Technical Note on NGMP Verification

Next Generation Multipurpose Microprocessor

Contract: 22279/09/NL/JK

TABLE OF CONTENTS

1 INTRODUCTION

1.1 Scope of the document

This document is a technical note describing part of the verification performed during the architectural design phase of the Next Generation Multipurpose Microprocessor (NGMP). The NGMP is developed within an activity initiated by the European Space Agency under ESTEC contract 22279/09/NL/JK.

The work has been performed by Aeroflex Gaisler AB, Göteborg, Sweden.

1.2 Overview of design

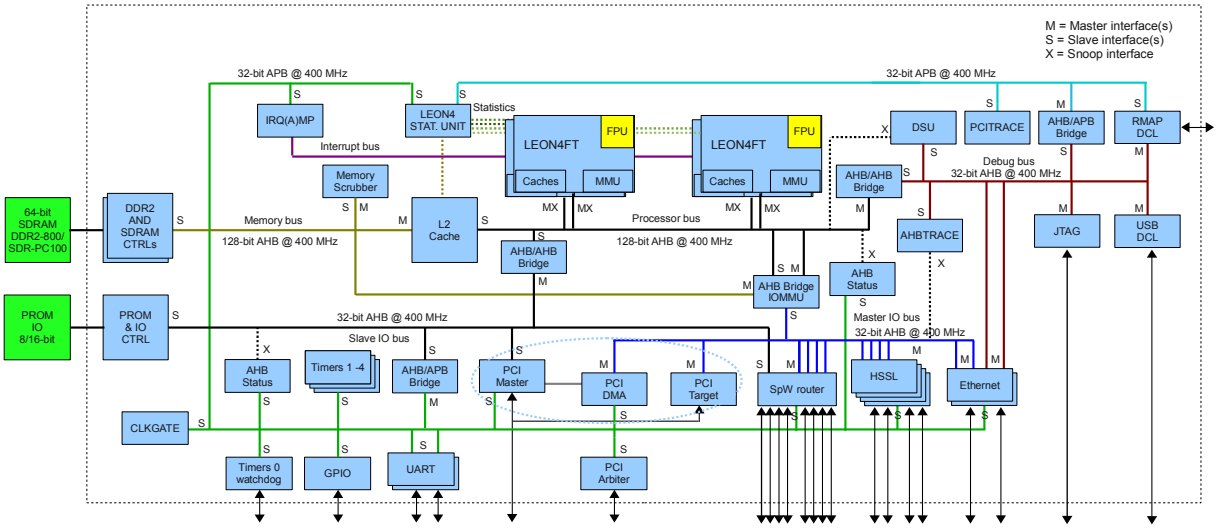

Figure 1: Overview block diagram

The system consists of five AHB buses; one 128-bit Processor AHB bus, one 128-bit Memory AHB bus, two 32-bit I/O AHB buses and one 32-bit Debug AHB bus. The Processor AHB bus includes 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 and SDRAM interfaces on shared pins, and includes a memory scrubber. The I/O bus has been split into two separate buses where all slave interfaces have been placed on one of the buses (Slave I/O AHB bus) and all master interfaces have been placed on the other bus (Master I/O AHB bus). The Master I/O AHB bus connects to the Processor AHB bus via an AHB/AHB bridge that provides access restriction and address translation (IOMMU) functionality. This AHB/AHB bridge also has a master interface connecting it to the Memory AHB bus. The AHB master interface to use when propagating traffic from a core on the Master I/O AHB bus is dynamically configurable.

The two I/O buses include all peripheral units such as PCI, Ethernet MACs, and SpaceWire interfaces. The dedicated 32-bit Debug AHB bus connects one debug support unit (DSU), JTAG, Ethernet, USB and SpaceWire debug links, AHB and PCI trace buffers and also provides a direct link to the LEON4 statistics unit that contains performance counters. The Debug AHB bus allows for non-intrusive debugging through the DSU and direct access to the complete system.

The target frequency of the LEON4FT processor cores and on-chip buses is 400 MHz, but depends ultimately on the implementation technology.

2 NGMP VERIFICATION

Variations of the NGMP design were implemented on several FPGA prototyping boards. The FPGA boards are off-the-shelf products from Aeroflex Gaisler, Xilinx and Synopsys.

The FPGA prototyping approach has been processor driven, based on the LEON3/4 software environment. This facilitates re-use of existing resources and ensures that the overall objective of a processor-controlled device is achieved.

The software development tools are based on the well-known LEON cross compilers, developed by Aeroflex Gaisler. The GRMON debug monitor was used for communication between the host system and target system.

2.1 Limitations

Since off-the-shelf development boards were used for FPGA prototyping it was expected that no board could be found that allowed using all interfaces of the NGMP design. In addition to this the NGMP design requires a large FPGA if the full design is to be implemented. The clock frequencies of the FPGA prototypes were also lower than the target frequency of the final chip. Other deviations from the final ASIC design include:

- Macros, such as PLL and DDR2 PHY
- DDR2 SDRAM and SDR SDRAM on shared pins was not possible prototype on FPGA
- Several of the core buffers will be implemented with flip-flops in the final design. On FPGA this may lead to the design growing to large and RAM blocks may be used instead.
- The SERDES (HSSL) link can not be verified.

2.2 Operating systems and drivers

Test suites and custom test application were run to demonstrate the hardware and to validate the porting work made to the operating systems. RTEMS 4.10, eCos 2.0, VxWorks 6.7 and Linux 2.6 were demonstrated to be operating correctly. Test were also run with the MKPROM2 bootloader creation tool.

2.3 Concurrent SMP and ASMP Configurations

The NGMP system has been designed with extended support for running ASMP configurations. The FPGA prototyping included two different tests, using two different operating systems; RTEMS and Linux. These tests demonstrate:

- SMP system: Linux SMP on three CPUs
- Heterogeneous AMP system: Linux on three CPUs and RTEMS on one CPU
- Homogeneous AMP system: RTEMS on two CPUs

The two tests successfully demonstrated how resources of the NGMP design can be shared in a heterogeneous AMP system, in a SMP System and in a homogeneous AMP system. It also demonstrated the flexibility of the extended multiprocessor interrupt controller controller: how different CPUs can be routed to a unique interrupt controller, and how multiple CPUs can be routed to the same interrupt controller.

2.4 I/O Performance Evaluation

Performance and throughput measurements are typically performed on FPGA prototypes due to the prohibitively long simulation runs required to measure throughput. The results from FPGA prototyping are then scaled if the FPGA prototype is not representative of the final design.

For NGMP, there are some difficulties involved in scaling the results from FPGA prototypes. Due to the placement of masters and main memory in the NGMP bus topology, masters on the Master I/O bus experience significant latencies on accesses made to main memory. In a typical LEON system, the AHB masters are placed on the same bus as the processors and main memory. In the NGMP system, masters must perform accesses over the IOMMU and possibly over the L2 cache in order to fetch data from main memory. Traversing the IOMMU and L2 cache adds latency cycles to each single access or to each block of burst accesses. These latency cycles, combined with latency from external memory is the main limitation to I/O throughput (apart from limitations in the bus fabric itself; arbitration cycles, data bus width etc.).

On a prototype system with a system frequency of 50 MHz, the latency clock cycles will give a latency that is eight times higher than the latency time in the final system. This is fine if input traffic can be scaled to be eight times slower. This is not always possible. One example applicable to the NGMP system is gigabit Ethernet. To test I/O throughput, one of the test cases would typically involve transferring data with TCP/IP. When the system cannot handle the stream of data, the Ethernet controller will experience overruns and packets will be dropped. This will lead to packets being re-sent, possibly in a slower rate. As soon as packets are being re-sent it becomes difficult to scale the results.

As the NGMP FPGA prototypes may not be representative for throughput tests, simulations were performed in order to build a view of the design's performance when running at the target frequency.

Traffic was generated via the SpaceWire router and Ethernet cores and tests were run on several configurations, including: L2 cache disabled, L2 cache enable, L2 cache with fault-tolerance enabled.

The results are summarised in the table below. The first column lists the configuration, the second column (1x Eth) lists the results for running the GRETH_GBIT throughput test on the first Ethernet controller. The column 2x Eth lists the results, per Ethernet core, from running the GRETH_GBIT throughput test on both controllers. Next (SpW) the results from running the SpaceWire router throughput test is shown divided per AMBA port and a total for all AMBA ports. The columns under combined test shows the results when running all tests simultaneously.

Configuration	**1x Eth**	**2x Eth**		**SpW**		**Combined test**			
		Eth 0	**Eth 1**	**Per port**	**Total**	**Eth 0**	**Eth 1**	**SpW/ port**	**SpW total**
L2 cache disabled	1.2 Gb/s	730 Mb/s	790 Mb/s	394 Mb/s	1.57 Gb/s	438 Mb/s	480 Mb/s	216 Mb/s	865 Mb/s
L2 cache enabled	1.7 Gb/s	1.7 Gb/s	1.7 Gb/s	1.56 Gb/s	6.25 Gb/s	1.4 Gb/s	1.5 Gb/s	1 Gb/s	4 Gb/s
L2 cache FT enabled	1.7 Gb/s	1.7 Gb/s	1.7 Gb/s	1.5 Gb/s	6.1 Gb/s	1.4 Gb/s	1.4 Gb/s	1.5 Gb/s	3.9 Gb/s

Table 1: Throughput of 400 MHz system

2.5 Memory Segregation Capability

The NGMP design provides separation capabilities via processor memory management units (MMUs) and a I/O Memory Management Unit (IOMMU). These capabilities were demonstrated with a test consisting of two RTEMS images, running in parallel trying to access allowed and protected memory areas using their respective CPU (load/store instructions). Attempts to access protected areas were also done using RMAP commands to the SpaceWire router's AMBA ports. The attempts to access protected areas were effectively blocked by the processor MMUs and the IOMMU.

2.6 Multi-core Debugging Support

To demonstrate multi-core debugging support, Aeroflex Gaisler's debug monitor GRMON is used to control and view the state of multiple CPU.

Four RTEMS images, one per CPU, are loaded to memory. Initialisation of CPU individual stack pointers and entry points are demonstrated using GRMON commands from a batch script. Once initialised and booted, the CPUs communicate with each other using shared memory. Each CPU spins in a tight busy loop waiting for a message to be passed on to the next CPU.

The test demonstrated:

- Stack pointer per CPU
- Entry point per CPU
- Viewing register file per CPU
- Viewing instruction trace per CPU
- Viewing bus transactions per CPU
- Instruction and bus transactions are sampled using the same counter making it possible to determine timing relative to each other
- Hardware watch points per CPU
- Hardware break points per CPU
- Continuing execution of all CPUs
- A practical example how to inspect D-cache snooping in action by inspecting the instruction traces of multiple CPUs

2.7 Sample benchmarks

To compare the performance or the NGMP to previous LEON2 and LEON3 systems, a small collection of benchmarks has been developed. These benchmarks can be compiled with the BCC tool-chain and run on systems without an OS and MMU. While not providing an exhaustive performance profile, these benchmarks still provide interesting compare points in the development of the LEON processor.

The benchmarks have been run on the following systems: AT697, UT699, GR712RC, NGMP.

The systems have the following processor configuration:

- AT697: LEON2FT, 32K + 16K cache, 5-clock MUL, load delay 1, Meiko FPU
- UT699, LEON3FT V1, 8K + 8K cache, 5-clock MUL, load del 2, GRFPU
- GR712: LEON3FT V2, 16K + 16K cache, 5-clock MUL, load del 1, GRFPU, branch pre.
- NGMP: LEON4FT, 16K + 16K L1 cache, 2-clock MUL, load del 1, GRFPU, 128K L2 cache

The following benchmarks will be run:

- 164.gzip (from the SPEC CPU2000 suite)
- 176.gcc (from the SPEC CPU2000 suite)
- 256.bzip2 (from the SPEC CPU2000 suite)
- AOCS benchmark
- Basicmath_large
- Coremark-1.0
- Dhrystone-2.0
- Linpack-DP
- Whetstone

All benchmarks have been compiled with gcc-4.4.2 tuned for SPARC V8.

All systems were clocked at 50 MHz during the tests, using 32-bit SDRAM (LEON2/3) or 64-bit DDR2 (NGMP). Table 2 shows the performance figures relative to AT697.

Benchmark	AT697	UT699	GR712RC	NGMP
164.gzip	1	0.94	1.1	1.31
176.gcc	1	0.79	0.97	1.3
256.bzip2	1	0.93	1.06	1.33
AOCS	1	1.2	1.52	1.79
Basicmath	1	1.3	1.46	1.62
Coremark, 1 thread	1	0.89	1.09	1.21
Coremark, 4 threads	1	0.89	2.05	4.59
Drystone	1	0.94	1.05	1.39
Drystone, 4 instances	1	0.94	1.61	4.81
Linpack	1	1.2	1.26	1.71
Whetstone	1	1.94	2	2.22
Whetstone, 4 instances	1	1.94	3.7	8.68

Table 2: Performance comparison

The table shows that the LEON4/NGMP system has approximately 30% better CPI than AT697 on integer benchmarks, and up to 100% better CPI on floating-point benchmarks. The Coremark benchmark can also be run multi-threaded, which shows on the high 4-thread results for GR712RC and NGMP. The benchmark will fit in the L1 cache, and therefore scales almost linearly with number of cores.

All benchmarks were run using the BCC runtime. Using the Linux SMP OS, multiple instances of Dhrystone and Whetstone was run on GR712RC and NGMP. It shows that performance scales better on NGMP than GR712RC, mostly due to wider data-paths and the presence of an L2 cache.

3 SUMMARY

The system has been verified by means of VHDL simulation and FPGA prototyping covering. The objective during the architectural design phase has been to verify the processing capabilities of the NGMP system. The system level tests of fault-tolerance capabilities have been deferred to the Preliminary SEE validation to be performed during the detailed design phase. Preliminary SEE validation is foreseen to include running the same set of tests, but also with error injection enabled.

Test results show that porting of all operating systems and driver development has been successful. Performance measurements indicate that the NGMP design is within specification. Providing enough bandwidth to satisfy four 6.25 Gb/s high-speed serial links will be challenging however.