TDT4295-Hardware Project course
Do's and Dont's

Konstantinos Antonakopoulos

Reviewer:
Marius Grannæs

Department of Computer and Information Science
Norwegian University of Science and Technology
Sem Sælandsvei 7-9, NO-7491, Trondheim, Norway
E-mail: kostas@idi.ntnu.no

September 21, 2009

Abstract

This document comprises the last lecture given for the TDT4295-Hardware Project course, of the Department of Computer and Information Science (IDI), Norwegian University of Science and Technology (NTNU), Trondheim, Norway. The course is being undertaken by students in their 4th year of undergraduate studies. The lecture describes some issues which students should pay special attention to, and refers to the design, construction, verification, and testing of PCB design along with some practical issues to make the project succesfull.

Contents

1 PCB **4**
- 1.1 Pull-ups/Pull-downs . . . 4
- 1.2 Decoupling capacitors . . . 4
  - 1.2.1 Decoupling . . . 4
  - 1.2.2 Switching subcircuits . . . 5
  - 1.2.3 Transient load decoupling . . . 5
  - 1.2.4 Placement . . . 5
- 1.3 Testpods . . . 5
- 1.4 Isolation, Jumpers . . . 6
- 1.5 XTAL/Resistor . . . 6
- 1.6 What can go wrong . . . 6
- 1.7 Footprints . . . 6
- 1.8 Test cards . . . 6
- 1.9 Voltage level conversion . . . 6
- 1.10 Overvolting . . . 7
- 1.11 Power lines . . . 7
- 1.12 PCB layers . . . 7
- 1.13 PCB cost . . . 7

2 FPGA **7**
- 2.1 Metastability . . . 7
- 2.2 Post route & route simulation . . . 9
- 2.3 High-speed busses . . . 9

3 AVR **9**
- 3.1 AVRFreaks . . . 9
- 3.2 Don't reinvent the wheel! . . . 9

4 Other issues **9**
- 4.1 Version control . . . 9
- 4.2 Latex format . . . 9

List of Figures

1 Crystal and resistor arrangement in an AVR 6
2 Example of Metastability output signals 8
3 Sample synchronization register chain 8

List of Tables

1 PCB

1.1 Pull-ups/Pull-downs

The main idea of using pull-up/down resistor is to avoid floating pins unconnected which can cause excessive power dissipation. Noise on input can cause the input stage to switch states at high frequency (oscillations). A pull-up resistor weakly "pulls" the voltage of the wire it's connected to towards 5V (or whatever voltage represents a logic "high"). However, the resistor is intentionally weak (high-resistance) enough that, if something else strongly pulls the wire toward 0V, the wire will go to 0V.

Similarly, pull-down resistors are used to hold the input to a zero (low) value when no other component is driving the input. They are used less often than pull-up resistors. Pull-down resistors can safely be used with CMOS logic gates because the inputs are voltage-controlled. TTL logic inputs that are left unconnected inherently float high, thus they require a much lower valued pull-down resistor to force the input low. This also consumes more current. For that reason, pull-up resistors are preferred in TTL circuits.

Pull-up resistors may be discrete devices mounted on the same circuit board as the logic devices. Many microcontrollers intended for embedded control applications have internal, programmable pull-up resistors for logic inputs so that minimal external components are needed.

Some disadvantages of pull-up resistors are the extra power consumed when current is drawn through the resistor, and the reduced speed of a pull-up compared to an active current source. Certain logic families are susceptible to power supply transients introduced into logic inputs through pull-up resistors, which may force the use of a separate filtered power source for the pull-ups.

1.2 Decoupling capacitors

A decoupling capacitor is a capacitor used to decouple one part of an electrical network (circuit) from another. Noise caused by other circuit elements is shunted through the capacitor, reducing the effect they have on the rest of the circuit.

An alternative name is bypass capacitor as it is used to bypass the power supply or other high impedance component of a circuit.

1.2.1 Decoupling

One common kind of decoupling is of a powered circuit from signals in the power supply. Sometimes, for various reasons, a power supply supplies an AC signal superimposed on the DC power line. Such a signal is often undesirable in the powered circuit. A decoupling capacitor can prevent the powered circuit from seeing that signal, thus decoupling it from that aspect of the power supply circuit.

Another kind of decoupling is stopping a portion of a circuit from being affected by switching that happens in another portion. Switching in subcircuit A may cause fluctuations in the power supply or other electrical lines, but you do not want subcircuit B, which has nothing to do with that switching, to be affected. A decoupling capacitor can decouple subcircuits A and B so that B doesn't see any effects of the switching.

To decouple a subcircuit from AC signals or voltage spikes on a power supply or other line, a bypass capacitor is often used. A bypass capacitor is used to shunt energy from those signals or transients past the subcircuit to be decoupled, right to the return path. For a power supply line, a bypass capacitor from the supply voltage line to the power supply return (neutral) would be used.

High frequencies and transient currents flow through a capacitor, in this case in preference to the harder path through the decoupled circuit, but DC cannot go through the capacitor, so continues on to the decoupled circuit.

1.2.2 Switching subcircuits

In a switching subcircuit switching noise must be suppressed. When a load is suddenly applied to a voltage source, the circuit tries to suddenly increase its current, but the inductance in the power supply line acts to oppose that increase. It opposes it by lowering the voltage the power line supplies. This is not just the voltage that the load in question sees, but the voltage that every other subcircuit that shares that power supply line sees. This is only temporary; the inductance ultimately loses the battle and the voltage comes back to normal. But even a temporary reduction in voltage can disturb other subcircuits.

To decouple other subcircuits from the effect of the sudden current demand, a decoupling capacitor can be placed between the supply voltage line and its reference (ground) next to the switched load. While the load is switched out, the capacitor charges up to full power supply voltage and otherwise does nothing. When the load is applied, the capacitor initially supplies demanded current. By the time the capacitor runs out of charge, the power supply line inductance cannot maintain the previous current, so the load can draw full current at normal voltage from the power supply (and the capacitor can recharge too). The voltage dip is reduced but not eliminated; i.e. the decoupling is not perfect.

1.2.3 Transient load decoupling

Transient load decoupling as described above is needed when there is a large load that gets switched quickly. The parasitic inductance in every (decoupling) capacitor may limit the suitable capacity and influence appropriate type if switching occurs very fast.

Logic circuits tend to do sudden switching (an ideal logic circuit would switch from low voltage to high voltage instantaneously, with no middle voltage ever observable). So logic circuit boards often have a decoupling capacitor close to each logic IC connected from each power supply connection to a nearby ground. These capacitors decouple every IC from every other IC in terms of supply voltage dips.

These capacitors are often placed at each power source as well as at each analog component in order to ensure that the supplies are as steady as possible. In these applications, the decoupling capacitors are often called bypass capacitors to indicate that they provide an alternate path for high-frequency signals that would otherwise cause the normally steady supplies to move. Those components that require quick injections of current can bypass the power supply by receiving the current from the nearby capacitor. Hence, the slower power supply connection is used to charge these capacitors, and the capacitors actually provide the large quantities of high-availability current.

1.2.4 Placement

A transient load decoupling capacitor should usually be placed as close as possible to the device requiring the decoupled signal. The goal is to minimize the amount of line inductance and series resistance between the decoupling capacitor and that device, and the longer the conductor between the capacitor and the device, the more inductance there is. A power supply decoupling bypass capacitor should be placed as close to the voltage/current source as possible. The idea is to minimize the line inductance and series resistance between the capacitor and the supplied devices. The guidelines for placing a high-speed decoupling capacitor on a multi-layer printed circuit board depend on whether the board has dedicated power distribution planes and how closely spaced those planes are.

1.3 Testpods

Do as many testpods as you can. With testpods you are able to analyze, test and debug any component using the logic analyzer.

1.4 Isolation, Jumpers

In addition you should make bridges (use pins for desoldering and debugging the components. Use jumpers to connect them back).

1.5 XTAL/Resistor

In case the crystal won't start, a resistor placed next to the crystal, will give a small current and a head start for the crystal. So the crystal will probably start giving frequencies.

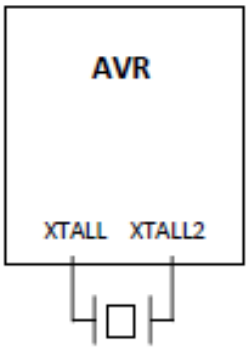

Figure 1. Crystal and resistor arrangement in an AVR

1.6 What can go wrong

Generally, you should always think of what can go wrong. Also, put as many pins as possible for debugging purposes.

1.7 Footprints

Almost always there is a problem with the footprints. A suggestion is to print out the footprints and check whether the parts match their pins on the printout. Also, check the ground and signal pins are in the correct orientation.

Attention: Don't let the PCB guys do that. They are probably already very tired!!! Let 2 other guys verify the footprints.

1.8 Test cards

Testcards are small cards representing the very same circuit found in the main board. Its a good idea to make as many small 2-layer designs to test the new ICs, if you don't have a clue how they work. This also enables easier software development.

1.9 Voltage level conversion

Check out the **latest** datasheets from the supplier, specially when it comes to voltages. You should verify that actually all parts should be able to receive and send current within the limits. Don't stress out the limits of any part, or don't expect the device to work efficiently on its limits for a long period of time.

1.10 Overvolting

In case of overvolting, you should consider ways to lower the voltage using resistors or use optoelectronic (i.e., with a photoresistor) packages.

1.11 Power lines

You should check that power lines are big enough to drive the load they're supposed to. If not, you may want to modify the lines' capacity throught the software.

1.12 PCB layers

In general, you should use 4 layers. The first and fourth layes are for routing, and the two inner layers become the power and ground layers.

1.13 PCB cost

PCB cost generally depends upon:

- Number of PCB layers
- Size of PCB
- Number of holes. The number of drill sizes should be reduced (2-3 drill sizes should suffice, max.4).
- Number of VIAS

Remember: Mounting holes don't need to be metal plated.

2 FPGA

2.1 Metastability

Metastability is a phenomenon that can cause system failure in digital devices, including FPGAs, when a signal is transferred between circuitry in unrelated or asynchronous clock domains. All registers in digital devices such as FPGAs have defined signal timing requirements that allow each register to correctly capture data at its inputs and produce an output signal. To ensure reliable operation, the input to a register must be stable for a minimum time before the clock edge (register setup time or t_{SU}) and for a minimum time after the clock edge (register hold time or t_H). The register output then is available after a specified clock-to-output delay (t_{CO}). If a data signal transition violates a registerŠs tSU or tH requirements, the output of the register may go into a metastable state. In a metastable state, the register output hovers at a value between the high and low states for some period of time, which means the output transition to a defined high or low state is delayed beyond the specified t_{CO}. In synchronous systems, the input signals must always meet the register timing requirements, so metastability does not occur. Metastability problems commonly occur when a signal is transferred between circuitry in unrelated or asynchronous clock domains. The designer cannot guarantee that the signal will meet t_{SU} and t_H requirements in this case, because the signal can arrive at any time relative to the destination clock. However, not every signal transition that violates a registerŠs t_{SU} or t_H results in a metastable output. The likelihood that a register enters a metastable state and the time required to return to a stable state vary depending on the process technology used to manufacture the device and on the operating conditions. In most cases, registers will quickly return to a stable

defined state. Figure 2 illustrates metastable signals. The input signal transitions from a low state to a high state while the clock signal transitions, violating a registerŠs tSU requirement. The data output signal examples start in the low state and go metastable, hovering between the high and low states. The signal output A resolves to the input dataŠs new logic 1 state, and output B returns to the data inputŠs original logic 0 state. In both cases, the output transition to a defined 1 or 0 state is delayed beyond the registerŠs specified t_{CO}.

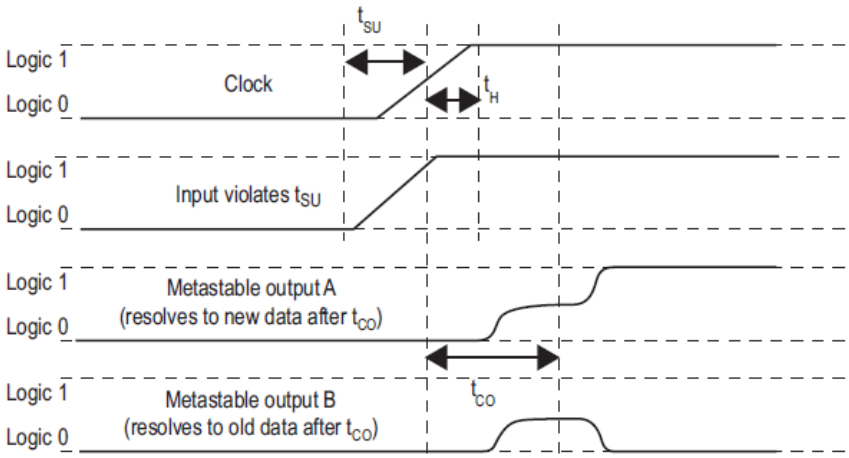

Figure 2. Example of Metastability output signals

To minimize the failures due to metastability in asynchronous signal transfers, circuit designers typically use a sequence of registers (a synchronization register chain or synchronizer) in the destination clock domain to resynchronize the signal to the new clock domain. These registers allow additional time for a potentially metastable signal to resolve to a known value before the signal is used in the rest of the design. The timing slack available in the synchronizer register-to-register paths is the time available for a metastable signal to settle, and is known as the available metastability settling time. A synchronization register chain, or synchronizer, is defined as a sequence of registers that meets the following requirements:

- The registers in the chain are all clocked by the same or phase-related clocks
- The first register in the chain is driven from an unrelated clock domain, or asynchronously
- Each register fans out to only one register, except the last register in the chain

Figure 3 shows a sample synchronization chain of length two, assuming the output signal feeds more than one register destination.

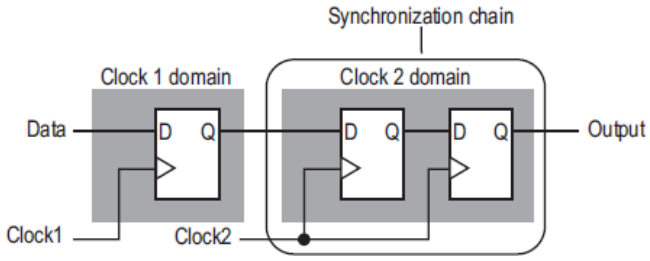

Figure 3. Sample synchronization register chain

2.2 Post route & route simulation

A way to solve the problem is to use post-place-and-route simulations in the simulation software.

2.3 High-speed busses

3 AVR

3.1 AVRFreaks

You can consult the AVRFreaks forum (*http://www.avrfreaks.net*) for acquiring as much information as you think would be benecifial for the project. Just remember if you acquire something, to give the right reference.

3.2 Don't reinvent the wheel!

You don't have to reinvent the wheel and consume your time when there is something you can find online. You don't get extra bonus but you get credit if you do it. Though, you should bear in mind that the project is very difficult and very time consuming itself. You need all the help you can get!

4 Other issues

Always document what you do, how you do and why you do it. This is not for the students to come and read your report in the future to get feedback on something, but is also very effective for your work, since there's high degree of difficulty and stress.

4.1 Version control

Its always a good idea to establish a version control system like SVN or CVS, in order to keep track of the changes you make.

4.2 Latex format

Its a good idea to work on multiple Latex files simultaneously, and then get a main file retrieve the part files. In that way, each group can work on its own pace and indepedently of the other groups (if that is possible).