

NEXT GENERATION MULTIPURPOSE MICROPROCESSOR

Jan Andersson⁽¹⁾, Jiri Gaisler⁽¹⁾, Roland Weigand⁽²⁾

⁽¹⁾*Aeroflex Gaisler, Kungsgatan 12, SE-411 91, Göteborg, Sweden, {jan,jiri}@gaisler.com*

⁽²⁾*European Space Agency, Keplerlaan 1 – PO Box 299, 2220AG Noordwijk ZH, The Netherlands, roland.weigand@esa.int*

ABSTRACT

The Next Generation Multipurpose Microprocessor (NGMP) is a SPARC V8(E) based multi-core architecture that provides a significant performance increase compared to earlier generations of European space processors. The NGMP is currently in development at Aeroflex Gaisler in Göteborg, Sweden, in an activity initiated by the European Space Agency (ESA).

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.

1 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 and low system cost. Another objective was to be able to manufacture in a Single Event Upset (SEU) sensitive semiconductor process. 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 LEON family includes the first LEON1 VHSIC Hardware Description Language (VHDL) design that was used in the LEONExpress test chip developed in 0.25 μm technology to prove the fault tolerance concept. The second LEON2 VHDL design was used in the processor device AT697 from Atmel (F) and various system-on-chip devices. These two LEON implementations were developed by ESA. Gaisler Research, now Aeroflex Gaisler, developed the third LEON3 design that is used in a number of avionics systems and also in the commercial sector. Aeroflex Gaisler recently announced the availability of the fourth generation LEON, the LEON4 processor.

Following the development of the TSC695 (ERC32) and AT697 processor components in 0.5 and 0.18 μm tech-

nology respectively, ESA has initiated the NGMP activity targeting an European Deep Sub-Micron (DSM) technology in order to meet increasing requirements on performance and to ensure the supply of European space processors. Aeroflex Gaisler was selected to develop the NGMP system that will be centred around the new LEON4FT processor.

2 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. The development work is scheduled to be finished by the 1st of December 2010.

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, High-Speed Serial Link, Ethernet MAC, and SpaceWire 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.

The target frequency NGMP design is 400 MHz, but depends ultimately on the ASIC technology.

The list below summarizes the specification for the NGMP system:

- 128-bit Processor AHB bus
 - 4x LEON4FT
 - 16 + 16 KiB write-through cache with LRU replacement
 - SPARC Reference MMU
 - Physical snooping
 - 32-bit MUL/DIV
 - GRFPU with 4-word instruction FIFO shared between pairs of LEON4FT
 - 1x 256 KiB Shared L2 write-back cache with memory access protection (fence registers), BCH ECC and cache-way locking
 - 1x 128-bit to 32-bit unidirectional AHB to AHB bridge (from Debug bus to Processor bus)
 - 1x 128-bit to 32-bit unidirectional AHB to AHB bridge (from Processor bus to slave I/O bus)
 - 1x 32-bit to 128-bit unidirectional AHB to AHB bridge with IOMMU (from master I/O bus to Processor bus)
- 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 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)
- 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

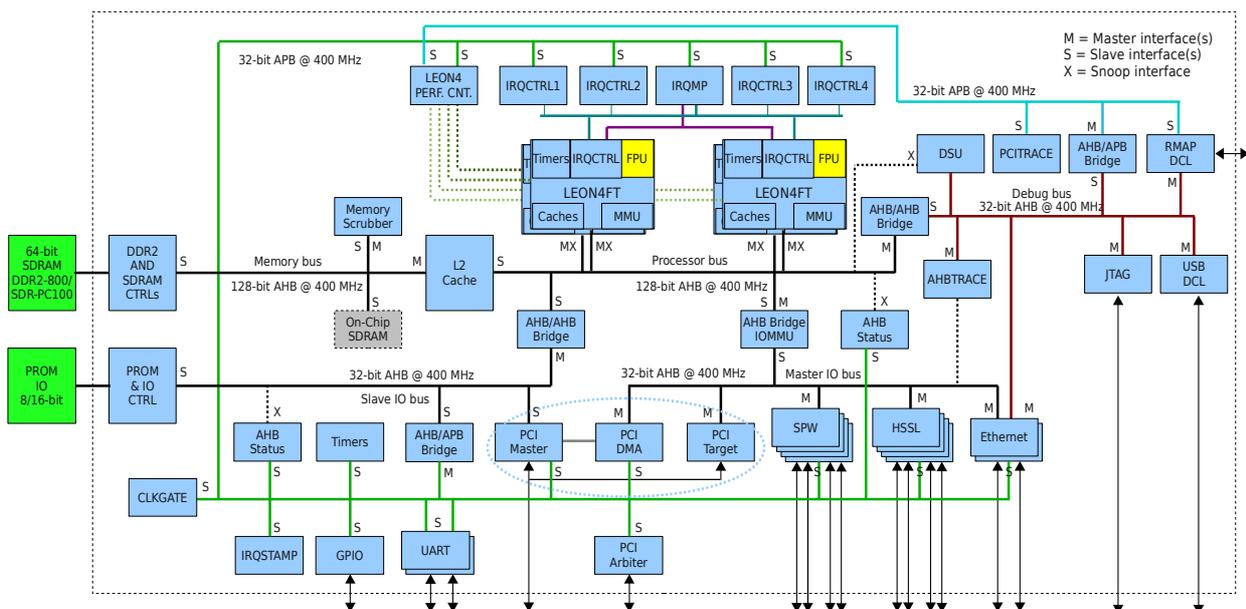


Figure 1: NGMP Block Diagram

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. The cores will buffer incoming 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 LEON4FT processor connects to an AMBA AHB bus with an 128-bit data width. This leads to a 4x performance increase when performing cache line fills. Single cycle load and store instructions increase performance and also take advantage of the wider AHB bus.

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. An 128-bit AHB bus will therefore be used to connect the LEON4FT processors. This will allow 32 bytes to be read in 2 clocks, not counting the initial memory latency. 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. The first read will have a delay of 3 clocks and the second word will be delivered after one additional clock. A cache line will thus be fetched in 4 clocks (3 + 1). Error correction will add an additional latency of 1 clock to all read accesses to allow time for checksum calculation.

2.2 Main Memory Interface

The baseline decision for the main memory interface is

to support 96-bit (64 data bits and up to 32 check bits) DDR2 and SDR SDRAM on shared pins. However the selection between DDR2 and DDR1 should be regarded as open. 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. This could potentially cause a timing problem since the Reed-Solomon codes are slow to generate, and the multiplexer tree would follow the codec directly. The decision on using column sparing has been deferred until the final technology and package has been selected.

The target is to use DDR2-800 and SDRAM PC100 memories. The SDR SDRAM interface will be able to run at the same or one fourth of the system frequency. The DDR2 interface will be run at the same or twice the system frequency. The clock scaling factor between the memory interfaces and the rest of the chip is selectable via an external signal.

Preliminary calculations that take memory controller latencies into account indicate that the memory interface will have the characteristics listed in tab. 1. The table also includes values for typical DDR-400 memories as this memory type could be a candidate for NGMP. It should be noted that these results are based on memories with characteristics that are deemed to be common for a wide range of SDR/DDR/DDR2 memory devices. The calculations are also based on the behaviour and latencies, with simplifications, of Aeroflex Gaisler memory controllers.

Memory interface	Min. time 32-byte fetch (ns)	Max. bw. 32-byte cache line (MB/s)	Min sys. freq. (MHz)	Max. sys. freq. (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

As seen in the second column of tab. 1, DDR-

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. The third column shows the maximum sustainable bandwidth when fetching several 32-byte cache lines back-to-back. 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 interfaces can be placed in a companion chip (FPGA/ASIC). The reason for this decision is that low-speed interfaces such as CAN, I²C, 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 system support is included on-chip. These include support for simple memory mapped I/O devices, two basic console UARTs, and one 16-bit I/O port for external interrupts and simple control.

The high-speed interfaces that are intended to be used in flight are four SpaceWire links, two 1000/100/10 Mbit Ethernet links, four high-speed serial links, and one 32-bit PCI 2.3 66 MHz master/target interface.

2.3.1 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 de-

mands 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's 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_GBITH core from Aeroflex Gaisler supporting 10/100/1000 Mbit/s operation. GRETH_GBITH 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 be 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_GBITH Ethernet cores. An extension to the GRETH_GBITH 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 a Ethernet debug link bandwidth of 100 Mbit/s.

A USB Debug Communications Link (USBDC) core provides a debug connection with relatively high bandwidth (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.

The JTAG communication link provides a link with a modest bandwidth of around 500 kbit/s, typically limited by the JTAG adapter. With modern USB JTAG adapters however, it is possible to run the JTAG communication link at 6 Mbit/s.

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 otherwise.

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 processes 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, 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-aside 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 allow non-intrusive accesses to the processors' debug support unit
- Hardware break- and watchpoints
- 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 resolve a problem.

2.8 Target technology

Baseline is the European ST 65 nm space technology. Possible backup options for target technology include UMC 90 nm with the DARE library and Tower (130 nm) with a library from Ramon Chips.

Power consumption of the NGMP ASIC core (without IOs) under worst case operating conditions and maximum software load is required to not exceed 6 W. Maximum power consumption in idle mode (no software activity, but conservation of status and SEE protection) is required to not exceed 100 mW.

2.9 Expected Performance

During the activity's definition phase a prototype system was developed based on existing components from Aeroflex Gaisler's GRLIB IP library. The prototype system was built on a Xilinx ML510 development board that has a Xilinx XC5VFX130T FPGA. To fit a prototype system on this FPGA, a reduced configuration of the system described in the NGMP specification was implemented.

For validation and performance measurements the SPEC CPU2000 and PARSEC 2.1 benchmark suites have been executed on a Linux Kernel (2.6.21.1) with Debian GNU/Linux 4.0 mounted using NFS over Ethernet.

To achieve a suitable execution time for each benchmark the SPEC CPU2000 benchmark suite was used with the TRAIN input set and the SIMSMALL (Facessim), SIMMEDIUM (X264), and SIMLARGE (Blackscholes, Ferret, and Fluidanimate) input sets were used for the PARSEC benchmark suite. With these input sets the benchmarks finish their execution within 2 to 40 minutes (depending on benchmark and number of cores). This gives a reasonable long execution time to sample benchmark statistics (MIPS, L2 cache hit-rate, and bus utilization are sampled every 5 seconds).

The average MIPS for all SPEC CPU2000 benchmarks is 123 for the quad core LEON4 system. When evaluating this number one has to take into account that the prototype system was operating at 70 MHz instead of 400 MHz which is the expected frequency of the final ASIC. If one would simply scale the measured MIPS with clock frequency the expected MIPS for the final NGMP system would be $(400/70 \cdot 123)$ 700. This is an overly optimistic figure since the DDR2 memory interface is operating at twice the clock frequency (140 MHz) of the prototype system (70 MHz). For the final system the DDR2 memory interface is expected to operate at the same frequency as the system (400 MHz). Another difference is that the read latency of the L2 cache will become longer due to the addition of error correcting codes. Taking these effects into account one

could expect an average performance of at least 500 MIPS for these benchmarks.

The average MIPS measured for all PARSEC benchmarks is 105 for the quad core LEON4 system. If the same frequency scaling as described for SPEC CPU2000 is used, the result is $(400/70 \cdot 105)$ 600 MIPS, which is a little less than what was achieved for the SPEC CPU2000 benchmarks.

3 INSTRUCTION SET SIMULATOR

The instruction set simulator for NGMP will build on the GRSIM multiprocessor simulator. The simulator consists of an AHB bus model with underlying event-driven simulation engine. C-models of IP cores are attached to the AHB model, and linked into a final simulator. The GRSIM library is re-entrant and thread-safe, and allows simulation of any number of buses and IP cores. It is therefore particularly suitable for simulating multi-processor LEON systems such as the NGMP.

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/watchpoint 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 performed 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.

4 SOFTWARE SUPPORT

The GRMON debug monitor from Aeroflex Gaisler will be extended to support all new functionality for debugging and profiling that will be included in the NGMP. The hardware platform will provide full backward compatibility with existing LEON3FT software and all standard compilers that can produce correct SPARC V8 code can be used.

Board support packages for the NGMP will be delivered for the following operating systems:

- RTEMS 4.8 and 4.10
- eCos
- VxWorks 6.7
- Linux 2.6

Other operating systems that are already ported to LEON3/4 include:

- LynxOs (LynuxWorks)
- ThreadX (Express Logic)
- Nucleus (Mentor Graphics)

Aeroflex Gaisler's bootloader creation tool MKPROM2 will be extended with additional support for booting AMP configurations.

5 CONCLUSION

The NGMP will be a SPARC V8(E) based multi-core architecture that provides a significant performance increase compared to earlier generations of European space processors, with high-speed interfaces such as SpaceWire and Gigabit Ethernet on-chip. The platform will have improved support for profiling and debugging and will have a rich set of software immediately available due to backward compatibility with existing SPARC V8 software and LEON3 board support packages. NGMP includes also specific support for AMP configurations and Time-Space Partitioning.

The NGMP is part of the ESA roadmap for standard microprocessor components. It is developed under ESA contract, and it will be commercialised under fair and equal conditions to all users in the ESA member states. The NGMP is also fully developed with manpower located in Europe, and it only relies on European IP sources. It will therefore not be affected by US export regulations.

The NGMP specification and other related documents are posted at the following link:

<http://microelectronics.esa.int/ngmp/ngmp.htm>