6.915 Digital Systems Engineering - 10/17/96

Signalling over Lossy Lines

1 Review: Lumped Signalling

There are two ways to transmit signals on-chip:

- Simple voltage-mode signalling, e.g. connecting logic gates together. Good for short distances when wires can be approximated as ideal.
- Transmission signalling, where a logic signal is carried by specialized transmitter and receiver circuits. Can implement reduced voltage swing to reduce power, and other improvements.

2 Implementing Reduced Swing Signalling

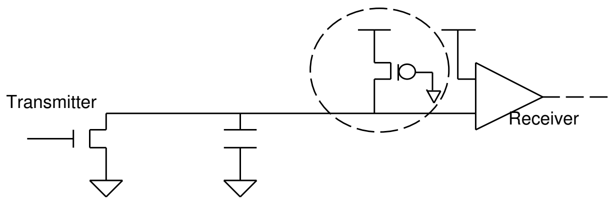

Figure 1: Possible implementation for reduced-swing signalling line

One way to implement reduced-swing signalling is to form a voltage divider including a pull-up (or pull-down) device, and the output device of the transmitter. A variety of circuits can be used for the pull-up network, including a feedback amplifier sensing the state of the receiver output, a resistor, or a MOS device in linear mode of operation.

Further improvements in performance can be made if more attention is given to voltage references. In the above example, the positive power supply was the voltage reference. An improved scheme is *differential signalling*, where a symmetric pair of wires are used for each signal. A signal and its complement are transmitted, which improves the voltage difference between signal and reference. If physical design is done carefully, the wires can be made symmetric, or in 'twisted-pairs' so that coupling noise becomes common-mode, and can be eliminated easily by the receiver.

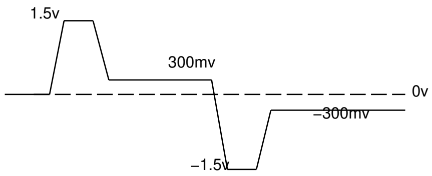

A compromise solution is pseudo-differential signalling, where a voltage reference is transmitted with every N signals. This is also known as *transmitter-generated reference*. The advantage is that it requires less wires per signal. The disadvantages are reduced voltage swing at the receiver, compared to differential; and that the wires and circuits are no longer symmetric, which implies less cancellation of noise.

3 Lossy Lines

Consider a wire as a series LRC circuit. The circuit will ring when there is energy in frequencies at or higher than the resonant frequency (1/sqrt(LC)). The value of Q = (sqrt(L/C) / R) can be used to estimate how long the ringing will persist. Ways to prevent ringing:

- Don't excite LC resonance; use slowly-changing signals that have little energy at or above resonant frequency.
- Reduce Q, so ringing stops quickly. This can be done by increasing output resistance of transmitter, a.k.a. using current sources.
- Add resistance at receiver, which is in series with capacitance of wire. This makes the RLC circuit even more lossy, to reduce resonant behavior.

Where does the L come from? There is significant inductance between a chip's internal power supplies and the rest of the system. On-chip power and ground lines also carry large currents, so there are V=Ldi/dt voltages.

Multi-level bit encodings can be used to reduce the power dissipated per bit transmitted. This is because of the reduction in average voltage swing per cycle. Challenges to this method are the sensitivity of the receiver, and the tradeoff of noise margins for more signal levels.

4 Driving Long Wires

Consider a long wire, with distributed resistance and capacitance. The propagation time for such a wire increases quadratically with length, for example: $t = 0.4l^2RC$, where 0.4 is an empirical factor, R and C are the resistance and capacitance per unit length of the wire.

This situation is prety grim, but there are ways reduce the propagation time over long wires:

- Use repeaters. Costs some area and power, but delay can become linear with length, instead of quadratic.
- Reduce capacitance by making wires wider. Resistance goes down, area capacitance goes up proportionally. But fringe capacitance as a percentage of total capacitance decreases.
- Reduce R while holding C constant.

The third item, implies the use of active circuits and signal-shaping. The RC line can be seen as a low-pass filter. So boost the high-frequency elements of a signal to compensate for the filtering action of the line. This leads to the idea of *overdriving*.

The high-frequency components of the signal are bosted greatly, for example by a factor of 5 in the figure. Thus, the voltage at the receiver end of the line will change much faster than when only +/-300mv drive is used. Issues with this scheme are: Increased crosstalk because of increased voltage swing, limitation of transmitter voltage range to power supply.

5 Fast Signals on Lossy Lines

Recall that a wire looks like a lossy LRC circuit, which attenuates high frequencies. At 1GHz, an on-chip wire might have 50% attenuation. We can transmit 300mV signals, which become 150mV at the receiver, which are easily detected by a receiver with 50mv sensitivity. Not quite.

The problem is that the line posesses hysteresis, or memory of bits/signals that it held previously. Consider a "lone pulse" scenerio, where a long string of 0's are followed by a single 1, then more 0's. The initial string of zeroes has an average DC value far below the threshold voltage. When the '1' is transmitted, it might not cause the receiver input to cross the threshold voltage.

The DC offset depending on previously sent data also introduces data-dependent jitter (timing uncertainty), because a signal will cross the threshold voltage at different times with respect to a clock signal, depending on the DC offset.

Ways to improve the situation include repeaters, passive or active circuits to equalize the line (boost high frequencies with respect to low frequencies).

6 Finally

Quick note: Voltage mode signalling wants a low impedance transmitter. Current mode signalling wants high impedance transmitter.

10/20/96 AC