Trade-offs in the Use of Commercial Off-the-Shelf Intellectual Property for Backplane Interfaces

A. Feizi[1] and R. Katz[2]
[1]QSS Group, Lanham, MD 20706
[2]NASA Goddard Space Flight Center, Greenbelt, MD 20771

Abstract

The use of commercial intellectual property is the trend for recent spacecraft electronics design. Two recent spacecraft projects had a common backplane requirement and both were to be PCI compatible. PCI was selected by management because it is a common commercial standard and cores were available. One project decided early to build a custom backplane chipset while the other completed the project with a set of commercial PCI cores. The development time, effort, and technical considerations of these two parallel efforts are compared and lessons learned are presented.

I. INTRODUCTION

The PCI bus is a popular backplane bus present in all modern PC's. It is a high-speed, parallel interface which, in the case of these two projects, implemented the 5V, 33 MHz standard. The technology used for these spaceborne processors was the Act 3 A14100A, 0.8 um FPGA. Consequently, a full-speed implementation was not possible. Fortunately, system requirements resulted in slower transfer speeds being acceptable. One bus was designed for the EO-1 WARP solid state recorder while the other was for the embedded computer in the GLAS Main Electronics Unit (MEU), part of a space science instrument. Both applications required the ability to move large blocks of data. Additionally, and a lower priority, was moving small amounts of data for command and telemetry tasks.

Ultimately, the EO-1 application, which was on a very tight development schedule utilized a custom parallel bus, the K-2, running at 8 MHz with no wait states. The GLAS MEU implemented a subset of the PCI protocol, running at 12 MHz with one state. Both busses were 32 bits. Both busses provided interfaces to the same processor, the Mongoose V.

II. WARP K-2 BUS

The initial implementation utilized PCI cores and a two chip set was planned, one for the initiator and the other for the target. An initial target design was completed but a low cost, affordable, production-quality core could not be delivered by the vendor. Clearly, the capabilities of the PCI bus were not needed for this design and the chipset, with or without cores, was more complex than it needed to be. At the time, with the available space-flight qualified microcircuit technology, PCI compliance was not possible.

An analysis of the PCI bus for this application showed certain limitations. First, the complexity of the bus added to the logic resources needed and the power consumed. The protocol, while flexible and powerful, is overkill for a simple application. Certain aspects of the PCI bus were found to be inconsistent with hi-rel applications. For example, the protection diodes in the I/O cells were used to clamp overshoot. The A14100A drivers were not PCI compliant. The control of tri-state buses was distributed, enabling the possibility of bus fights if there was a loss of control as a result of an electrical transient or single event upset (SEU). Another concern was the tight requirements on clock skew, as all flip-flops were clocked with a single edge. Indeed, this could present problems when extender cards were. For a tight development schedule and no breadboards or engineering models, it was critical that the design work the first time. Therefore, the PCI bus was not used for this project and the K-2 bus, designed for space-flight use, was designed and built.

Without the restrictions imposed by commercial standards, the K-2 bus design was developed for high-reliability space-flight applications. For example, all signals, with the exception of the data bus, were made unidirectional and always on, eliminating concerns with bus fights and simplifying the protocol. The PCI specification requires that a signal, before it is tri-stated, first be driven high. Another major change was the clocking scheme. For this bus, opposite edge clocking was used eliminating clock skew as a concern. Control of all tri-state buffers on the data bus was embedded in a single sequencer in the master, eliminated the possibility of bus fights as a result of either electrical transients or SEUs. Parity was not included in this version of the K-2 bus but features such as bus time-outs is supported as well as transfer word counts, etc.

The performance of the K-2 bus, at 8 MHz, was satisfactory for the WARP application. For the lowest speed-grade parts, worst-case analysis showed that 16 MHz was not achievable, using the SEU-tolerant design approach. The limiting factor was the clock to output parameter. While the I/O module flip-flops in the Act 3 architectures offer high-speed, deterministic timing, they are SEU sensitive. For this application, since the speed was not required, all flip-flops are C-C constructs, which have a speed penalty when driven off chip. Higher performance could have been easily obtained by using the I/O module flip-flops but with increased SEU sensitivity. Since reliability was the concern, we chose safety over unneeded speed. For a next generation bus, the SX or SX-A architecture would be optimal and the design easily ported. This architecture is ideal for mapping the K2 logic design since there are no I/O module flip-flops. Additional considerations will be presented in the full paper.

The resources required for the developing this bus was minimal. Concept and design was accomplished in three man months. There was no debugging time as the chip set worked, without failure, upon the initial application of power.

III. MEU PCI BUS

The GLAS MEU decided to continue with the PCI-based approach. This project felt that the benefits of commercial-off-the-shelf standards and test equipment would be cost-effective. There was also additional time for the development of the MEU.

Some of the considerations for PCI are described in §2 above and will not be repeated here.

The use of off-the-shelf core required customization to adopt to the project unique requirements. The initiator was required to interface with a Mongoose V, a radiation-hardened processor compatible with R3000 code. The target was an interface to a four different interface cards.

The performance of the core was limited to 12 MHz with wait states required for data transfers, for each cycle. The design of the core would not allow faster operation. The target core also did not include DMA transfer capability.

The breadboard systems initially had problems with clock skew. Proper clock distribution on the flight initiator board was implemented and acceptable.

Additionally, the design was difficult to modify as a result of the lack of vendor documentation, one of the factors that led to the discontinue of PCI use for the WARP. The vendors code had to be examined, in many cases, to understand the behavior of the core. During development, which took approximately two years, was complicated by one of the key designers resigning from service. Without the personnel change, the task might have been completed in approximately one year. Most of the problems encountered were logic design errors. The target, which was used on a previous project and modified for the GLAS requirements, was operational after a few revisions of the FPGA design. The initiator, a new design, took many iterations to achieve working hardware.

IV. CONCLUSIONS

The use of COTS IP cores appear attractive and, according to the trade magazines and popular belief, are the way to rapidly design reliable hardware. We have found, in this case study, that that is not always the case when designing space qualified space-flight electronics on a tight schedule.

The limitations encountered are poor documentation, non-compliant hardware and microcircuit limitations, as well as complexity. The complexity makes customization difficult and limits performance. The PCI bus is designed for commercial applications. This is a concern for reliability with diodes clamping signal overshoot and a distributed control of tri-state drivers.