© 2014 IJIRT | Volume 1 Issue 5 | ISSN : 2349-6002

ADVANCED MULTIPURPOSE MICROPROCESSOR

Pankaj Gupta, Ravi Sangwan
IT Department, mdu university

***Abstract-* the Next Generation Multipurpose Microprocessor (NGMP) is a SPARC V8 (E) there have been three major revisions of the architecture. The first published revision was the 32-bit *SPARC Version 7* (V7) in 1986. *SPARC Version 8* (V8), the main differences between V7 and V8 were the addition of integer multiply and divide instructions, and an upgrade from 80-bit "extended precision" floating-point arithmetic to 128-bit "quad-precision" arithmetic. This paper describes the baseline architecture, points out key choices that have been made and emphasises design decisions that are still open. The software tools and operating systems that will be available for the NGMP, together with a general overview of the new LEON4FT microprocessor, are also described.**

I. BACKGROUND

The LEON project was started by the European Space Agency in late 1997 to study and develop a high-performance processor to be used in European space projects. The objectives for the project were to provide an open, portable and non-proprietary processor design, capable to meet future requirements for performance, software compatibility to maintain correct operation in the presence of SEUs, extensive error detection and error handling functions were needed. The goals have been to detect and tolerate one error in any register without software intervention, and to suppress effects from Single Event Transient (SET) errors in combinational logic.

The "Scalable" in SPARC comes from the fact that the SPARC specification allows implementations to scale from embedded processors up through large server processors, all sharing the same core (non-privileged) instruction set. One of the architectural parameters that can scale is the number of implemented register windows; the specification allows from 3 to 32 windows to be implemented, so the implementation can choose to implement all 32 to provide maximum call stack efficiency, or to implement only 3 to reduce cost and complexity of the design, or to implement some number between them. Other architectures that include similar register file features include Intel i960, IA-64, and AMD 29000, recently announced the availability of the fourth generation LEON, the LEON4 processor.

-

ESA has initiated the NGMP activity targeting a European Deep Sub-Micron (DSM) technology in order to meet increasing requirements on performance and to ensure the supply of European space processors.

II. ARCHITECTURAL OVERVIEW

It should be noted that this paper describes the current state of the NGMP. The specification has been frozen and the activity is currently in its architectural design phase. SPARC machines have generally used Sun's SunOS, Solaris or Open Solaris, but other operating systems such as Next STEP, RTEMS, FreeBSD, OpenBSD, NetBSD, and Linux have also been used.

Fig. 1 depicts an overview of the NGMP architecture. The system will consist of five AHB buses; one 128-bit Processor bus, one 128-bit Memory bus, two 32-bit I/O buses and one 32-bit Debug bus. The Processor bus houses four LEON4FT cores connected to a shared L2 cache. The Memory bus is located between the L2 cache and the main external memory interfaces, DDR2 SDRAM and SDR SDRAM interfaces on shared pins, and it will include a memory scrubber and possibly on-chip memory. As an alternative to a large on-chip memory, part of the L2 cache could be turned into on-chip memory by cache-way disabling.

The two separate I/O buses house all the peripheral cores. All slave interfaces have been placed on one bus (Slave I/O bus), and all master/DMA interfaces have been placed on the other bus (Master I/O bus). The Master I/O bus connects to the Processor bus via an AHB bridge that provides access restriction and address translation (IOMMU) functionality. The two I/O buses include all peripheral units such as timers, interrupt controllers, UARTs, general purpose I/O port, PCI master/target, and High-Speed Serial Link, Ethernet MAC, and Space Wire interfaces.

The fifth bus, a dedicated 32-bit Debug bus, connects a debug support unit (DSU), PCI and AHB trace buffers and several debug communication links. The Debug bus allows for non-intrusive debugging through the DSU and direct access to the complete system, as the Debug bus is not placed behind an AHB bridge with access restriction functionality.

Tagged add and subtract instructions perform adds and subtracts on values checking that the bottom two bits of

both operands are 0 and reporting overflow if they are not.
The list below summarizes the specification for the NGMP system:
The SPARC architecture has been licensed to many companies who have developed and fabricated implementations such as:

- Afara Websystems
- Bipolar Integrated Technology (BIT)
- C-Cube
- Cypress Semiconductor
- Fujitsu and Fujitsu Microelectronics
- HAL Computer Systems
- Hyundai
- LSI Logic
- Magnum Semiconductor
- Meiko Scientific
- Metaflow Technologies

- 128-bit Memory AHB bus
  - 1x 64-bit data DDR2-800 memory interface with Reed-Solomon ECC (16 or 32 check bits)
  - 1x 64-bit data SDRAM PC133 memory interface with Reed-Solomon ECC (16 or 32 check bits)
  - 1x Memory scrubber
  - 1x On-chip SDRAM (if available on the target technology)
- 32-bit Master I/O AHB bus
  - 4x SpaceWire cores with redundant link drivers and RMAP @ 200 Mbit/s
  - 4x High-Speed Serial Link, exact definition TBD
  - 2x 10/100/1000 Mbit Ethernet interface with MII/GMII PHY interface
  - 1x 32-bit PCI target interface @ 66 MHz
- 32-bit Debug AHB Bus
  - 1x Debug support unit
  - 1x USB debug link
  - 1x JTAG debug link
  - 1x SpaceWire RMAP target
  - 1x AHB trace buffer, tracing Master I/O bus
  - 1x PCI trace buffer

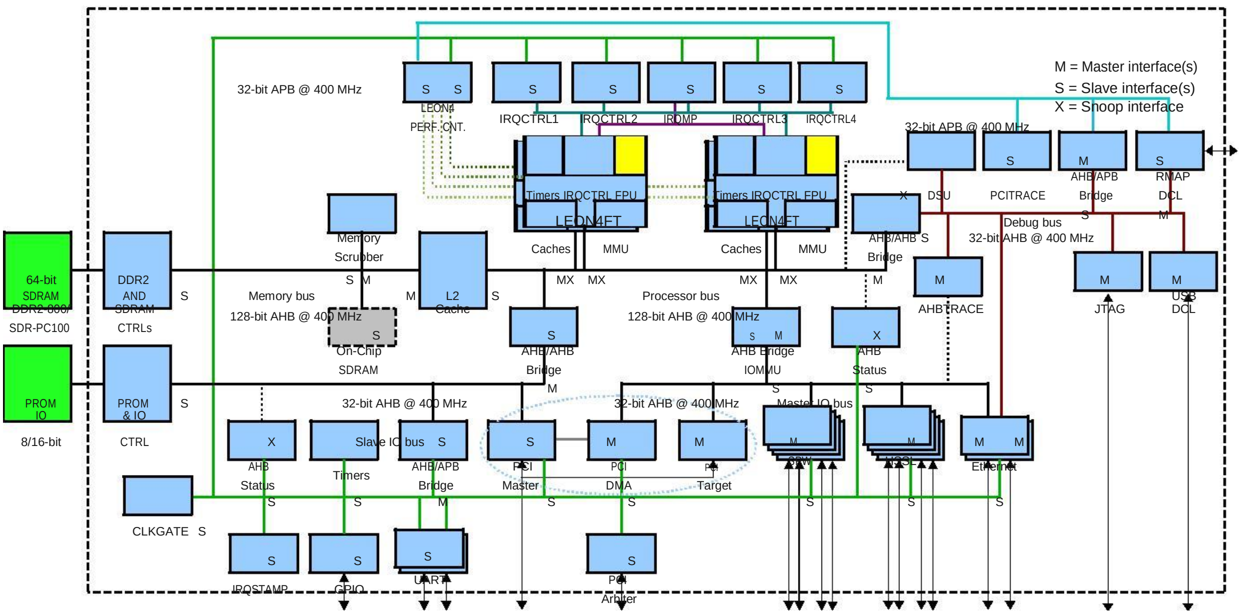

Figure 1: NGMP Block Diagram

- All I/O master units in the system contain dedicated DMA engines and are controlled by descriptors located in main memory that are set up by the processors. Reception of, for instance, Ethernet and SpaceWire packets will not increase CPU load. 32-bit Slave I/O AHB bus
  - 1x 32-bit PCI master interface @ 66 MHz with DMA controller mapped to the Master I/O bus
  - 1x 8/16-bit PROM/IO controller with BCH ECC
  - 1x 32-bit AHB to APB bridge connecting:
    - 1x General purpose timer unit
    - 1x 16-bit general purpose I/O port
    - 2x 8-bit UART interface
    - 1x Multiprocessor interrupt controller
    - 4x Secondary interrupt controller
    - 1x PCI arbiter with support for four agents
    - 2x AHB status register
    - 1x Clock gating control unit
    - 1x Interrupt time-stamp unit
    - 1x LEON4 statistical unit (perf. counters)

The cores will buffer in-coming packets and write them to main memory without processor intervention.

2.1 LEON4 Microprocessor and L2 Cache

The LEON4 processor is the latest processor in the LEON series. LEON4 is a 32-bit processor core conforming to the IEEE-1754 (SPARC V8) architecture. It is designed for embedded applications, combining high performance with low complexity and low power consumption. LEON4 improvements over the LEON3 processor include:

- Branch prediction
- 64-bit pipeline with single cycle load/store
- 128-bit wide L1 cache

The LEON2 is a synthesisable VHDL model of a 32-bit processor conforming to the IEEE-1754 (SPARC V8) architecture. It is highly configurable, and was designed for embedded applications with the following features on-chip:
Static ("always taken") branch prediction has shown to give an overall performance increase of 10%. The LEON4 also has support for the SPARC V9 compare and swap (CASA) instruction that improves lock handling and performance.
The L2 cache uses LRU replacement algorithm. It is a 256 KiB copy-back cache with BCH Error Correcting Code (ECC). One or more cache ways can be locked to be used as fault-tolerant on-chip ('scratchpad') memory.
An important factor to high processor performance and good SMP scaling is high memory bandwidth coupled with low latency. A 128-bit AHB bus will therefore be used to connect the LEON4FT processors. To mask memory latency, the GRLIB L2 cache will be used as a high-speed buffer between the external memory and the AHB bus. A read hit to the L2 cache typically requires 3 clocks, while a write takes 1 clock. A 32-byte cache line fetch will be performed as a burst of two 128-bit reads.
Atmel has manufactured an ASIC version of the LEON2-FT in the ATH18RHA rad hard process, available through their catalogue as part number AT697F. The AT697F is qualified according to QML-Q.

2.2 Main Memory Interface

The baseline decision for the main memory interface is Fujitsu's K computer ranked #1 in TOP500 - June 2011 and November 2011 lists. It combines 88,128 SPARC64 CPUs, for a total of 705,024 cores—almost twice as many as any other system in the TOP500 at that time.
The K Computer was more powerful than the next five systems on the list combined, and had the highest performance-to-power ratio of any other supercomputer system.
The flight models of the NGMP are scheduled 4 to 5 years into the future. At that time there may be additional information available regarding memory device availability. Availability of I/O standards on the target technology may also impact the final decision.
The width of the SDR SDRAM data interface could potentially be made soft configurable between 32 and 64 data bits (plus check bits), allowing for NGMP systems with a reduced width of the memory interface to support packages with low pin count. This is not considered technically feasible for DDR/DDR2.
To further improve resilience against permanent memory errors, a scheme with a spare device column could be envisaged. If a permanent error occurred in one memory device, the spare column would replace the faulty one. The decision on using column sparing has been deferred until the final technology and package has been selected.
The LEON2 (non-FT) model is no longer maintained. It is superseded by LEON2-FT, and the subsequently released LEON models (LEON3, LEON4). The ASIC implementation of the LEON2-FT is available from Atmel, as part number AT697F.
The calculations are also based on the behaviour and latencies, with simplifications

Memory interface	Min. time 32-byte fetch	Max. bw. 32-byte cache line	Min sys. freq.	Max. sys. freq.

	(ns)	(MB/s)	(MHz)	(MHz)
SDR PC100	100	320	-	400
DDR-400	50	533	86	400
DDR2-800	42.5	512	62.5	400

Table 1: Memory interface alternatives

400/DDR2-800 memories require half the time, or less, to deliver 32 bytes of data compared to PC100 SDRAM. 32 bytes is the cache line size that will be used by the L2 cache. In this case DDR memories offer better performance compared to DDR2 memories. This is due to the cache line size, with longer cache lines a longer burst memory burst length can be used and DDR2 memories will eventually outperform DDR memories since the actual data will be fetched at a higher clock frequency when using DDR2 SDRAM.

The importance of the minimum time required to fetch one single cache line versus the maximum sustainable bandwidth when fetching several cache lines is highly application dependent and depends on parameters such as memory footprint and data access patterns. One observation that can be made is that L2 cache hit rate can indeed be key to high performance, especially in MP systems. The target clock frequency of the NGMP is 400 MHz, which gives a clock period of 2.5 ns. One cache line fetch from DDR2-800 memory will in other words take 17 clock cycles. A hit in the L2 cache means that the data will be delivered in less than one third of the time required to access external memory.

2.3 I/O Interfaces

An early design decision was to only include high-speed I/O interfaces on-chip, while legacy low-speed inter-faces can be placed in a companion chip (FPGA/ASIC). The reason for this decision is that low-speed interfaces such as CAN, I^{2}C, 1553, UARTs etcetera do not generate enough data rates to require DMA capabilities, and can easily be implemented off-chip and connected to the NGMP using one of the high-speed interfaces. However, a set of standard peripherals required for operating sys-tem support is included on-chip. These include support for simple memory mapped I/O devices, two basic con-sole UARTs, and one 16-bit I/O port for external interrupts and simple control.

2.3.1 The LEON4 is the latest implementation of the SPARC V8 architecture by Aeroflex Gaisler, in the form of a synthesizable VHDL model of a 32-bit microprocessor. As was the case with the previous LEON models, the LEON4 is also highly configurable, and particularly suitable for system-on-chip (SoC) designs. The LEON4 extends the LEON3 model with support for an optional Level-2 (L2) cache, a pipeline with 64-bit internal load/store data paths, and an AMBA interface of either 64- or 128-bits. Branch prediction, 1-cycle load latency and a 32x32 multiplier results in a performance of 1.7 DMIPS/MHz, or 2.1 CoreMark/MHz.**PCI Interface**

The currently used AT697 processor and several LEON3FT designs have a 32-bit PCI interface. This makes a 32-bit PCI format a suitable candidate for the local backplane, since it will make the NGMP backward compatible with existing backplanes. The downside with the PCI interface is that it requires many I/O pins and is relatively slow. However, selecting a more modern interface, such as PCI express would increase demands on companion chips that could prevent the use of many types of currently available programmable logic devices as companion devices.

2.3.2 SpaceWire Links

Considering the wide adoption of SpaceWire links, the NGMP system will implement four SpaceWire link interfaces directly on-chip. The links will be based on the GRSPW2 core from Aeroflex Gaisler IP Library GRLIB, and include support for the Remote Memory Access Protocol (RMAP). The maximum link bit rate will be at least 200 Mbit/s. The GRSPW2 core will have its redundant link capability enabled, with two link drivers per core for redundancy.

2.3.3 1000/100/10 Mbit Ethernet

The Ethernet interfaces will be supplied by the GRETH_GBIT core from Aeroflex Gaisler supporting 10/100/1000 Mbit/s operation. GRETH_GBIT has internal RAM that allows buffering a complete packet. Support for multicast will also be included to allow reception of multicast packets without setting the interface in promiscuous mode.

2.3.4 High-Speed Serial Links

The availability and specification of the High-Speed Serial Link (HSSL) IP cores to be integrated within the European DSM ASIC platform is at the time of writing very limited. Aeroflex Gaisler is working with the European Space Agency in order to be able to provide, at the minimum, a descriptor based DMA to control the SerDes macros that are expected to provide 6.25 Gbit/s of bandwidth per link.

2.4 Debug Communication Links

The NGMP has a wide range of debug links; JTAG, SpaceWire RMAP, USB and Ethernet. The controllers for the first three links are located on the Debug bus and will gated off in flight. The controllers for the two Ethernet debug links are embedded in the system's Ethernet cores.

The two Ethernet debug links use Aeroflex Gaisler's Ethernet Debug Communication Link (EDCL) protocol that is integrated in the GRETH_GBIT Ethernet cores. An extension to the GRETH_GBIT core allows users to connect each Ethernet debug link either to the Debug bus or to the Master I/O bus. The Ethernet cores' normal function is preserved even if the debug links are active. The debug traffic is intercepted and converted to DMA at hardware level. The selected buffer size for debug traffic in the NGMP gives an Ethernet debug link band-width of 100 Mbit/s.

A USB Debug Communications Link (USBDCL) core provides a debug connection with relatively high band-width (20 Mbit/s). The wide adoption of USB will allow the NGMP system to be debugged from nearly any modern workstation without the need for configuration that is normally required when using an Ethernet Debug Communications Link.

Microprocessors are a core component of modern electronics and on-board computers do not escape this rule. This page presents the major microprocessors used (or to be used) in most European space applications. A dedicated SpaceWire RMAP target was included on the Debug bus in order to accommodate users who use the NGMP in SpaceWire networks. With a dedicated SpaceWire debug link it becomes easy to use existing infrastructure to control the NGMP system. The SpaceWire RMAP target will typically provide a debug link bandwidth of 20 Mbit/s.

2.5 Fault-tolerance

The fault-tolerance in the NGMP system is aimed at detecting and correcting SEU errors in on-chip and off-chip RAM. The L1 cache and register files in the LEON4FT cores are protected using parity and BCH coding, while the L2 cache will use BCH. External SDRAM memory will be protected using Reed-Solomon coding, while the boot PROM will use BCH. Any RAM blocks in on-chip IP cores will be protected with BCH or TMR. Flip-flops will be protected with SEU-hardened library cells if available, or TMR other-wise.

2.6 Improved Support for Time-Space Partitioning and Multi-Processor Operation

Beyond support for standard symmetric multiprocessing (SMP) configurations, e.g. with a central multi-processor interrupt controller, NGMP will also support asymmetric multiprocessing (AMP) configurations: Per CPU-core dedicated timer units and interrupt controllers allow running separate operating systems on separate processor cores.

Each processor core has a dedicated memory management unit (MMU) that provides separation between pro-cesses and operating systems. The system also includes an IOMMU that provides access restriction and address translation for accesses made by the DMA units located on the Master I/O Bus. The MMU and the IOMMU provide access restriction and address translation to blocks of memory divided into, a minimum size of, and 4 KiB pages. In order to be able to grant selective access to the registers of one and only one peripheral core, all core register addresses are 4 KiB aligned.

This has the side effect of theoretically increasing the miss rate of the MMU Translation Look-a-side Buffer (TLB) as more pages will be required in order to map all registers of peripheral units. The performance impact of the potentially increased TLB miss rate is expected to be low, as measurements have shown that TLB misses when accessing peripheral registers are also common in systems that map all peripheral registers within one page.

In addition to the MMU's in each of the CPU cores and the IOMMU, memory read/write access protection (fence registers) are implemented in the L2 cache. This functionality is primarily intended to protect backup software but can also be used to add another layer of protection with regard to time-space partitioning.

2.7 Improved Support for Performance Measurement and Debugging

The NGMP will include new and improved debug and profiling facilities compared to the LEON2FT and LEON3FT. The selection of available debug links has previously been discussed, additional debug support features of NGMP include:

- AHB bus trace buffer with filtering and counters for statistics
- Processor instruction trace buffers with filtering

- Performance counters capable of taking measurements in each processor core
- Dedicated debug communication links that al-low non-intrusive accesses to the processors' debug support unit

- Hardware break- and watch points
- Monitoring of data areas
- Interrupt time-stamping in order to measure interrupt handling latency
- PCI trace buffer

All performance counters and trace buffers are accessible via the Debug bus. The processors can also access the performance counters via the slave I/O bus.
The rich set of debugging features gives users the ability to quickly diagnose problems when developing systems that include the NGMP. Some features, such as the PCI trace buffer could easily be replaced by external units or by performing measurements in simulation.
However, experience among Aeroflex Gaisler engineers has shown that on-chip debugging resources that are readily available, and supported by a debug monitor, can significantly shorten the time required to diagnose and re-solve a problem. GRSIM can be run in stand-alone mode, or connected through a network socket to the GNU GDB debugger. In stand-alone mode, a variety of debugging commands are available to allow manipulation of memory contents and registers, breakpoint/watch point insertion and performance measurement. Connected to GDB, GRSIM acts as a remote target and supports all GDB debug requests. The communication between GDB and GRSIM is per-formed using the GDB extended-remote protocol. Any third-party debugger supporting this protocol can be used.
The overall accuracy will depend on the accuracy of the simulation models for the DDR2 memory controller and the L2 cache. The target is a maximum error of 10% during an extended period of simulation, this level of accuracy is considered challenging but will be the goal during development.

III. SOFTWARE SUPPORT

LEON, a 32-bit, SPARC Version 8 implementation, designed especially for space use. Source code is written in VHDL, and licensed under the GPL.
OpenSPARC T1, released in 2006, a 64-bit, 32-thread implementation conforming to the UltraSPARC Architecture 2005 and to SPARC Version 9 (Level 1). Source code is written in Verilog, and licensed under many licenses. Most OpenSPARC T1 source code is licensed under the GPL. Source based on existent open source projects will continue to be licensed under their current licenses. Binary programs are licensed under a binary software license agreement.
S1, a 64-bit Wishbone compliant CPU core based on the OpenSPARC T1 design. It is a single UltraSPARC v9 core capable of 4 way SMT. like the T1, the source code is licensed under the GPL.

IV. CONCLUSION

During the activity's definition phase a prototype system was developed based on existing components from Aeroflex Gaisler's GRLIB IP library. The prototype sys-tem was built on a Xilinx ML510 development board that has a Xilinx XC5VFX130T FPGA. To fit a proto-type system on this FPGA, a
Reduced configuration of the system described in the NGMP specification was implemented.
As a result of SPARC International, the SPARC architecture is fully open and non-proprietary. Thus SPARC is an open set of technical specifications that any individual or entity can license and use to develop microprocessors and other semiconductor devices based on Published industry standard