Payload processing based on LEON3FT SMP architecture

Gaisler, J.
Gaisler Research

"1. Introduction Payload processing has traditionally been performed by dedicated signal processing units, either in form of a DSP processor or by custom designed hardware. With the availability of LEON3 SMP architectures, it has become feasible to implement certain signal processing algorithms on a more general purpose processor such as LEON3. This presentation will outline some possible LEON3 multi - processing architectures, and discuss both the performance that can be reached and the limiting factors. 2. LEON3 multi - processor support The LEON3 multi - processor suppor t features allow to attach up to 16 processor to the same AMBA AHB bus for increased performance. Cache snooping, processor enumeration and interrupt steering allows the hosting of SMP- capable operating systems and applications. The limiting factor in SMP systems is mostly the bandwidth of the AHB bus, and attaching more than four processor s is rarely recommended as it does not bring any further performance gains. 3. LEON3 GINA architecture To provide a baseline architecture for payload processing, the LEON3 'GINA' architecture has been developed under an ESA contract. This architecture includes four LEON3FT processor s with one fully pipelined FPU each, four Spacewire links, dual 100 Mbit Ethernet MACs, PCI, and a number of additional on- chip peripherals. Targeted for a 130 or 90 nm process, this architecture can provide a peak performance of 1 GIPS or 1 GFLOPS when clocked at 266 MHz. The average throughput depends strongly on the application and data rates, but a sustained performance of 500 MIPS /MFLOPS has been simulated. 4. Limiting factors To analyze the applicability of GINA for a certain payload application, the limiting factors needs to be understood. In data - intensive applications, the overall memory bandwidth will become the bottle - neck. Using space- qualified SRAM, the external memory bandwidth is in the range of 500 – 750 Mbit / s (using a 32- bit bus). Using SDRAM, the practical bandwidth is ~ 1.5 Gbit / s using 32- bit bus and ~ 3 Gbit / s using 64- bit bus. Data from instrument s will pass the main memory at least twice, once when received from the instrument and once when read by the processor(s) into the cache. This means that already two Spacewire interfaces at 200 Mbit / s will consume a major part of the available bandwidth. A second limiting factor is the computational throughput . To achieve maximum performance, the inner loops of the computational algorithm should be coded in assembly. Even under such conditions, it is difficult to reach more than ~ 50% utilization of the floating- point unit due to data latencies and loading and storing of operands. The GINA architecture is thus suitable for payload processing where data rates are below ~ 500 Mbits / s and the aggregated throughput requirement is less than 500MIPS/MFLOPS."