

Next Generation Microprocessor Functional Prototype SpaceWire Router Validation Results

SpaceWire Components, Short Paper

Jonas Ekergrarn, Jan Andersson, Andreas Larsson,
Daniel Hellström, Magnus Hjorth
Aeroflex Gaisler AB
Göteborg, Sweden

Roland Weigand
Microelectronics Section
European Space Agency
Noordwijk, The Netherlands

Abstract—The Next Generation Microprocessor is a quad-processor system-on-chip that contains a SpaceWire router with eight external SpaceWire links and four on-chip AMBA ports. This paper describes the validation work done for the SpaceWire router within the Next Generation Microprocessor functional prototype development.

Index Terms— SpaceWire, Networking, Spacecraft Electronics.

I. INTRODUCTION

The Next Generation MicroProcessor (NGMP) is a quad-processor system-on-chip currently being developed by Aeroflex Gaisler. The design includes four LEON4 SPARCV8+ processors with a shared Level-2 cache, DDR2-800 SDRAM main memory interface, a SpaceWire router with eight external SpaceWire links and four internal AMBA ports, two 10/100/1000 Mbit Ethernet MACs, 32-bit 66 MHz PCI interface and other interfaces.

The SpaceWire router allows the NGMP to act both passively and actively in a SpaceWire network. The target frequency for the NGMP device is 400 MHz. Preliminary results for this target frequency show that, using only internal

routing, the architecture is able to sustain a data throughput of 1.5 Gb/s per SpaceWire AMBA port. In a scenario where the two full-duplex Ethernet links and all SpaceWire AMBA ports are run at full speed, the sustainable throughput is roughly 1.5 Gb/s for the Ethernet links and 1 Gb/s per SpaceWire AMBA port. In addition to this, the SpaceWire router will also be able to simultaneously route packets at maximum speed.

The implementation of NGMP in rad-hard technology was put on hold in April 2011, pending advances in the development of a suitable Deep-Sub-Micron technology for space. Development has instead progressed in the development of a NGMP functional prototype (NGFP) device targeting eASIC Nextreme2, a structured ASIC technology based on a 45 nm process. Silicon was received in August 2012 and an evaluation board has been manufactured.

One of the primary goals of the NGFP development is to allow use of the architecture at higher clock frequencies than what is attainable with FPGA prototype implementations. The prototype devices do not reach the full target frequency of the final device (400 MHz) but can be run at a system frequency of 200 MHz with the same clock frequency used for the SpaceWire router.

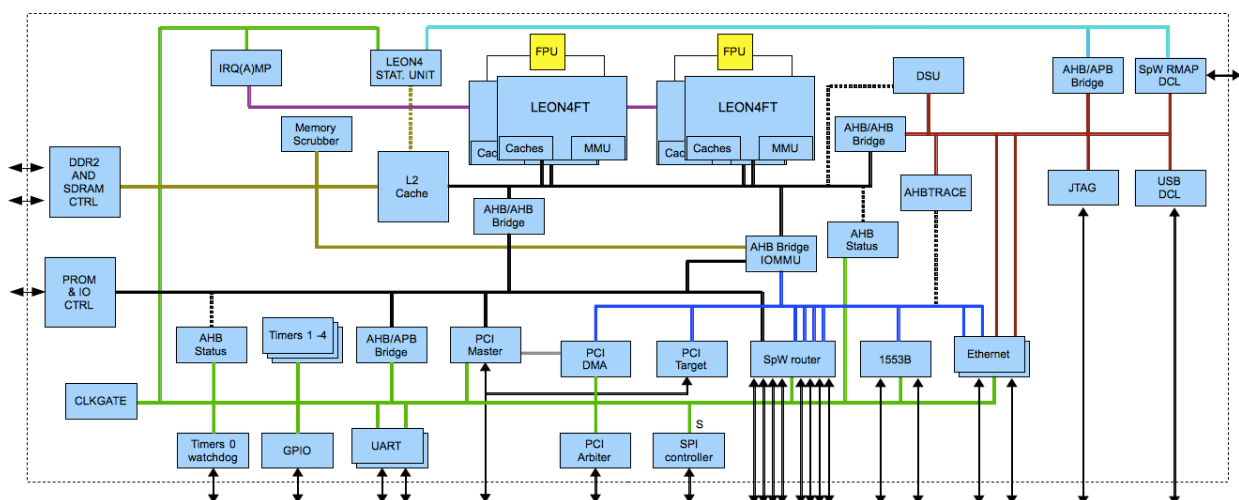


Fig. 1 Block diagram

II. DESCRIPTION OF NGMP FUNCTIONAL PROTOTYPE ARCHITECTURE

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 connects four LEON4 processor cores connected to a shared Level-2 (L2) cache. The Memory bus is located between the L2 cache and the main external memory interfaces, DDR2-600 SDRAM and PC100 SDRAM, and also connects a hardware 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 connect all the peripheral cores. All memory-mapped interfaces of peripheral cores that can be directly accessed by the processors 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 timer units, interrupt controller, UARTs, general purpose I/O port, PCI master/target, Ethernet MACs, MIL-STD-1553B, Serial Peripheral Interface bus and SpaceWire router. 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, as an example, Ethernet and SpaceWire packets will not increase the 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 bridge connecting the Debug bus to the Processor bus allows unrestricted access to the memory space.

The NGMP architecture has been designed to provide a significant performance increase compared to earlier generations of European space processors. The platform has 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. The design also includes specific support for asymmetric multi-processing configurations. Five memory management units (MMUs), one per CPU core, and the IOMMU provide access protection. Several dedicated interrupt controllers allow interrupt steering to a specific CPU and duplicated timer units allow to run one operating system per CPU core with full space-partitioning.

III. SPACEWIRE ROUTER IP CORE AND CONFIGURATION

The design includes Aeroflex Gaisler's GRSPWROUTER SpaceWire router IP core. The IP core implements a SpaceWire routing switch as defined in the ECSS-E-ST-50-12C standard. It provides an RMAP target for configuration port 0 used for accessing internal configuration and status registers. In addition to this, the implementation described by

this paper implements two different port types; external SpaceWire links and on-chip AMBA interfaces.

One AMBA AHB slave interface is also provided for access in the port 0 registers from the on-chip AMBA bus. Group-adaptive routing and packet distribution are fully supported.

The GRSPWROUTER was implemented with the following characteristics:

- 64 entries per 9-bit receiver FIFO (N-Char FIFO)
- 32 entries per 32-bit AMBA port FIFO
- Four DMA channels per AMBA port
- Hardware RMAP target in each AMBA port

IV. SPACEWIRE ROUTER ROLE IN SYSTEM-ON-CHIP DESIGN

The system-on-chip architecture is a multi-processor 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 NGMP was initially specified to include for SpaceWire codecs with AHB host interfaces and hardware RMAP targets. The four SpaceWire codecs would use redundant ports giving a total of eight external SpaceWire links, where four of the links could be used separately.

The four SpaceWire codecs were later replaced by the SpaceWire router. The register interface of the AMBA ports of the SpaceWire router are software compatible with the register interface of the previously used codecs, giving little or no overhead for software implementations. While the SpaceWire router supports redundant ports the choice was made to implement the eight ports as separate links and instead to recommend group-adaptive routing as an alternative to the redundant port feature. This gives users the alternative to forego redundancy and instead use all eight available links simultaneously.

The NGMP is targeted at general payload processing with the main design goal of increasing the average processing performance. The main use of the SpaceWire router within this context is not to route SpaceWire traffic from external entities but instead to provide the same functionality as the previously included SpaceWire codecs.

The inclusion of the SpaceWire router provides more options to system designers. The device can be used to both provide SpaceWire connectivity to the on-chip processing components while also acting as a router for external entities. Protection mechanisms in the architecture also allow the use of the SpaceWire router to be completely separate from the rest of the design. In effect packaging the router and microprocessor components together with the gain of reducing the number of required devices.

V. TRAFFIC ROUTING

The first use of the SpaceWire router within the NGFP validation effort was to study the effects of routing AMBA traffic either through or behind the system's Level-2 cache.

The test consists of an RTEMS application that is transferring data over the four AMBA ports simultaneously.

The router is configured to route the SpaceWire packets from AMBA port 0 to AMBA port 1, then AMBA port 1 to port 2 and so on back to AMBA port 0. This means that every packet will exercise eight DMA operation channels every time one packet makes one round-trip. The number of round-trips is counted and performance figures are calculated based on these counts.

The packets are marked with a unique sequence number and contain 16-bit incremented data. This is done to be able to verify packet receive/transmit ordering and data correctness. The packet sequence is verified for every received packet. The data is verified after all transmissions are finished.

The tests were performed with the internal SpaceWire fabric and AMBA system running at 200 MHz.

The test was run in several different configurations, of which three are considered here:

- CFG2 – Cache-coherent system with Level-2 cache, caching all traffic
- CFG5 – Cache-coherent system with Level-2 cache, SpaceWire DMA buffers and traffic not cached
- CFG10 – System with Level-2 cache. SpaceWire DMA buffers not cached by Level-2 cache. SpaceWire DMA traffic does not pass through Level-2 cache. In this configuration the cache coherency of the L1 cache cannot be maintained through bus snooping. The processor MMU is used to mark the DMA buffers as noncachable to solve the coherency issue.

The results of the tests showed that the highest performing configuration is CFG2 where the Level-2 cache caches all DMA traffic. This is expected as the software execution causes little interference and the Level-2 cache is essentially dedicated for DMA buffers. In a configuration where additional software instances made use of the Level-2 cache it is expected that the SpaceWire traffic throughput and software application performance would be negatively affected due to the shared resource in the Level-2 cache. The combined throughput for all DMA ports was measured to 1.54 Gbit/s.

The test case CFG5 showed that SpaceWire throughput is more than halved when marking the DMA buffers as uncached in the Level-2 cache. In this configuration the Level-2 cache does not add any benefit when fetching data from external memory, instead the cache only adds latency on each DMA access.

The CFG10 test case showed the effects of bypassing the Level-2 cache and routing traffic directly to the main memory controller and was completed using a FPGA prototype due to NGFP IOMMU silicon errata. The test showed that the throughput decreases with 18% in CFG10.

While the data throughput for this particular test is lower when bypassing the Level-2 cache it is important to recognize the effects on the processor system. When bypassing the Level-2 cache the DMA traffic will have negligible, if any, impact on software instances with high Level-1 and Level-2 cache hit rates. This allows large amount of data to be transferred to main memory without processor intervention and without impacting performance of software.




Fig. 2 Example of test rig setup

VI. SPACEWIRE ROUTING TESTS

The traffic routing test described in the previous section studied the effects of using the AMBA ports to transfer large amounts of data. The second set of validation tests performed on the functional prototype device that involved the SpaceWire router focused on the routing capabilities of the router. The tests were divided into four major groups:

- All SpaceWire ports – Exercise the router by generating traffic on all ports
- Group-adaptive routing
- Packet distribution
- Priority routing
- Packet timeout

All tests described below were performed with the internal SpaceWire fabric running at 200 MHz and the AMBA system running at 200 MHz. All SpaceWire links were configured to operate at a bitrate of 200 Mbit/s.

A. All SpaceWire ports

To validate that all the SpaceWire ports of the SpaceWire router can handle both receive and transmit at a rate of 200 Mbit/s, each SpaceWire port was connected to another SpaceWire port. 4 MiB packets were then sent from an AMBA port, routed out onto a SpaceWire port, received at another SpaceWire port, and then routed to an AMBA port where the data was validated. This test was repeated so that all SpaceWire ports were utilized, and both path addresses and logical addresses were used for the packets.

B. Group adaptive routing

The SpaceWire router supports group adaptive routing for all path addresses and logical addresses. Group adaptive routing means that packets can be routed through the network over different paths depending on which of the router's ports that are available when the packet arrives. For example, a packet with address 0x40 arrives at SpaceWire port 1 of the router, and address 0x40 is configured with group adaptive routing to SpaceWire port 2 and 3. The router will then route the packet to either port 2 or port 3 depending on which port becomes available first. If both ports are available, the router will send the packet on the port with the lowest port number. The group adaptive routing mechanism was validated by connecting four SpaceWire ports together and then sending packets from an AMBA port where the address byte of the packets were configured with group adaptive routing to two of the four ports. When the packets arrived at the router again they were routed to another AMBA port. It was then verified that the packets arrived correctly as long as one of the two SpaceWire used as output ports were connected to another port. If none of the two SpaceWire ports used as output ports were connected then the packet was not received at the AMBA port used as destination. Group adaptive routing was also verified further in the packet distribution validation (see below).

C. Packet distribution

Packet distribution - which means that data arriving at a input port is sent to multiple ports simultaneously - is supported by the SpaceWire router for both path addresses and logical addresses. This feature was validated by connecting four SpaceWire ports to each other and then sending a packet with two address bytes from an AMBA port. The first address byte was configured with header deletion and packet distribution out on the four SpaceWire ports, and the second address byte was configured with group adaptive routing to AMBA ports 0-3. When the packet was sent from the AMBA source port the first address byte was removed by the use of header deletion, and the packet was routed out onto the four SpaceWire ports. It was then verified that the four packets, arriving at one SpaceWire port each, was routed to one AMBA port each (because group adaptive routing was used for the second address byte). This test also adds additional validation of group adaptive routing since the test validates that group adaptive routing works when the destination ports are busy with transmitting data. The validation of group adaptive routing described above only validated the case when the destination links were not running.

D. Priority routing

When packets are to be routed, each destination port is arbitrated individually using a two level priority. The priority is based on the first address byte of the incoming packet, and all path addresses and logical addresses can be assigned either a high or low priority. Round-robin is used when one or more packets with the same priority competes about the same destination port. The validation of the priority routing mechanism was done by enqueueing four different packets, each one from a different AMBA port, where all packets were to be routed out on the same SpaceWire port. Three of the packets contained an address that had been assigned a low priority, while the fourth packet contained an address with high priority. The SpaceWire port that the packets would be routed out onto was connected to another SpaceWire port of the router, and the second address byte in all packets was the path address of one of the AMBA ports (same for all packets so that the order could be observed). The three low priority packets were sent slightly before the high priority packet, and it was then validated at the destination AMBA port that the first packet received was the first low priority packet, followed by the high priority packet, and then followed by the two remaining low priority packets. It was also validated that if the high priority packet was instead changed to low priority it was received last of the four packets.

E. Packet timeout

The SpaceWire router implements packet timers in order to prevent situations where the ports become blocked for ever if, for example, a source stops sending data without terminating the packet with an end of packet marker (EOP and EEP). In such a situation the router will detect that no data has been sent for a certain amount of time (configurable), and the packet will then be spilled and the destination port released. Connecting two of the router's SpaceWire ports to each other has validated

the packet timeout feature. Then a packet containing two address bytes was sent from an AMBA port. The first address byte made the router route the packet out onto the first of the two mentioned SpaceWire ports. The second address byte made the router try to route the packet to a SpaceWire port that was not connected to anything. After the first packet was sent another packet was sent from a second AMBA port. The first address bytes of the second packet informed the router to route the packet out onto the same SpaceWire port as the first packet, and the second address byte was the address of a third AMBA port. In the case that the SpaceWire routers timers were not enabled it was verified that the second packet never reached its destination (because the first packet blocks the outgoing SpaceWire port for ever). If timers were enabled it was verified that the second packet eventually reached its destination, since the first packet was spilled by the router after a timeout period when it failed to route it.

VII. CONCLUSION

The parts of the NGFP validation effort that included the SpaceWire router aimed to prove the design decision to allow AMBA traffic to be routed so that it bypasses the Level-2 cache and to demonstrate core functionality of the router.

The functionality to bypass the Level-2 cache was successfully demonstrated using the SpaceWire router AMBA ports and as a side effect also verified high-speed communication between the router's AMBA ports.

Core functionality of the router was also demonstrated by generating traffic on all ports and execution of test cases using group-adaptive routing, packet distribution, priority routing and packet timeouts.

The NGMP is part of the ESA roadmap for standard microprocessor components and it will be commercialized under fair and equal conditions to all users in the ESA member states. The NGMP is 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.

ACKNOWLEDGMENT

The Next Generation Microprocessor is developed in activities commissioned and funded by the European Space Agency under contracts 22279/09/NL/JK and 18533/NL/JD.