Best Practices for Motherboard / Embedded System Design

Annie Drees

Introduction

Turning an idea into a functional circuit board is a complex and time consuming process that can prove to be error-prone. The procedure of creating the schematics includes both the circuit design and the component selection. This is a lengthy process in itself. Then the board has to be laid out in CAD. Once this is done the board design is sent out to board manufacturers who will take additional time to build the boards and then deliver them to the designer. After all these steps are complete, the testing can finally begin. During the test phase, issues are discovered requiring hardware or software changes. The complexity of the schematic design and the process itself increases the potential for these bugs.

Any electrical changes required to fix a circuit board issue will hamper the delivery schedule. Not only will boards need to be modified or built, but the new design will have to be re-verified. Occasionally a bug can be fixed by replacing an existing component with another that is in the same package, which is generally referred to as a BOM or build of materials change. A BOM change is the least painful scenario as a new schematic and layout are avoided. Typically it is a very pointed change so verification on the new component is minimal. However, BOM changes can require new parts to be procured and depending on the part, this can add additional delay to acquire sufficient quantity. The worst case of all would be a circuit change, which would mean a new PCB and all the build and additional validation delays. Reworked boards have questionable health and reliability. New circuits also affect layout which could mean effecting circuitry unrelated to the original bug, extending the amount of validation needed to be rerun on the new boards.

Given that schematic design is complicated and mistakes are costly, the more systematic the design and test process the more reliably products will be delivered to schedule. This paper highlights three key practices: circuit reuse, design reviews and thorough testing that will develop a robust design and validation plan and will minimize schematic issues to insure a faster product delivery.

Circuit Reuse

Circuit reuse refers to using the same circuit across multiple designs. For example, a voltage regulator circuit could be used across different products with the same voltage and current requirements. Once a circuit has been implemented and fully validated in a design, it is considered "known good". Using this same circuit reduces the risk that would be generated from a new design. Another source of circuit reuse is another design that is undergoing testing. This may be beneficial in a different way from post-production reuse. Differences prompt conversation; is there a need for a difference? Concurrent testing will give further confidence in the robustness of the circuit. From the validation angle, using a common schematic will allow you to compare test results to diagnose a failure in a more meaningful way. If circuit board A exhibits a failure, but circuit board B does not, the engineer can quickly identify differences between the two boards. This can mean a quicker issue resolution. In general, reusing a circuit will speed up the schematic design. Reusing known good parts will minimize the test and evaluation time.

To enable circuit reuse, schematics must be available and readable. Previous generations of a design must be available to the current designers. A shared database for design schematics and layout will allow engineers to review content from others, minimizing unnecessary differences of design. Also, the schematics themselves should be well organized so that designers unfamiliar with the design can locate different subsystems quickly. For more elaborate schematics, a table of contents is highly recommended. Block diagrams with page designators and part numbers can also be used so that the design can be readable at a glance. These practices will facilitate the sharing of designs across multiple teams and continents.

Of course, a circuit should not be reused lightly. It must be examined thoroughly and evaluated for its usage in the new design. An IC used on a previous platform may not have the proper feature set for the new design. A new chip that operates at a lower voltage for lower power consumption shouldn't be used with the voltage regulator used on a previous higher power version. Reusing circuits can be done to a fault, too. Designs can reuse an IC generation after generation, using the part long after it makes sense. Silicon vendors are constantly redesigning parts to be smaller and consume less power. Avoiding a redesign now could eventually force one when a part reached obsolescence.

Reviews

Holding a peer review among engineers is beneficial on many levels. This review will catch human errors and schematic errors. It also encourages a sharing of knowledge and expertise across the engineering community. If a circuit is reused from a previous design, there is always a possibility that an issue came up post-production which required a schematic or layout change. Having all the engineers together talking through the schematic page by page facilitates these discussions and discoveries. In this way, issues can be identified and fixed prior to the board's being sent out to the manufacturer to be built, or gerber. A community effect where schematic designs are reused and reviewed as a team will have a collaborative effect that will foster a better design.

IC vendors typically review schematic pages with their parts or provide reference schematics. This should be seen as an opportunity for discussion with the vendor. For example, an IC reference schematic uses a specific oscillator to drive the IC. On discussions with the vendor, it might turn out that that was the only oscillator that had been validated with the IC, or perhaps there is a known issue with another oscillator. The IC vendors have a wealth of knowledge on the part and how it was verified. Sending schematics to them or questioning their reference schematics opens a dialogue that can highlight areas of risk in the circuit board design.

The productivity of a review can vary greatly. The following points will maximize the effectiveness of the meeting. The invite list should be extensive and inclusive. Managers who encourage attendance at design reviews from their people will create higher attendance. It is imperative to dwell on the circuitry and encourage questions and concerns. The more eyes on the design and discussion generated, the more likely potential bugs are to be identified. Publishing a roster and issues discovered at the review will keep the designers accountable for following up on any items.

Design reviews are an essential part of a successful design. Some engineers can be hesitant to open their designs to criticism, but even the best designer can benefit from a second opinion. A more comprehensive design reviews up front will catch a problem prior to gerber. Schematic reviews can be held in tandem with the board layout to avoid taking any additional schedule time.

Test Plans

The more thorough the test plan, the better the product. These test plans should be completed before the boards come back from the manufacturer. Fundamentally, a circuit board test plan includes two major categories: electrical analysis (or signal integrity), and functionality. It is not uncommon for emphasis to be put on the functional testing of the product, but both are needed to insure the design's robustness.

Electrical analysis focuses on the signal level electrical requirements of the design. The electrical requirements of any ICs must be evaluated to confirm accordance with the specification of the manufacturer. These parameters may include power up sequencing, maximum and minimum voltages on each pin, etc. Most circuit boards have at least one standard interface used to communicate between two ICs, such as a memory or I2C bus, etc. Any interface must be tested for adherence to the standardized spec. Rise and fall times, frequency measurements, voltage levels, etc. make up the signal integrity aspect of the test plan. A signaling failure in the interface might cause an issue at either end of the communication; both ICs are vulnerable. If the signal level specifications are not being followed, the behavior of the IC is unknown. It might intermittently fail; it might fail only in a particular scenario. Functional testing alone will most likely not catch a failure of the IC specification that will become a problem. Signal integrity testing is usually done before functional testing since it typically requires a circuit change to correct an issue, which means any functional testing would have to be rerun.

Functional testing focuses on what the circuit board is designed to do. Does it work as designed? Are all the features fully functional? There will be different corner cases to be run that will test the full functionality of the board. Every failure mode should be validated. These tests also need to consider the end customer and how they will use the circuit board. If the board will come to the customer in a specific mode of operation, it should be tested that way. Tests must be carefully written to address how the design will be used by its customer.

Over time, test plans are made more robust and thorough by incorporating past experiences. If a design has an issue post-production then the existing test plans need to be modified for the new designs. Identifying how the issue made it through the validation process is essential to avoid it in the future. Building this recursiveness into the test methodology should be done in a way that all circuit board designs will benefit wherever they are tested. Once a certain level of quality is in place, the addition of new features adds minimal churn to the test plan.

Beyond Electrical Engineering Test

Functional testing will identify issues at nominal conditions; the next hurdle is to validate the system operating within real world scenarios. Temperature, voltage, and frequency can vary from part to part within a range of expected behavior. Pre-production builds do not reflect the same material that will be used in production. Typically the circuit board will have some pre-production silicon on it that will not be finalized until close to production. Even if the silicon doesn't change prior to production, undoubtedly something will be different. Parts may be manufactured in a new location or even a new variation that develops between lots. Sensitivities within the silicon that are not manifested in typical conditions pre-production will suddenly assert themselves during production.

One of the problems with a circuit or silicon reliability issue is it frequently manifests itself in a singular way during pre-production. The circuit board may fail once and never again, for example. Or perhaps, a single unit will repeatedly fail; however, no additional units can be made to fail. This often leads to the failure being dismissed – that circuit board must have been damaged, or maybe the original failure was user error since it cannot be repeated. These issues should not be overlooked or dismissed lightly.

In Conclusion

A systematic schematic design and validation process will produce a robust design with fewer iterations. Although these ideas outlined in the paper are not novel they are easily overlooked. Embedding these practices in the process will create a strong foundation of design. Once a mature process is in place, issues will be found earlier in the process, some corrected before the schematics are completed. The next challenge is to develop a reliability test plan that will encompass the real world environment and insure functionality through the life of the product.