

A very High Performance Multi Purpose Computing Card for TM/TC and Control Systems

Markus Wittkamp¹, Anderson Zigiotta²

1: German Aerospace Centre, Mobile Rocket Base, Wessling, Germany

2: Instituto de Aeronáutica e Espaço – Divisão de Eletrônica, São José dos Campos, Brazil

Abstract: For future sounding rocket and balloon missions, a powerful processing system is needed for ambitious rate and attitude control algorithms as well as data acquisition for telemetry with high data rates (e.g. digital video).

The card comprises a dual core DSP together with a large FPGA to collect and process experiment data, build telemetry data streams and decode commands for the onboard experiment module network.

Due to the use of FPGA controlled periphery the Euro sized card offers the flexibility to meet the various requirements on sounding rockets and balloons.

This paper describes the requirements for on board computing in service modules and how a multi function card can solve them.

Keywords: control system, computing card, TM/TC, FPGA, DSP

1 Introduction

In order with increasing requirements of upcoming projects the presently used TM/TC and controlling card gets overloaded in many situations. In some projects several of these cards are needed simultaneously to satisfy the requirements of the actual projects. Now a state of the art processing system is to be developed as a replacement for the present system.

2 Requirements for a Computing Card

The most important requirement of a TM/TC card is to handle the communication between the experiments and their experimenters as safe and fast as possible. The data bandwidth needed for this tends to increase continuously. A new card should be able to generate and read analogue waveforms which are common to RF components like Biphase for telemetry or GMSK for telecommand without additional hardware.

Many experiments need special environmental conditions like low gravity, attitude pointing or navigation solutions.

The card should perform the necessary algorithms to setup these conditions. It should also drive appropriate actuators directly or indirectly such as solenoid valves of cold gas thruster systems or electric motors.

The service module is a kind of supervisor to the payload. Typical components of a service module are power supplies, RF systems, GPS receiver and various housekeeping- and other sensors. The module needs to control the environment for the experiments and monitor itself and the payload.

Flight on rockets or balloons requires electronics which can perform over a large temperature range and need to withstand high vibration and shock levels. The card should also fit in our existing electronic box concepts. This defines a Euro card with 5V power supply as well as a standard 96 pin connector.

A computing card for service modules should be as flexible as possible. Each project has different requirements regarding data formats and bandwidth, power supply and environmental conditions.

3 Concept of the Card

A way to meet the requirements is to design a multi function card (MFC) which is built from generic, programmable functions together with a powerful and flexible computing architecture. With this approach a MFC is usable for a wide spectrum of applications by adaptation of the FPGA design and the CPU firmware to the needs of the actual project.

The block diagram of the MFC concept is shown in Figure 1.

Processing Units

Our concept uses a symmetric multi processing (SMP) dual core Blackfin 561 DSP from ADI as CPU together with a large Cyclone 3 FPGA made by Altera.

By following the approach of generic functions, the FPGA controls all the peripheral and interface parts of the MFC which are as simple and universal as possible. This central device is shown in the middle of Figure 1.

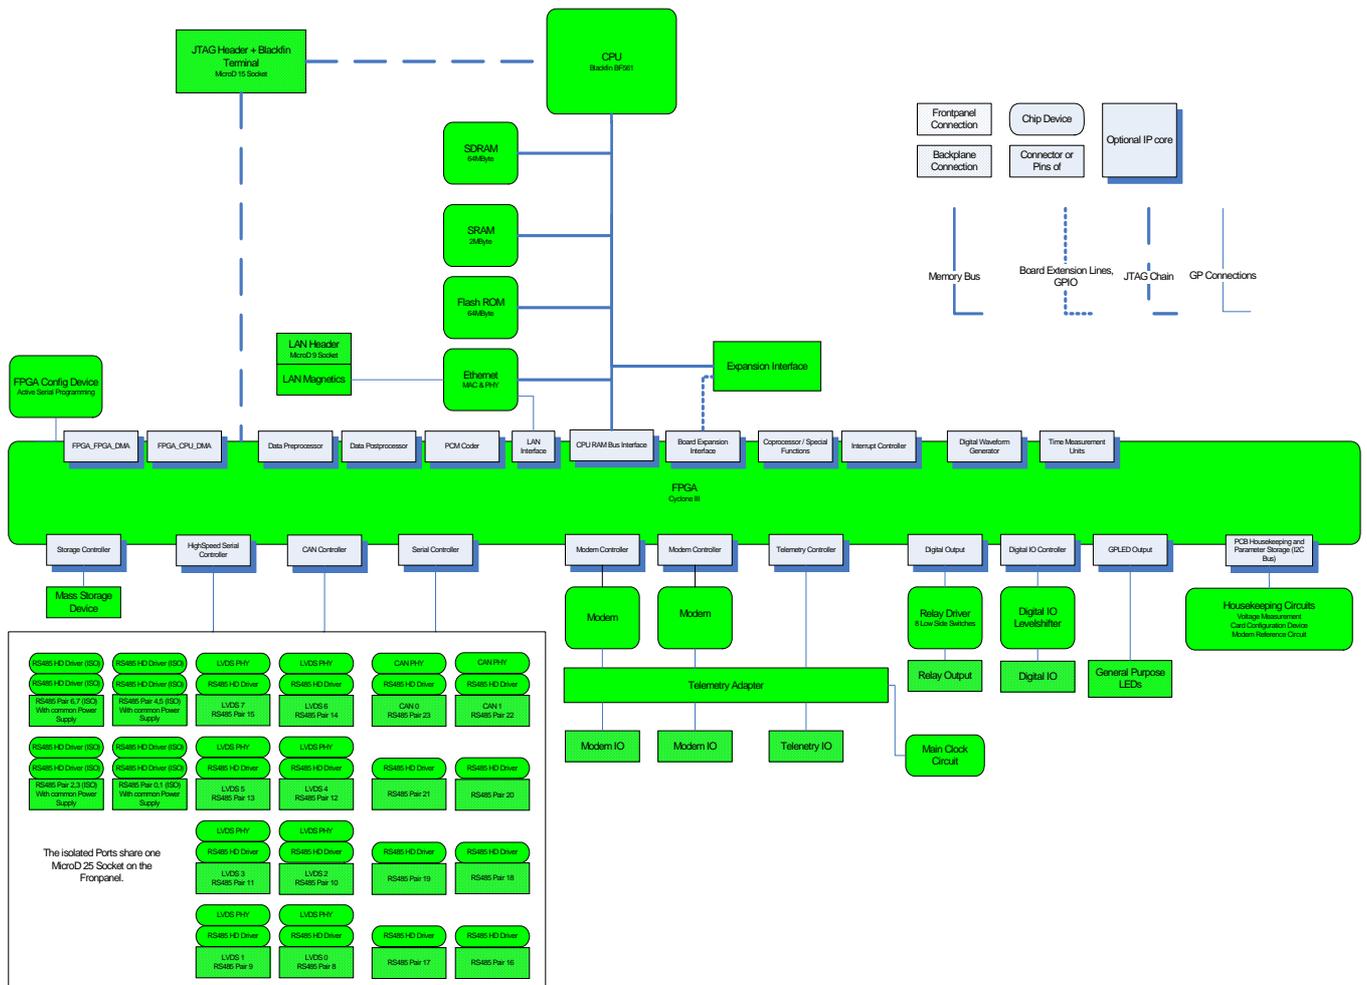


Figure 1: Block Diagram of the MFC Concept

The FPGA can be used to implement IP-Cores for more complex functions e.g. UART's or communication buses. The chip is large and fast enough to run additional processors, DMA controllers and pre- or postprocessors for the different types of data used by an experiment. The advantage of a large FPGA is the possibility to run many tasks in parallel.

The chip contains embedded multipliers to accelerate the calculation of mathematical functions. There are 20 global clock lines available and 4 PLL's for clock management and high-speed I/O. The MFC printed circuit board (PCB) layout is designed to fit with both top models, the EP3C80 and EP3C120 of the Cyclone 3 device family. The device family is produced using a 65nm process.

To estimate the logic usage of a complex state machine on the EP3C120 variant, a test design with a 32 Bit Altera NIOS 2 IP processor was made for example. The result was a usage of about 2% for logic which leaves enough space for running several complex tasks on this device in parallel. The test was compiled to run at 50 MHz.

The FPGA will boot from its own dedicated serial Flash ROM in less then one second. The exact boot time depends on the size of the compressed image stored in the configuration device. This device is larger then needed for

just booting the chip. The spare memory space can be used by the application to store additional data.

The CPU will be used for calculation intensive parts of the project with large data buffers like compression algorithms for video or other forms of digital signal processing. It will run with an real-time kernel for encapsulation of different processing threads and easy programming of the system. The chip will boot its firmware from a dedicated NOR Flash ROM which is connected to the main memory bus together with a 64 MByte SDRAM, 2 MByte asynchronous SRAM and an Ethernet controller. Main memory of the CPU is the SDRAM which can be controlled directly by the Blackfin hardware. The SRAM can be used to share large processing buffers between the main parts of the MFC. The Blackfin CPU generates chip select signals for several synchronous and asynchronous memory banks. Size of the asynchronous banks is 64 MByte and the synchronous ones can be up to 128 MByte. These signals are assigned to the memory chips and all of them are connected to the FPGA to map register files of IP-Cores to the CPU address space.

The memory bus is the backbone of the card. All data intensive components are connected to this bus. To extend the functionality of the card, an additional Expansion Interface is connected to the same bus. Expansion cards

can also share the bus and become part of the MFC structure. It's also possible to use the Interface for debugging purposes to monitor the bus and the running software. For further extension possibilities, 32 signal lines of the FPGA are also routed to the Expansion Interface connector.

The Blackfin needs a single clock source for generation of its core clock (CCLK) by an internal PLL using multiplication. The periphery of the chip runs with the system clock (SCLK) which is derived from CCLK by a divider. Both multiplier and divider can be set by software. The memory bus and all the periphery of the chip are running at SCLK. The theoretical bandwidth (BW) of the memory bus can be calculated as

$$BW = f_{SCLK} \cdot WL_{CPU} \quad (1)$$

With a CPU word length (WL) of 32 Bit and SCLK nominal speed of 133 MHz, the theoretical Bandwidth is 4256 Mbit/s. Due to bus sharing and wait states the real value will be much smaller.

Further, the Blackfin has two parallel interface ports for high speed data transfer to special chips. This is called the "Parallel Peripheral Interface". It is a synchronous, half duplex, 16 Bit parallel interface to connect with e.g. A/D converter chips which support this interface. At the MFC, it is connected to the Expansion Interface Connector.

Another feature of the CPU is to have two universal serial ports which are programmable to operate with various synchronous and asynchronous serial communication protocols, with up to half the speed of the system clock.

There are several internal cache memories, 100 kByte assigned to each core and 128 kByte of shared memory for both cores. Depending on the selected chip derivate, the cores can be clocked with up to 600 MHz. Several MMU and DMA controllers support the cores with data transportation. The chip is able to control its own core voltage with a buck converter design with a few external components. With this feature and several run modes, it is possible to influence the total power consumption of the chip.

Debugging

It is possible to load firmware to the MFC and debug the card when it is plugged into an electronic box. The JTAG chains of both Blackfin and FPGA are connected to a dedicated debug header. The single UART of the Blackfin is also connected to this header. The firmware can use this function to transmit status data or an implementation of terminal software to log in. The Flash ROM of the card can be programmed indirectly by the CPU. Therefore a loader firmware is necessary to get "write access" to the chip. In principle it is also possible to use this technique to write data during runtime.

Four LED's with different colours can be programmed to show status information during the test phase of a project. They are visible at the front panel using light pipes.

Mass Storage

In case of loosing the TM link or just as backup, data can be stored on the MFC to a mass storage device. This would be a NAND flash based device like e.g. a Multi Media Card. The host controller to this device resides inside the FPGA fabric like all other peripheral controllers.

Serial Interfaces

Many experiments are using serial communication like asynchronous RS422 or synchronous RS485. These and other standards use the same physical drivers. They can be built up from generic half duplex line driver chips together with project dependent logic inside the FPGA. In this way the card is able to implement all communication standards using RS422 levels by just exchanging the FPGA logic and setting some control signals of the drivers. The selected devices are able to drive data rates up to 30 MBit/s in a 100 Ohm system and can be terminated according to the project needs. Other termination levels are also possible. The outgoing pins of these 16 drivers are located on the backplane of the MFC for short connection to other cards inside the electronic box or the service module.

The eight isolated serial drivers of the design use their own independent power supply connections. Groups of two share one of these supplies because most communication standards need at least two drivers. To be more flexible it is also possible to supply the chips with internal power by setting a jumper. The isolated drivers work up to 35 MBit/s. They are placed at the front panel of the MFC connected to a 25 way Micro D socket.

For even higher data rates up to 200 MBit/s, eight RS422 drivers are sharing pins on the back plane with LVDS type drivers. LVDS driver chips working in the same way as the RS422 ones but at higher data rates with a smaller signalling level. As above, the FPGA controls the behaviour of the drivers – direction, enable, etc. – and the logic is project dependent. The RS422 drivers are disabled in case of LVDS usage and vice versa. Another advantage of LVDS type signalling compared with RS422, is less power consumption due to the lower signalling level.

The third type of driver implements the popular CAN interface. It also shares pins on the back plane with RS422 level drivers. If a project needs CAN, the respective CAN controller will be implemented as part of the FPGA fabric. Some resistors can be used for slew control to reduce EMI. The resistors are driven directly by FPGA. Inside the FPGA fabric the controller will be connected to a state machine or directly to the CPU bus for sending and receiving of messages on a higher processing level.

For more complex payloads or EGSE usage, an Ethernet interface is connected to a header mounted to the front panel of the MFC. There are vibration insensitive connectors available on the market so it is possible to use them with high data rates like 100 Base-T Ethernet during sounding rocket flights. We have chosen a 10/100 Base-T LAN9218i Ethernet controller chip made by SMSC. This chip uses a standard asynchronous memory bus interface

for communication with the CPU and several additional control lines connected to the FPGA. The memory interface of the Ethernet controller is part of the memory bus except of the chip select generation, which is done by the FPGA and is part of one of the FPGA internal memory banks. In this way the chip does not occupy an entire address bank of the CPU. The output signals of the LAN chip run through a special transformer chip for DC decoupling of the Ethernet signals. The Ethernet controller uses its own oscillator for clock generation.

Digital Interface

Some functions of a service module need to be switched by the MFC e.g. relays. A chip with eight low side switches can be used to switch currents of half an Ampere and voltages up to 50 Volts. An SPI bus with additional status lines is used for controlling and monitoring of the device. These high current signals are connected to the back panel connector.

Usually a MFC needs additional PCB's to interface with sensors or external devices. Therefore general purpose digital I/O lines are implemented. They are controlled by an FPGA IP Core and connected to level shifters with built-in ESD protection. It is possible to select the direction and voltage level of these lines by groups of eight bits. The voltage levels can be set by jumpers. There is no dedicated function assigned to these lines. The only restriction is to take care of the direction and the signalling voltage. All functions might be generated inside the FPGA fabric and connected to higher instances of control logic like event generators or a CPU interrupt controller.

Telemetry Adapter

To generate TM signal waveforms like Biphase or NRZ, a special board can be plugged to a Telemetry Adapter connector. This PCB became part of the MFC because it can be soldered to this socket. It is intended to implement project specific filters or amplifiers. The connector is placed between the FPGA and the back plane and protected against ESD from outside. The connection to the back plane is realized by ten lines. Their usage will only be defined by the Telemetry Adapter board. The "inside" eight lines end directly at the FPGA, they are also project dependent. One of them is able to drive the internal clock network of the card.

Two CMX909 GMSK Modems made by CML are used for receiving telecommands from independent receivers during flight. For EGSE or very slow telemetry usage, these modems are also able to send data streams. The modems are able to pre process MOBITEX compatible frames with forward error correction and interleaving of bits. Because this technique is used to secure TC links in almost all our projects, we decided to implement them directly, not as IP-Core into the FPGA. It is also possible to use the modems without these features just for generation or decoding of GMSK signals. The chips are able to work up to 38.4 kBit/s. They are connected to the FPGA by an 8 Bit parallel bus with address and chip select lines. An oscillator generates the clock signals for

both devices. A filter network is necessary to connect the GMSK I/O signals of both modems to the Telemetry Adapter. The adapter can be used to implement amplifiers or additional filters.

Housekeeping

An internal I²C Bus connects a housekeeping sensor to the FPGA. This multi channel A/D Converter is connected to the different power supply voltages for monitoring purposes. It is also capable to measure the temperature of the card. Therefore it is physically located next to the Cyclone which is expected to be the hottest single device on the card.

The GMSK modems use a software adjustable resistor network with I²C interface for calibration of their gain values.

A serial ferromagnetic non-volatile RAM (FRAM) from RAMTRON which is connected to the same bus can be used for storing of constants as well as board or project specific values.

Reliability

The usage of highly integrated devices like the SRAM based FPGA or RAM chips leads to the question of radiation hardness of the design. A sounding rocket flight doesn't take longer than half an hour. The card itself will be placed inside an electronic box - made of aluminium - which is mounted in the payload structure. That should reduce the risk of failure due to radiation effects. Additionally there is the opportunity that FPGA and CPU monitor each other. This can be done by calculating known values or checksums by a CPU task which are compared with values calculated by the FPGA. In case of an error, the FPGA can restart the CPU. On the other hand an IP-Core can calculate such values which are monitored by the CPU. The decision of what to do must be made during the design phase of each project. The designer has to select only one of these opportunities because of a missing third calculation source. In a worst case, the entire card has to be restarted as a result of radiation effects.

A dedicated watchdog timer chip is implemented as part of the global reset system. It needs reset pulses to restart the internal timer. If these pulses are missing, the timer will run to its end and will then reset the card. The second function of the chip is to compare the supply voltage with an internal reference. If it is too low, the chip will also reset the card until the voltage is stabilized again. An additionally reset pin can be used to restart the card from an external device or a user.

As mentioned above the FPGA will start in less than one second. The CPU will also need time until it is ready for operation after a restart. To minimise this application dependant time the Flash ROM is connected by cost of PCB area to the parallel memory Bus instead of using slower serial versions.

Global and Local Clocks

There are several clock sources implemented on the MFC which are selectable by jumpers. The main oscillator feeds a zero delay clock buffer with internal PLL. It works between 10 and 133 MHz which allows a wide range of main oscillators. The buffer distributes the clock signal to CPU and FPGA. Both FPGA and CPU contain PLL's to adjust this source to their own needs. The Telemetry Adapter can also feed a clock signal to the buffer. This can be used to implement PCM data streams with high accuracy for slant range measurement. The source is selectable by a jumper. Usually there is a high precision quartz oscillator for slant range measurement implemented in the service module. With the global clock it is possible to generate a TM data-stream with the accuracy of this oscillator.

4 Implementation

Estimation of Power Consumption

Especially the RS422 drivers will consume high power and get warm if they are used at such high data rates. A worst case estimation of this power is about 6W for all drivers, estimation for the whole card assuming bad conditions is about 12W. The power supply is designed to supply up to 16W to provide safety margin. A realistic value depends on the actual FPGA layout and the CPU firmware which are different for each project.

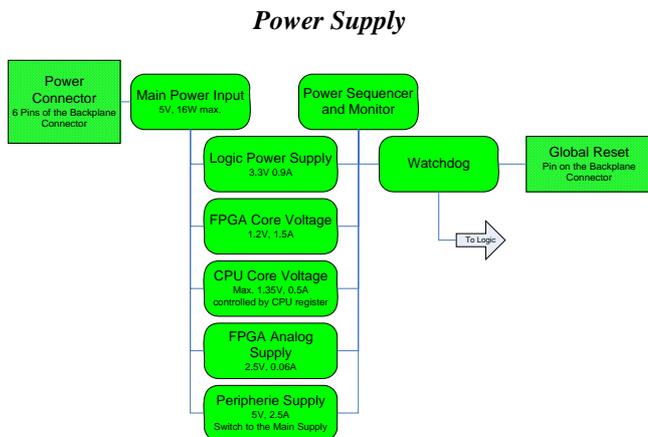


Figure 2: Diagram of the MFC Power Supply Structure

Figure 2 shows a block diagram of the Power Supply and the reset logic. Supply of the MFC is a 5V DC/DC converter in the service module. This rail is the main supply voltage of the card. Some components like RS422 drivers are connected to this by a switch.

The logic voltage for Cyclone - and Blackfin I/O's is 3.3V. It is derived from the 5V by a buck converter. But other chips use this rail as well, e.g. Clock circuits, LAN and memory chips, modems and mass storage. The core voltage of the Blackfin is generated by itself. A register controls the buck converter which is partly implemented inside the chip with the use of a few external parts.

A third buck converter produces 1.2V from the 5V rail which is used for the Cyclone core voltage. A linear converter regulates 2.5V for the analogue circuits of the Cyclone.

PCB Layout

Due to the high speed buses on the card, most of the signals must be designed to be impedance controlled. The 12 layer design is made of standard FR4 epoxy material. The memory bus and most other signals need termination for correct working. The nominal impedance for most signals was selected to be 100 Ohm. With the given geometry of the PCB, this results in narrow tracks which are a challenge for production.

Highly integrated devices usually came in ball grid array (BGA) packages. The size of the board limits the useable area for placement of parts. This leads to small surface mount devices (SMD) in general. The standard size of resistors and capacitors was chosen to be 0402 and all selected parts are as small as possible. Figure 3 and Figure 4 show pictures of the partly assembled first prototype.

On the left side, the Micro D connectors of front panel are located and on the right side, the 96 way back plane connector.

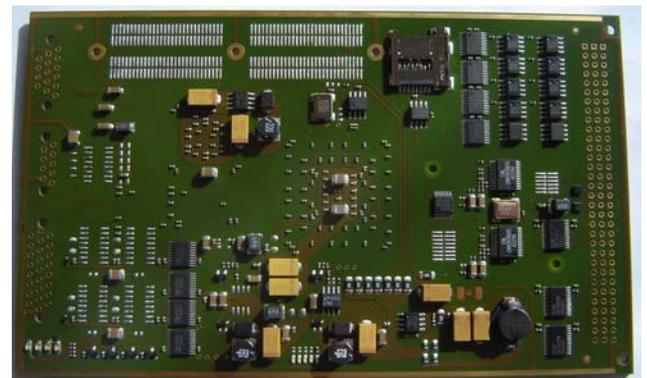


Figure 3 Top View of the MFC

The SMD connector on the upper side of the Top View is the Expansion Interface. Its post connectors are not soldered at this prototype. The power supply is located on the lower side.

Central part on the Bottom View in Figure 4 is the Cyclone 3 FPGA with the Blackfin CPU to its left. The smaller BGA chips left of the FPGA are RAMs and the Flash ROM. The Quad Flat Pack on the left side is the Ethernet controller.



Figure 4: Bottom View of the MFC Prototype

The MFC should fit to the existing electronic box design and it consumes several Watts of power, depending on the tasks which it is to perform. Therefore a box is proposed for additional encapsulation of the card inside the electronic box. This works as a heat spreader and additional shield against radiation and EMI.

5 Concept of Usage

There are two fundamental different ways to use this card. First of all the FPGA can be used as very large glue logic that just connects all parts to an internal bus. This bus must be visible to the CPU by using a bridge to the memory bus of the MFC. In this case the CPU needs hardware drivers for all implemented IP-Cores and their register files. The CPU needs to control everything. This approach maximises the debug capabilities of the design. It is easy to stop the CPU at all time and read register and memory values, or use the so called background telemetry transfer system to read data without stopping the execution of software. The disadvantage is that the CPU needs to calculate everything in a sequential way, which takes time, maximises load and power consumption of the chip by unloading the FPGA.

The second way is to idle the CPU completely and compile the entire logic with excessive use of state machines to the FPGA. In this case the FPGA needs to do all the work. The parallel architecture of the FPGA will reduce the calculation time for many algorithms at the cost of power consumption due to the logic usage.

Most projects will use a compromise of both ways to get advantages of both possibilities. There can be special processors inside the FPGA fabric for supporting the CPU with data handling and pre- or post processing of data. A hardware library can be used with tools like Alteras SOPC-Builder to generate custom FPGA images. A device driver library for a real-time kernel running on the CPU, eases the implementation of software.

6 Conclusion

The paper describes the requirements of a computing card for sounding rocket and balloon service modules. The concept and parts of the implementation for a multi function card were shown. With the approach of universal

functions together with high performance computing chips, it is possible to have a powerful tool for upcoming projects using advanced algorithms for data compression or navigation together with high bandwidth on TM links at the same time.