
TECHNOLOGY DESCRIPTION

The EE CMOS technology used by Vantis in programmable logic is a single-poly, double- or triple-metal process. It has been optimized for high-speed programmable logic devices, which do not have the same density constraints of memory devices. The basic characteristics of the EE process are:

- ◆ CMOS
- ◆ Grounded substrate
- ◆ Single-poly, dual metal or single-poly, triple metal
- ◆ Five generations in production: EE4, EE5, EE6.5, EE7, EE8
- ◆ 1.3 μm , 0.9 μm , 0.7 μm , 0.45 μm , or 0.35 μm minimum feature
- ◆ 0.85 μm , 0.65 μm , 0.5 μm , 0.35 μm , or 0.25 μm gatelength (L_{eff})
- ◆ 150 \AA or 80 \AA gate oxide thickness
- ◆ 90 \AA tunnel oxide thickness

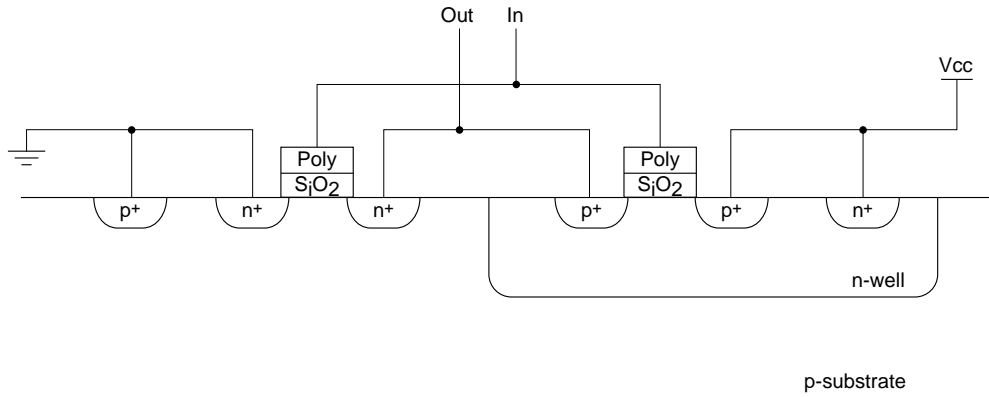
CMOS PLDs use standard CMOS logic internally, with the addition of a programmable array. The output buffers of most devices are designed to be compatible with TTL circuits, and therefore have n-channel enhancement pull-up transistors.

Vantis' EE CMOS process for programmable logic is simplified by the absence of standard depletion-mode transistors in the more advanced processes. Depletion mode transistors are a vestige of NMOS design, and are not really needed. This results in the elimination of a mask and implant step, reducing the process cost and simplifying the structure.



Transistor Cross-Section

Figure 1 shows a cross-section of a basic inverter. This is a very straightforward structure. The gates consist of poly-silicon; the other connections are made with metal.




16507C-001

Figure 1. CMOS Inverter Cross-Section

Erasable Technology

Any erasable CMOS technology is based upon the concept of stored charge. The charge is stored on a transistor with a **floating gate**—that is, a gate that has no connection. The transistor actually has two gates: one that floats, and one that acts as a control gate. The control gate is used to establish the field across the floating gate (see Figure 2).



16507C-002

Figure 2. Floating-Gate MOS Transistor

In the programmed state, there is a net deficit of electrons in the floating gate. The resulting positive charge turns the transistor ON. In the erased state, there are enough electrons on the control gate so that the negative charge turns the transistor OFF.

There are two basic ways of transferring the charge onto the floating gate: a) hot electron injection, and b) tunneling. Electrically erasable devices rely on tunneling.



Electrically Erasable Technology


Electrically erasable devices use **Fowler-Nordheim** tunneling as the mechanism for getting charge onto the floating gate. This is defined roughly as tunneling that occurs as a result of a field placed across the barrier that the electrons tunnel through.

Some amount of **direct** tunneling, or tunneling that occurs without an applied field, is always possible through any energy barrier. It may be extremely small or significant, depending on the width of the barrier. Since tunneling electrons are going through the barrier instead of over it, the height of the barrier does not affect the amount of tunneling.

For an electrically-erasable cell, the tunnel oxide is about one third the thickness of the oxide of a UV-erasable part; therefore, tunneling occurs at relatively low fields. Even so, the field used to cause tunneling is about five times the field used to cause hot-electron injection for UV parts. Note that tunneling is theoretically possible on a UV part using a very high field, and the normal electron injection would swamp out any tunneling that would occur.

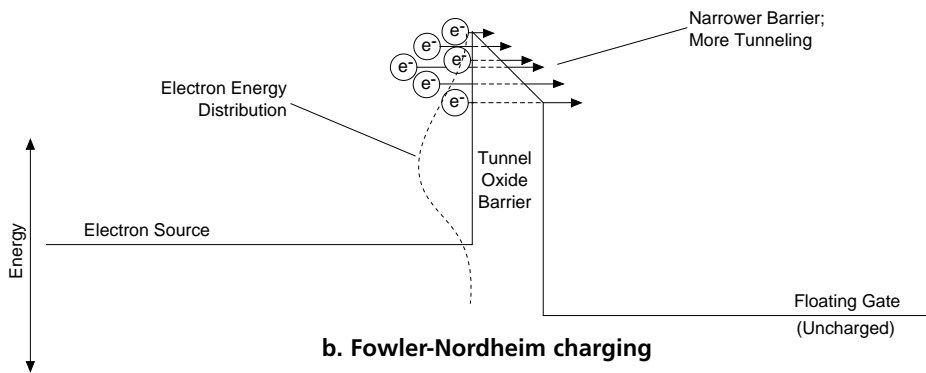
Fowler-Nordheim tunneling involves placing a potential across the barrier which distorts the band diagram as shown in Figure 3. The “angle” caused by the applied potential effectively thins part of the already-thin barrier, making tunneling easier. It is this tunneling under bias that is used to program electrically-erasable devices. Note that by reversing the bias, the tunneling can occur just as well in the opposite direction. This is what makes electrical erasure possible.

Electrical erasure has advantages over UV erasure both in cost and quality. Because the erasure is electrical, no expensive window is required in the package. This makes erasability cost-effective even in high-volume production quantities. In addition, the fast erasure allows Vantis to reprogram the device many times, allowing many more paths to be tested than can be tested in a UV part. This provides much higher quality, especially in higher-density devices.




a. Direct tunneling

16507C-003



b. Fowler-Nordheim charging

16507C-004



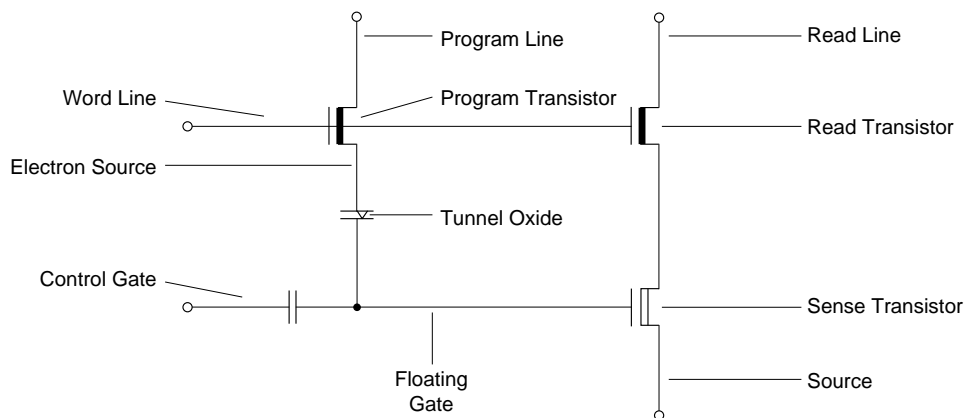
c. Fowler-Nordheim discharging

16507C-005

Figure 3. Energy Band Diagrams

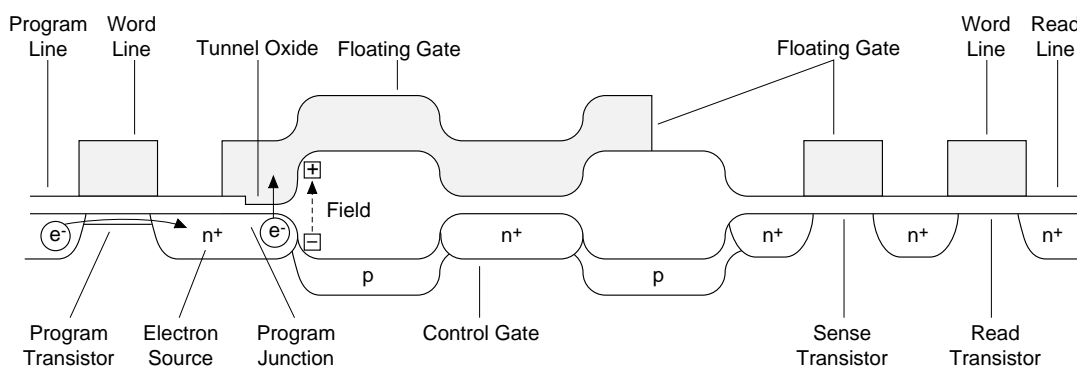
Cell Configuration and Programming

The programming cell is shown in Figure 4. To improve device speed, the programming cell has been divided into the programming portion and the data path portion. In addition to speed, there are a number of other benefits to this approach. At the most basic level, this eliminates a poly-silicon layer, simplifying the process. This reduces costs and improves reliability



a. Circuit

16507C-006



b. Cross-section

16507C-007

Figure 4. EE PLD Programming Cell

The programming half requires long-channel transistors capable of sustaining high electrical fields; the data path requires short-channel transistors that are fast. Note that this does take more space, but in PLDs, the size of the cell is not a limiting factor as it is in memories. In a PLD, the programming array can take up as little as 10% of the die area, while a memory typically uses more than 90% of the die area for the programming array.

Programming and erasure are complementary procedures in EE technology. However, the sense of programming and the sense of erasing are perhaps opposite to what one might assume. A cell is considered to be programmed if there is a charge deficit on the floating gate, providing a positive voltage; it is erased if there is excess charge on the floating gate, generating a negative voltage. This means that programming a device only requires turning ON those cells that are needed, rather than turning OFF all of the cells that are not needed.

A cell fresh from wafer fabrication has no net positive or negative charge on the gate. To balance the threshold of the transistor for reliable turn-on and turn-off, a **cell implant** is used to center




the threshold voltage near 0 V. Programming and erasing involve either removing electrons from the conduction bands of the poly-silicon gate or adding excess electrons, providing a net charge that will move the gate voltage solidly to one side or the other of the threshold voltage.


When programming or erasing the device, a voltage is applied between the program and control gate nodes. The direction of the voltage determines whether the cell is erased or programmed.

When erasing, the control gate is given a positive voltage, and the program node is grounded. This attracts electrons from the program transistor across the tunnel oxide to the floating gate, turning the read transistor OFF (see Figure 5a).

When programming the cell, the program node voltage is elevated, and the control gate is grounded, reversing the electron flow, as indicated in Figure 5b. Enough electrons flow off the floating gate to leave a net positive charge; this turns the transistor ON.




16507C-008



16507C-009

Figure 5. Cell Biasing

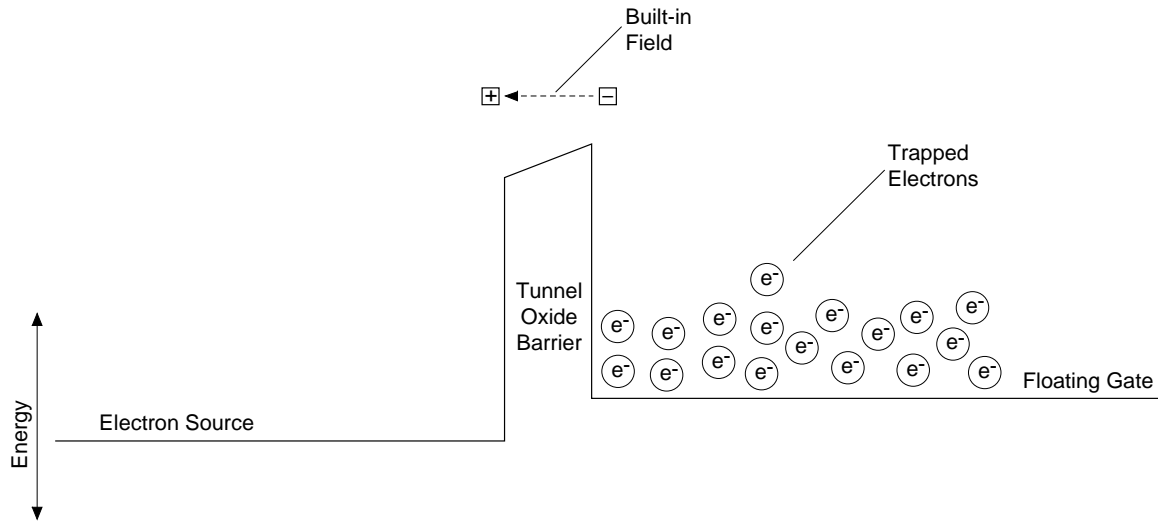
Vantis has modified the programming cell to increase programming efficiency and has a number of patents on the resulting circuit. On traditional devices, the source node is grounded during programming. On Vantis' devices, the source node is raised to the same potential as the control gate, as shown in Figure 6. This increases the **coupling ratio** of the cell. The coupling ratio is the percentage of the applied field that appears across the tunnel oxide. When the source is grounded, the field across the tunnel oxide is reduced (since there is another capacitor in parallel with the tunnel oxide). By raising the source voltage, more of the field is available for programming. The coupling ratio can therefore be thought of as a measure of the programming efficiency; since the efficiency is higher, lower voltages are required for programming.



16507C-010


Figure 6. Source is at Same Potential as Control Gate to Improve Coupling Ratio

The split-cell configuration also allows a simpler programming algorithm, since the programmer can take advantage of the self-limiting nature of programming and erasure. The split cell places the read cell gate and the floating cell gate in “parallel” with each other. Therefore, the floating cell can be either completely charged (with a net excess of electrons) or completely discharged (with a net deficit of electrons, or an excess of holes), as shown in the energy diagrams in Figure 7. This is simple to do, since the electrons that have crossed the barrier set up a field that opposes further tunneling. As more electrons cross the barrier, the opposing field grows strong enough to block more electrons from tunneling (Figure 8). Regardless of the state of the floating cell, the select line will turn on or off the read transistor; the cell will only be read, however, if both the read transistor and the floating transistor are ON.



a. Charged (erased)


16507C-011



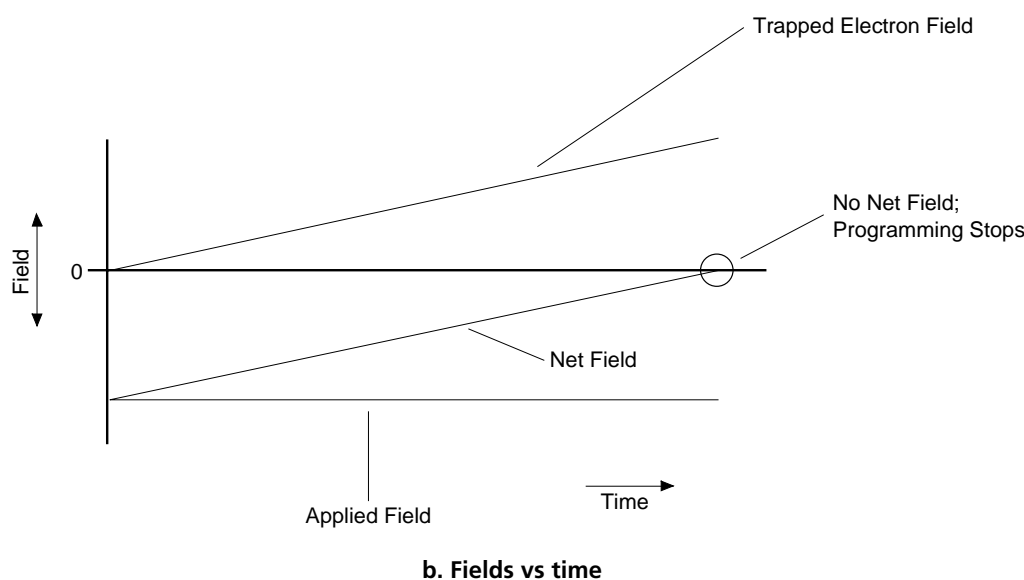
b. Discharged (programmed)

16507C-012

Figure 7. Stable EE Cell




16507C-013



16507C-014


Figure 8. Self-Limiting Programming and Erasure

In standard one-transistor cells, the two gates are actually in “series”. If the floating gate is charged, then the transistor is OFF, regardless of the state of the select line. In order to read the cell, the floating gate has to be neutralized so that the select line controls the transistor (Figure 9). If the floating gate were completely discharged, then the transistor would be ON regardless of the state of the select line. The programming algorithm is therefore more complicated, since the amount of charge removed must be monitored to ensure that just enough charge is removed to neutralize the floating gate.



a.

16507C-015



b.


16507C-016

Figure 9. UV Cell: a. Charged (Programmed) State; b. Neutral (Erased) State

Array Configuration


The discussion above focused on individual cells. These cells must be joined together to form a complete array that is driven by input lines and drives product terms.

There are two configurations used in Vantis' EE CMOS devices (see Figure 10). The configuration in Figure 10b provides some benefit to device speed because the parasitic capacitor does not couple the input line and the product term, but both approaches are actively used in designs.



a. Bit line at the drain

16507C-017



b. Bit line at the source

16507C-018

Figure 10. Two Array Configurations

PROGRAM INTEGRITY

Reliable programming of PLDs requires the use of well-calibrated, quality programming equipment. To ensure that the device is correctly programmed, the correct voltages and times must be applied.

As discussed above, it is impossible to over-charge or over-discharge the programming cell since the mechanism is self-limiting. This provides more leeway and makes the programming algorithms less sensitive to programmer variations. This ultimately provides higher, more consistent programming yields under real-life production programming conditions.

However, if the cell is under-programmed or under-erased, an insufficient amount of charge might transfer onto or off of the floating gate. When programming, this might not turn the read cell ON sufficiently, potentially slowing down the device. In the case of erasure, the read cell might be partially ON if it is not completely erased. This may cause “disconnected” inputs to appear partially connected. Thus, it is important to ensure that the programming pulsewidths are long enough to provide adequate programming.

If the programming voltages are slightly inaccurate, CMOS devices often can still be programmed correctly. However, excessive voltage might cause device damage if breakdown voltages are exceeded. Extremely low voltages might fail to engage the programming circuitry completely.

Because of the need for accurate programming, and for ensuring that the programming algorithms are up-to-date, we certify programmers that meet strict criteria for all products.




Data Retention

In an electrically-erasable device, the floating cell is programmed by forcing electrons to tunnel through the tunnel oxide into the floating gate. Ideally, these trapped electrons mean that the device remains programmed indefinitely. Actually, the charge cannot remain indefinitely, but its lifetime is normally extremely long. The stability of the program charge is called **data retention**, the ability of the device to retain its charge as programmed.

There are two basic leakage mechanisms: direct tunneling and thermal leakage. These mechanisms occur independent of whether the cell was programmed by electron injection or tunneling. The amount of direct tunneling is a function of the potential across the tunnel oxide, and is generally very low. Leakage is normally dominated by thermal charge decay.

On one side of the energy barrier, there are electrons with a distribution of energies (see Figure 11). Some have enough energy to escape over the top of the barrier. As the temperature is raised, more electrons achieve the energy required to overcome the barrier.



16507C-019

Figure 11. Data Loss Mechanisms



The tendency of the gate to leak can be modeled as an Arrhenius function, which means the formula for the programming “decay time” t_d has the form:

$$t_d = Ke^{E_a/kT}$$

where

E_a is the intrinsic **activation energy**

T is the temperature in Kelvin

k is Boltzmann’s constant

K is a scaling constant.

If we can measure the rate at two known temperatures, then:

$$\frac{t_{d1}}{t_{d2}} = \frac{Ke^{E_a/kT_1}}{Ke^{E_a/kT_2}}$$

Note that the constant K drops out, so we need not be concerned with its specific value. From this we find that:

$$E_a = \frac{T_2 \ln(t_{d1}) - T_1 \ln(t_{d2})}{k(T_1 - T_2)}$$

This lets us measure E_a , which should be constant for a given process. The higher the value of E_a , the longer the decay time will be. This is because E_a roughly represents an energy “barrier” that must be overcome for an electron to leak away. The higher the barrier, the fewer electrons have the energy to overcome E_a .

Charge leakage can be aggravated by poor quality tunnel oxide. Defects in the oxide provide a lower energy path for discharging, effectively lowering E_a . Baking a device accelerates this leakage and identifies devices with weak oxide. Vantis uses a bake for all EE products to ensure that the production devices have a high E_a and therefore good data retention. The average E_a for all devices, including those with weak oxide, is about 0.8 eV. After eliminating the weak devices by a 250°C 24-hr bake, the average E_a is about 1.8 eV.

Data retention time depends on the temperature to which the devices are exposed. The higher the temperature, the shorter the decay time because the electrons have more energy, and more can leak off the gate.

There are two temperatures that may be of concern for different reasons: the maximum device storage temperature (150°C) and the maximum operating temperature (125°C for military). In the first case, the idea is to know that if a programmed part sits on a shelf for some period of time before being used, the program will remain intact for that time. The second case is intended to give an idea of how long a device will remain operational in-system.



Using the equation above to solve for the decay time at these temperatures, the result is several decades for the storage temperature, and even longer for the operating temperature. For room temperature, the exponential nature of the function makes the decay time increase to centuries.

Vantis specifies 10 years at the maximum storage temperature (an industry standard for EPROMs and EEPROMs), and 20 years in-system under worst-case military conditions. That the calculated numbers are so much higher builds confidence in the numbers specified. In general, the typical end-of-life failure mechanisms that affect all devices (and which are unrelated to the EE cells) will cause device wear-out before the program data is lost.

The integrity of the charge in the electrically-erasable cell also stands up to any electrical fields that exist in surrounding equipment. For charge to be transferred off of, or onto, the floating gate, a field must be placed across the oxide. Such a field cannot be generated outside the programming mode; an external field, no matter how strong, cannot set up the programming mode.


The charge might also be pulled through some other oxide if the field were large enough. However, to remove the charge through anything but the tunnel oxide requires an external field so high that the rest of the device would break down before any cell charge were ever lost. This would occur on any device, programmable or not. Therefore, any external field strong enough to remove charge from a floating cell will destroy the rest of the device first.

CELL ENDURANCE

Another factor that affects long-term data retention is the **cell endurance**. The endurance is the number of times the device can be erased and reprogrammed. Over time, the oxide can wear out, resulting in a gradual reduction in E_a . This occurs as defects are created in the oxide. These defects trap electrons, which then oppose the field that is required for programming. Given enough trapped charges, the established potentials will be insufficient for programming. This typically happens after hundreds of thousands of reprogramming cycles.

The ability to charge up a cell with good data retention can be measured by the *margin voltage*. This is the voltage that must be applied to the control gate to counteract the charge on the floating gate. If the gate is highly charged, a larger margin voltage is needed to overcome the charge. Thus, put simplistically, a higher margin voltage indicates better cell charging.

Figure 12 illustrates measurements of the margin voltage as the number of program/erase cycles is increased. By 100,000 cycles, the margin voltage still is greater than 4 V; for the cell to fail, the margin voltage must fall to below about 1 V.



16507C-020

Figure 12. Cell Endurance: Margin Voltage Solid After 100,000 Program/Erase Cycles

For EEPROMs, which often are reprogrammed in-system, it is important to know how many thousands of times the device can be reprogrammed. However, most EE PLDs are not intended to be programmed in-system, and probably are programmed very few times. Most production units are programmed only once by the user. Prototypes might be programmed tens of times at most. Therefore we specify a maximum number of 100 erase/reprogram cycles.


This does not imply that the devices are weaker than EEPROMs; it is just that more extensive testing would have to be done to justify specifying a larger number. Since this larger number is not needed, a cost savings is realized because of the test simplification. Note that the devices are actually programmed hundreds of times in testing before they are shipped out, giving outstanding programming and functional yields; however, the number of erase/reprogram cycles specified refers only to programming done by the user.

DEVICE CHARACTERISTICS


Power Dissipation

CMOS technology is associated with low power consumption, and indeed, all CMOS PLDs use lower power than their bipolar counterparts. However, most PLDs do not provide the zero-standby power that standard CMOS logic parts do.


The basic CMOS inverter lowers operating power because at any given time, only one of the two transistors can be fully ON. The other is OFF and blocks the flow of DC current. Thus, when the device is in a stable state, no current can flow. While the device is switching, both transistors are partially ON, allowing for a transient current spike. This means that power is consumed only when the device switches. Because a spike occurs for each transition, the average power consumption is affected by the frequency of operation (see Figure 13).



a. $V_{OUT} = \text{Static LOW}$ 16507C-021



b. $V_{OUT} = \text{Static HIGH}$ 16507C-022



c. **Dynamic power dissipation**

16507C-023

Figure 13. CMOS Inverter Power Dissipation

This type of circuitry is found throughout most of a PLD circuit. However, one portion of the PLD circuit does not use a standard CMOS inverter: the programmable array. One of the necessary elements of zero-power operation is that the output of the inverter have a voltage swing from ground to V_{CC} , so-called **rail-to-rail** operation. In the array, such a wide swing makes the propagation delays too long. To speed up the device, the sense amps that determine the state of a product term are designed to have a much more limited swing. This means the sense amps are constantly drawing power, even when not switching. These are the half- and quarter-power CMOS PLDs; their power consumption is still less than that of a bipolar PLD. Since most CMOS PLDs are used in TTL sockets, the CMOS PLDs work well.

I_{CC} vs V_L and Loading

The greatest external contributors to I_{CC} are the input HIGH level (V_{IH}) and the output load.

As the V_{IH} drops from its ideal level of V_{CC} , the inverter starts to draw current. The worst case scenario would be a V_{IH} at the minimum of 2.0 V, which could contribute some 5 mA per input buffer to the power consumption.

The output load can also have a dramatic effect on power dissipation, especially on devices that have many I/O pins. For an output driving a purely capacitive load, the power dissipation contributed by the load for one output is determined by the load capacitance, the frequency at which the output is switching, and the output voltage swing (V_S). The output stage will go through a process of repeatedly charging and discharging the capacitor. Although the direction of charge flow reverses itself every other transition, the relative voltage change does too, so that the power contribution is the same for a charge and a discharge.



If we consider the case of charging the capacitor, we will be placing a charge, Q_L , on the capacitor that is determined by:

$$Q_L = C_L V_O$$

where C_L is the load capacitance and V_O is the output voltage. The current contribution from this is:

$$\begin{aligned} i &= \frac{d_L Q}{dt} \\ &= \frac{C_L d_{V_O}}{dt} \end{aligned}$$

In one half the output transition period t_p , the change in output voltage will be equal to the output swing V_S . This means that:

$$\begin{aligned} i &= C_L \frac{V_S}{\frac{t_p}{2}} \\ &= 2C_L \frac{V_S}{t_p} \\ &= 2C_L V_S f_o \end{aligned}$$

where f_o is the frequency at which the output is switching.


The power dissipation is the product of the current and the voltage. Since the voltage is changing during the time that the power is being dissipated, we can approximate by dividing the voltage swing by 2. This gives:

$$\begin{aligned} P &= i v \\ &= 2C_L V_S f_o \frac{V_S}{2} \\ &= 2C_L V_S^2 f_o \end{aligned}$$

This means that for a 100-output device (PLD or any other device) with each output driving 35 pF loads, where the output swing is 3 V and the output frequency is 50 MHz, the power dissipation contributed only by the load will be about 1.6 W regardless of the power dissipation of the chip itself.

I_{CC} vs Frequency

The operating current increases with frequency for standard CMOS devices. The difference is the current at low frequencies. A standard device typically can draw 35 mA at 0 MHz. Figure 14 shows a typical curve for standard devices.



16507C-024


Figure 14. I_{CC} vs Frequency, Standard Device

I_{CC} vs Number of Product-Terms

The number of product terms switching can sometimes affect I_{CC}. On standard devices, however, the design of the particular sense amp determines whether the I_{CC} will increase or decrease with more product terms. Therefore, it cannot generally be predicted. From a practical standpoint, the change in I_{CC} due to different numbers of product terms is negligible.

I_{CC} vs Temperature

The amount of current drawn by a device depends on how much current can pass through the transistors. Simplistically speaking, the channel of a transistor can be modeled as a resistor. The resistance is affected by temperature since this affects the mobility of electrons. Electrons are less mobile in a hot device. The hotter the device, the more the molecules will vibrate and the harder it is for electrons to pass through without a collision with a molecule. This means that the resistance of the channel is higher, which in turn means that the device conducts less current. Therefore, I_{CC} is greatest when the device is cold, and is minimized when the device is hot. A typical curve is shown in Figure 15. This curve has been generalized by normalizing the current to the room temperature current.



16507C-025


Figure 15. I_{CC} vs Temperature, Normalized to Room Temperature

I_{CC} vs V_{CC}

The variation of I_{CC} with changes in V_{CC} should come as no surprise; as V_{CC} increases, so does I_{CC} . This means that the power consumption actually increases roughly as the square of V_{CC} , since power consumption can be expressed as:

$$P = V_{CC} I_{CC} = \frac{V_{CC}^2}{R_{eff}}$$

where R_{eff} is defined as V_{CC}/I_{CC} . This is a simplification, of course, since R_{eff} is non-linear, and varies with V_{CC} . A typical I_{CC} vs V_{CC} curve is shown in Figure 16.



16507C-026

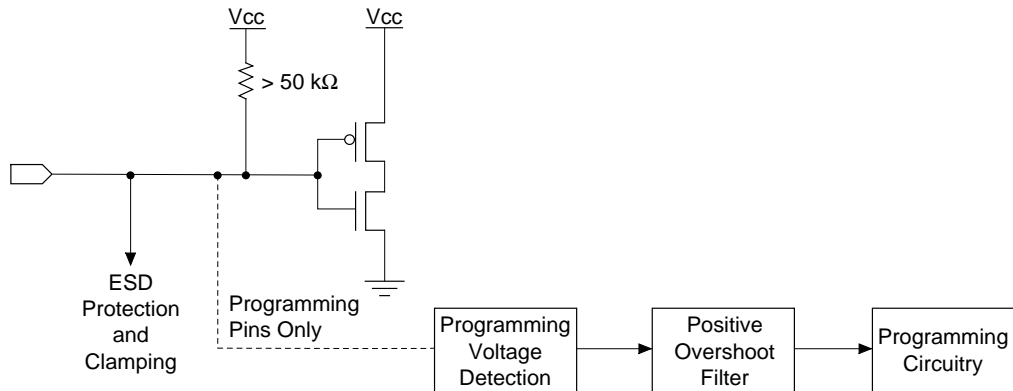
Figure 16. I_{CC} vs V_{CC}



Input/Output Structures


The basic input and input/output structures are shown in Figure 17. The ESD circuits and the programming voltage detection circuits will be discussed in more detail later.

Newer devices have pull-up resistors as shown below. In these devices, there is also a transistor in series with the resistor, which reduces die size and improves testability of the I/O circuit.



a. Input with pull-up resistor and overshoot filter

16507C-027



b. Output with pull-up resistor

16507C-028


Figure 17. Equivalent input/output schematics



I-V Curves

Figure 18 shows a typical I-V curve for an input buffer. Within the range of normal input signals, the input buffer has extremely high impedance, with diodes and MOS transistors that turn on when the input is below ground. On higher speed devices, this has the effect of a high-speed diode capable of clamping negative overshoot on noisy signals.

Since the input is effectively a capacitor, the impedance has no real component; the imaginary portion falls with increasing frequency. A typical device has an input capacitance of 8 pF at 1 MHz. Assuming a capacitance around 8 pF at higher frequencies, this yields a capacitive reactance of 2.5 K Ω at 50 MHz.




16507C-029

Figure 18. I-V Curve for an Input with No Pull-Up Resistor




Figure 19 shows typical I-V curves for high and low TTL-style outputs. The impedance of a low output is about $10\ \Omega$; a high output has an impedance of about $30\ \Omega$. The fact that the impedances are somewhat more symmetric than those found on a bipolar device makes it easier to terminate long traces accurately.



a. Output LOW

16507C-030




b. Output HIGH

16507C-031

Figure 19. I-V Curves for a TTL-Style Output with No Pull-Up Resistor




Figure 20 shows the curves for rail-to-rail switching outputs. The p-channel impedance, when the output is HIGH, ranges from 200 Ω when extremely heavily loaded to about 50 Ω when lightly loaded. The n-channel impedance is lower, at about 10 Ω .



a. Output LOW

16507C-032



b. Output HIGH

16507C-033

Figure 20. Curves for a CMOS-Style Output

General Information



Output Drive vs Temperature and V_{CC}

The output drive varies with temperature just as I_{CC} does. As the temperature increases, electron mobility decreases, cutting the drive. Likewise, the drive increases as temperature decreases. For example, at 75° C, I_{OL} decreases by about 18% from its room temperature value; I_{OH} decreases by about 7%.

The drive also varies directly with V_{CC} , although the effect is most pronounced on I_{OH} ; it increases by about 18% when taken from 5.0 V to 5.25 V. Because a low output transistor is already ON hard, the little extra bit of drive that its gate gets as V_{CC} goes to 5.25 V only increases I_{OL} by about 3%.

There is no explicit current-limiting resistor on the pull-up. The resistance of the pull-up channel limits the current. The fact that this resistance is smaller than what one might find in a bipolar device contributes to the more symmetric impedances, but also gives a higher short-circuit current I_{SC} . The slew-rate-limiting circuit also limits the drive; slew-rate-limiting is discussed below.

AC Parameters

AC parameters vary with a number of conditions. The data sheet specifications designate one set of conditions that act as a benchmark for confirming the guaranteed performance, but as the application changes the conditions, the actual system performance may change.

t_{PD} vs Temperature

Propagation delays decrease (that is, they speed up) at colder temperatures for the same reasons that I_{CC} increases. In general, devices at 0°C operate about 15% faster than those at 75°C.

t_{PD} vs V_{CC}

As V_{CC} is increased, more power is available, and the device can operate faster. However, the effect is less pronounced than with temperature. A device operating with a 5.25 V supply runs about 4% faster than one running with a 4.75 V supply.

t_{PD} vs Loading

The t_{PD} increases as the device load increases, although much of this results from the increase in rise and fall times of the outputs. For every 50 pF change in load, roughly a 2- to 5-ns change in the rise and fall time can be expected. In addition, as the load increases, more transient current is switched, creating more internal noise. This can slow the speed path inside the chip.


Power-Up Reset

Power-up reset is a feature that forces a device to power up into a known state. Without this feature, the power-up state is not known. Power-up reset helps make system initialization and testing simpler.

The ramp rate of V_{CC} is not critical to the power-up reset function. However, there are two other requirements: the supply ramp must be monotonic, and the clock must be suppressed until power-up is complete.

The monotonicity requirement basically says that there should be no low-going glitches in the power-up ramp (Figure 21). The danger in such glitches is that if the timing and voltage are just right, the registers themselves may think that the device powered down temporarily, causing them to lose their state. If the glitch is fast enough, however, the power-up reset circuit may not

notice the glitch and may think that everything is proceeding just fine. At the end, the registers may be in a random state. Even if the power glitches low enough for long enough time to shut down all circuits, the power-up timing must be restarted from the end of the glitch.




16507C-034

Figure 21. Non-Monotonic Power-Up Can Cause Power-Up Reset To Fail

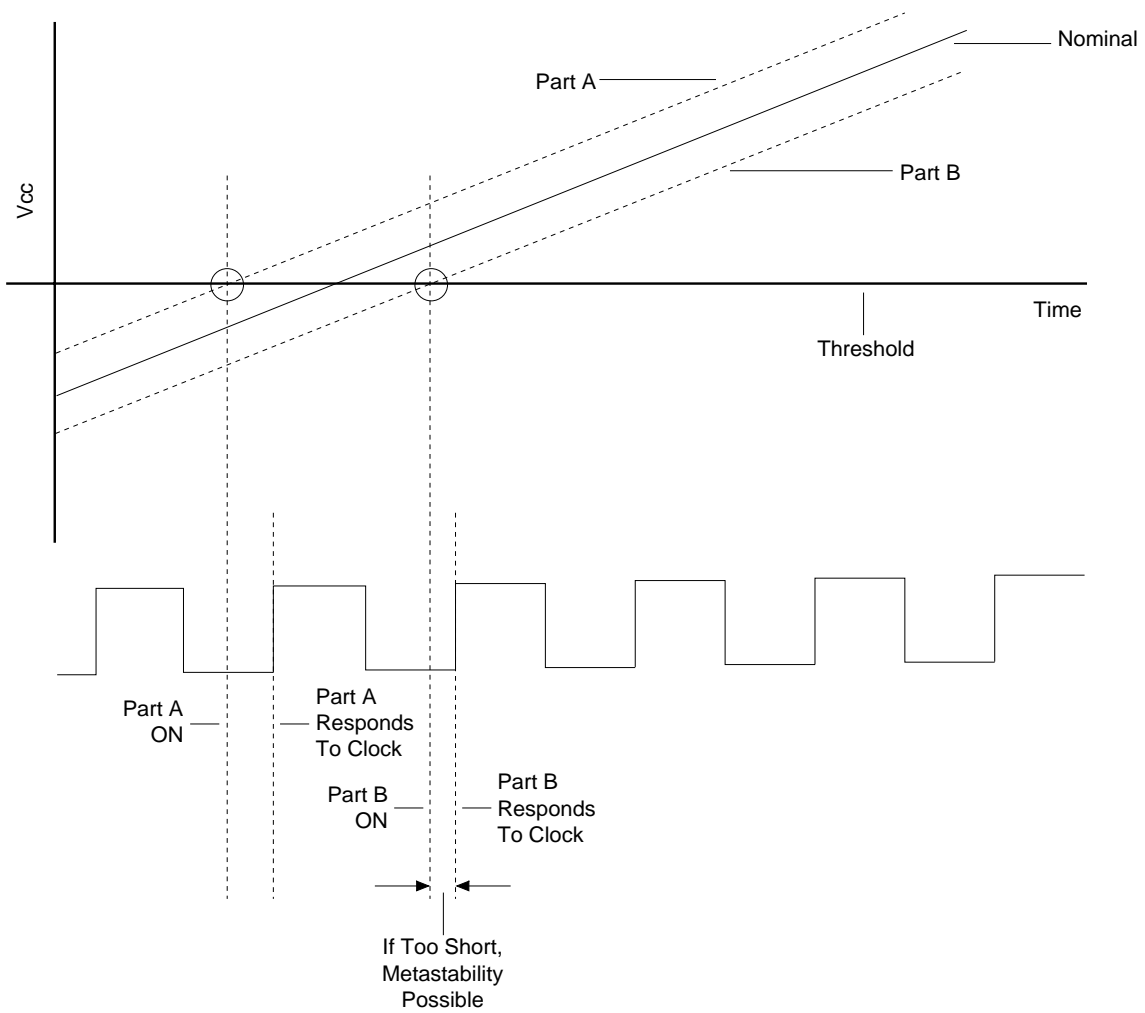
There is also a requirement that the clock not be running during power-up (Figure 22a). If the clock is running while the device is powered up, then as different parts of the device—and, indeed, the whole circuit board—turn on, parts of a single device, or different devices, may be out of synchronization with each other (Figure 22b). At some point, a part of a device will be ON enough to start recognizing the clock. It will then start to sequence as per the inputs it sees. If the inputs are not stable, the sequence may not be correct. In addition, if not all parts of the circuit or board recognize the clock at exactly the same time, some parts will start cycling before others, and the whole system will be out of synchronization.

The other potential (although remote) problem with clocking during power-up is metastability. If a register powers ON in time to see the clock edge, its setup time might have been violated, making the results at the output unpredictable.



a. Correct operation

16507C-035



b. Free-running clock places part B one clock cycle out of sync with part A

*2.7 V for Low-Voltage MACH

16507C-036

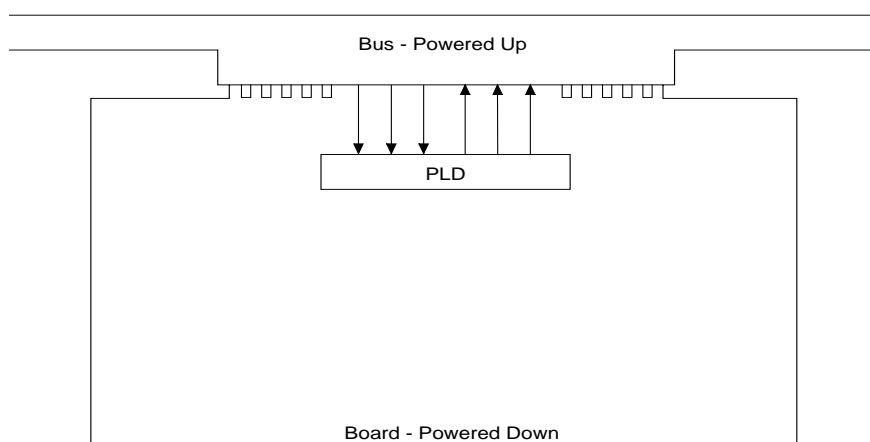
Figure 22. Clocking During Power-Up Reset

Powered-Down Characteristics

Some applications place the CMOS PLDs in a situation where it is itself powered down, but it is driving or is driven by other devices that are still powered up. This is especially typical of devices which talk directly to a bus (Figure 23).


The characteristics of the device in such a condition depend on how the power was removed. There are two ways of removing power:

- ◆ Opening up the V_{CC} line (e.g., if V_{CC} is fused, and the fuse blows; Figure 24)
- ◆ Grounding V_{CC} (Figure 25)




16507C-037

Figure 23. Powered-Down Device with Active Inputs and Outputs



16507C-038

Figure 24. Power Down with V_{CC} Open




16507C-039

Figure 25. Power Down with V_{CC} Grounded




It is important to know whether, for a given device, there is some kind of path from the pin to V_{CC} when V_{CC} is lower than the pin voltage. If any current can flow, it is not necessarily catastrophic, but there can be some effect. If V_{CC} is grounded, then there is a direct path to ground for any current flowing from the pin to V_{CC} (Figure 26a). If V_{CC} is open, then the only path from V_{CC} to ground is through the device itself, and through the V_{CC} lines of any other devices on the same V_{CC} line (Figure 26b). In the latter case, the pin is essentially powering up the device(s) itself; realistically, it cannot provide enough power to drive the chip, and this could result in the pin being loaded down.



a. V_{CC} grounded

16507C-040




b. V_{CC} open


16507C-041

Figure 26. Powered-Down Current Paths with P-Channel Pull-Up

Most of Vantis' CMOS PLDs have no such path when powered down. Figures 27 and 28 show the I-V curves of inputs and I/O pins while V_{CC} is open and V_{CC} is grounded. Figure 27 is for TTL-compatible devices which have n-channel pull-ups on the outputs. Figure 28 is for the HC/HCT-compatible zero-power devices and others which have p-channel pull-ups.




16507C-042




16507C-043

Figure 27. Power-Down Characteristics of TTL-style CMOS Inputs and Outputs



16507C-044




16507C-045

Figure 28. Power-Down Characteristics of CMOS-Style Output

Note that for most of the TTL-compatible devices, there is no leakage on the pins. This means that signals on a pin are not affected by the powered-down device. Therefore, it can be safely connected to an active bus. It also allows for safe **hot insertion or “hot socketing,”** where the device (or the board that contains the device) is plugged into a socket that has V_{CC} applied.

As a result of one of the ESD structures (which are discussed below), some devices do conduct some current when V_{CC} is powered down (Figure 28b). Newer devices do not have this characteristic. This is described in the following ESD section.

With the HC/HCT-compatible devices, the input structures are the same as for TTL devices, but the outputs conduct because of the p-channel pull-up. There is a parasitic diode between the output and V_{CC} (Figure 29). This can cause latch-up if the output voltage is higher than V_{CC} . Thus, it is not recommended that devices with p-channel outputs be directly connected to a bus if the device will be powered down while the bus is active. **Hot insertion** of these devices should also be avoided.



16507C-046

Figure 29. Parasitic Diode in CMOS-Style Outputs

DEVICE INTEGRITY AND ROBUSTNESS

The reliability of Vantis' CMOS processes is documented in product and process qualifications. For EE process products, generally the extended life FIT rate is under 100. The rate for the devices with 2000 hours burn-in is around 30; similar devices with only 1000 hours burn-in have a FIT rate closer to 100. With more burn-in experience, the FIT rate will decline even further due to the statistics used to calculate FIT rates. The FIT rate calculation is such that with fewer burn-in hours, a lower confidence factor is applied, giving higher FIT rates on newer products even when there are no failures.

ESD

Every pin on the devices is protected against electro-static discharge (ESD), a formal name for static electricity shocks. Output pins rely on the large output drivers as protection. Inputs normally do not have large drivers, so a circuit must be added for input protection. These input protection circuits also provide clamping against negative overshoot.

All new devices make use of the structures in Figures 30a and 30c for ESD protection. Most input pins use the circuit in Figure 30a. On pins requiring high voltages, the circuit has been modified as shown in Figure 30c. Some older devices have the configuration shown in Figure 30b. Because the active pull-down transistor is not ON when V_{CC} is disconnected, it cannot necessarily hold off the ESD transistors; this causes the current seen in Figure 27b. This circuit is no longer being used in new devices.

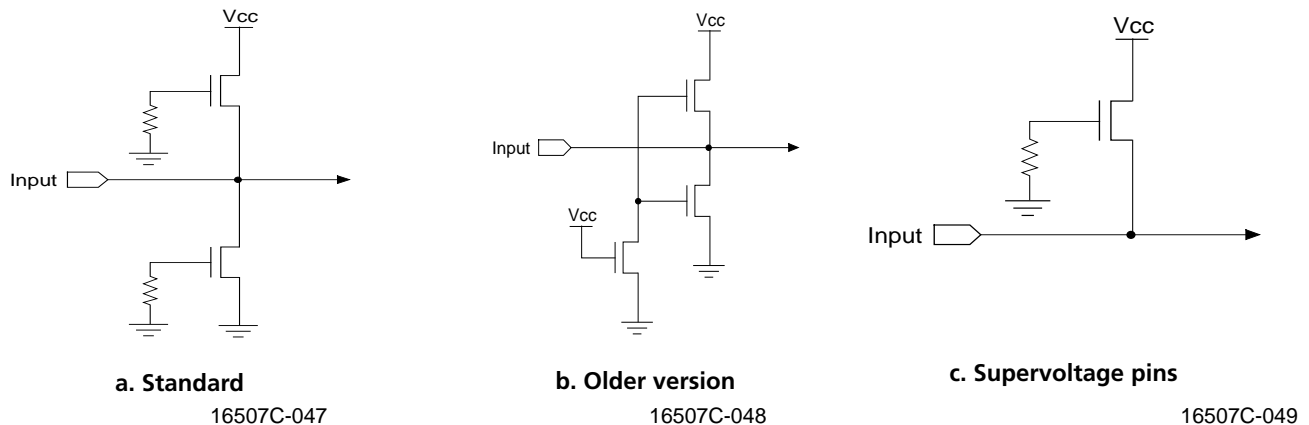


Figure 30. ESD Protection

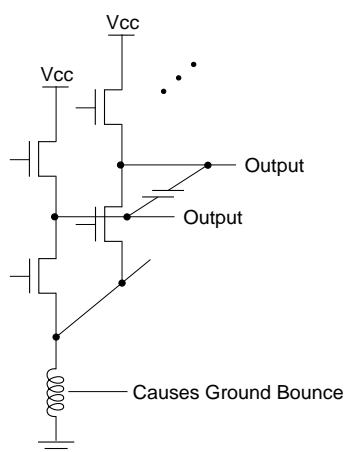
NOISE GENERATION AND SENSITIVITY

Vantis' CMOS PLDs are designed with noise concerns in mind. This affects both the amount of noise generated by the devices and the way in which the devices react to externally-generated noise. As more is understood about the nature of system-level noise, new design techniques are being used to make the devices quieter and more robust.

Ground Bounce

Ground bounce occurs when many outputs simultaneously switch from HIGH to LOW. This occurs because of the fact that CMOS devices generally have outputs that switch very quickly. If left uncontrolled, ground bounce can make a device with many outputs unusable.


Ground bounce is generated by the natural parasitic inductance in the ground lead (see Figure 31). When a large current surge goes through the inductor, the high di/dt induces a voltage that puts the ground level on the chip at a higher voltage than the ground level seen on the board.



16507C-050

Figure 31. Origins of Ground Bounce

Any output that is at a static LOW level maintains a V_{OL} with respect to the chip ground. If the chip ground is bouncing with respect to the board ground, the LOW output will track the moving chip ground and will also appear to bounce (see Figure 32). This is sometimes seen as a glitch by the next device. Even if there is no output glitch, instances of high ground bounce can slow the performance of the internal circuits by temporarily starving them of power. In extreme cases, this can interrupt the internal circuits.




16507C-051

Figure 32. Symptoms of Ground Bounce

Excess ground bounce can be handled in two ways: by limiting the amount of ground inductance and by reducing the di/dt . Inductance can be reduced by improving the configuration of the ground pin.

Ground bounce is also controlled by limiting the slew rate of all the output drivers (see Figure 33). This slows down the fall time and reduces the rate of current change by as much as 25%.



16507C-052


Figure 33. Output Drivers with Slew-Rate Control

Overshoot Sensitivity

Overshoot is a form of noise usually generated when signal traces act as transmission lines but have not been adequately terminated. The resulting reflections can cause significant overshoot, with as much as double the intended swing applied to the input in the negative or positive direction.

Negative Overshoot


Negative overshoot (Figure 34) poses no problems for a device that has been carefully designed. There is no detrimental effect as long as no unexpected parasitic behavior occurs due to the fact that ground is no longer the most negative voltage. However, the ringing that usually follows overshoot can slow down system performance, since the system has to wait for the ringing to subside.



16507C-053


Figure 34. Definition of Negative Overshoot and Undershoot

Clamp diodes are useful for stealing the energy present in the ringing and cutting the ringing short. A fast clamp reacts to the overshoot as it occurs, cuts the amplitude of the overshoot, and reduces or eliminates ringing. Figure 35 shows the ESD protection circuit used on most input pins. Parasitic p-n junction diodes exist between the substrate and the n-type source and drain, although these diodes are relatively slow. Faster reaction is provided by the n-channel devices themselves. When the input is too negative, the gate-to-drain voltage is positive. If the drain is more negative than the threshold voltage, the transistors turn on in the reverse direction, with the drains acting as a sources. This happens very quickly and acts as a clamp. This will also happen on an I/O pin, with the low output driver acting as the clamp.



a. Circuit diagram

16507C-054




b. Cross-section

16507C-055

Figure 35. Negative Overshoot Clamping

While it might appear that parts with negative substrate bias can “tolerate” more negative overshoot, it is really more accurate to say that these parts **allow** more negative overshoot, since there is no clamping. If there are effective input clamps, which are possible with a grounded substrate, then it will look like the part never gets as much negative overshoot. This does not mean it cannot handle the overshoot; it means that it is clamping the overshoot. If you take the part out of the socket, you will see that when unclamped, the overshoot will increase dramatically, as illustrated in Figure 36. Since Vantis’ devices have a grounded substrate, they are inherently better equipped to handle negative overshoot.




Ch. 1 = 2.000 volts/div

Timebase = 50.0 ns/div

a. Signal driving empty socket

16507C-056



Ch. 1 = 2.000 volts/div

Timebase = 50.0 ns/div


b. Signal driving same socket with CMOS PLDs in it

16507C-057

Figure 36. The Effect of Clamping: a. Signal Driving Empty Socket

Positive Overshoot

Large amounts of positive overshoot (Figure 37) can be a problem on most PLDs, regardless of technology or vendor. This is because most PLDs are programmed using **supervoltages**, and the pins therefore have supervoltage detectors that turn on the programming or test circuits, and potentially disable parts of the normal operating circuitry.



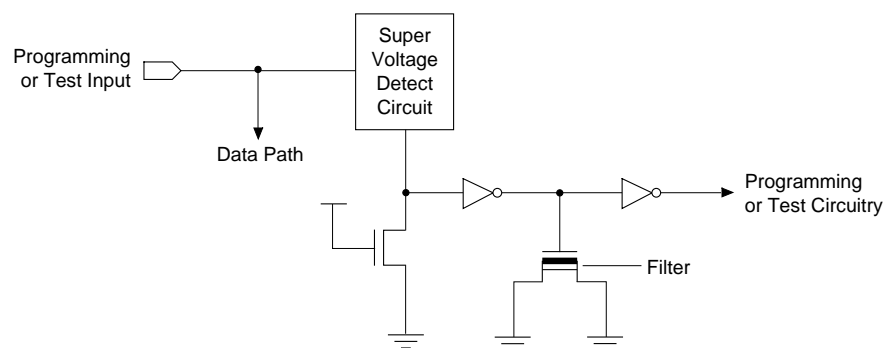
16507C-058

Figure 37. Definition of Positive Overshoot and Undershoot

If there is too much positive overshoot, the signal can travel into the programming voltage range, briefly activating the programming circuitry. This can result in functional interruptions, such as outputs momentarily starting to disable or going from HIGH to LOW.

For earlier devices, the problem can only be avoided by revising the design to reduce the overshoot. A particular design in a particular device may work, but this may be because that device has no supervoltage function on that particular pin. However, if an alternate source with different supervoltage pins is used, the design may not work.

New Vantis CMOS devices incorporate a filter, or delay circuit, that delays the reaction of the programming circuit for about 100 ns. This is enough to reject overshoot signals, which usually last for less than 30 ns. Positive overshoot will not cause any functional interruptions on devices with this protection (see Figure 38).



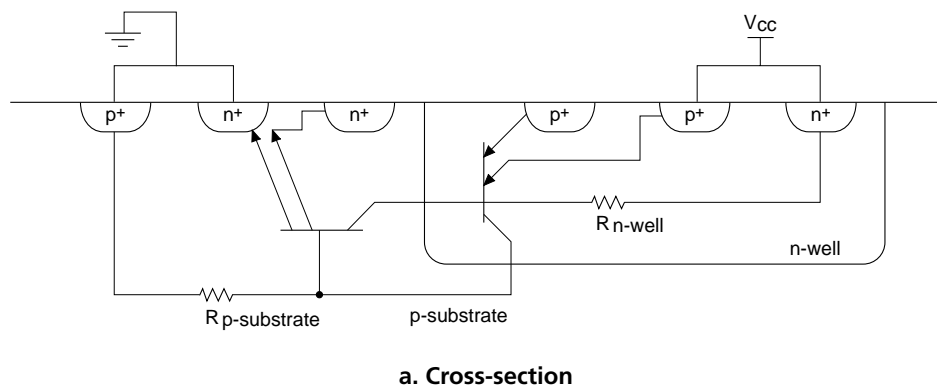
16507C-059

Figure 38. Positive Overshoot Filter


LATCH-UP

Latch-up occurs as a result of parasitic bipolar transistors between the n-channel and p-channel devices (see Figure 39a). These transistors form a parasitic SCR (see Figure 39b), which turns ON when triggered, conducting large amounts of current. It is usually impossible to shut OFF without removing all power from the device. The amount of current drawn is so high that it can either overload a power supply or, if the power supply can supply huge amounts of current, destroy the device.

Latch-up is normally triggered by an input or output at a voltage significantly above V_{CC} or below ground, with enough current drawn to cause the SCR to turn on. This condition usually occurs when hot socketing a vulnerable part; i.e., plugging a part into a powered-up board or inserting a board into a powered-up system. When this happens, the inputs and V_{CC} power up uncontrolled, and there is a risk of latch-up.



16507C-060



16507C-061

Figure 39. Latch-up Mechanism



TTL-compatible outputs are intrinsically less susceptible to latch-up, since they have no p-channel pull-up. This accounts for nearly all of Vantis' CMOS PLDs; these devices can be used for hot-insertion.

For true CMOS outputs, the SCR is an intrinsic part of the CMOS structure and cannot be eliminated. The SCR must be made as difficult as possible to turn ON by using guard rings and very carefully laying out input and output circuits. All of Vantis' CMOS devices are guaranteed to endure a current pulse of 100 mA into or out of the pin without inducing latch-up; most devices can actually withstand over 500 mA.

GROUND BOUNCE

Because CMOS devices generally have higher output slew rates, designs having many outputs switching at the same time (particularly if the outputs are heavily loaded) can cause more ground bounce than that generated by a comparable TTL device. It is important to use devices with output slew rate control.

The slew-rate-limiting circuits help minimize the occurrence of conversion problems, but even when the output slew rate is limited, the signal still can switch more quickly than that from a TTL output. If a design cannot be modified to accommodate the faster edge rates, this ground bounce may make a conversion unfeasible. If design changes are possible, any of the following can be tried:

- ◆ Limit the number of outputs that can switch at once.
- ◆ Reduce the loading on the outputs.
- ◆ Go to a lower-lead-inductance package (like a PLCC).
- ◆ Ensure that the ground path on the circuit board has low inductance.

OVERSHOOT

The other possible problem when converting from bipolar to CMOS is reaction to signal overshoot in a noisy system. This is only an issue if the CMOS device has no overshoot protection. Overshoot sensitivity is not specifically related to CMOS, but results from programming algorithms being different between the technologies. This also can occur when changing between bipolar vendors, or when changing between CMOS vendors. If the noise on a signal can disturb supervoltage circuitry, this can be troublesome.

Different devices have different sensitivities; this accounts for some of the apparent incompatibility. However, the culprit usually is the fact that supervoltages appear on different pins for different devices, and the supervoltage functions vary. Thus, overshoot on one pin of a particular bipolar device might have had no effect. Once that device is changed (whether to CMOS or any other device that has no overshoot filter), the new device might react to the overshoot and cause problems.

The solution is to ensure that all signals are clean and have minimal overshoot, making them compatible with any device. Signal noise reduction can be accomplished most effectively by controlling the impedance of the signal traces and terminating correctly. As an alternative, if the driving device has extremely fast edge rates, it can be replaced with a device that has better controlled output slew rates.



SUMMARY

By concentrating on the needs of CMOS PLD users, Vantis has developed industry-leading CMOS technology that can provide cost-effective PLDs of unequalled quality, reliability, and performance. Vantis provides value through:

- ◆ State-of-the-art fabs, for better control of quality, reliability, volume, and costs
- ◆ Electrical erasure, for higher quality and lower cost
- ◆ The highest performance available
- ◆ Robust technology that is quiet and yet tolerant of noise
- ◆ An extremely broad offering of products; low and high density, low and zero power

This information has detailed many of the aspects of the technology that make it superior to any alternatives.

