LEON SVF, a hardware accelerator for LEON2-FT software tools

10th International Workshop on Simulation for European Space Programmes

7-9 October 2008
at ESTEC, Noordwijk, the Netherlands

Mauro Caleno[(1)], Frederic Pasquier[(2)]

[(1)] ESA/ESTEC
Keplerlaan 1 - 2200 AG Noordwijk ZH - the Netherlands
Email: mauro.caleno@esa.int

[(2)] Astrium SAS
31, Avenue des Cosmonautes - 31402 Toulouse Cedex 4 -France
Email: frederic.pasquier@astrium.eads.net

SUMMARY

With the LEON2-FT processor being released for space applications, the frequency of the clock of the processor has gone up from the region of 25MHz to 100MHz. The software simulators cannot yet keep up with these frequencies. ESA/ESTEC has initiated a study at a consortium lead by Astrium-SAS including TERMA, SAAB & Astrium-Gmbh for the development of a versatile LEON Software Validation Facility (LEONSVF) for this new target processor. Versatile means that the LEONSVF can be configured either as a full software simulation tool or with a LEON accelerator board running the LEON2-FT VHDL code. An SVF running the LEON2-FT VHDL offers higher performance and representativity compared to the software-based instruction-set simulators. It offers an accurate timing and functional behaviour especially for the caches, FPU and fault handling. The study is currently on-going and a prototyping phase including the LEON accelerator implemented on a Commercial Off-The-Shelf Field Programmable Gate Array board (COTS FPGA) with the Peripheral Component Interconnect (PCI) interface has been completed. It shows promising results in accuracy with respect to the LEON2-FT functional behaviour and associated timings, and the capability to interact with software simulations running on a host platform. There are still optimisation possibilities to improve the overall execution speed to reach real-time or above performance while simulating an on-board computer clocked at 80-100MHz. This paper will present the current status at the end of the prototyping phase and the anticipated way-forward to a full system.

INTRODUCTION

The issue of On-Board Software Validation Facilities (SVF) for new generation of computers (ERC32 or LEON based) is being addressed since a while by the European Space community.

Two main families of solutions exist:

- Full software simulation facilities
- Hybrid facilities with the processor hardware in the loop (HIL)

Full software simulation facilities based on a software instruction-set simulator of processors are fully operational for ERC32 based computers (e.g. SIMERC32 in Pléiades, Galileo IOV, Gaia, Bepi-Colombo programs, TSIM in Herschel/Planck program). Their adequacy in simulation accuracy while bringing full observability of on-board software (OBSW) execution has been demonstrated. The execution speed and representativity of the LEON CPU simulators are limited (e.g SIMLEON by Astrium, TSIM by Gaisler Research, LSVE by Spacebel). The increase in the clock rate of the flight CPUs from 25MHz to 100MHz, does not yet allow the software-based instruction-set simulators to simulate flight computers at real-time while maintaining full representativity.

Hybrid facilities based on processor Hardware-in-the-loop (HIL) were developed for the previous generation of ESA micro-processors. The LEON processor is available as VHDL (Very High Speed Integrated Circuits Hardware Description Language). This opens the possibility of synthesising the LEON onto a FPGA and thereby building a LEON-HIL emulator based on a COTS FPGA board. This will be referred to as the LEON-accelerator. Availability of the target CPU in VHDL eases the design of the interface logic and upgrade to new FPGAs. So, this set-up ensures accuracy with respect to the LEON ASIC while delivering a performance comparable to the LEON ASIC or better as faster FPGAs will become available.

Several SVFs and post launch on-board software maintenance facilities with the HIL emulators were deployed in the last 15 years for several ESA missions. These facilities were based on Simulation Handling Modules (SHAM) which are generic reusable boards specifically designed for a given CPU and equipped with the ASIC of the target CPU. SHAM boards were developed for the MAS281, Ma3-1750, Pace1750, the ERC32 3-chip and single chip and the TSC21020. The SVFs and maintenance facilities of XMM, Envisat, SOHO, Artemis, Rosetta, Rosetta Star Tracker were based on these boards.

OVERALL OBJECTIVE AND SCOPE OF THE PROJECT

The objective of the LEONSVF is twofold:

- Demonstrate that the LEON accelerator can be integrated with software simulations while bringing full fidelity and an emulation speed approximately real-time compared to an on-board computer (OBC) equipped with the Atmel Leon chip clocked in the range 80-100 MHz.
- Demonstrate that the LEON accelerator and the software simulators are interchangeable with no impact on the rest of the simulations such as the computer interfaces (1553, Spacewire, etc) and the units. This is achieved by a common interface called FCPP (Functional Component Plug-in Protocol) which is implemented over the ESS (Emulator Support Software) in the hardware-based platform.

The scope of the project is to exercise the LEONSVF in four configurations (see Fig.1):

- Two configurations will be based on the ESA VSRF (Virtual Spacecraft Reference Facility) including the Eagle Eye OBC Model upgraded for the LEON2-FT. One configuration is based on the TSIM software simulator and the other configuration is based on the LEON accelerator.
- Two configurations will be based on the Astrium facility for the Galileo-IOV software validation. This simulation facility has been upgraded with LEON2-FT computer features and will be referred to as GaliLeon. It is configured to run fully in software using the SIMLEON as well as with the hardware in the loop using the LEON accelerator.

THE ARCHITECTURE OF THE LEON ACCELERATOR BOARD

The Fig. 2 represents the top level architecture of the LEON-HIL emulator implemented in the FPGA. The LEON accelerator implements the Leon2-FT core with additional logic to interface the processor with the simulation software running on the host. The FPGA board plugs into a PCI or PCIe (PCI express) slot in the host.

The memories on the FPGA board support the Leon memories. The block “SRT counter” (Simulated Real Time) keeps track of the time elapsed for the Leon by counting the clock ticks. When the block “IO Detection” detects that the Leon accesses the I/O space (input-output), the block “Suspend Controller” suspends the clock to the Leon, stops counting the time and requests the host to provide the value to feed to the Leon or to store the value produced by the Leon. When the host completes the simulation of an I/O access, the “Suspend Controller” resumes the clocks. Thanks to this suspension mechanism, the on-board software experiences time the same way as it does in the flight context.

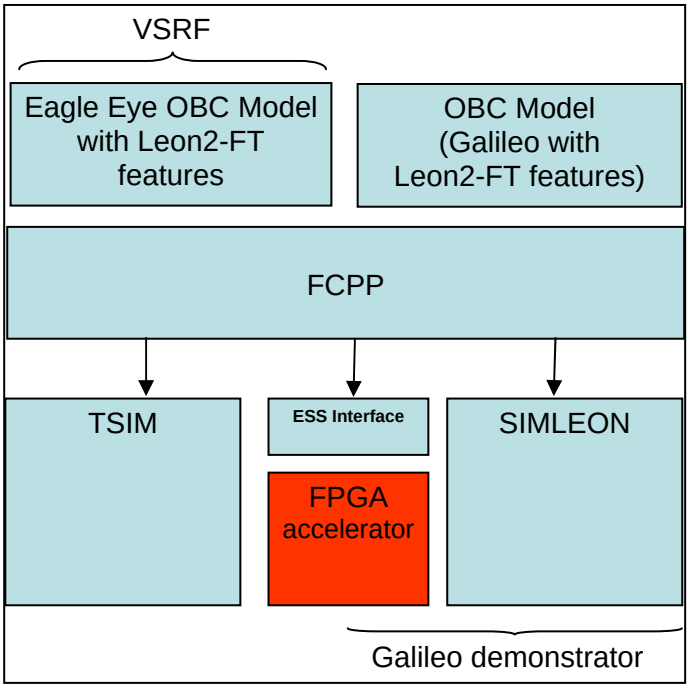

Fig. 1. Scope of the LEONSVF project.

The block “Event Handler” enables the simulation to suspend the Leon after a given interval of SRT time. The host can use this facility to simulate periodic signals or delayed responses from the OBC.

In general, the glue logic enables monitoring of all outputs and controlling all inputs of the Leon CPU, e.g. interrupts.

The purpose of the block “exclusion memory” is to maximise the overall execution speed of the SVF. This block supports the dual-ported memories used by OBCs to buffer data exchanged between the on-board software and the OBC interfaces. For example, they can store the TM packets sent over a Spacewire link or the TC packets from the Packetwire links. The simulations can prepare TM/TC packets in the “exclusion memory” using fast block-oriented transfers. This increases the overall speed of the simulator, because no PCI transaction is required for each individual I/O access.

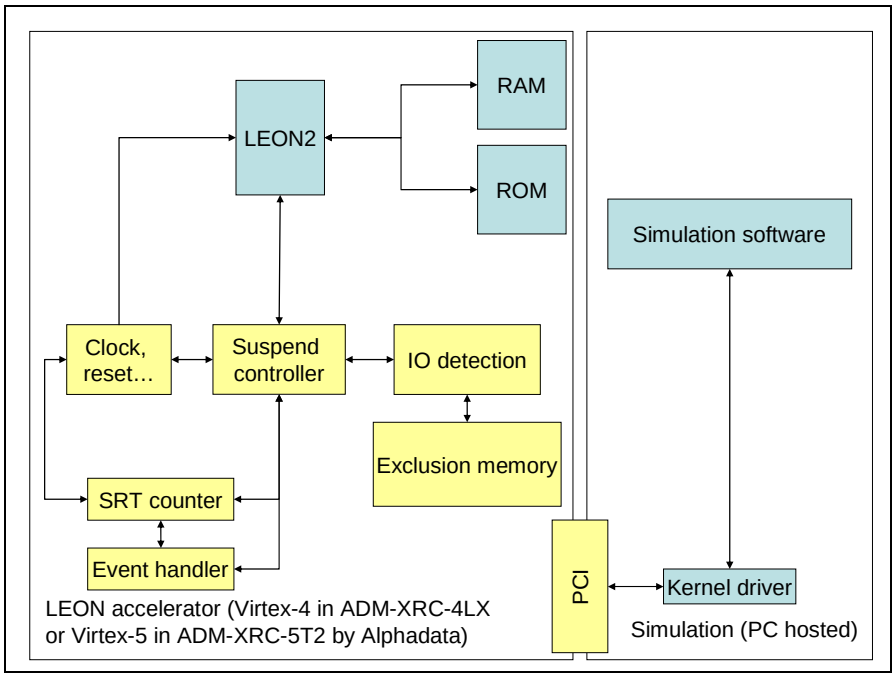

Fig. 2. Top-level architecture of the LEON accelerator board within the host

THE PROTOTYPING PHASE

The prototyping phase focussed on the implementation of the LEON accelerator in the FPGA and on the operability of the accelerator board within the overall HIL simulator. This phase started in 2007 on a COTS board featuring a Virtex-4 FPGA and moved on to a Virtex-5 board in 2008 when this became available on the market. The Virtex-5 board enables clocking the Leon2 core around 100MHz whereas the Virtex-4 board allowed achieving only 66MHz.

The accelerator design was completed for the Virtex-4, the board was integrated in the host and the individual accelerator functionalities were successfully tested. The Leon execution was successfully confidence-tested using general benchmarks, like the Stanford and Paranoia programs.

A test on-board software was developed which uses the resources of the GaliLeon on-board computer: 1553 interface, the Spacewire and the Packetwire interfaces. It was executed on the full LeonSVF software facility based on the SIMLEON instruction-set simulator and the I/O trace was collected. The software and the trace were then executed open loop on the accelerator. This set-up demonstrated that the accelerator board works correctly under

sustained I/O load and enabled to characterise the timing overhead due to the PCI bus and software layers. This overhead amounts to less than 5 micro seconds per I/O access.

The accelerator board was then fully integrated in the GaliLeon Software Validation Facility and the test on-board software was executed with the model of the Galileo on-board computer in closed loop. The software was tuned to generate an amount of I/O accesses which is representative of a real mission (5000 I/O per second). This set-up demonstrated that the HIL LEON accelerator board can simulate a representative on-board computer in real-time up to a clock frequency of 60MHz on the Virtex-4 board. The Virtex-4 FPGA is clocked at 66MHz and the overhead of the PCI transfers and the closed loop simulations accounts for the loss of about 10%.

Measurements on the full software-based simulator show that this can simulate an OBC clocked at 30-40MHz in real-time. So, the HIL LEON accelerator based on the Virtex-4 FPGA already offers better performance than instruction-set software simulators.

In order to speed up the simulation further to the 80-100MHz OBCs, the design was ported to a Virtex-5 board. In order to maximise the reuse of the accelerator design in the FPGA, the manufacturer of the Virtex-4 board was selected also for the Virtex-5 board. Unfortunately the Virtex-5 board features 8 MB of static RAM which are barely sufficient to support all Leon memories. To work around this, the dynamic RAM of the board was used to support all Leon memory areas. This required adapting the memory controller in the accelerator so that the timing of the dynamic RAM on the board looks the same as the static RAM of the Leon. This is achieved by suspending the Leon clock when the transactions on the dynamic RAM begins and releasing it as soon as the data is available. As a result, the performance of the accelerator depends on how frequently the on-board software accesses the external RAM compared to the internal Leon caches.

Preliminary measurements and estimations of the performance show that the Virtex-5 accelerator can emulate an 80MHz Leon-based OBC with a speed varying from 100% to 50% of real-time depending on the cache hit rate.

CURRENT STATUS

The prototype phase focused on the configuration with the HIL LEON accelerator board, concentrating on two main issues, namely the accelerator implementation and the control and interaction of the accelerator together with simulations running on a host PC.

The main results achieved follow.

- The functional and timing accuracy of the accelerator was verified versus the software simulation
- The operability of the accelerator within a larger simulator was achieved
- Software debug features have been successfully exploited

- The execution speed was measured on the Virtex-4 and shows that a flight-representative on-board computer up to 60Mhz can be simulated in real-time
- The overhead for I/O simulation amounts approximately to a 10% due to synchronisation, PCI transactions, operating system and simulations under a realistic I/O traffic profile.
- The Virtex-5 FPGA can speed-up execution of the Leon core up to 100MHz

LESSONS LEARNED ON USING COTS

The procurement of the COTS FPGA board and the development of our application in the VHDL has suffered from a number of set-backs, some expected and inherent to the use of COTS and others unexpected.

The statement of work required to use a COTS FPGA board for the project. The purpose was to reduce the custom design to the minimum (e.g. skip the design and validation of the PCB) thus shortening the development time, effort and cost. However, there are also some drawbacks in this approach.

The validation status (quality) of a COTS board is not guaranteed. It is proportional to the number of users or 'flight hours'. The Virtex-4 board design was already about 1 year old when we purchased it and was stable since we began using it. However, this was not the case for the Virtex-5 board.

Problems which affect a COTS board, need technical support from the supplier. This is likely to have a longer turn-around time compared to internal developments.

An inherent disadvantage of a COTS board is that the requirements of our application had to be adapted to the boards available on the market. The Virtex-4 board fit our requirements well. However, the amount of static memory available on the Virtex-5 board was limited and we had to use the dynamic RAM of the board to simulate the static RAM of the Leon. This approach forced us to a sophisticated design and brings a penalty in the performance which we could quantify reliably only after prototyping.

We planned to upgrade the design of the memory step by step, first using the limited amount of static RAM and then migrating to the dynamic RAM. Unfortunately, the Virtex-5 board was a new product which suffered from a number of problems, e.g. the memory controller did not compile and the static RAM was not operational. After spending time in troubleshooting the static RAM without success, we migrated to the dynamic RAM. Eventually, this prototype showed that the level of performance which depends on how the on-board software uses the Leon cache's, is sufficient only in the best case. So it is the intention to move back to the static RAM with the limitation of 8MB.

Another inherent disadvantage of COTS components is that they often have a shorter life in the market compared to the duration of space projects. When this is the case, porting is

needed and the effort required depends on the similarity between the old and new generation boards.

FUNCTIONALITY OF THE HIL LEON ACCELERATOR

The main functionalities of the HIL LEON accelerator are:

- the accelerator enables to validate unmodified software developed for the Leon2FT in a simulated and representative context
- the accelerator enables real-time simulation of flight on board computers which run at 80MHz
- the accelerator is fully representative of the Leon2FT including cache's, FPU and 'processor bugs'
- the accelerator is a PCI or PCIe board easy to accommodate in common workstations
- the accelerator enables to save/restore the status of the Leon
- the accelerator enables to connect a debugger
- the HIL accelerator and the software simulator can be exchanged thanks to a common interface

CONCLUSIONS AND WAY FORWARD

The hardware-in-the-loop LEON accelerator board can be integrated in larger simulators to enable validation of unmodified on-board software developed for the Leon2FT (Atmel 697) with full representativity and performance. The LEONSVF project has demonstrated operation of the HIL simulator in the representative context of the Galileo on-board computer (OBC) adapted for the Leon2-FT (GaliLeon). It has also demonstrated performance ranging between 60 and 100% of real-time for an 80MHz OBC depending on the way the on-board software uses the cache. In addition, it is the intention to modify the accelerator design to achieve 100% real-time simulation for an OBC at 90MHz regardless of cache use. The current level of performance of the LEON accelerator is already attractive compared to the software-based instruction-set simulators which enable simulating an OBC in real-time only up to 40 MHz.

Furthermore, Terma is also integrating the HIL LEON accelerator in the Virtual Spacecraft Reference Facility at ESTEC.