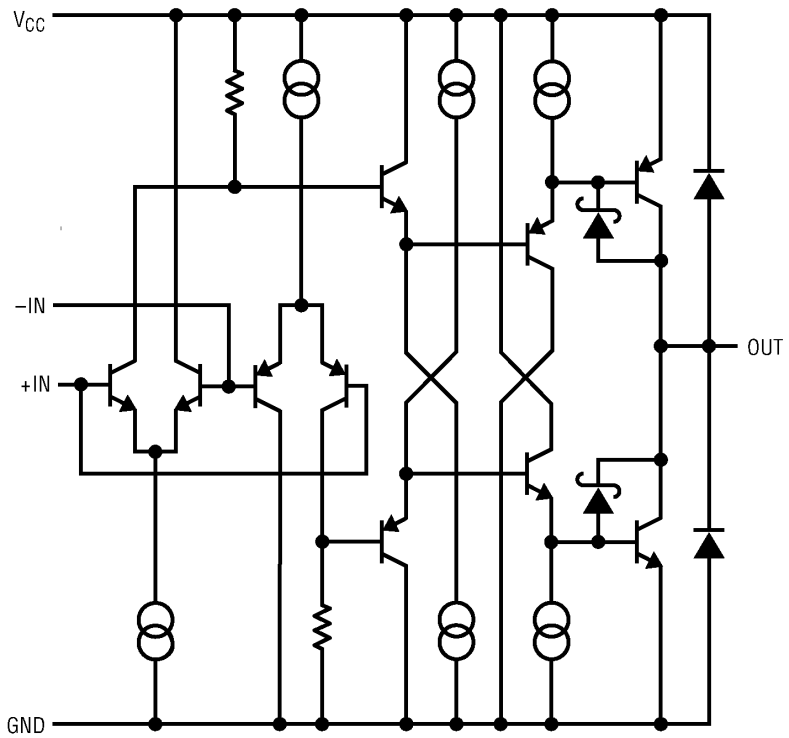

Figure 2-4: Output Stage Topology

2.3 LT1719 Differences

The LT1719 is a single comparator version of the LT1720 with some additional modifications. The LT1719 adds a shutdown control to reduce power consumption when the comparator is not in use. This feature actually led to some circuit redesign and since it will be left out of the dual comparator described in this report, the details will be ignored.

The LT1719 has another feature, however, that would be very important for the new comparator. The LT1719 has separate input and output supplies to accommodate separate analog input ranges and output logic levels. The input stage and biasing uses the input supplies exclusively, since the common mode and input voltage range should be decided entirely by the range of the input supplies. The gain stage is connected to the bottom input supply (V_{EE}) and to the top output supply ($+V_S$). This allows the gain stage to act as a level shifter between the input supplies and the output supplies. The output stage, while biased mainly by levels in the input supply's bias, runs entirely on the output supplies. The range of supply possibilities is illustrated below, with the connotations: V_{CC} and V_{EE} are the input supplies, and $+V_S$ and 0V (ground) are the output supplies.

$$(\text{Gnd} + 2.7\text{V}) \le +\text{Vs} \le (\text{Gnd} + 6\text{V})$$
$$\text{Vee} \le \text{Gnd}$$
$$(\text{Vee} + 2.7\text{V}) \le \text{Vcc} \le (\text{Vee} + 10.5\text{V})$$

Some examples of possible power supply configurations can be seen in Figure 2-5. The new comparator will use the same power supply implementation.

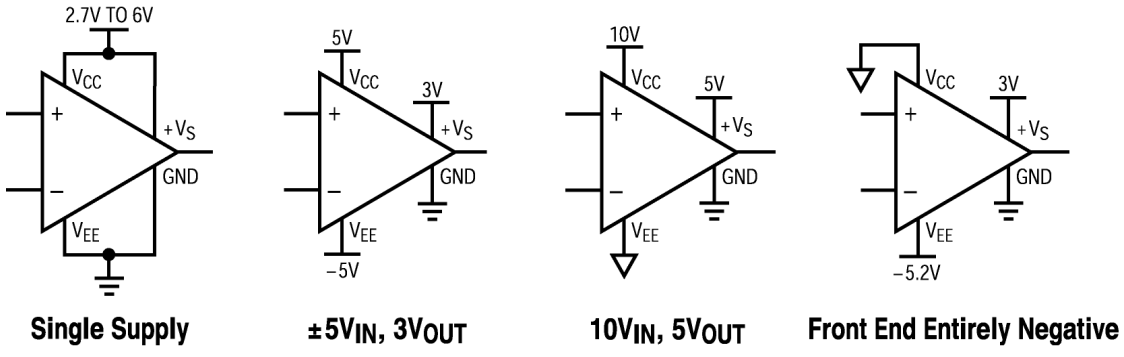

Figure 2-5: A Variety of Possible Supply Configurations

3 Understanding Toggle Rate Limitations

3.1 Expectations For Toggle Rate

While the LT1720 and LT1719 both excel in high-speed applications, they are both constrained by their toggle rates. Toggling is defined as slewing the output between some output-high voltage (V_{OH}) and some output-low voltage (V_{OL}) in response to an input sinusoid of some selected amplitude. The LT1720 is specified to run at 62.5MHz typical at room temperature with 50mV of overdrive. This number may seem suspicious since the rise and fall times of the output waveform on a 5V supply are just 2.5ns and 2.2ns, respectively. Therefore, a complete toggle should be possible in just 2.5ns plus 2.2ns, or 4.7ns, which would allow a maximum toggle rate of 212MHz. No LT1720 or LT1719 could possibly toggle even half that fast.

The input and gain stages of the LT1720 family have no problem with high-speed inputs. The on-chip parasitic capacitances are relatively low since devices are small and metal routing short and narrow, and the signal amplitudes are generally only plus or minus one base emitter voltage to switch differential pairs on and off. Even in the output stage, signals are buffered at low voltages all the way to the output devices. However, it is the common-emitter output device that causes our limitations. Due to its large size, the large amount of metal routing required to connect it, and the significant package parasitics and output load at its collector, the output device is the slowest part of the comparator.

3.2 Toggle Rate Limitations Explained

3.2.1 Output Device Near Saturation

The greatest limitation to toggle rate performance is caused by the very low V_{CE}'s on the output devices required by rail-to-rail operation. While the Baker clamp and another, proprietary clamp keep the collector voltage from falling a Schottky diode

drop below the base voltage, limiting V_{CE} to $V_{BE} - V_D$, or generally .3V or more when the device is on. However, although the device does not fully saturate, it still does not react as quickly as a device performing in the forward active region. This is because, although the device is not fully saturated, its base-collector epitaxial region is becoming forward biased, and charge is accumulating. The accumulation of charge reduces turn-on and turn-off speed, since the charge needs to be removed before the output device can switch itself on or off.

The output device is able to turn on very quickly since the emitter follower leading to the output device can supply as much current as the low output impedance emitter follower device is limited to. Ordinarily, when the voltage applied to the base of the device decreases, the current decreases and the device turns off. However, if the device was saturated, or in this case, quasi-saturated, turning it off would require extracting the excess charge in the base accumulated from operation in that region. Therefore, there must be a turn off current. A minimal current was in the original design in the LT1720 output stage because this effect was not recognized (or modeled in PSpice, at the time), and the sole intention was to keep the emitter follower in the forward active region in the absence of a drive signal.

Once the quasi-saturation problem was understood, the solution was simple. Due to the operating region of the output devices, they were guaranteed to gain or lose excessive amounts of base charge when their collector-to-emitter voltages were minimal. The NPN output common-emitter, for example, while slewing the output low, would go into the base push-in region and would gather a substantial additional base charge. Then, if the complimentary PNP output device started to slew the output high again, the bottom NPN would not be immediately turned off by decreasing its base voltage, and it would begin to carry some substantial fraction of the PNP's slew current, preventing the output itself from slewing high. Figure 3-1 illustrates the shoot-through currents in a PSpice simulation. The bottom plot displays the output voltage waveform while the top waveforms are the output device currents. The 1ns-long current spike after the slew current has ended is the substantially minimized shoot-through current simulated in my final circuit. In the original LT1720, this shoot-through approached a substantial fraction of the slew current and lasted for several nanoseconds.

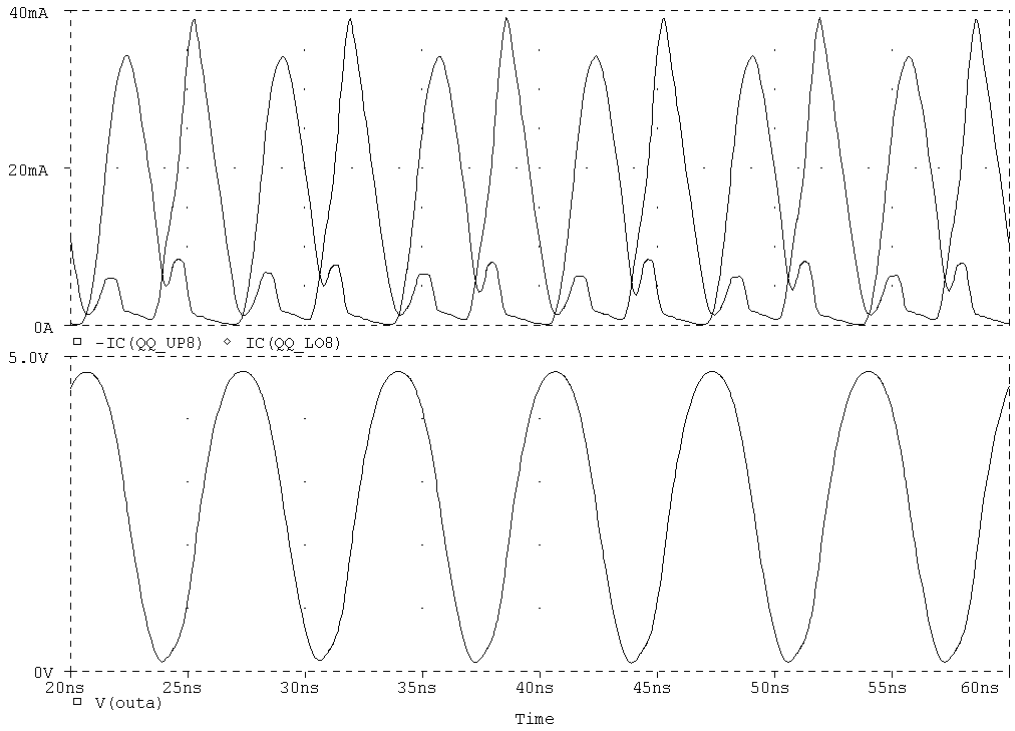

Figure 3-1: PSpice Plot of Shoot Through Currents

3.2.2 Base Drive

Another toggle limitation is slew current. The output can be summarized as a symmetric (top and bottom) Darlington configuration with the output device Baker clamped. Therefore, from the input of the Darlington to the collector of the output device, there is β-squared current gain.

However, the circuit is not quite that simple. The current source at the base of the output device will be removing (in the NPN case) some of the current from the emitter follower, thereby reducing the current gain. When circuitry is added later in this paper, the current source is significantly increased in magnitude, and a substantial amount of current that would be multiplied by β to the output disappears. So if there was β times the 50uA or 100uA available at the base of the emitter follower, but several milliamps taken away at the base of the output device, there would be minimal base current to β multiply, especially if β is small.

Furthermore, β may drop considerably with process variations and under extremely cold operating conditions. For the part to maintain its low ~5mA quiescent operating current per comparator, the stage before the Darlington was originally running at only

around 100uA. Since the output node contains parasitics from the large output devices, the large output trace, and the bond inductance, 25-40pF is representative of the capacitive load presented to the output devices. In order to drive a 40pF load 4V within 2ns (the rise time for around a 100MHz toggle rate in the finished design), I=C*dV/dt would indicate that the output device would need to carry a continuous 80mA. This requires a substantial amount of current gain in lieu of the losses to the current source mentioned earlier. Also, without positive feedback or another boosting mechanism, it is hard to regulate how much V_{BE} is required to turn the Darlington on hard enough to get that. The V_{BE}'s of the previous stages are dependent on at least one IR drop, and across some parameters, there isn't enough base voltage drive to drive the Darlington input to its current saturation limit. Modifications would be made to eliminate the problem.

3.3 Explaining Quasi-Saturation

3.3.1 Conceptual Understanding

In order to increase the collector to emitter breakdown voltage (VB_{CEO}) of a bipolar transistor for use in circuits with large supply voltages and therefore large V_{CE} potentials, the epitaxial collector region is made lightly doped. However, at high injection levels the internal base-collector junction may begin to forward bias although at the external base and collector connections, the junction will look reverse biased. This region with saturation-like effects, unseen from the transistor's outer terminals, is termed quasi-saturation (soft saturation). Although external voltages imply that the transistor is operating in the forward active region, additional base current is being conducted through the base collector diode, decreasing the current gain of the device. In a DC I_C vs. V_{CE} plot, this leads to a smooth rounding of the curve between the forward active region and the saturation region, and is especially noticeable at high current levels where the collector region has even larger internal voltage drop.

Using a charge-storage model understanding of the bipolar, the problem's effect in the rail-to-rail output stage is more apparent. In order to turn on a silicon junction, some amount of charge needs to be stored at the junction. The depth of the depletion region and the width of the junction, as well as the relative doping levels determine the amount of charge stored in forward biased junction. In order to push a transistor

into the forward active region, some amount of base charge is applied. However, in the quasi-saturation region defined above, there is now another junction being forward biased, therefore requiring some amount of charge. Due to the low doping level of the collector, this charge turns out to be fairly significant, much like the very large amount of accumulated charge (q_s) when the transistor is saturated (large compared to the charge required (q_f) for the base emitter junction to be forward biased).[10] In the AC f_T vs. I_C plot, this is especially apparent as the curves become more significantly V_{CE} dependent, and f_T is generally much lower for small collector to emitter potentials than parameters would otherwise indicate.

The quasi-saturation region is a problem for the circuit designer because for a device to be turned off it is not enough just to short the base-emitter junction to set V_{BE} to zero. The internal base voltage decay time-constant between the BE and BC junction capacitance and the base resistance will take a substantial amount of time compared to the turn off time in the normal forward operating region, where only the base emitter junction has stored charge. To increase the speed of a transistor in this region, a current should be applied to the base of the transistor to extract the excess charge.

3.3.2 Quasi-Saturation Spice Modeling

In PSpice, quasi-saturation is modeled with the inclusion of several new circuit elements. The collector epitaxial region resistance (R_{CO}) is used to define a current through the layer, I_{EPI}. The amount of charge stored in the layer is divided into two distinct charge storage elements, Q_O and Q_W, related by scaling factor Q_{CO} and γ (gamma), the doping factor of the collector region.[11] The math is complex, but the result in the PSpice model is an additional pair of capacitors, C_X and C_I, from the internal base node (inside R_B) to the collector region, separated by a current source, I_{EPI}, connecting the external collector junction to the internal collector junction, where the Gummel-Poon transistor model is used.[12]

---

[10] For more description of the charge control model, look to: H. K. Gummel and H. C. Poon, "An integrated charge control model of bipolar transistors," *Bell Syst. Tech. J.*, Vol. 49, May 1970:827-852.

[11] For additional information on the mathematical elements in the Spice quasi-saturation model, see: G. M. Kull, et al., "A Unified Circuit Model for Bipolar Transistors Including Quasi-Saturation Effects," *IEEE Transaction on Electron Devices*, Vol. 32, No. 6, June 1985:1103-13.

[12] The Spice transistor model is well described both mathematically and schematically in "Star HSpice manual", Release 1998.2, *Avant! Corporation*, 1998.

4 Design Solutions

4.1 Tools for Circuit Design

In order to simulate circuit ideas, MicroSim PSpice version 8 and Orcad PSpice version 9 were used to run hand-written Spice scripts. Engineering Capture System by the CAD/CAM group was used to enter schematics and produce PSpice netlists. All layout I completed on the PC was done using LEdit.

4.2 Considerations and Limitations

Many factors constrained the design of a new output stage topology. The silicon used a predefined device layout, so the number of devices and the space to route metal between devices was limited. Additionally, very wide metal busses would be required for both the power supplies and the output channels in order to limit metal electro migration as well as trace resistance, further constraining the layout possibilities. To minimize power and output coupling, important signals and devices could not be placed under these traces, also limiting device options.

Secondly, there were many performance constraints. The comparator needed the rail-to-rail output stage with similar drive and output levels, similar or better propagation delay with as symmetric as possible rise and fall times, and minimally increased supply currents. Next, the output circuit needed to work regardless of process variations (β, early voltage, resistor sheet rho), and across all operating temperatures as well as with all resistive and capacitive loads.

4.3 Failed Circuit Ideas

The motivation of any new circuitry in the output stage would be to speed up the turn on and turn off speed of the output devices by extracting more charge from the base of the NPN

devices in order to turn it off, and by adding charge to the base of the PNP to do likewise. The easy way to do this would be to place a current source at the base of each output common-emitter. In the case of the bottom NPN, the current source would extract charge and force the device off when the NPN follower at the base was off. Otherwise, the current source would pull current from the emitter follower also attached to the base. Obviously, the larger the current source, the faster the output device could be turned off. However, a larger constant current source would both increase quiescent current when it was not needed, and would decrease the amount of base current into the output device thereby decreasing the current gain of the Darlington output configuration and the slew current. Both of these effects are highly undesirable, so all of my circuits attempted to minimize current draw when unnecessary while maximizing it during turn off periods.

4.3.1 Mirrored Output Currents

One of my earliest efforts to solve the problem was to mirror the slew current in the output device over to the opposing device's base. The simplified circuit can be seen in Figure 4-1. A Schottky diode was added to prevent saturation of the new transistor used for base charge extraction. By using the mirrored slew currents, I would only

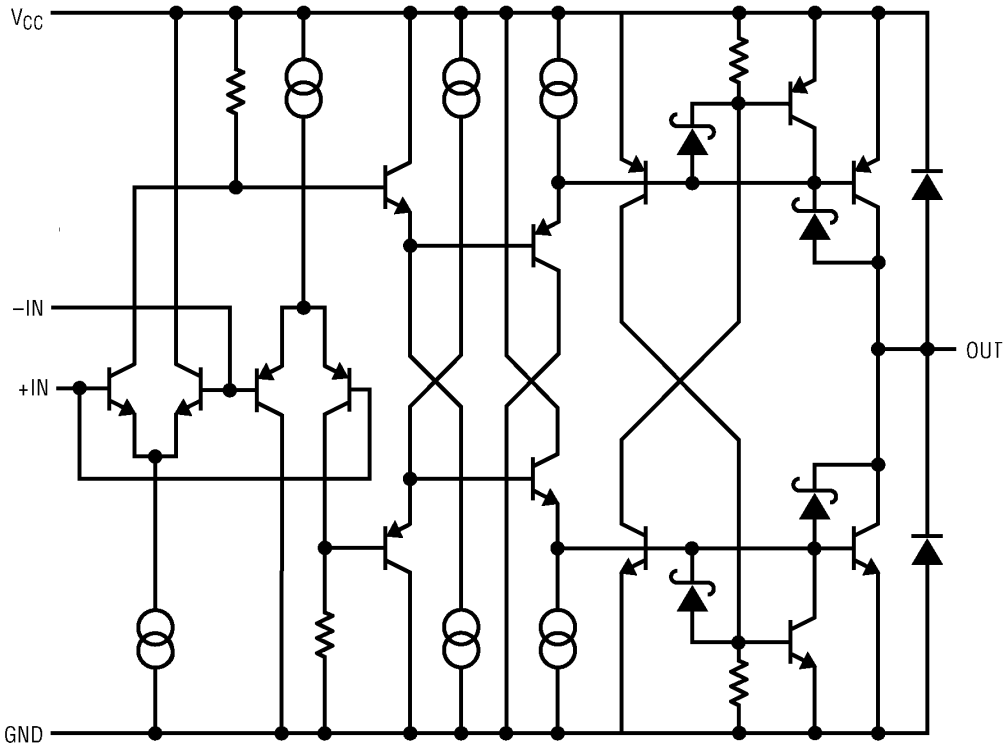

Figure 4-1: Mirrored Slew Currents to Extract Base Charge

consume extra current when toggling leaving quiescent current unaffected, and when one output device was on and slewing, it would not lose any base current to my added circuitry.

There were too few transistors available in the physical area to use a transistor mirror below the base of the output device, so a resistor to transistor mirror was chosen. This was acceptable since I was activating the mirror only under peak currents. The voltage dropped across the resistance would be fairly well set by the maximum slew current and only variations in Is (and temperature) would change the necessary V_{BE} to turn the transistor on hard. By selecting a midpoint of mirror gain on a curve of process parameters and temperature, acceptable performance could always be achieved, and ideal performance could be selected for nominal parameters and temperature.

However, the circuit did not work as well as anticipated. The output device is very large, so to mirror an acceptable amount of current the device in parallel with the output device needed to be fairly large. This proved difficult due to my device and layout constraints. For the small amount of mirrored slew current I needed a fairly sizeable resistor, and I did not have one available in my predefined layout. The compromises I made were not ideal. Additionally, slew currents vary drastically with process parameters, temperature and load, so selecting an ideal mirror gain between the resistor and the base current extraction transistor was difficult.

The greatest problem with the circuit was its long reaction time. As the opposite device began to slew, the device intended to be off would still be "on" and would see a large shoot-through current before the base charge extraction transistor turned on. Therefore, often the circuit did not do enough to curb the shoot-through before a toggle cycle was complete. At some process corners and temperatures, the base extraction transistors would quasi-saturate and stay on. This was acceptable in that it continued to extract near constant current from the base of the output device, but of course this would limit slew currents and therefore the amount that the opposing device was on. This meant that the first several cycles of toggling might look different than later cycles since the base current extraction device would be turning on during that time. After that, although the output device would be able to switch

on and off more quickly, the slew currents of the device would be substantially lowered.

4.3.2 Feed Forward From "Earlier" Nodes

The results of the base charge extraction approach above were not a complete failure. The idea worked to turn off the output device that was meant to be off, reducing shoot-through current, and allowing higher speed toggling. The next approach, shown in Figure 4-2, used the same resistor to transistor mirror at the base of the output common emitter, but the source of the current came from elsewhere in the output stage. The main flaw from the previous circuit was the delay between when an output device "needed" to be turned off, and when the mirrored slew currents actually turned it off. The maximum slew current doesn't occur until some time after the output begins to slew, and then there is a short delay before the extraction transistor turns on, and another delay before enough charge is extracted from the base of the output device and it actually powers down.

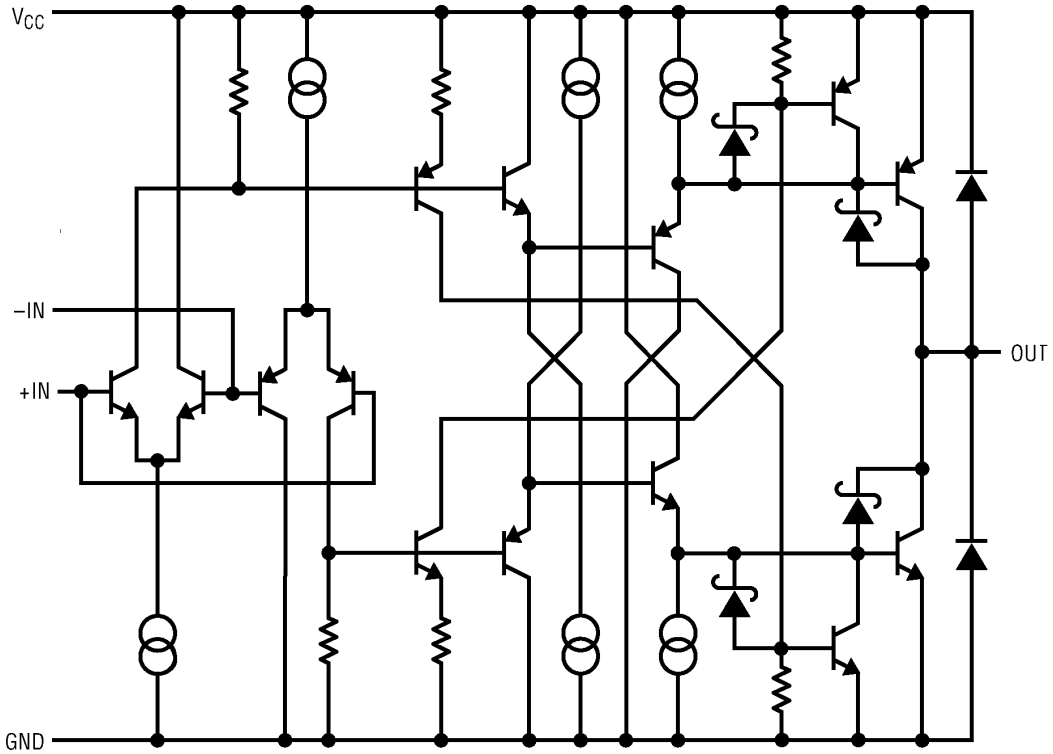

Figure 4-2: Feed Forward From "Earlier" Nodes

To avoid some of the delay, I chose a high-speed node earlier in the propagation of the output signal. Since the earlier nodes were still moving by about a V_{BE}, the earlier nodes are capable of switching a current to the output turn-off mirror from my previous circuit. The slowest nodes in the output stage by far are the bases of the output device due to the very large capacitance of the large devices as well as the their Millered C_π's. By using a signal 'earlier' in the output stage, I would be able to turn on my extraction mirror half a nanosecond to one nanosecond earlier than before.

This circuit also was not successful. The delay was not well matched between the top and bottom current sources and the outputs. Additionally, the delay varied with output load, process parameters and temperature. Of course, these variations would be of minimal concern so long as the base extraction device did not turn on until it was intended to. The real problem was that in some cases the extraction device at the base of the output device would turn on hard before that output device had finished slewing. When this happened, the output device would be turned off quickly (it had not yet reached quasi-saturation), and the output would not finish slewing before changing direction. Since there was a skew between the top and bottom output signals, this would sometimes only occur on just the PNP output device and not the bottom NPN. In this case, the output would toggle at low amplitude, starting at the bottom output supply rail.

4.4 Successful Circuits

4.4.1 The Turn-Off Current Boost

With the failure of the two previous attempts at providing an instantaneous current to the output devices in order to extract charge from their bases and turn them off, a compromise was invented. The primary reasons for not merely adding current sources to the bases, as discussed earlier, were a decrease in maximum output current due to limitation of β current gain, and increased quiescent current.

A compromise circuit that adds more constant current to the device switched off is shown in Figure 4-3. At DC when there is no switching, one output device – for example the bottom NPN – will be either on or diode clamped and therefore forced off. In the opposite (top) differential pair, the current will be switched from the resistor (that feeds the series of transistors to the output device) to the newly added

diode connected transistor. That transistor mirrors the current to the base of the output transistor. By doing so, the current when the transistor is intended to be off is increased, and by using various degeneration and device sizing, the current can be tweaked to optimize the compromises mentioned earlier. While this circuitry is not enough to meet my goals for increased turn off current and therefore turn off speed by itself, it was a compromise that could be dialed in to provide additional base charge extraction at a minimal cost to quiescent current and at a cost to slew currents.

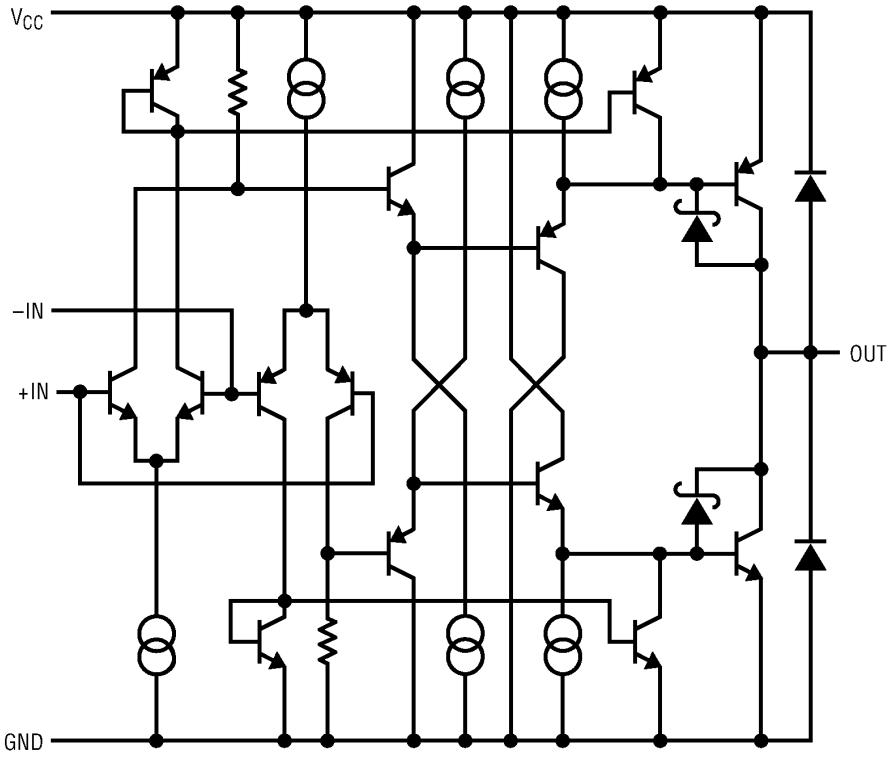

Figure 4-3: Turn-off Current Boost

4.4.2 Boosting on Indirect Slew Current Sensing

The real accomplishment of the new output stage modifications all lay in the following circuitry. After the failure of switching the base charge extraction current from earlier nodes as discussed in section 4.3.2, no new voltage nodes were found to regulate the additional base extraction current. However, I soon noticed that the collector current of the emitter follower before the output device, ordinarily tied to the rail, would be the output slew current divided by β plus the current from the extraction current sources. I could easily use this current to switch on additional base extraction devices at the right time. Figure 4-4 illustrates the circuit.

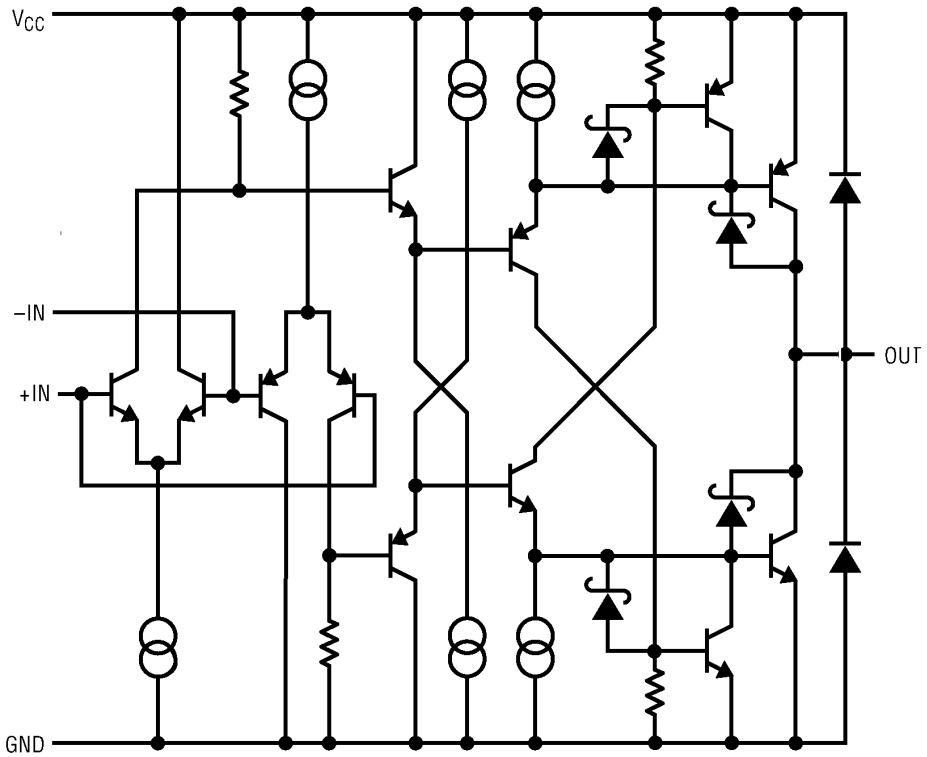

Figure 4-4: Extraction Current Created From Indirect Slew Sensing

The best attributes of this circuit were that it did not compromise any other performance specifications. It did not add any additional devices in the signal path, and did not add any significant slow down to the emitter follower since the Cπ capacitance would see only a small modulated voltage across the mirror resistor.

4.4.3 A Dreaded Latch Up

The schematic in Figure 4-5 shows an SCR-like latch formed by the previous circuit, caused if the current gain is greater than one through the cascode with the extraction transistor on. Since the gain of the current mirror with the resistor and the transistor is exponential, for a large enough instantaneous slew current, the IR drop of the mirror resistor would turn on the extraction transistor very hard, which would in turn pull more current from the opposing mirror resistor. Decreasing the gain of the resistor to transistor mirror would be a potential solution, but unfortunately, reducing the gain adversely affects performance, and exponential gain is difficult to keep below one. Although the chances of slew currents reaching the levels required for the latch up were minimal in a revision of the circuit with minimized gain, the chance

of having even a single part latch up in an application was too great of a risk, and the circuit idea was dropped.

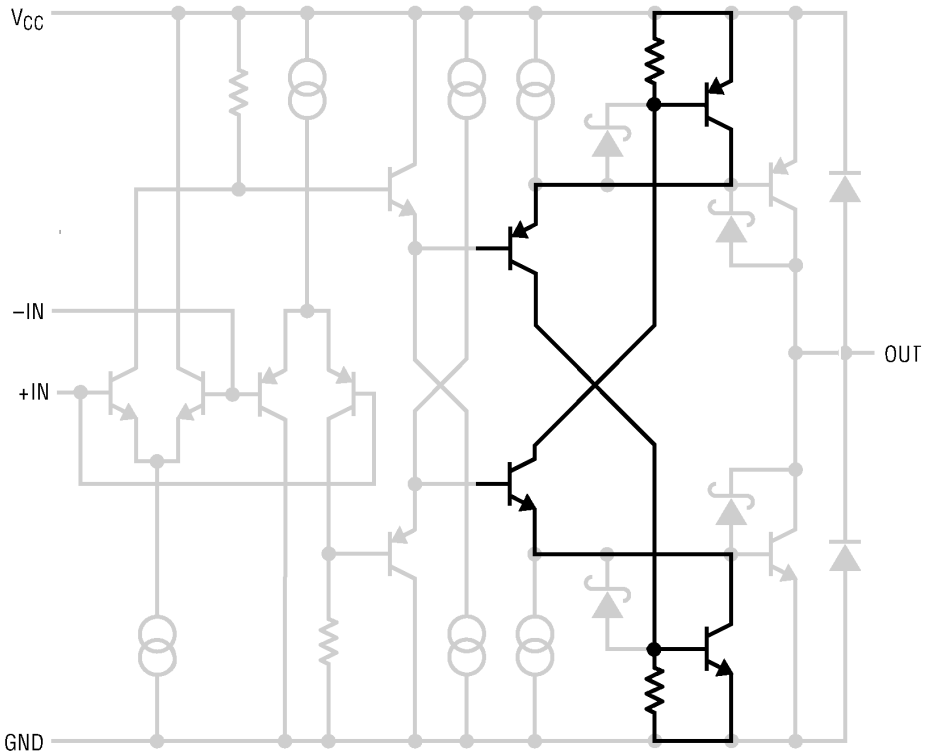

Figure 4-5:The Potential Latch-Up Circuit

4.4.4 Final Design

The schematic of Figure 4-6 shows the simplified final design for the LT1715. The current boosting when off is added to both the top and bottom devices, with the bottom NPN receiving substantially more quiescent current due to the asymmetry described below. The indirect slew current sensing was eliminated from the lower NPN output device since the PNP device would benefit more from a well-optimized instantaneous boost. The PNP's potentially lower β makes the constant current boost less appealing due to the current gain losses described previously. The NPN is only marginally affected. Without the original symmetry, there is obviously no chance of latch up. Additionally, The asymmetry leads to different bias currents depending on the state of the output, but still results in symmetric rise and fall waveforms with similar drive abilities in both directions.

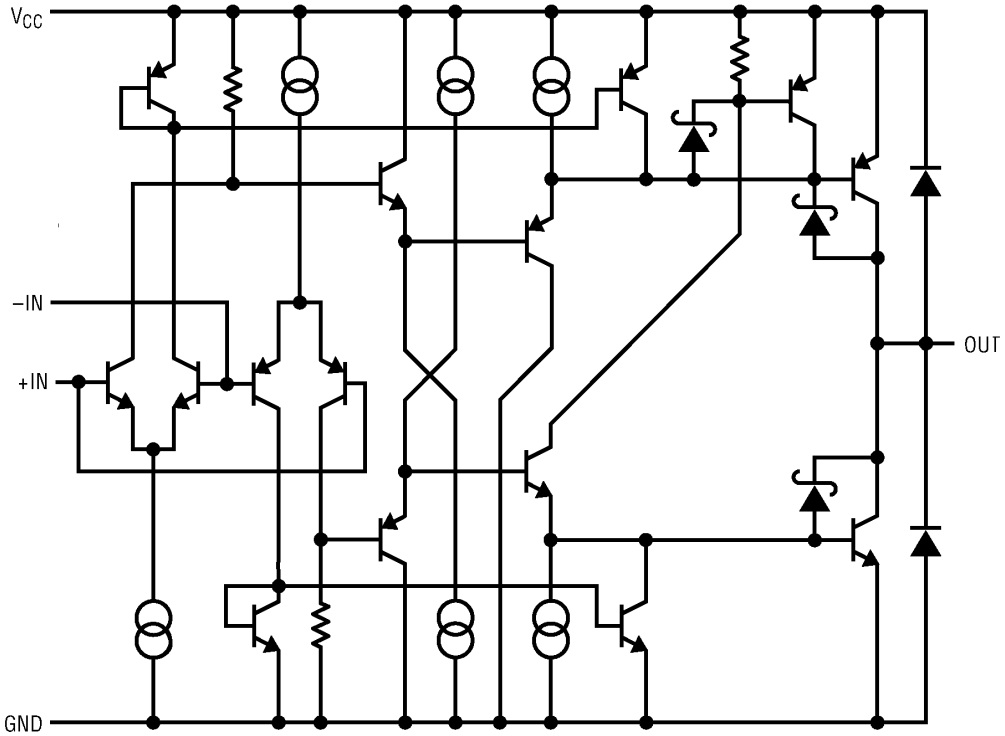

Figure 4-6: The Final LT1715 Output Stage Topology

5 Other Circuit Design Issues

5.1 Increasing Slew Rate

In order to raise the toggle rate of the new comparator above 150MHz, the drive ability of the PNP output would need to be improved. The output 10-90% rise times for the LT1719 and LT1720 were stated as 2.5ns and the fall time 2.2ns. However, outside these typical numbers, the rise time could easily increase above 3ns – especially when cold. This would prevent valid toggling at frequencies above 100 MHz even without quasi-saturation problems because the output would have insufficient time to pull high before the NPN driver would be turned on slewing the output downward.

The PNP device was almost at its current saturation limit even when the it wasn't β-squared limited, so I first increased the size of the device by 50%, increasing the maximum current it could carry by around 50%. The current boost at turn off, described in the previous chapter, was minimized to prevent excessive loss of base current while slewing, and the indirect slew current sensing was maximized. At room temperature with nominal device parameters, my PNP and NPN output devices were precisely balanced in current carrying capability. Additionally, a relatively substantial excess current was available in case β's were low or temperatures were cold, so even at the far corners of performance, the PNP device could keep up at the output.

The expected result of this modification, beyond allowing an increase in toggle rate, would be to make the rise and fall times of the comparator symmetric, leading to a symmetric propagation delay under nominal conditions. However, with the additional drive currents added to both the NPN and PNP output device Darlingtons, the fall time in the final product was still faster than the rise time.

5.2 Increasing Input Offset Accuracy

A good measure of comparator quality is the input offset voltage. Variations in processing can cause device sizes to be mismatched, and resistances to be slightly varied on the same die. Even with the extremely tight layout of the input stage of the LT1715, input offset voltages could be in the millivolt range, reducing the ability of the part to resolve small signals, or at least reducing the user's ability to use the part for small signals.

To combat this inherent defect, binary weighted trims were added to the input stage to adjust the current sources feeding the input devices. Since the circuitry and the layout of the input stage was nearly untouched from the previous LT1719 and LT1720, data from their offset variations was used to determine step sizes for optimal trimming. The trims themselves change the currents enough to vary the V_{BE}'s of the input devices by some number of millivolts thereby decreasing the offset between the two inputs. In production, this trimming could be done automatically at wafer sort by machine, ensuring incredibly accurate offset voltages, in this case less than half a millivolt plus or minus. In Chapter 9, I discuss how a defect in the hysteresis circuitry as designed undermined the very high accuracy of the input trims, forcing the datasheet limit to a guaranteed maximum of 3.5mV when it could have been drastically lower. Luckily, the primary appeal of the LT1715 should be its high speed and not necessarily it's high accuracy – engineers looking to resolve very tiny signals might want to use a comparator without internal hysteresis.

6 Testing Methods and Procedures

6.1 Testing First Silicon

While the LT1715 wafers were in the fab, a test fixture was constructed to evaluate the DC performance of the initial wafers. The bonding and probing pads on the individual LT1715 die could be connected to a test board by using a probe insert with fine, aligned tungsten leads. The board and insert are clamped into a test station and a microscope is used to align the probe insert before lowering it to contact the wafer.

The test fixture features extensive bypassing of supplies both on the probe card insert as well as on the board itself. However, AC characteristics were believed completely immeasurable due to both the distance of the bypassing from the die and the series inductance of the probe leads on the supplies. Both factors together would likely cause large V=L*di/dt fluctuations in supply voltages as the output devices switched on and off while slewing the output. The bypassing was adapted from what was required on the LT1720 to prevent any oscillation, including 1000pF shunting the inputs together. This capacitor makes high frequency noise common mode on each comparator input to eliminate spurious toggling, while having almost no effect on the slow and DC inputs used to test DC characteristics.

Several jumpers on the board for each channel allow for easily changing setup between different tests. A 49.9 **Ω** resistor from the board input to ground allows for termination of a signal generator as well as reduction of input signals through a voltage divider with a series 49.9 k**Ω** to measure hysteresis and offset. Each input can also be jumpered on the board to ground to be used as a common mode, or it can be shunted directly to the input. The option to shunt a .01uF capacitor to ground was also included to filter out medium frequency noise on the input, potentially useful if the input was otherwise connected to a DC voltage.

An HP33120A Arbitrary Signal Generator and an HP5461 150MHz oscilloscope with standard 10M, 10pF, 10X probes were used to work on the initial silicon. A variable triple power supply was used to provide variable supplies for the input supplies and the output supply.

Trip points were measured by applying a 1KHz triangular wave to the inputs through a 49.9kΩ resistor with the jumper set to terminate into 49.9 Ω. This created a 1000:1 divider with the signal generator, which was set to output 10Vpp, which into this non-50□ Ω termination, turned out to be 20Vpp. The inverting input was tied directly to ground and the input supplies were split +/-5V with a 5V output supply. As the input ramp crossed each trip point, the output would switch. Since the ramp is very slow, the comparator delay is negligible. A 10pF probe on the output line then relayed the signal to the scope to show the output signal in reference to the input ramp. By looking at the time after which the comparator trips, and by assuming a perfectly linear voltage slope, the trip point voltage can be determined. In this case, the input went up 20V in 500us and back down again in another 500us. Therefore, every small ($1/5^{th}$) notch in a division with the scope set to 100us per division was .8mV of input voltage. I was then able to vary the common mode as well as the power supplies to find CMRR and PSRR by varying the three supplies to obtain the desired supply conditions to watch the effects on the inputs. DC results are shown in Chapter 8.

While I had the bench set up running, I decided to crudely get an idea of the toggle rate. For this setup, I terminated an HP 8446A sinusoidal signal generator into the 49.9 Ω on the test board while shorting the opposing input to ground, maintaining the +/-5V input supplies with a 5V output supply. While the 1nF capacitor at the inputs could have an effect at very high frequencies, I did not anticipate interference at my testing frequencies since my series resistance was low, pushing $1/(2\pi RC)$ toward 1GHz assuming some tenths of an ohm series resistance. Results of the initial toggle rate tests are reviewed in Chapter 7.

6.2 Wafer Sort of Initial Silicon

Soon after basic functionality was verified on the probe station, wafers were passed through an automated test program to sort the individual die based on functionality and performance. Furthermore, the parts were automatically trimmed to minimize their input offset. Data for each wafer was combined into a database so that individual parameters could be plotted by location on the wafer, plotted by value, or plotted in reference to each other. All characteristics tested at this level were DC, measured with automated but roughly similar

methods to the silicon on the probe station explained above. Graphs of compiled data were produced to show the statistical distribution of performance among a huge sample of test die. This data shows real world performance of the comparator with variation in processing.

6.3 AC Packaged Part Testing

A copper board was designed, routed and stuffed to evaluate the AC performance of the LT1715 packaged parts. A simplified schematic with the details of one drive channel is shown in Figure 6-1, and the copper routing which contains circuitry for inputs to all four channels (comparators A and B, both + and -) is shown in Figure 6-2 and Figure 6-3. Do note that in the routing photos that all the blank white area is continuous ground plane.

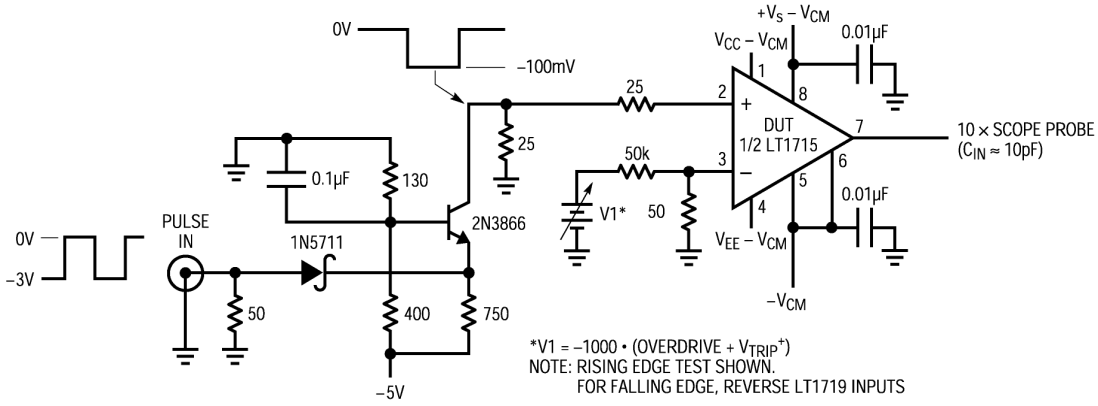

Figure 6-1: Single Channel of AC Test Fixture

On the test board, each supply is bypassed with a .1uF and a 1000pF ceramic surface mount capacitor. For toggle rate testing, the NPN transistors in the input section are removed. This leaves the inverting input tied to ground through 50 Ω with an overdrive or offset pin connected through 4.99 kΩ to a banana jack for changing the common mode. The non-inverting input with the transistor removed sees a 50 Ω input impedance and a 2:1 divider with a 50 Ω termination from the BNC monitor connection where the collector of the transistor was, previously. At the output of the comparator, mounting pads are in place for various terminations as well as R and C loads.

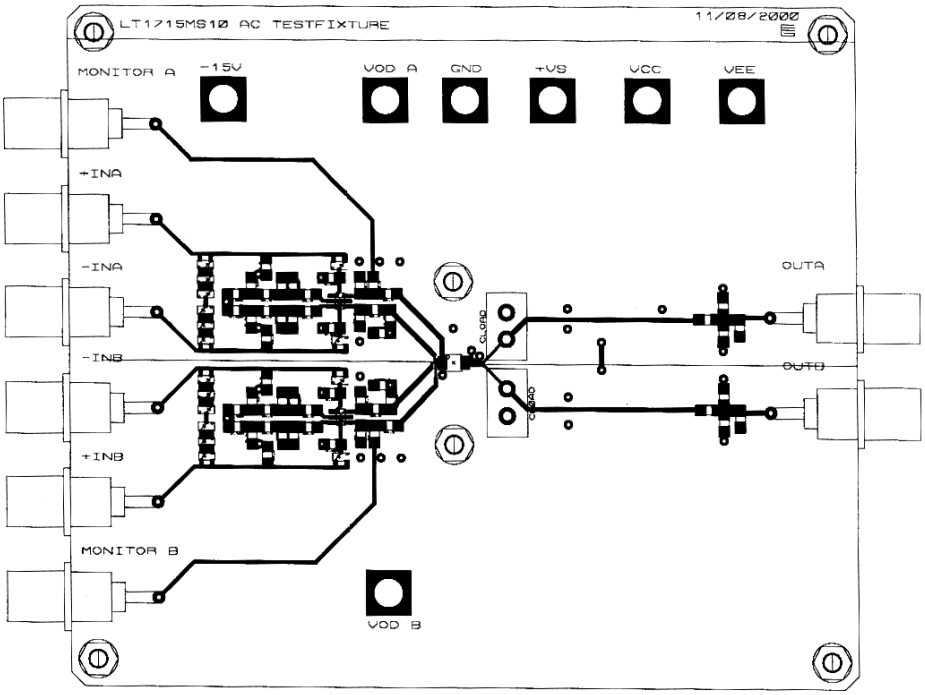

Figure 6-2: AC Test Fixture Front Side Copper Routing

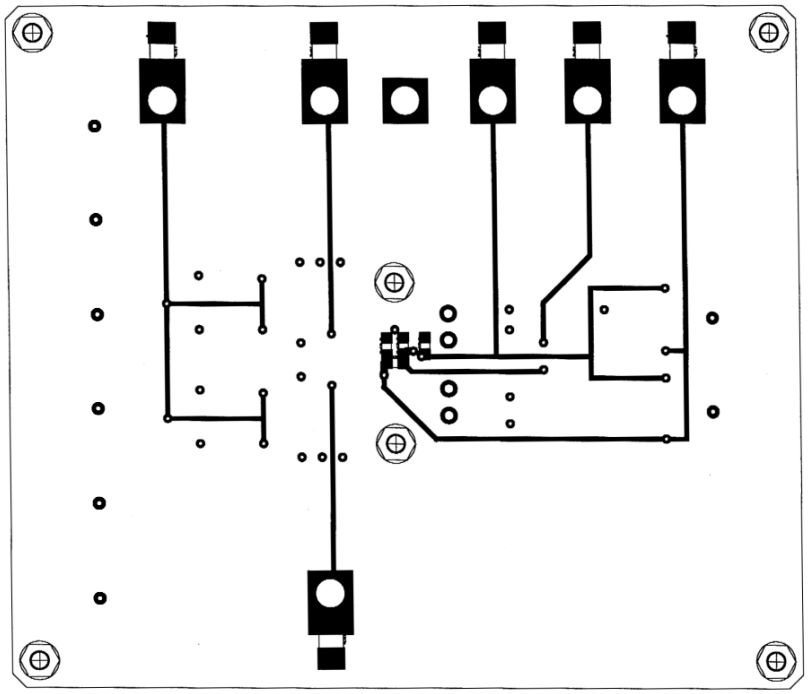

Figure 6-3: AC Test Fixture Back Side Copper Routing

The transistor and other components at the input are used for creating accurate and fast voltage steps at the comparator inputs. The circuitry is duplicated on each of the four input pins and is used to square up the input pulses for propagation delay at the input of the device under test (DUT). Jumpers and resistors are soldered to connect the comparator for a desired test, and since the inverting and non-inverting channels have essentially identical performance, the board was generally left the same. The base of the transistor sees a thoroughly bypassed –1.2V bias voltage. The BNC input onto the board is diode clamped to prevent destroying the transistor, and the diode into the emitter of the transistor prevents overdriving the common base. When the input to the circuit falls below –.6V minus two diode drops, the transistor is pulled out of cut-off and the diode turns off. The 3V at the emitter divided by the 750 **Ω** resistor creates 4mA of collector current. Dropped across 25 **Ω**, that current creates a fast, –100mV step at the input of the DUT. As long as the input passes quickly from –1.9V to -.6V, regardless of the shape, speed or amplitude of the input signal otherwise, the DUT input will see a 100mV square step. To test different overdrives, offset is summed to the inverting pin through a 1000:1 divider with the 50 k**Ω** resistor.

6.3.1 Accurate Measurements

In order to get accurate measurements of toggle rate, I would need to get an accurate picture from my measuring equipment. I expected to see relatively fast edges on the oscilloscope since PSpice predicted slew rates of 2V to 3V per nanosecond, and the initial tests had shown actual silicon to be even faster. Therefore, a 3V step, the minimum size for a valid toggle, would occur in as little as a nanosecond.[13] I would need a very fast scope to pick up these transitions. For a first order system, as a scope with almost no overshoot would be, the rise time to a step input would be $.35/f_h$, where f_h is generally the bandwidth displayed on the scope. I therefore would need .35/f to be equal to 1ns, making f_h= 350MHz. To eliminate any loss of rise time, I wanted to use an even faster scope, and a 500MHz unit was available. It would have a step response rise time of approximately 700ps. The initial toggle rate measurements were made using an Agilent 54610B 500MHz sampling scope and the Tektronix P6201 FET probe.[14]

However, after a couple hours of measurement, I was quite bothered by enormous overshoot on the rising and falling edges of my output waveforms. With a 5V output

[13] Toggle rate requirements are specified in Chapter 8 as they were selected for the LT1715 datasheet.

single supply, I was seeing 1V of overshoot in each direction. This seemed highly unlikely given the low inductance of the LT1715's MS10 package and bonding, and at .6V above or below the output supplies my ESD diodes should have clamped the voltage. The LT1715 was performing very close to what PSpice predicted with the sole exception of this overshoot response, leading me to believe that either the long output traces or the probe was causing anomalous readings.

Jim Williams offered a fast test to see if my scope or probes were to blame – an HP213B 100ps 10-90% time Tunnel Diode Step Generator. With an input that fast, my measurement on the scope would be only the step response characteristic of the equipment taking the measurements. On several other scopes, I observed the output of the generator to have no discernable overshoot. However, when the step generator was connected from it's BNC output directly into my Agilent scope internally terminating with 50 **Ω**, there was almost 10% overshoot (over 20mV overshoot with a 240mV step)! Additionally, the active FET probe, coupled through a BNC adapter and with a 50 **Ω** termination at the probe, increased overshoot to 60mV for the 240mV step – a total of 25% overshoot. A photo of the two overshoots is shown in Figure 6-4. This overshoot could easily lead to 1V of overshoot in a compromise between the slower slewing across 5V and the faster response in around a nanosecond. This error was likely to invalidate my previous measurements since the V_{OH} and V_{OL} levels I had been measuring were drastically exaggerated by the

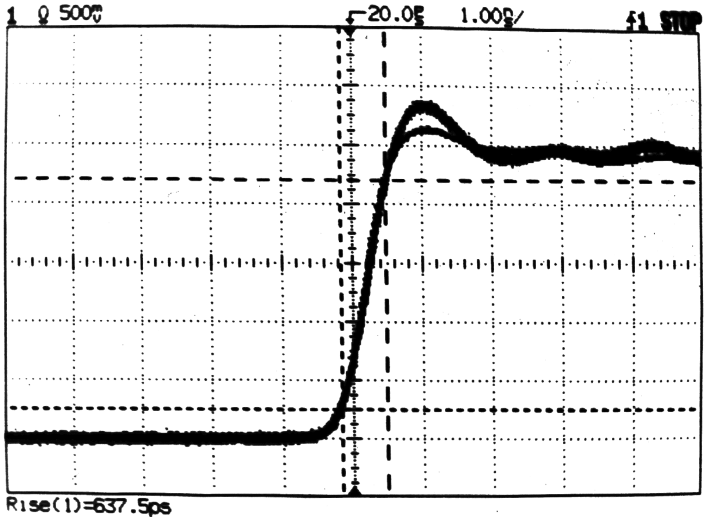

Figure 6-4: Overshoot Response to 100ps Step With and Without FET Probe

[14] Specifications for the Tektronix P6201 active FET probe can be found at http://www.tektronix.com .

overshoot. Most likely the amplitude had rolled off considerably at the time that the exaggerated waveform was just reaching the selected limits for valid toggling, so my numbers would be substantially optimistic.

I upgraded to a Tektronix TDS 640A 500MHz, 2GS/s scope and used only the Tektronix P6156 500 Ω probes or the BNC's with series 450's into the scope termination for the remainder of all testing.[15] The 100ps step generator showed the new scope to have only 3% overshoot, though a slower rise time was shown. The 500 Ω probe had negligible effect on both overshoot as well as rise time. A scope photo of the results is shown in Figure 6-5.

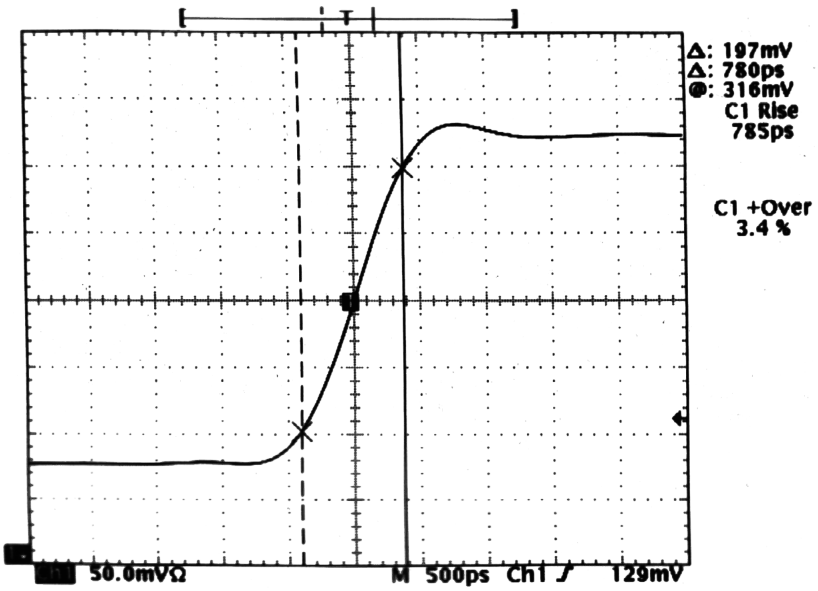

Figure 6-5: Overshoot Response of TDS640 Scope and 500-ohm Probe to 100ps Step

6.3.2 Ground Plane and Oscillation

The initial design of the AC test fixture seemed effective, but a flaw was soon uncovered. When one channel was being driven to test toggle rate, the other comparator would be left with both of its inputs tied to ground through 50 Ω. Quite often, the unused comparator would also toggle along with the driven comparator. First efforts attempted to better bypass the supplies by placing 500pF at the pins to reduce any high-speed spikes caused by the fast slew currents in the output stage. There was almost no effect from this effort.

[15] Specifications for the Tektronix P6156 500 ohm probe can be found at http://www.tektronix.com , although similar probes in the P615x line may soon be replacing it.

I decided to inspect the inputs to investigate where noise was coupling into the comparator. The inverting and non-inverting traces are adjacent and bordered on opposing sides with separate ground planes. Each input is terminated to ground through some 25 **Ω** or 50 **Ω**, but each connection to ground occurs on the opposing ground plane. I took a capacitor of .1uF to connect any combination of the two ground planes and the inputs while observing the effect on cross talk. Although most combinations of connections made noticeable but minimal changes to the spurious toggling on the undriven comparator, connecting the inputs to the ground under the DUT had a drastic effect. When placing the capacitor between the non-inverting, outside trace and ground, the cross talk became much worse – matching the output of the driven comparator almost exactly, regardless of frequency. When connecting the capacitor between the inside, non-inverting input and the ground plane under the DUT, the spurious toggling almost completely disappeared. My theory was that the ground plane under the LT1715 part was being contaminated with high frequency noise from the ground pin, while the other ground plane was free of noise. By coupling one input to one ground plane, and the other input to a different plane, I was creating some differential noise, in this case enough to toggle the comparator.

I connected the termination resistors together on the inside ground away from the DUT and also placed more vias between the frontside and backside ground planes hoping to lessen the noise as well as to eliminate any ground noise being differential. With these changes, the communication between the driven and the unused comparators was drastically lessened. However, the problem was not eliminated and while the trick with the capacitor still had some effect, occasionally when the undriven comparator oscillated, I could not stop it.

I believed that my problem was still ground noise, only now I envisioned it capacitvely coupling from the output supplies into the outermost, non-inverting input. To lessen the coupling, I cut the ground plane copper between the inputs and the output supply and outputs, basically straight down the underside of the chip. By separating the input and output ground planes on both the front and rear of the board, I imagined that the ground noise would be isolated to the output side of the board, while the minimal currents on the input side would create no problems. The ground plane on the input side would still be low enough impedance that minimal noise would couple to it, and the capacitance between the output ground and the inputs

would at least be halved with two trace-to-trace capacitors in series instead of one. With this modification, the oscillations on the unused comparator channel were completely eliminated. Now I was ready to begin solving other problems.

6.3.3 Output Trace Capacitance

As seen in Figure 6-2, the output traces on the AC test jig are fairly long. In hindsight, this was not efficient or necessary. Ideally, they would have been 50 **Ω** strip lines leading to a 50 **Ω** BNC. However, I didn't have time to make a new fixture, and my best option was to measure the output capacitance that I had created between the output trace and BNC connector to ground while correlating any effects from wave reflections in the uncompensated trace impedance.

To determine the trace capacitance, I planned to study supply current while toggling with changing output capacitance, expecting the current to increase as more load needed to be slewed. I measured the output supply current with one comparator toggling at 50MHz into the trace and BNC, and the other comparator's output forced low. I then placed a FET probe with a nominal 1.5pF capacitance on the pin to examine the waveform as well as to provide another data point. After these initial data points, I lifted the output pin from the trace on the test fixture, and began soldering capacitors from the output directly to ground. The following table expresses the results of various output capacitors, all of which were 1206 surface mount ceramics with relatively tight tolerances. My table of values is found in Figure 6-6.

Measure	**Capacitance (pf)**	**Supply Current (mA)**
Pin floating	~0	11.08
5.6pF 1206	5.6	12.70
5.6pF 1206 and probe	7.1	13.23
10pF 1206	10	13.95
10pF 1206 and probe	11.5	14.65
Pin on trace	x	13.85
Pin on trace and probe	x+1.5	14.70

Figure 6-6: Estimating Output Capacitance by Experiment

From this data, it appears that 10pF rather accurately estimates the capacitance of the trace and BNC to ground. The waveforms looked different on the scope but maintained very similar rise times for the similar conditions. The difference in the waveforms could be characterized by noisy ("bumpy") high and low times between transitions with the trace and BNC touched down, instead of smooth and high and low times as observed with the probe at the raised but loaded output pin. In addition to these bumpy peaks and troughs, the overshoot was larger when the output was on the trace. All of these effects are likely from wave reflections in the output trace. The impedance is unspecified and the end of the long copper trace unterminated (since the 1715 cannot drive low impedances). These results are as expected since the long line was not carefully designed.

While the line was now characterized for driving capacitive loads, I still needed more data on the parasitics of the output trace. High-speed probes could not be used in the oven due to the high temperatures. Therefore, BNC cabling would be run from the oven. In order to use 50 **Ω** BNC cable, I would use a 450 **Ω** series resistance on the board before the BNC, and 50 **Ω** termination at the scope. This minimizes the load current from the DUT while also minimizing reflections in the cabling. However, to do this I would place the series resistance on the board before the BNC connector, thereby changing the capacitance of the trace. For this test, I first used a 500 **Ω** probe to examine the output with the nominal 10pF trace. I separately saved the output on a digital sampling oscilloscope for both channels independently toggling at 125MHz and 155MHz and then soldered in the series 453 **Ω** resistor and reconnected the outputs to the 50 **Ω**-terminated scope with BNC's. At both frequencies, I attempted to match the saved waveforms rising edge from the initial probing by adding various load capacitance at the output of the DUT. I tried 3, 5 and 8pF on the output, and found that 5pF looked most similar at the higher frequency, but 3pf showed more accurate rise and fall times at both frequencies while looking most similar on the 125MHz signal. I arbitrarily decided that 3pF would be added to the board for all tests using a 500 **Ω** load with BNC's.

6.4 DC Packaged Part Testing

Testing bias, input current and output levels with differing output currents were easy to perform on the AC test board, and would similarly be verified in the testing department using digital current and voltage meters. The DC input specifications can be made by

measuring the input trip points to derive offset and hysteresis and were made again on the AC board in the same manner explained above for the initial silicon tests. A large triangular wave was divided down to the inputs and observed on the scope with the output signal. By scaling the input voltage slope to time, the input trip points could be read off the scope by where they crossed the x-axis. A scope photo of these waveforms (with the results of one channel's response saved and restored with the other channel now displayed, simultaneously) is shown in Figure 6-7.

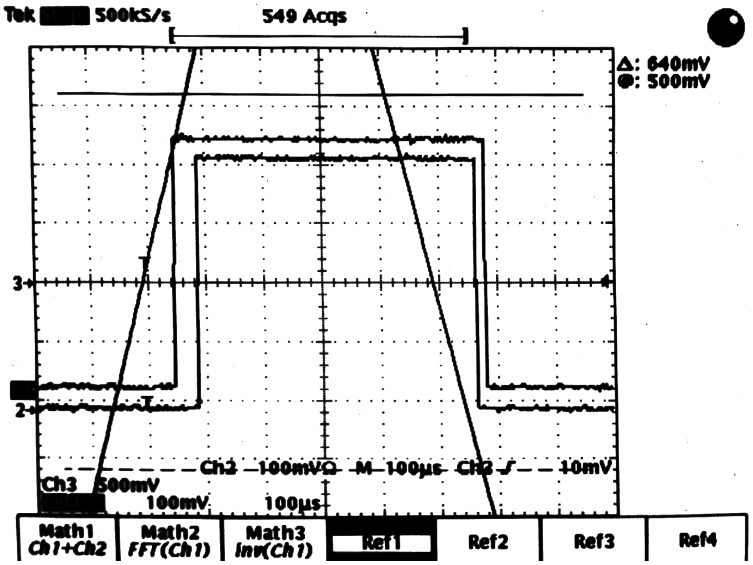

Figure 6-7: Two Hysteresis Plots in Response to Input Triangle

However, this method is only accurate to the resolution of what can be read on the scope. Additionally, noise at the inputs is not averaged and the output waveform has significant jitter. For these reasons, a more accurate measurement board was designed to accurately measure DC input performance. A simplified schematic of the board is shown in Figure 6-8, and the copper clad layout is shown in Figure 6-9 and Figure 6-10.

The board uses a 1KHz triangular wave divided down to an amplitude of just millivolts to trip the comparator under test. An LT1638 dual op-amp is used to create the triangular wave by charging an RC centered around the common mode voltage with a range of V_{cm} +/- 7.5V.[16] By using positive feedback from the output to the input, the op-amp rails low each time it reaches the trip point of $15/2+V_{cm}$ and rails high at $-15/2+V_{cm}$. Independent switched capacitors sample the amplitude of the triangular wave when the comparator changes output

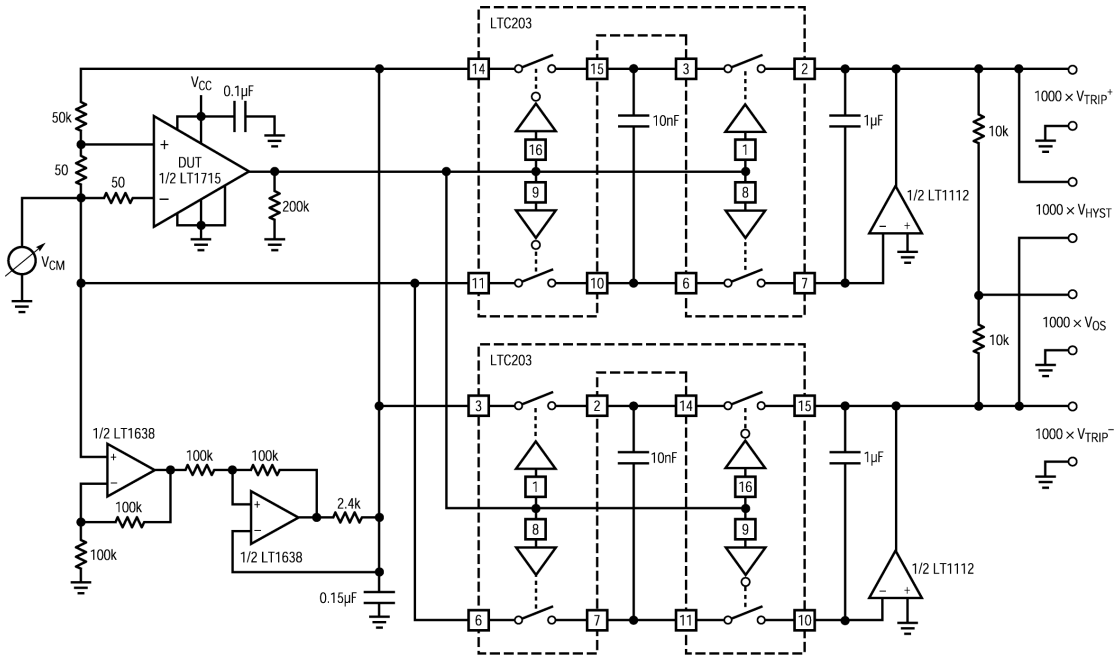

Figure 6-8: Simplified DC Input Test Fixture Schematic

state in either direction. When the LT1715 output changes polarity, the voltage sample is charge-pumped into a buffered output capacitor. The output is averaged over time since the sample capacitor is 1/100th the output capacitor and can therefore only make incremental changes to the amount of charge on the output stage. Because the input waveform is attenuated 1000:1 into the DUT and the output is a sample of the original triangle, each trip point output is expressed as 1000 times V_{TRIP}. The trip points were read with a Fluke 83 DVM. V_{OS} can be measured by resistively summing the trip points, and V_{HYST} is just the differential voltage between the two trip voltages.

[16] The datasheet for the LT1638 Over-the-Top op-amp can be found at http://www.linear.com/prod/datasheet.html?datasheet=444 .

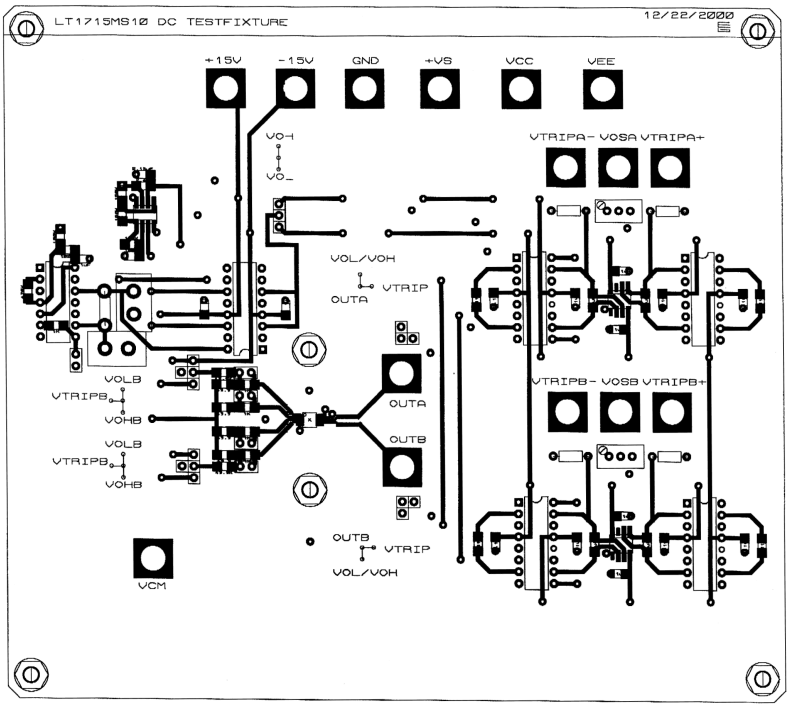

Figure 6-9: DC Test Fixture Front Side Copper Routing

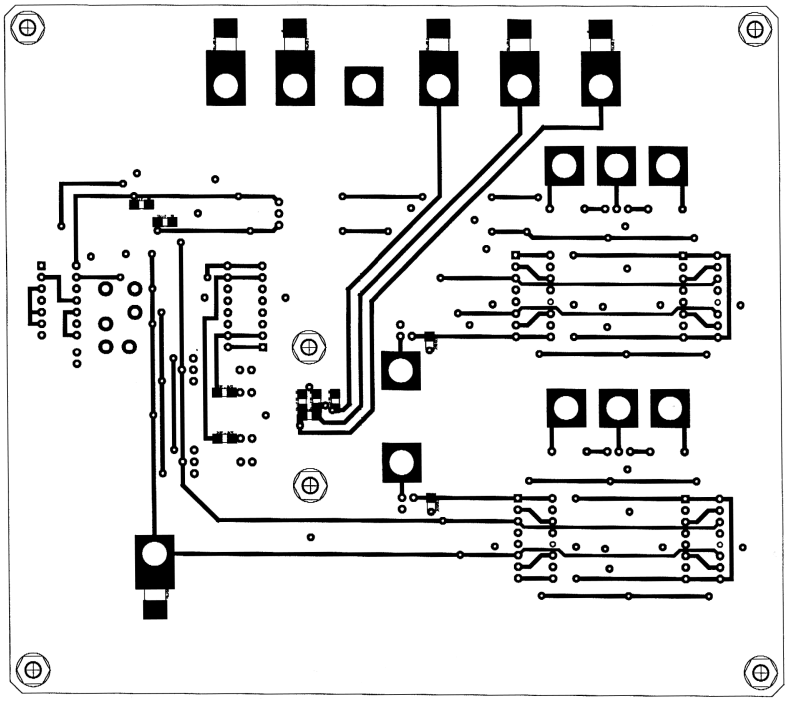

Figure 6-10: DC Test Fixture Back Side Copper Routing

7 Toggle Rate Testing

7.1 Toggle Rate Specifications

While toggle rate is not always a concern in comparator circuits, it was the driving goal for my entire design due to recent demand from customers hoping to interface with digital logic. Therefore, selecting what was required from the comparator's toggling was very important not only in terms of design, but in terms of user specifications.

For the design, I looked at the valid logic output levels for different TTL and CMOS families. The chart in Figure 7-1 shows the results and compares the desired output in the LT1715. All numbers are specified for a 5V single supply and sinking or sourcing around 4mA of output current.

	74LS series	**74HC series**	**74LVC1GU series**	**LT1715 DC spec**	**LT1715 toggle spec**
V_{OH}	< 2.7	4.5	> 4.3	***4.6***	***4***
V_{OL}	.5	.4	.45	***.4***	***1***
V_{IH}	2	3.5	3.75	N/A	N/A
V_{IL}	.8	1.5	1.25	N/A	N/A

Figure 7-1: Levels in Different Logic Families

TTL logic like the 74LS series is not capable of large output amplitude and is only capable of sourcing .4mA at its output. However, the low input level creates some serious restrictions. Luckily, it is a fairly dated logic family that is very rarely used and sometimes even difficult to obtain. CMOS creates the upper end restriction. While HC logic presents a

rather typical CMOS requirement, the LVC1GU low voltage, sub-micron, very fast logic family requires a stiff 1.25V minimum low, and 3.75V minimum high.

With these numbers in mind, a reasonable compromise for output current was defined. Due to the requirements of CMOS, and the essentially symmetric output of the LT1715, a symmetric 1V and V_S – 1V was chosen as the valid logic output levels of the part. When driving signals in the tens of megahertz, the DC output levels are fairly accurate. Additionally, the limit to toggle rate is usually V_{OH}, especially with any non-negligible DC output current (this is due to the higher β and I_S of the NPN which drives low opposing the PNP output device that drives high). Finally, the 5V supply was chosen since high toggle rates are more difficult to obtain on the larger supply due to slew current limits. Few, if any, users would use a 6V supply, and the 3V supply performance is generally superior to 5V performance. For this reason, most lab and production testing would occur with a 5V supply testing for V_{OL} and V_{OH} limits of 1V and 4V. 3V performance would initially be characterized by hand to ensure that it did, in fact, outperform the 5V numbers, after which performance would be guaranteed by design.

7.2 Results of Hand Testing

Both initial wafers and final packaged parts were tested in the lab. One lot of silicon was used on the probe station that the fab had indicated to have nominal tolerances on all design parameters, with the exception of NPN and PNP β's that were high. That lot was tested in packaged part form as well as an additional lot with entirely nominal parameters. The nominal lot was the focus of almost all of my packaged testing since it had the lowest β's and therefore was the most pessimistic. Occasionally results were verified on the samples with higher β, but results were generally improved and optimistic. If the results to the same experiment improved, I ignored the better results and chose the more pessimistic figures for datasheet and thesis information.

7.2.1 Testing Toggle on the Probe Station

Using the probe board described in Chapter 7, I looked at the toggle rate of the initial silicon. Using a sinusoid with an amplitude of several volts on one channel and ground on the other, I drove a single die and watched the output with a 10pF probe clipped on to the end of the probe card. In hindsight, I realize that I was likely driving 20+pF, but at this point I was looking for crude verification of functionality, and not a real measure of performance. What I observed was exciting. Although the

output would cease to toggle around 100MHz, it came back before the 110MHz limit of the sinusoidal generator. I tried switching the driven and grounded inputs and had identical results. With the other comparator, I observed perfect toggling all the way to the limit of the generator. Several other die were quickly tested at 110MHz with the same successful results.

There was one interesting fact, however. The output at these high frequencies was around 8V, from –1.5V to 6.5V, while the output supplies were only ground and 5V. However, there are plenty of good hypotheses for this dramatic display. First, there are very sizable inductances from the supply bypassing on the test board to the die itself due to the long traces to the probe tips in addition to the half inch probe tips themselves. Therefore, both the output supplies as well as the output itself may bounce when the currents increase or decrease dramatically. Since V=L * di/dt, the inductances are on the order of 3nH or so, and the di/dt can be as high as 30mA in .3ns when the circuit clamps itself, the voltage bounce could be .3V on both the supply and the output. Additionally, a huge amount of overshoot likely existed in the measurement due to wave transmission problems with the unterminated, non-strip lined 6 inches of output trace on the debug card. Needless to say, the part was definitely toggling at 110MHz, and there was no way that the output on the die was really going a volt and a half beyond the rails; the ESD Schottky diodes would have prevented that on both the supplies and the outputs. I was satisfied to see preliminary evidence that the part would be fairly functional at high speeds.

7.2.2 Testing Toggle of Packaged Parts

Several weeks after testing the initial silicon, I received packaged parts in MSOP-10. With the AC board ready and waiting for their arrival, I immediately began my testing. Using around 100mV of peak-to-peak input amplitude, the parts easily reached 180MHz. This was slightly better than what PSpice had predicted nominally, but when the high PNP and NPN β's of the initial silicon were modeled, PSpice roughly agreed that 170MHz to 190MHz was reasonable. Later on, I received packaged parts with all process parameters measuring nominally. These nominal parts toggled to only 165MHz or 175MHz. A quick sampling of a dozen parts confirmed that no parts would fail to reach toggle specifications to at least 160MHz.

Unfortunately, I did run into some difficulties while toggling. Those problems will be discussed as error sources and discrepancies in Chapter 9.

A substantial log of data was taken for toggle rate generalization in the LT1715 datasheet. Due to a badly compensated FET probe, a large amount of the data turned out very pessimistic. At higher frequencies, the probe rolled off the amplitude so much that I only saw maximum toggle rates of around 155MHz. I retook a series of data for toggle rate versus supply voltage and found 20MHz and greater of a discrepancy with the probe properly calibrated. However, since my intention with the graphs was solely to create a general picture of their performance with respect to a single changing parameter, the pessimistic downward scaling was advantageous as a safety margin. The data was converted to scatter plot graphs in Excel, and low-order quadratic fit lines were made to better shape the slightly spurious points. All of these graphs can be found on pages 4 and 5 of the LT1715 datasheet in Appendix A.

7.2.3 Toggle Rate Correlation

The measured data for toggle rate matches very well with PSpice data. Failure modes and general performance often look exactly like what I saw in PSpice. Real world toggle rate figures could be found by tuning parasitic capacitances and bond wire inductances in PSpice and comparing. For example, Figure 7-2 and Figure 7-3 show output toggling with .2nH of added supply bond wire inductance both in reality and in simulation. When leaving the parasitic capacitances relatively optimistic, the numbers match up perfectly, and vary with β as they did in actual silicon.

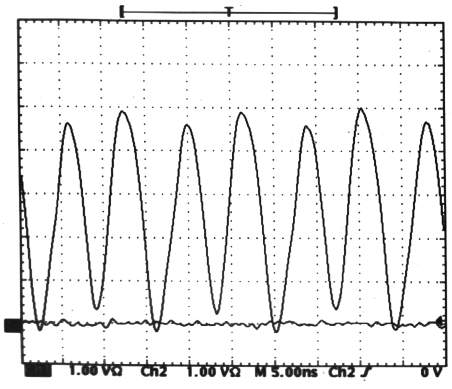

Figure 7-2: Actual Output Toggling

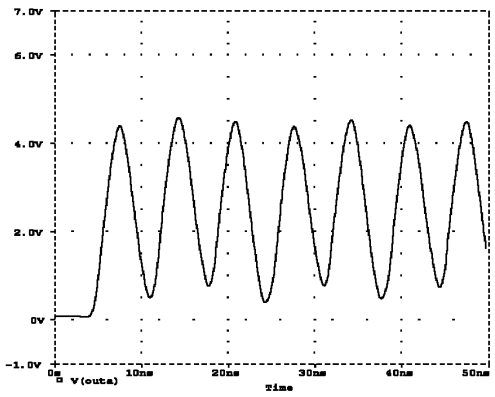

Figure 7-3: Simulated Output Toggling

7.3 Production Toggle Rate Testing

While my data indicated that all of the parts would perform fairly well, the part was conceived to have a guaranteed toggle rate. If the specification could be guaranteed to customers, it needed to be tested or guaranteed by design. The toggle rate capabilities of the part would be highly dependent on β and especially upon β mismatch, as well as on sheet resistances of the various resistor materials. Since toggle rate performance was dependent upon so many interdependent variables, it was decided that parts should be individually tested across temperature to determine whether they could be guaranteed for a given toggle frequency.

7.3.1 Measuring High Speed Pulses

While I knew and could identify what specifications were required for a valid toggle rate, I did not immediately know how to test for them. My first challenge would be just to catch the fast edges at the peaks and troughs of the toggling output. The applications engineers at Linear Technology believed that 1ns was about the smallest pulse width that could be detected, and even that would require 100mV or more of overdrive.[17] This would not be practical for my application, since at the breakdown of valid toggling the peak or trough could be well under 1ns in width, and I wanted to accurately detect amplitude to within less than 100mV of overdrive.

Luckily, Signal Processing Technologies makes very fast comparators. For example, the SPT9689 has less than .7ns of propagation delay with a meager 10mV of overdrive.[18] The SPT9689 also boasts a 900 MHz toggle rate, but luckily for the LT1715, the SPT part is a 20mA quiescent current, ECL-output comparator, and not a competitor.

A circuit to catch these narrow edges and regenerate a fast square wave to logic levels was devised as seen in the schematic in Figure 7-4. A window comparator is made from two comparators to test for signal crossing the high and low thresholds. When the toggle rate output reaches the high threshold on comparator A, the comparator pulses high until the toggle output falls. When the toggle output goes below comparator B's low threshold, the output of comparator B pulses high until the

[17] One example of a peak level detector is shown on page 34 of Linear Technology Application Note 72, available online at http://www.linear.com/pdf/an72f.pdf .
[18] The datasheet for the SPT9689 can be found online at http://www.spt.com/datasheets/products/9689.pdf .

toggle output goes high again. These two pulses, one occurring at the high threshold, and one occurring at the low threshold, are fed to an incredibly fast R-S latch, in this case, an MC100EL31 2GHz ECL input and output D flip-flop with asynchronous set and reset. The asynchronous set and reset allow the dual comparators to convert their two short pulses into a full square wave.

This threshold detector circuit was built and tested. The layout is shown in Figure 7-5 and Figure 7-6. Trace lengths for high-speed signals were kept to a minimum although they were not strip-lined to have a 50 **Ω** impedance. The digital logic would be fairly immune to transmission line defects, and the only possibility for degradation to the circuit's performance would be in accurate sensing of input amplitude. However, the transmission line impedance mismatching would be minimal on the board, and the large BNC would likely contribute more than the less than a centimeter of trace.

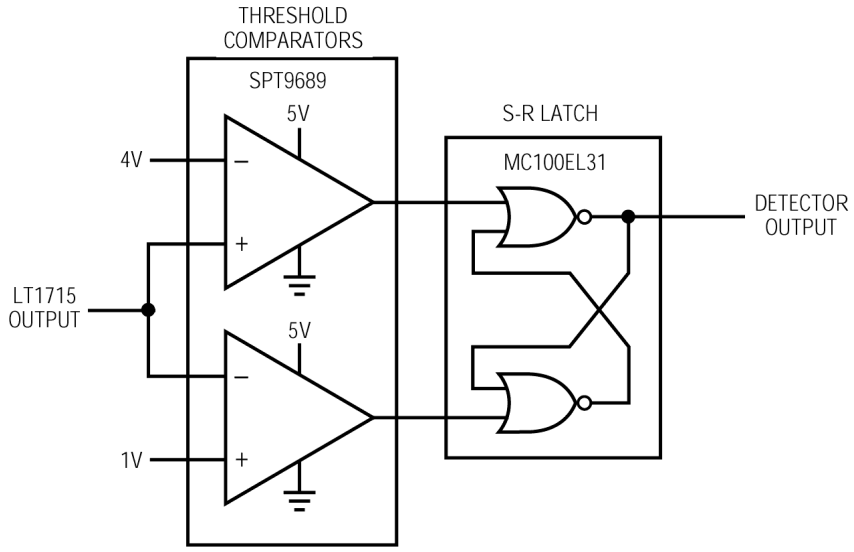

Figure 7-4: Threshold Detector for Toggle Rate Testing

On the test bench, the circuit was capable of almost 200MHz operation using a sinusoidal input with $3.02V_{pp}$ input, allowing a maximum of 10mV overdrive on the comparator inputs. By increasing input amplitude to 3.5Vpp, the circuit easily reached 600MHz. Clearly the circuit would be well suited to accurately sense the peak and minimum levels of my LT1715 at high frequency. The board would be implemented, along with the rest of the toggle rate test circuit to be discussed in the next section, in the final packaged part test board for the LT1715.

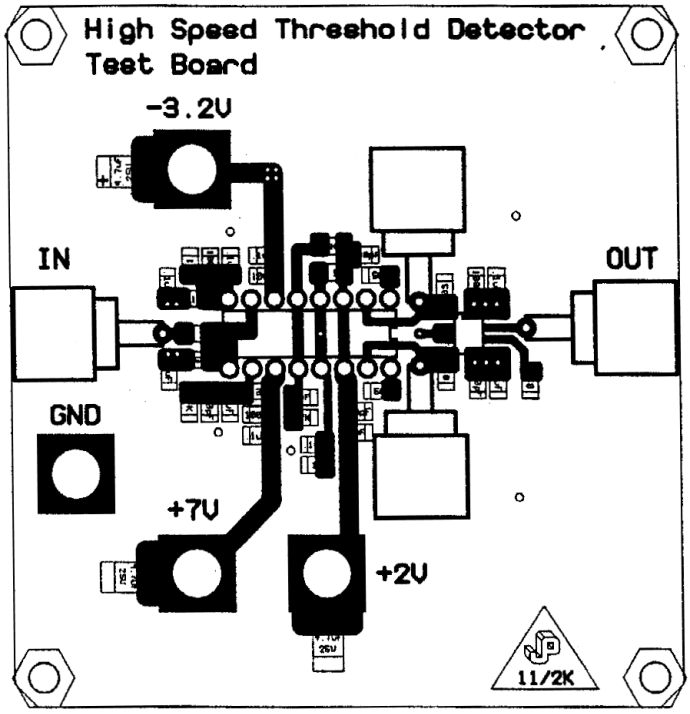

Figure 7-5: High Speed Threshold Detector Front Side Copper Routing

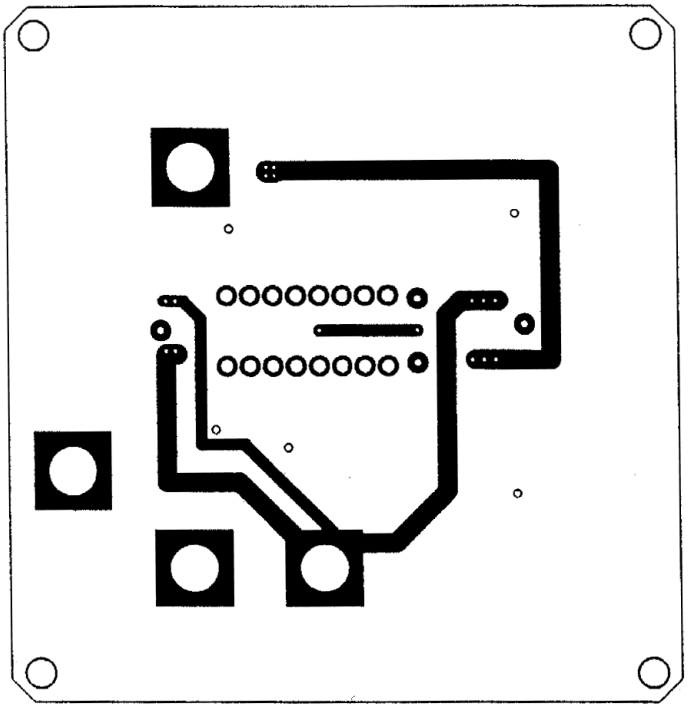

Figure 7-6: High Speed Threshold Detector Back Side Copper Routing

7.3.2 Variable PLL Test Board

A maximum toggle rate test circuit was devised that placed the LT1715 under test in the feed forward path of a phase-locked loop. A block diagram is shown in Figure 7-7 (the full schematic has been omitted due to its size and complexity). The toggle frequency of the PLL is slowly increased until the LT1715 fails to toggle and the loop falls out of lock. A threshold comparator trips when the loop loses lock, disabling the counter and signaling the tester to read the frequency count. The final count is read such that it will be one increment below the current toggle frequency that resulted in failure, and therefore, the maximum successful toggle rate.

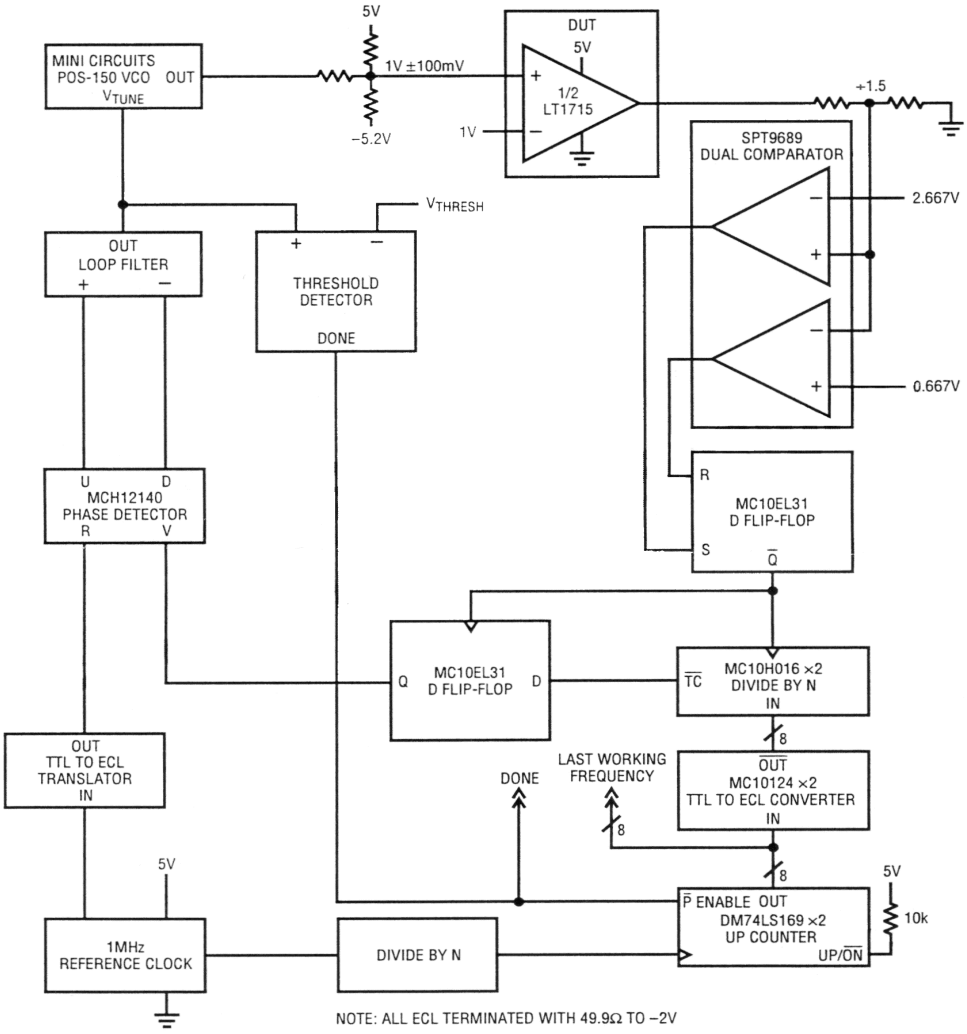

Figure 7-7: PLL Toggle Rate Test Circuit

The heart of the circuit is a counter that sets a divide-by-N to divide the LT1715 toggle frequency down to match the reference 1 MHz oscillator frequency before the phase detector. The phase detector's output feeds through an active loop filter and into the VCO to set the oscillation frequency. The output of the VCO is then buffered and resistor divided to provide a common mode voltage of 1V with approximately 200mV peak-to-peak amplitude. The output of the LT1715 under test is then fed into the two SPT9689 threshold comparators and ECL R-S latch described in Section 7.3.1. The square output of the latch ensures that the setup and hold times of the divide-by-N are not violated. Finally, the feedback loop is closed with the connection of the latch back to the divide-by-N chip.

The 1MHz crystal oscillator provides both a reference clock to the phase detector and, after being divided down, a lower frequency clock to increment the counter. Incrementing the counter increases the divisor by which the LT1715 toggle frequency is reduced down to 1MHz. As the counter increments, the loop is forced to increase toggle frequency in order to maintain stability.

When the LT1715 does not toggle to the limits required at the threshold comparators, the latch will not toggle and phase error will increase. This increases the phase detector output voltage, causing the VCO input voltage to increase, thereby increasing the oscillation frequency of the VCO. As the VCO's toggle frequency increases, the LT1715' s toggle capability is further exceeded and positive feedback ensures that the loop will lose lock. Since the bandwidth of the loop is sufficiently higher than the clock rate of the counter, the counter will not increment during this time.

Once the voltage at the input of the VCO passes through the reference voltage at the lock sensing comparator, a hold signal is triggered. The lock reference voltage limits the maximum testable frequency to 160MHz, where the comparator will be tripped and the counter held – even assuming perfect toggling – in order to shorten testing time. With the lock-sensing comparator triggered, the counter is held at its last count and the hold signal passes to the tester, indicating that the test is complete. The final count on the counter is the maximum toggle rate of the part, expressed in megahertz.

7.3.3 Results From Toggle Rate Test Board

Unfortunately, the PLL test board has not yet been built and tested. It will be completed at a future date.

8 Characterization and Correlation

8.1 DC Performance

DC characteristics important in the design of a comparator include quiescent bias currents, input offset voltage, input bias currents, PSRR, CMRR, output swing voltages and maximum output source and sink currents. The primary focus of my project was to improve AC specifications, so only DC characteristics I modified in my design – bias currents, output characteristics, and input trip points – are described in this section

8.1.1 Biasing, DC Output Levels

An initial verification of the LT1715's functionality was the supply current levels. I was immediately satisfied to find out that my initial silicon biased up to approximately correct levels. The PSpice high-β figures for both comparators powered by input supplies of +5 and –5 and an output supply of +5 and 0 with both outputs driven high were I_{CC} = 2.1mA I_{EE} = 5.8mA and I_S = 9.1mA, and my initial silicon with similarly high β's showed values of I_{CC} = 2.2mA I_{EE} = 5.9mA and I_S = 10.6mA. Additionally, I could test that my output stage current boosting modifications were working as planned on a DC level by switching each of the outputs high or low. Doing so changed the + V_S current from 5.3mA to 4.8mA for each comparator, similar to PSpice's simulation that I_S would change from 4.6mA to 4mA for each comparator. Setting comparator A's output high and B's low also led to the exact same currents as setting A low and B high, so my output stage was working symmetrically, also.

After wafers were sorted, I was given data on their test performance from which I could find the mean values for the bias currents. In addition, I could correlate that data to PSpice measurements since I had approximate values of the device

parameters as determined when the wafers themselves were tested. Almost all device parameters were nominal on one wafer, and the average of thousands of die was only different by a worst case of 6.4%. With the higher PNP and NPN β's on another wafer, I had around 15% error compared to the average of all the parts. This error was easily accountable by the inaccuracies between the model parameters and the device parameters of the wafer under test. I did not think it was necessary to design my PSpice device models to exactly match one batch of silicon. Packaged parts and final specs were also quite close for the LT1715; the simulation tools I used are incredibly accurate, at least at DC. See Figure 8-1 for a summary of these values.

	I_{CC} (mA)	I_{EE} (mA)	I_S (mA)
Initial Wafer	2.2	5.9	10.6
Sort Nominal Mean	2.05	5.80	9.22
PSpice Nominal	2.08	5.54	8.63
Error, Nominal Cases	1.5%	4.5%	6.4%
Sort High-β Mean	2.30	5.92	10.74
PSpice High-β	2.11	5.77	9.14
Error, High-β Cases	8.3%	2.5%	14.9%

Figure 8-1: Bias Currents in Simulation and in Lab

Another useful verification of the part's functionality was the V_{OH} and V_{OL}, in this case with no output current. While I didn't measure this on my initial wafers, the numbers I received during wafer sort indicated that my output stage specifically was biasing up as expected. The mean numbers for V_{OH} and V_{OL} for sorted wafers with nominal parameters were 4.85V high and .15V low on supplies of +5, -5, +5 and 0, and I had expected 4.84V and .15 nominally from PSpice.

8.1.2 Input Offset and Hysteresis

Shortly after I had confirmed the functionality of my initial silicon, I tested the offset and hysteresis using the methods described in Chapter 6. I checked the functionality of the input trimming I had designed for the part, and found that it would work near perfectly to trim the input offset to nearly zero millivolts. I also saw between 3.5mV and 4.5mV of hysteresis band on my parts while PSpice had predicted 4mV to 6mV. This was low, but my sample size and range of device parameter values for one die were quite limited.

Wafer sort validated the early input functionality data, as expected. There was an unfortunately low mean value of 3.6mV for the hysteresis band, but on a positive note, the average offset of all trimmed die was well within +/- .5mV. In fact, the entire distribution fit within +/- 1mV, even better than anticipated. The careful layout and trace routing in addition to the accurate simulation capabilities had proven very successful. However, there was one big problem. The two comparators on any given die did not have the same amount of hysteresis. In fact, comparator A averaged .25mV more hysteresis than comparator B. The cause of this problem is discussed in detail in Chapter 9.

In order to find DC performance variations with varying supplies, temperature and common mode, I decided to wait until a DC test fixture was assembled. Unfortunately, testing was not started until after my departure from Linear, and problems were discovered that would go unsolved. Pages 4 and 5 of the LT1715 datasheet in Appendix A illustrate DC performance. All but one curve – the input trip points versus temperature – performed almost exactly as planned. The deficiencies in input performance versus temperature are discussed in Chapter 9.

8.2 AC Performance - Propagation Delay

The good DC performance, though slightly refined by tweaks I did to the original circuitry, is not essentially my work. However, the AC performance (measured only with packaged parts) was quite directly correlated to my efforts. Delay on the LT1715 was decreased (improved) slightly, as expected. This was in part due to the increased drive current capability in the PNP side of the output, and additionally from the smaller package and shorter bond wire and lead inductance at the output. This allowed a front-page specification for propagation delay to be 4ns, specified for dual supply 5V operation into 10pF. PSpice

had predicted something over 4ns for the rising and falling propagation delays in response to a 100mV step, with a good deal of variation with process parameters and simulated parasitics. In the nominal case, PSpice simulated a 4.4ns propagation delay going from low to high and 4.6ns delay going from high to low. The average measured data, collected from a wide sampling of different packaged parts, showed an average of 4.1ns rising and 3.7ns falling.

I had expected the propagation delay to be significantly faster than the simulations due to my inaccurate parasitic modeling. In order to simulate the effects of metal-to-metal and metal-to-substrate capacitance, a huge array of capacitors were extracted from the layout information. However, to ensure that the parasitics accounted for fringing and lateral capacitances, I doubled the parasitics capacitances. On the majority of the routing in the LT1715, this would be accurate. However, on the very wide output traces and power traces which steer the large slew currents of the part, the capacitances were already large, and now doubled. By cutting them down my propagation delays fell about .4ns each. However, the exchange of rising being faster than falling to the opposite situation is strange. Since the comparators were merely working better than expected, no investigation of this discrepancy was attempted.

9 Error Sources and Discussion

9.1 Hysteresis Mismatch

The first defect found in the LT1715 was that the hysteresis band was smaller than desired. When wafer sort data arrived, one of the causes became apparent. Comparator 1 has .2mV to .6mV less hysteresis than comparator 2. As explained in Chapter 2, hysteresis is provided by switching a current with a differential pair and feeding it into the emitter followers that buffer the input signal into the gain stage (see Figure 9-1 for a reminder of the circuitry). The matching of currents through the emitter followers with the additional current switched to provide the hysteresis is critical. Careful circuit design was intended to make this matching nearly perfect regardless of process parameters, power supplies, or temperature.

During layout, however, the extreme vulnerability of the circuit was not fully realized. Although the emitters of the mirrors shown in the top of the schematic were connected perfectly and immune to any I*R drops in the routing metal, the current sources shown at the bottom of the schematic were not. This is a problem for several reasons. First, the two current sources at the top of the schematic and to the left, shown linked to the diode-connected transistor in the top right corner, are actually dependent on the current through the main differential pair that provides voltage gain. All of the combined currents pass through a degeneration resistor used to compensate for early voltage effects with varying supply voltages. Furthermore, the current source for the main differential pair uses the same bias voltage as the current sources in the bottom right corner of the schematic.

The mismatch is caused by a large I*R drop between the emitters of these current sources mentioned above. Figure 9-2 shows the path of the V_{EE} trace from the bond pad to the emitters of the current sources. Although the metal resistance of the V_{EE} trace is relatively small, all the current from each comparator is conducted through it, and the voltage drop

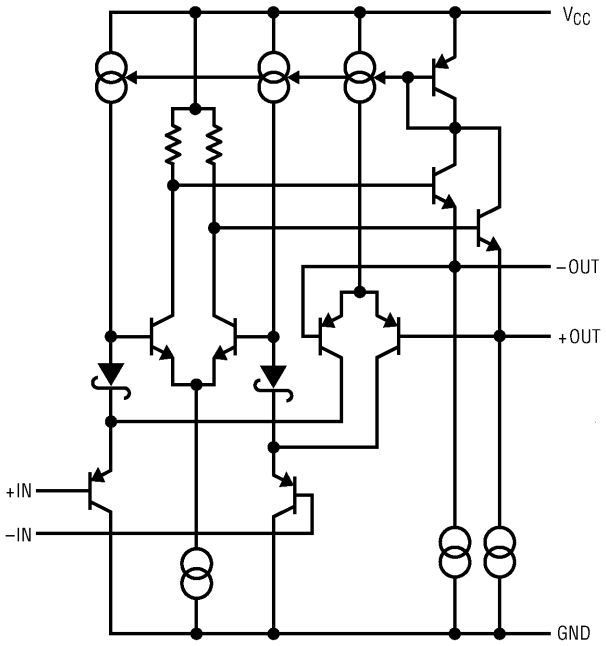

Figure 9-1: The Gain and Hysteresis Stage From Chapter 2

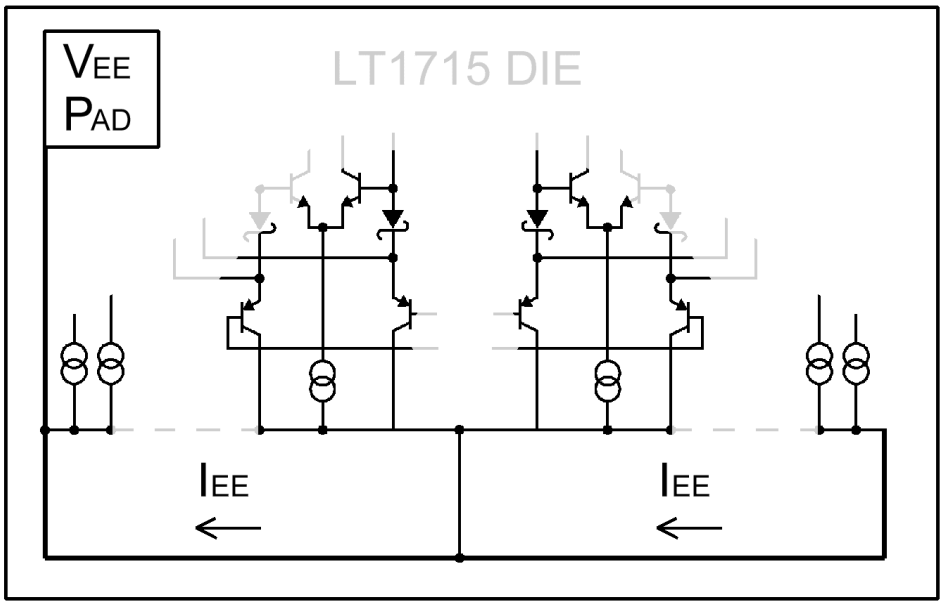

Figure 9-2: The I_{EE} Current Path on the Die

across it is 10mV. A PSpice simulation was run with the voltage drop across the V_{EE} routing included. In the right-hand comparator, the I*R drop decreases the current through the right-hand current sources by about 7%, which changes the switched hysteresis current by the same 7%. On the other hand, the gain stage differential transistor currents only change by about 3% in response to the current source change because

they use much more degeneration and also depend on the gain differential pair's current. These miniscule differences in currents turned out to cause a significant change in input trip voltages. Figure 9-3 shows the hysteresis of comparator B (COMP2) plotted against the hysteresis of comparator A (COMP1) for individual die, illustrating the average difference of 15% less hysteresis on comparator B. As unattractive as it looks, the variation in hysteresis falls within the specified limits, so no changes were made to the layout to eliminate the problem. Easy solutions include thickening the metal V_{EE} trace that is currently very narrow and moving the V_{EE} bond pad on the die to make the current paths shorter and closer to symmetric.

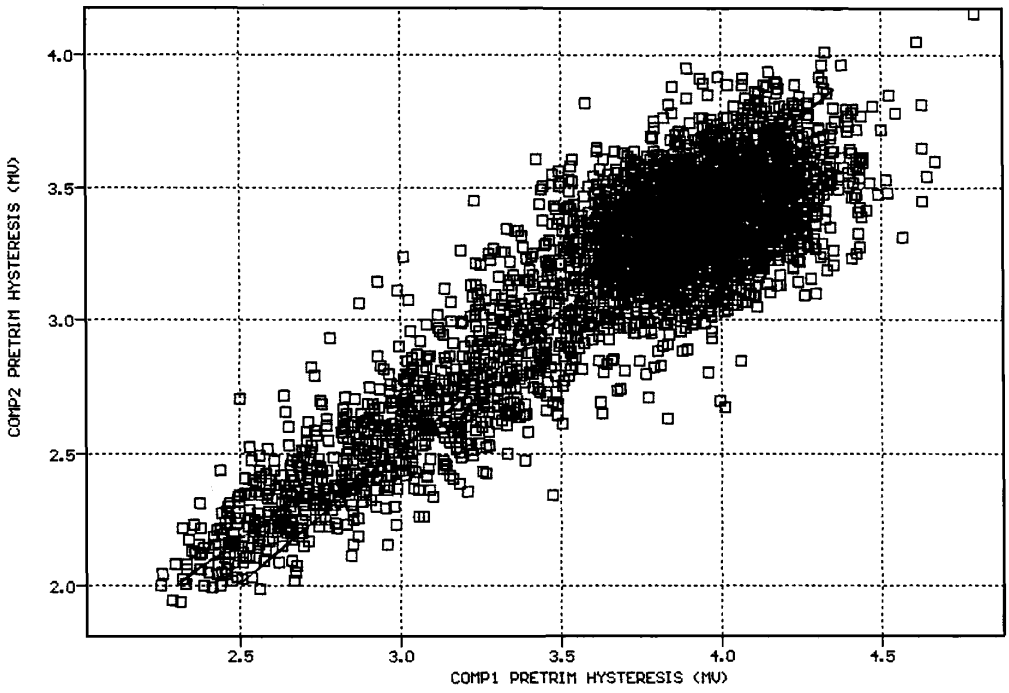

Figure 9-3: Cross Plot of Hysteresis Between Comparators

9.2 Channel Interactions

With the DC specifications looking very good (with the exception of the hysteresis mismatch), AC testing was begun. All initial toggle rate testing was done with only one channel driven and simultaneously measured. However, several PSpice simulations from the design process had shown an interaction between the comparators when they toggled simultaneously, and now that variables could be tweaked and results viewed in real time (instead of running overnight simulations), I was ready to see just what would happen.

I was not prepared for the severity of the results. The LT1715's two channels are designed to be entirely independent, however, at frequencies approaching and exceeding 100MHz the output signals began interfering with each other, causing jitter or even failure to toggle. Figure 9-4 shows one of the LT1715's outputs toggling at 100MHz with the other output constant (not driven) on a sampling scope set to display infinite persistence. Figure 9-5 shows the same output channel in response to the same input signal when the second comparator is toggled across an array of frequencies. At frequencies well above 100MHz but within the operating toggle rate of the LT1715, there is a distinct possibility that any signal on the opposing channel could cause failure on the higher frequency channel.

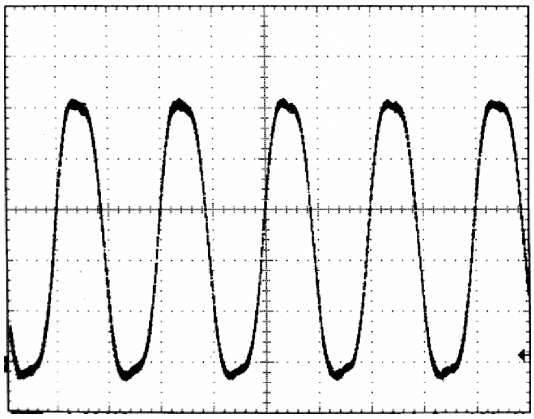

Figure 9-4: Clean 100MHz Toggling

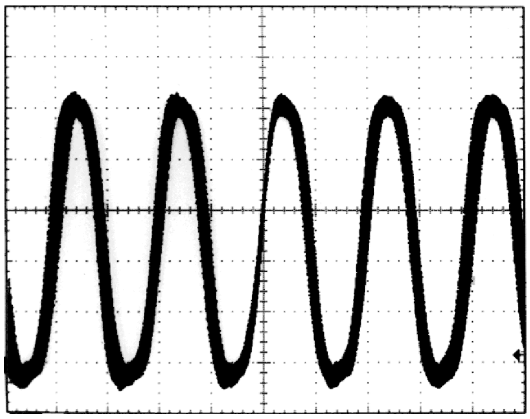

Figure 9-5: Jitter With Both Channels Driven

The believed cause of the problem is oscillation in the power supplies caused by the inductance of the supply bond wires and package leads when currents turn on and off very quickly. The ESD devices at the bond pads look like big capacitors at high frequency, and the rapid current steps cause some oscillation. This oscillation is coupled to other sections of the chip and can affect the output swing of the other comparator. If the bond wire inductance is significant enough, the V=L*di/dt term could also cause capacitive voltage coupling into the more sensitive output circuitry, reducing toggle capability.

The problem is very complex, however. Changing the external bypassing of the LT1715 has an effect, likely by minimizing the inductance of supply traces. Purposely adding inductance with an inch of wire to the ground pin on the LT1715 caused the interaction to become drastically worse, reinforcing the idea that power supply inductance coupling to the internal voltage nodes through the ESD capacitances was to blame. When the rising or falling edges of the output signals align, the supply currents are drastically increased and the oscillation

becomes significant. The best solution for the problem was to layout an application PC board with as low impedance of a ground connection as possible and with superior bypassing as close to the LT1715 as possible. However, even my best efforts could not eliminate the high frequency cross talk.

9.3 Hysteresis Variation with Temperature

Another problem with the LT1715 was discovered very late in its development. PSpice simulated that the newly designed hysteresis circuitry performed flawlessly over temperature and with supply variation. Input offset voltage and hysteresis width were perfectly constant – improved over previous designs. The nominal figures on the initial batches of parts looked to be well related to the PSpice numbers, and easily measured supply variations had no effect. However, when final testing began to look at the performance of the LT1715 over temperature, the DC input specifications wandered drastically. A plot of the LT1719's offset and hysteresis compared with the LT1715's versus temperature can be seen in Figure 9-6.

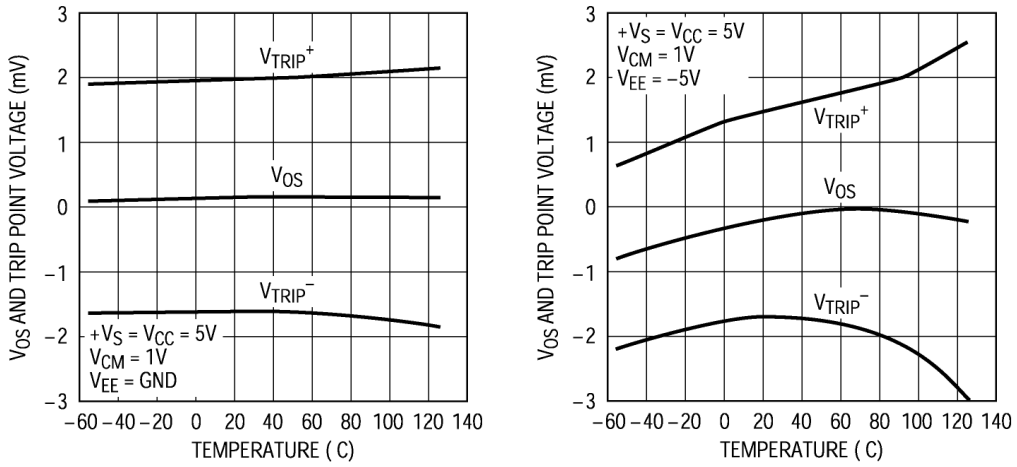

Figure 9-6: Input Offset and Hysteresis of the LT1719 (left) and the LT1715 (right)

A likely cause of the problem is apparent. The current sources at the top of the schematic in Figure 9-1 use varying amounts of emitter degeneration to create different temperature coefficients so that the relative voltage differences are unchanged with temperature. The diode-connected transistor in the top right corner that sets up the currents in the top sources has only 30 **Ω** of degeneration and less than 1mA of current, creating only tens of millivolts of degeneration. In final silicon, the emitter is tied through a substantial length of minimum width metal to the degeneration and to the supply. The extra resistance of this metal, although only ohms, is enough to throw the very sensitive current mirroring out of match as

temperature varies, and is also likely the cause of the slightly smaller than desired average hysteresis, even at room temperature.

The variation forced a widening of the temperature-guaranteed data sheet specifications for offset voltage and hysteresis and is clearly shown in the datasheet charts. Further causes of this particular defect may also exist, and time to investigate them has been set aside starting in July of 2001, after the release of this paper, as the mistakes made in the LT1715 should not be repeated.

10 Comparisons and Conclusions

10.1 Rail-to-Rail Output Stage Speed

The title of this paper mentions toggle rate in a rail-to-rail comparator output stage, but a large amount has been written about the overall performance of the comparator that the output stage was implemented for. Looking at the performance of just the output stage, and more specifically, just the output Darlington, shows the real success of the design. The final design has enough output source or sink current to slew 3V/ns in either direction, allowing the LT1715 to boast the lowest propagation delay in its class - 4ns. Additionally, the output still remains close to the output supply rails even while supporting up to 20mA or 30mA of constant sink or source current. However, the real beauty of the circuitry is its ability to toggle – to turn the output devices on and off remarkably quickly.

In my final testing, I noticed that the LT1715 toggle limitation would always be failure to either reach V_{OH}, or failure to reach V_{OL}. On a 5V output supply, the comparator would be toggling from .3V to 3.9V or from 1.1V to 4.7V, and never from 1.1V to 3.9V. As the comparator passed its maximum toggle frequency, it was the skew in propagation delay at frequency that stopped the comparator from going faster.

To test this theory, a small amplitude sinusoid with a varied DC offset was used to change the effective duty cycle of the output, thereby compensating for different delays through the comparator to the output devices. With the offset perfectly tuned, the results were impressive; with completely nominal silicon, 300MHz was normal fare as shown in Figure 10-1, and frequencies approaching 350MHz were possible with the higher β silicon. This was the final proof that quasi-saturation had been completely overcome, and that the output stage was a true success. Nevertheless, the 300MHz toggle speed was only for show; the causes of skew are too varied and interdependent to ever be compensated for in circuitry.

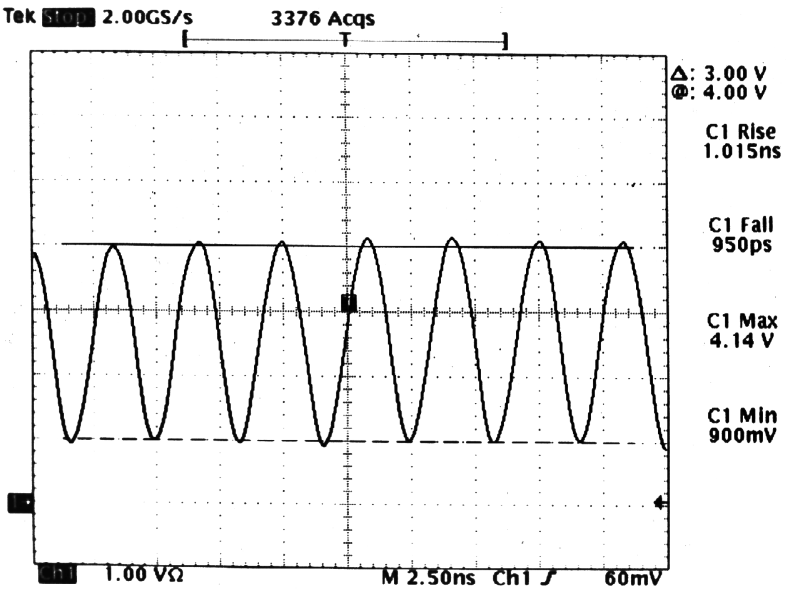

Figure 10-1: Compensating for Skew to Allow 300MHz Toggling

10.2 Comparisons

The LT1715 fits into the Linear Technology comparator family exactly as intended. It is significantly faster than its peers, but similarly versatile with the exceptions of the shutdown of the LT1719 and the latches, rail-to-rail inputs and complimentary outputs of the LT1714. However, with separate input and output supplies, controlled internal hysteresis, an exceptionally fast propagation delay, and the ability to toggle at high frequencies, the LT1715 still stands out.

The LT1715 continues to excel in comparison with the products of competing companies. Maxim Integrated Products has a similar large family of comparators (MAX 96x's) boasting 4.5ns propagation delays. However, they specify propagation delay into 5pF instead of 10pF and do not test it in production other than at large overdrive. Additionally, the parts have insufficient gain to handle low overdrive, have less refined hysteresis and input offset specifications, and don't offer as much output drive or as little output resistance as the Linear products. At the same time, the Maxim parts all have rail-to-rail input stages and many are full featured with a latch, shutdown or both, so they are still worthy competitors.

Another direct competitor, Analog Devices, produces the AD8611 and AD8612 with impressive propagation delay (the front page claims 4ns but the typical number on a 5V supply is 5ns) and a typical 100MHz toggle rate. However, the small common mode range (0 to V_{CC}-2V), lack of internal hysteresis, and omission of a rail-to-rail output stage limit the

usability of these parts. Analog Devices successfully markets the comparators as high-speed pin for pin replacements for the older LT1016 and LT1394, but the LT171x and LT172x families would likely find their way into new designs over the older Linear parts, anyway.

To conclude, the LT1715 design is a success. The improvements made to the output stage allow high-speed toggling well above 100MHz – a speed completely unmatched in the market. The extremely fast propagation delay afforded by the same modifications is equally noteworthy. By overcoming soft-saturation in the output stage and tuning an already high-speed design, the final circuitry in the LT1715 has placed it among the best TTL-output comparators that money can buy. Engineers will no longer be limited by the speed of available comparators; the LT1715 will fulfill all their needs for speed.

Appendix A

LT1715 Final Datasheet

(16 Pages)