

Power Management for FPGAs

By **Frederik Dostal**

Share on   

There have been many technical discussions about designing a good power management solution for an FPGA application, as it is not a trivial task. One aspect of this task involves finding the right solution and selecting the most suitable power management product, while another is how to optimize the actual solution for use with FPGAs.

Finding the Right Power Supply Solution

Finding the best possible solution to power FPGAs is not simple. Many vendors market certain products as suitable to power FPGAs. What makes the selection of dc-to-dc converters specific to powering FPGAs? Not much. Generally, all power converters can be used to power FPGAs. Recommendations for certain products are usually based on the fact that many FPGA applications require multiple voltage rails, such as for the FPGA core, the I/Os, and possibly an additional rail for DDR memory termination. Often PMICs (power management integrated circuits), where multiple dc-to-dc converters are all integrated into one single regulator chip, are preferred.

One popular way to find a good solution for powering a specific FPGA is to use pre-existing power management reference designs, which many FPGA vendors offer. This is a good starting point for an optimized design. However, modifications of such designs are often necessary, since a system with an FPGA usually requires additional voltage rails and loads that also need to be powered. Additions to the reference design are also often necessary. Another thing to consider is that the input power of FPGAs is not fixed. The input voltage heavily depends on the actual logic levels and the design that the FPGA is implementing. After completing modification to the power management reference design, it will look different from the reference design's original suggestion. One can argue that the best solution is to not even bother with a power management reference design, but to enter the required voltage rails and currents straight into a power management selection and optimization tool such as LTpowerCAD from Analog Devices.




Figure 1. LTpowerCAD tool to select the right dc-to-dc converters to power FPGAs.

LTpowerCAD can be used to come up with a power solution for individual voltage rails. It also offers a collection of reference designs, providing designers with a good starting point. LTpowerCAD can be downloaded for free from the Analog Devices website.

Once a power architecture and individual voltage converters have been selected, we need to choose suitable passive components and design the power supply. When doing this we need to keep the special load requirements of FPGAs in mind.

These are:

- ▶ Individual current requirements
- ▶ Voltage rail sequencing
- ▶ Monotonic rise of voltage rails
- ▶ Fast power transients
- ▶ Voltage accuracy

Individual Current Requirements

The actual current consumption of any FPGA depends heavily on the use case. Different clocking and different FPGA content requires different amounts of power. Because of this, the final power supply specification for a typical FPGA design is bound to change during the FPGA system design process. FPGA manufacturers supply power estimation tools that help to calculate the kind of power level that the solution will need. This information is quite useful to have before actual hardware is built. Still, the design of the FPGA needs to be final, or at least close to final, to get meaningful results with such power estimators.

Often, engineers design the power supply with the maximum FPGA current in mind. Then, if it turns out that the actual FPGA design requires less power, they scale down the power supply.

Voltage Rail Sequencing

Many FPGAs require different supply voltage rails to come up in a specific sequence. Often times the core voltage needs to be supplied before the I/O voltages come up. Otherwise some FPGAs will be damaged. To avoid this, the power supply needs to be sequenced in the correct order. Simple up-sequencing can easily be done by using enable pins on standard dc-to-dc converters. However, controlled down-sequencing is usually also required. It is difficult to achieve a good result when only enable pin sequencing is performed. A better solution is to use a PMIC with advanced integrated sequencing features, such as the [ADP5014](#). The special circuit block that enables adjustable up and reverse order down-sequencing is indicated in red in Figure 2.

Figure 3 shows the sequencing done with this device. The time delay for up- and down-sequencing can be easily adjusted with delay (DL) pins on the ADP5014.

If individual power supplies are used, an additional sequencing chip can take care of the required on/off sequencing. One example is the [LTC2924](#), which can control either the enable pins of dc-to-dc converters to turn on and off the power supplies or it can drive high-side N channel MOSFETs to attach and detach an FPGA to a certain voltage rail.




Figure 2. ADP5014 PMIC with integrated support for flexible up- and down-sequencing.




Figure 3. Start-up and shutdown sequencing of multiple FPGA supply voltages.

Monotonic Rise of Voltage Rails

Besides the voltage sequencing, a monotonic rise of the voltages during startup may also be necessary. This means that a voltage will only rise linearly, as shown by Voltage A in Figure 4. Voltage B in this plot shows an example of a voltage not rising monotonically. This can happen when the load starts pulling large currents at a certain voltage level during startup. One way to prevent this is to allow for a longer soft start of the power supply and to choose power converters that can quickly supply high amounts of current.




Figure 4. Voltage A rising monotonically, with Voltage B not rising monotonically.

Fast Power Transients

One other FPGA characteristic is that FPGAs very quickly start drawing high currents. They cause high load transients on the power supply. For this reason, many FPGAs require extensive input voltage decoupling. Ceramic capacitors are used very closely between the V_{CORE} and the GND pins of the device. Values up to 1 mF are quite common. Such high capacitance helps to reduce the demand on the power supplies to deliver very high peak currents. However, many switching regulators and LDOs have a maximum output capacitance specified. The input capacitance requirement of the FPGA can exceed the maximum allowed output capacitance of the power supply.

Power supplies do not like huge output capacitors since, during startup, this capacitor bank looks like a short circuit on the output to the switching regulator. There is a solution for this problem. A long soft start time can allow for the voltage on the large capacitor bank to come up reliably without the power supply to go into short-circuit current limit mode.




Figure 5. Input capacitor requirement of many FPGAs.

Another reason why some power converters do not like excessive output capacitance is that this capacitance value becomes part of the regulation loop. Converters with integrated loop compensation do not allow for excessive output capacitance to prevent loop instability of the regulator. Often there are ways to influence the control loop by using feedforward capacitance across the high-side feedback resistor, as shown in Figure 6.




Figure 6. Feed forward capacitor to allow for control loop adjustment when no loop compensation pin is available.

For the load transient and start-up behavior of a power supply, the development tool chain including LTpowerCAD and especially LTspice is very helpful. One effect that lends itself well to modeling and simulation is the decoupling of the large input capacitors of the FPGA from the output capacitors of the power supply. Figure 6 shows this concept. While the POL (point-of-load) power supply tends to be located close to the load, often there is some PCB trace between the power supply and the FPGA input capacitor. When there are multiple FPGA input capacitors next to each other on the board, the ones that are furthest away from the power supply will have a smaller effect in the power supplies' transfer function, since there is some resistance

but also parasitic trace inductance between them. These parasitic board inductances can allow for the input capacitance of an FPGA to be larger than the maximum limit of output capacitance of the power supply, even though all the capacitors are connected to the same node on the board. In LTspice, parasitic trace inductances can be added to the schematic and such effects can be modeled. Simulation results are close to reality when adequate parasitic components are included in the circuit modeling.




Figure 7. Parasitic decoupling between the power supply output capacitors and the FPGA input capacitors.

Voltage Accuracy

The voltage accuracy of an FPGA power supply usually needs to be quite high. A variation tolerance band of only 3% is quite common. For example, keeping a Stratix V core rail at 0.85 V within a 3% voltage accuracy window requires a complete tolerance band of only 25.5 mV. This small window includes voltage variation after load transients, as well as dc accuracy. Again, the available power supply tool chain including LTpowerCAD and LTspice is essential in the power design process for such strict requirements.

One last piece of advice is in regards to the selection of FPGA input capacitors. For them to quickly deliver large currents, ceramic capacitors are usually chosen. They work well for this purpose, but they need to be selected so that their true capacitance value does not drop with dc bias voltage. Some ceramic capacitors, especially the Y5U types, change their true capacitance value down to only 20% of the nominal face value when they are biased with a dc voltage close to their maximum rated dc voltage.

Frederik Dostal [frederik.dostal@analog.com] Frederik Dostal studied micro-electronics at the University of Erlangen-Nuremberg, Germany. Starting work in the power management business in 2001, he has been active in various applications positions, including four years in Phoenix, Arizona, working on switch-mode power supplies. He joined Analog Devices in 2009 and works as a power management technical expert for Europe.



Frederik Dostal

Also by this Author:

[Intermediate Voltage to Increase Power Conversion Efficiency](#)
Volume 52, Number 1