

DEVELOPMENT OF A FUNCTIONAL PROTOTYPE OF THE QUAD CORE NGMP SPACE PROCESSOR

Jan Andersson, Magnus Hjorth, Sandi Habinc, Jiri Gaisler

Aeroflex Gaisler, Kungsgatan 12, SE-411 91, Göteborg, Sweden

Tel: +36 31 7758650 Fax: +46 31 421407

E-mail: {jan,mhjorth,sandi,jiri}@gaisler.com

ABSTRACT

The Next Generation 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 Göteborg, Sweden, in activities initiated by the European Space Agency (ESA).

The presentation and paper describes the baseline NGMP architecture and the functional prototype and validation board development.

This paper describes an ongoing development, at the time of writing the devices and validation board has not yet been manufactured. Validation boards with devices are expected to be available during the third quarter of 2012.

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 has subsequently 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 technology 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 is centred around the new LEON4FT processor.

The main NGMP activity was put on hold in April 2011 pending advances in the development of the European DSM technology. NGMP development has instead progressed in the development of a functional prototype targeting eASIC Nextreme2 technology.

2 ARCHITECTURAL OVERVIEW

Figure 1 shows an overview of the NGMP architecture as implemented in the functional prototype. The system consists of five Advanced High-performance Buses (AHB); 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 Level-2 (L2) cache. The Memory bus is located between the L2 cache and the main external memory interfaces, DDR2 SDRAM and SDR SDRAM, and also connects a memory scrubber. As an alternative to a large on-chip memory, part of the L2 cache can 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 controller, UARTs, general purpose I/O port, PCI master/target, Ethernet MAC, MIL-STD-1553B, Serial Peripheral Interface bus and SpaceWire interfaces. 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.

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 baseline target frequency for all on-chip buses in NGMP is 400 MHz. The functional prototype has a reduced frequency of 150 MHz.

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
 - FPU shared between pairs of LEON4FT
 - 1x256 KiB Shared L2 write-back cache with memory access protection (fence registers), 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 (with 16 or 32 check bits)
 - 1x 64-bit data SDRAM PC133 memory interface with Reed-Solomon ECC (with 16 or 32 check bits)
 - 1x Memory scrubber
- 32-bit Master I/O AHB bus
 - SpaceWire router with eight external ports and four AMBA ports, with RMAP @ 200 Mbit/s
 - 2x 10/100/1000 Mbit Ethernet interface with MII/GMII PHY interface
 - 1x 32-bit PCI target interface @ 66 MHz
 - MIL-STD-1553B interface
- 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
 - 2x 32-bit AHB to APB bridge connecting:
 - 5x General purpose timer unit
 - 1x 16-bit general purpose I/O port
 - 2x 8-bit UART interface
 - 1x Multiprocessor interrupt controller
 - 1x PCI arbiter with support for eight agents
 - 2x AHB status register
 - 1x Clock gating control unit
 - 1x LEON4 statistical unit (perf. counters)
 - 1x SPI master/slave controller
 - PLL and pad control units
- 32-bit Debug AHB Bus
 - 1x Debug support unit
 - 1x USB debug link




Figure 1: NGFP Block Diagram

- 1x JTAG debug link
- 1x SpaceWire RMAP target
- 1x AHB trace buffer, tracing Master I/O bus
- 1x PCI trace buffer

2.1 Deviations From NGMP Baseline Specification

The architecture of the functional prototype described above is intended to match the final NGMP design. However, due to target technology and target device constraints, some deviations have been made from the NGMP baseline specification in the development of the functional prototype.

The subsections below describe aspects where the functional prototype is different from what is specified as the NGMP baseline design.

2.1.1 Radiation tolerance

The target technology does not allow implementation of a radiation hardened device, and as the functional prototype (FP) is intended as a platform for software development and to further evaluate the architecture, this has been omitted from the prototype.

2.1.2 High-Speed Serial Links

The FP does not contain the four High-Speed Serial Links (HSSL) present in the NGMP baseline design. The HSSLs are 6.25 Gbit/s serial links that may be used via a DMA engine based on the GRETH_GBIT IP core, from Aeroflex Gaisler, or potentially via SpaceFibre controllers.

2.1.3 Main Memory Interfaces

The baseline NGMP design has DDR2 SDRAM and (SDR) SDRAM on shared pins. This design has separate pins for the two primary memory interfaces. The change has no functional impact visible to software.

The maximum operating mode for the DDR2 SDRAM interface has been decreased from DDR2-800 to DDR2-600.

2.1.4 Fault-Tolerance in LEON4FT Cores

Integer and floating-point unit register files in the processor cores are not implemented using triple-module redundancy. This has no impact for software as data errors would have been corrected transparently without impact at application level.

The cache memories have been implemented without byte parity, this prevents error injection in cache memories.

2.1.5 Fault-Tolerance in Level-2 Cache

The RAM cells that would hold BCH checkbits in the Level-2 cache have not been implemented. This means that the check bits cannot be read out using diagnostic accesses and it also prevents error injection in the Level-2 cache.

Despite not physically storing the checkbits, the Level-2 cache has been configured to have the same latency as if fault-tolerance had been enabled. This is to ensure the Level-2 cache's impact on benchmarking results are representative for the final NGMP.

2.1.6 Ethernet and SpaceWire cores

The Ethernet and SpaceWire cores have been implemented without fault-tolerance for internal buffers. This has no impact for software.

2.1.7 Clocking

The maximum frequency that should be used for the AMBA system has been decreased from 400 MHz to 150 MHz.

The NGMP features bootstrap signals for selecting the frequency of the AMBA system and the frequency, and clock source, for the main memory interfaces. The bootstrap interface has been left intact.

2.1.8 Dynamic PLL and Pad Control

The target device has PLLs that support dynamic configuration also has support for dynamically controlling the drive strength and slew of selected pads. In order to take advantage of this, an interface has been developed so that the clock frequencies used in the design can be configured by software. Dynamic configuration of the PLLs can be interesting for measurements on, for instance, memory bandwidth as a function of memory interface frequency and system frequency, and can also make it easier to over-clock the device.

Simple interface cores have been developed in order to allow software to configure clock frequencies and pad characteristics. These interfaces have been mapped in memory areas that are reserved in the NGMP baseline design and should not cause any issues for software development.

2.2 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 is therefore used to connect the LEON4FT processors. This allows 32 bytes to be read in 2 clocks, not counting the initial memory latency. To mask memory latency, the GRLIB L2 cache is used as a high-speed buffer between the external memory and the AHB bus. A read hit to the L2 cache typically requires 5 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 5 clocks and the second word will be delivered after one additional clock. A cache line will thus be fetched in 6 clocks (5 + 1).

2.3 Shared Floating-Point Units

The design features two floating-point units (FPU) that are shared between pairs of processors. Benchmarking has shown that the performance effects of FPU sharing are low for many applications. For instance the floating-point benchmarks from the SPEC CPU2000 benchmarks show a speed-up of between 1.00 and 1.08 for a system with dedicated FPUs compared to a system where each FPU is shared between two processors.

In benchmarks specifically designed to target the worst case behaviour of shared FPUs a speed-up of 1.35 can be observed for a system with dedicated FPUs.

As benchmarks of what is regarded as typical applications has shown a performance decrease of less than 10% for systems with dedicated FPUs the baseline specification for NGMP is to have FPUs shared between pairs of processors. The advantage of FPU sharing are savings in area and power. It should be noted that the

decision on FPU sharing will be re-evaluated for future implementations of the NGMP as the savings depend on the target technology.

2.4 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. As mentioned earlier the functional prototype implements these interfaces on separate pins.

For future NGMP developments, the selection between DDR2 and DDR1 should be regarded as open. The flight models of the NGMP are scheduled several 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 will also impact the final decision. The functional prototype will support DDR2-600 SDRAM and PC100 SDRAM.

The width of the memory interface is configurable via a bootstrap signal between 32 and 64 data bits (plus check bits). This allows both a reduced mode where a board is not fully populated with memory devices and also makes it possible to support packages with lower pin count.

To further improve resilience against permanent memory errors, the architecture supports a scheme where the interleaving mode of the error correcting code can be switched on-line. Software can initiate a switch between a mode that uses 32 check bits to a mode that makes use of 16 check bits. Combined with multiplexing this allows any faulty memory device to be removed from operation while keeping the system operational during the mode switch.

2.5 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. After feedback from early adopters the design was also extended to include a MIL-STD-1553B controller and a SPI controller.

The high-speed interfaces that are intended to be used in flight are the eight SpaceWire links, two 1000/100/10

Mbit Ethernet links, four high-speed serial links (high-speed serial links not included on functional prototype), and one 32-bit PCI 2.3 66 MHz master/target interface.

2.5.1 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 makes 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.

2.5.2 SpaceWire Links

Considering the wide adoption of SpaceWire links, the NGMP system will implement eight SpaceWire link interfaces directly on-chip. These links are connected to the GRSPWROUTER SpaceWire router 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 design also includes a GRSPW2 SpaceWire IP core that acts as a RMAP target on the Debug AHB bus and provides a SpaceWire debug link. This link has a separate SpaceWire link interface.

2.5.3 10/100/1000 Mbit Ethernet

The Ethernet interfaces will be supplied by the GRETH_GB1T core from Aeroflex Gaisler supporting 10/100/1000 Mbit/s operation. GRETH_GB1T 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.

2.6 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_GB1T Ethernet cores. An extension to the GRETH_GB1T 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.

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

Beyond support for standard symmetric multiprocessing (SMP) configurations, e.g. with a central multiprocessor interrupt controller, NGMP also supports asymmetric multiprocessing (AMP) configurations through the inclusion of additional timer units and an interrupt controller with AMP extensions.

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.

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.

2.8 Improved Support for Performance Measurement and Debugging

The NGMP includes 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 and 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.

3 TARGET TECHNOLOGY

The selected ASIC technology is Nextreme² from eASIC. The Nextreme² technology was announced by eASIC in 2008 and is a family of devices, manufactured on a 45 nm CMOS process, and built using eASIC's single-via customization technology. The device is fabricated by Chartered Semiconductor Manufacturing, who are part of Global Alliance.

The choice for target technology stood between a selection of large high-end FPGA devices, a multi-project-wafer (MPW) run or eASIC Nextreme². To manufacture the device on a MPW run could result in a higher target frequency. The main issue with MPW is that the number of delivered devices is typically low and that the delivered devices are not tested. Using a high-end FPGA device may have allowed the same maximum frequency as Nextreme² but at a higher cost per device

Nextreme² was considered a good trade-off based on the number of delivered, tested, devices and cost. Also, the manufacturing time advertised by eASIC between tape-out and delivery of tested devices is eight weeks.

4 VALIDATION BOARD

The baseline specification for the validation board is to support all implemented functions and interfaces of the design. The validation board, shown in figure 2, comprises a custom designed PCB in a 6U Compact PCI format, making the board suitable for stand-alone bench top development, or if required, to be mounted in a 6U CPCI Rack.

The board contains the following main items:

- LEON4 NGMP functional prototype device
- CPCI interface




Figure 2: NGMP Validation Board

- Memory
 - DDR SDRAM: 1 bank, 96 bits wide, DDR2-SODIMM (DIMM:s for data- and checkbits on opposite sides)
 - PC100 SDRAM 1 bank, 96 bits wide, discrete chips
 - Parallel Boot Flash 128 Mbit
- Power, reset, clock and auxiliary circuits
- Interface circuits required for the remaining interfaces

The principle interfaces and functions are accessible on the front and back edges of the board, and secondary interfaces via headers on the board. The interface connectors on the front edge of the board provide:

- Dual RJ45 10/100/1000 Mbit GMII/MII Ethernet interface
- 8 port SpaceWire interface (8 x MDM9S)
- SpaceWire debug communications link (MDM9S)
- USB debug communications link (USB-Mini-AB interface)
- 16-bit general purpose I/O
- MIL-STD-1553B interface (DE9 connector)
- FTDI serial-to-USB interface (provides access via USB cable to JTAG debug link and UARTs)

The interface connectors on the back edge of the board provide:

- Compact PCI interface (32 bit, 33/66 MHz), configurable for host or peripheral slot.
- Input power connector (+5V nominal) for stand-alone use when not powered via PCI.

To enable convenient connection to the interfaces, most connector types and pin-outs are compatible with the standard connector types for these types of interfaces.

Additionally, on board headers and components provide access to functions and features such as: DIP switches for GPIO signal configuration, LED indicators connected to GPIO signals, SPI interface on 0.1" header, two serial UART interfaces with DE9 female connectors, push buttons etc.

5 FUNCTIONAL PROTOTYPE VALIDATION

A validation plan for the NGMP devices has been established as part of the initial NGMP development. The validation plan covers verification of all IP cores and four multi-processor operating systems (eCos, Linux, RTEMS and VxWorks). The functional prototype validation is to repeat what has been done previously by

means of FPGA prototyping, but at higher operating speed. In addition to the tests previously performed on FPGA prototypes several performance/stress tests will be performed on the functional prototypes.

6 SOFTWARE SUPPORT

The GRMON debug monitor from Aeroflex Gaisler has been extended to support all new functionality for debugging and profiling included in the NGMP. The hardware platform provides 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 has been extended with additional support for booting AMP configurations.

7 INSTRUCTION SET SIMULATOR

The instruction set simulator for NGMP builds 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 final 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 is the goal during development.

8 CONCLUSION

The NGMP is 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 has improved support for profiling and debugging and a rich set of software immediately available due to backward compatibility with existing SPARC V8 software and LEON3 board support packages. NGMP also has specific support for AMP configurations and Time-Space Partitioning.

The development of a NGMP functional prototype is done to further evaluate the design and to provide a software development platform that is close to the baseline NGMP design. Previously, software efforts have used FPGA prototypes that only provide downsized configurations of the NGMP design running a considerably lower clock frequencies compared to the functional prototype.

The NGMP is part of the ESA roadmap for standard mi-

croprocessor 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>