

Current and Next Generation LEON SoC Architectures for Space

Flight Software Workshop 2012
November 7th, 2012

www.aeroflex.com/gaisler



Outline



- History of LEON Project
- LEON3FT Overview
- LEON4 Overview
- LEON3FT systems: UT699, AG RTAX, GR712RC
- LEON4 Next Generation Microprocessor
- Summary

LEON Processor History



- Project started by European Space Agency in 1997 with the objectives:
 - To Provide an open, portable and non-proprietary processor design → Based on SPARC architecture
 - To be able to manufacture in SEU sensitive semiconductor process and to maintain correct operation in presence of SEUs
- LEON1 VHDL design (ESA) → LEONExpress test chip (2001, 0.35 um)
- LEON2(FT) VHDL design (ESA / Gaisler Research) → AT697 (2004)
- LEON3(FT) VHDL (Gaisler Research) → UT699 (Aeroflex CoS)
- LEON4(FT) (Aeroflex Gaisler) → Next Generation Microprocessor

LEON3FT

- IEEE-1754 SPARC V8 compliant 32-bit processor
 - 7-stage pipeline, multi-processor support
 - Separate multi-set caches
 - On-chip debug support unit with trace buffer
 - Highly configurable
 - Cache size 1 - 256 KiB, 1 - 4 ways, LRU/LRR/RND
 - Hardware Multiply/Divide/MAC options
 - SPARC Reference Memory Management Unit (SRMMU)
 - Floating point unit (high-performance or small size)
 - Fault-tolerant version available
 - Register file protection
 - Cache protection



Certified SPARC V8 by SPARC international

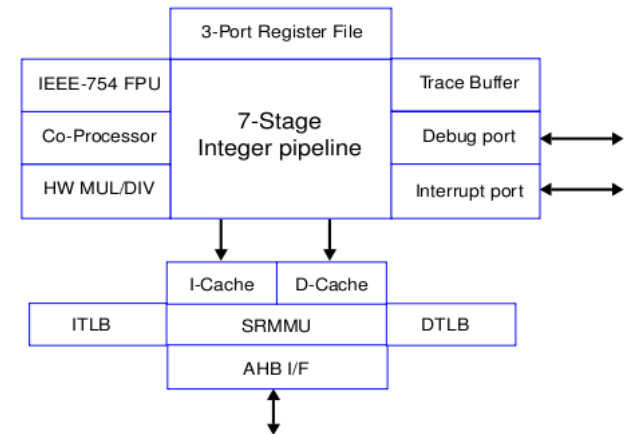
Suitable for space and military applications

Baseline processor for space projects in US, Europe, Asia

LEON4

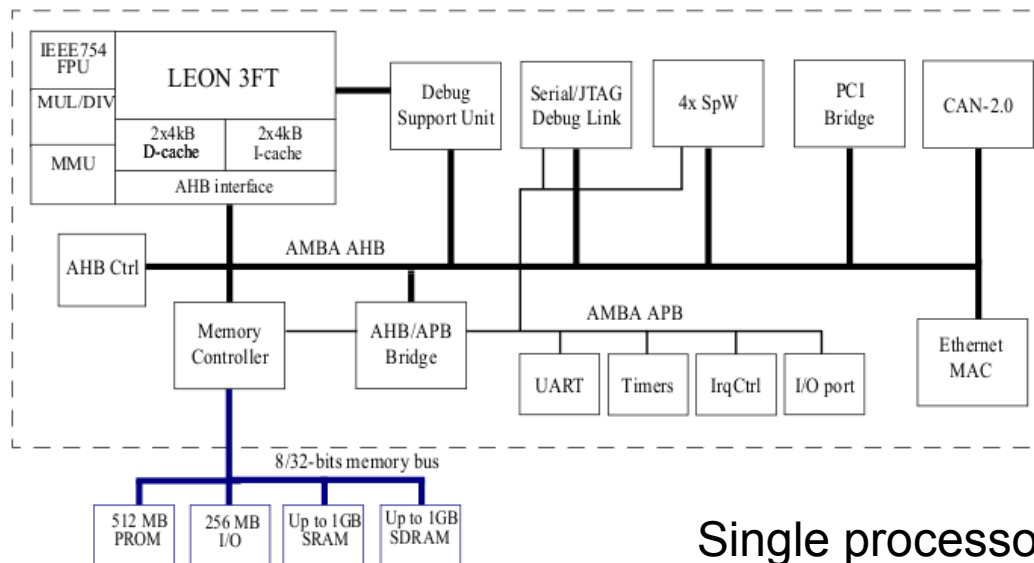


- IEEE-1754 SPARC V8 compliant 32-bit processor
- Same features and architecture as LEON3 (7-stage pipe, MP support etc.)
- Offers higher performance compared to LEON3
 - 64-bit single-clock load/store operation
 - 64- or 128-bit AHB bus interface
 - Branch prediction
 - SPARC V9 Compare-and-Swap instruction
 - Performance counters
 - 1.7 Dhrystone MIPS/MHz
 - 0.6 Whetstone MFLOPS/MHz
 - 2.1 CoreMark/MHz (comparable to ARM11)
 - Typically used with L2 cache to reduce effects of memory acc. Latency
- Currently in use for commercial projects and also for ESA's "Next Generation Microprocessor"



Traditional LEON SoC Architecture

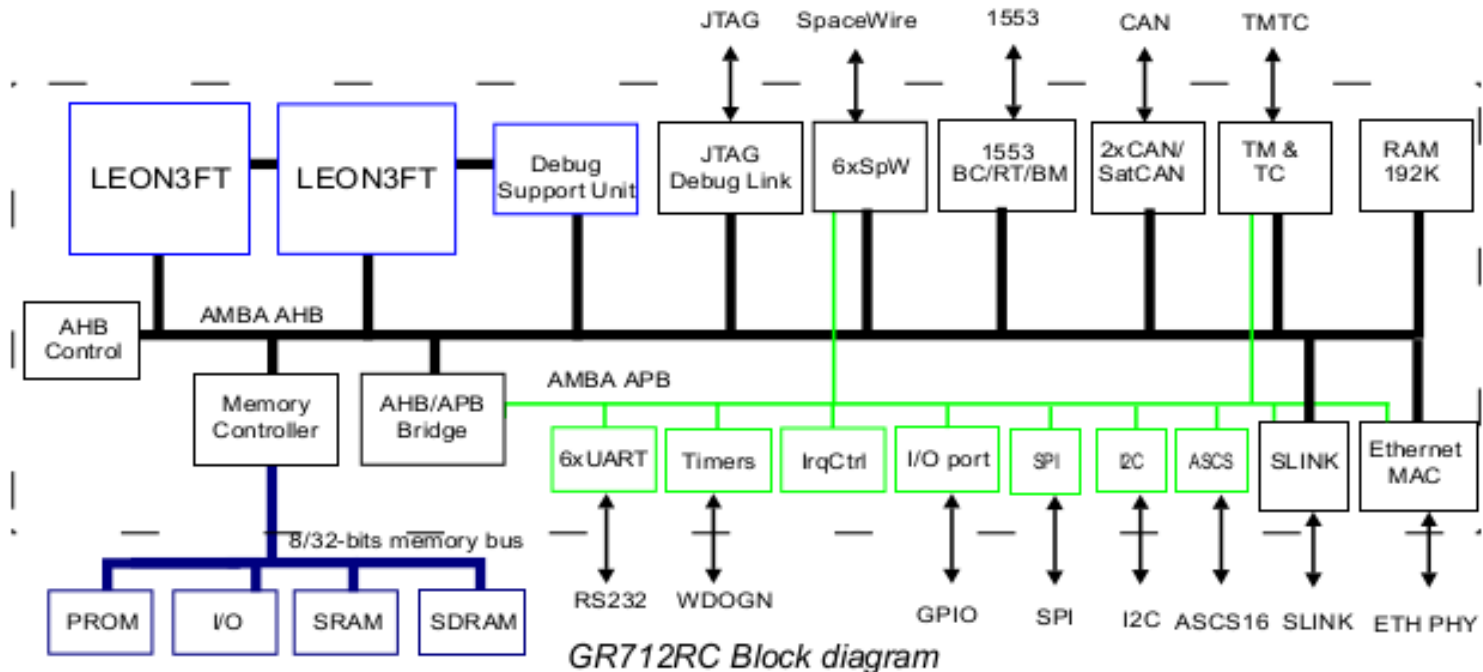
- Aeroflex Gaisler LEON3FT-RTAX, Aeroflex CoS UT699, etc.



Single processor core
AMBA AHB 2.0 with one or
several AHB/APB bridges
DMA capable peripheral units

Current Multi-core LEON: GR712RC

- GR712RC: Dual core LEON3FT



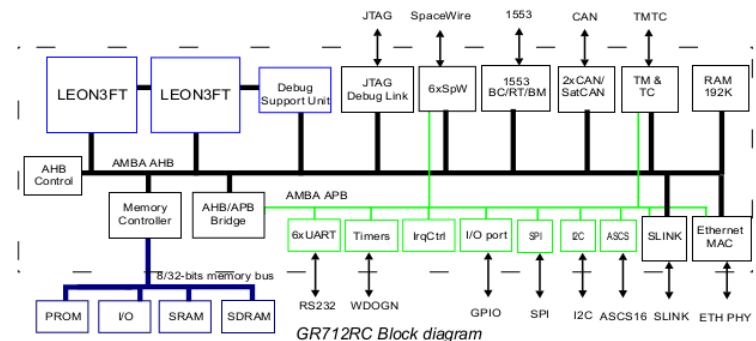
Traditional topology
Two processor core
connected to same bus

GR712RC Features



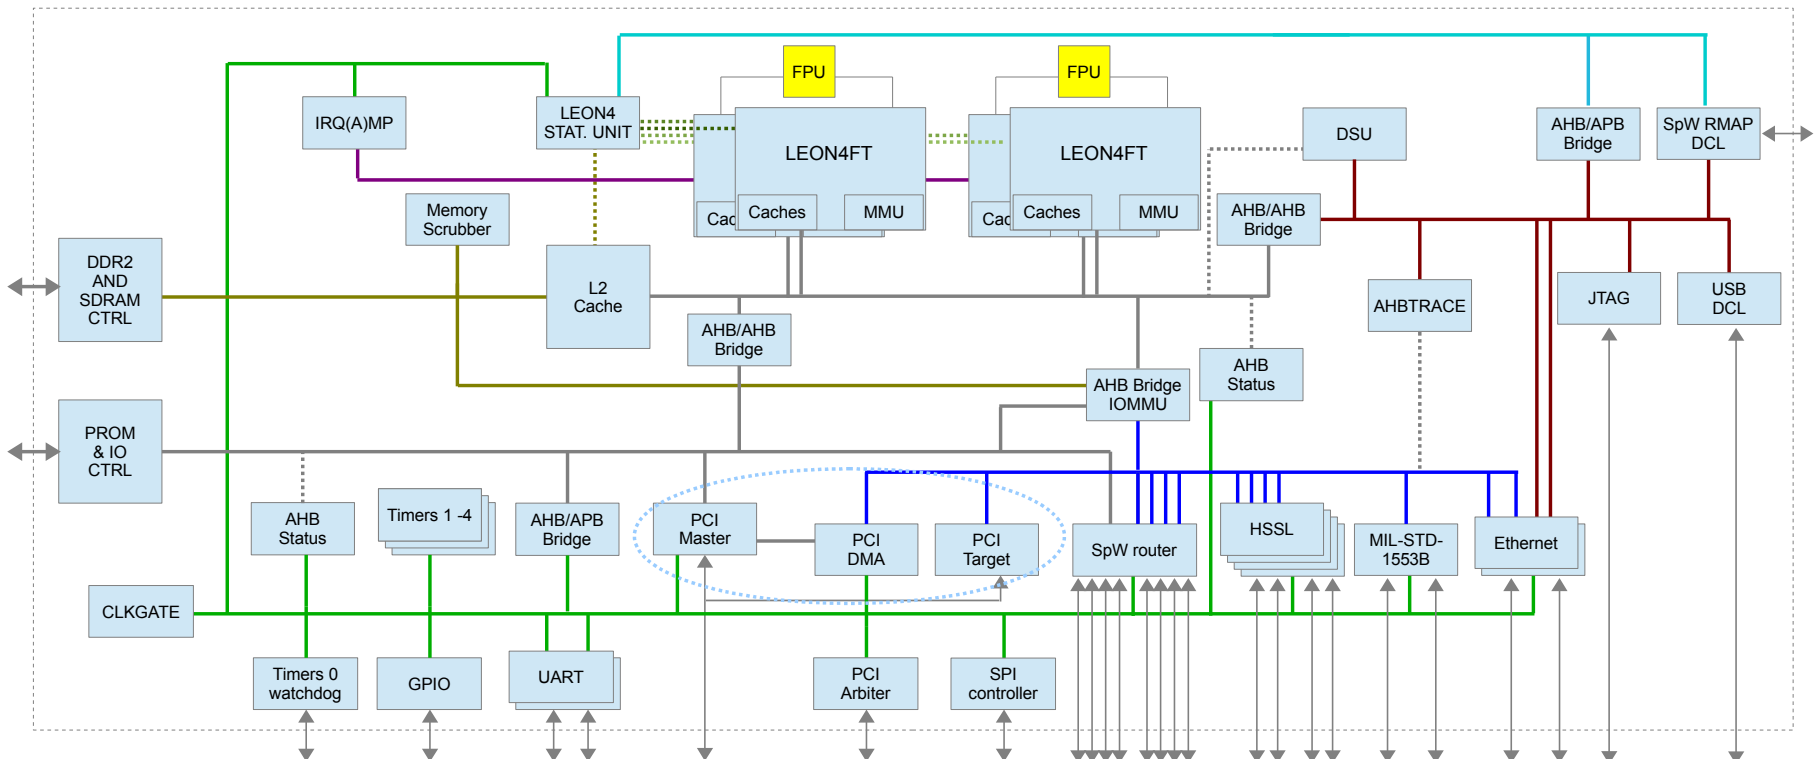
- Two LEON3FT SPARC V8 processors
 - SRMMU, IEEE-754 FPU, 16 KiB I-cache, 16 KiB D-cache
- On-chip 192 KiB SRAM with EDAC
- Off-chip memory types: SDRAM, SDRAM, Flash PROM / EEPROM

- I/O matrix connecting subset of:
 - Four 200 Mbit/s SpaceWire ports
 - MIL-STD-1553B (BC/RT/BM)
 - Ethernet MAC, 10/100 Mbit/s
 - Two CAN 2.0B bus controllers and one SatCAN controller
 - Six UART port, SPI master, I²C master, ASCS16 and SLINK serial ports
 - CCSDS/ECSS 5-channel TC decoder, 10 Mbit/s input rate
 - CCSDS Telemetry encoder, 50 Mbit/s output rate
 - 26 input and 38 input/output GPIO ports



Next Generation LEON SoC: NGMP

- Quad core LEON4FT connected to shared L2 cache

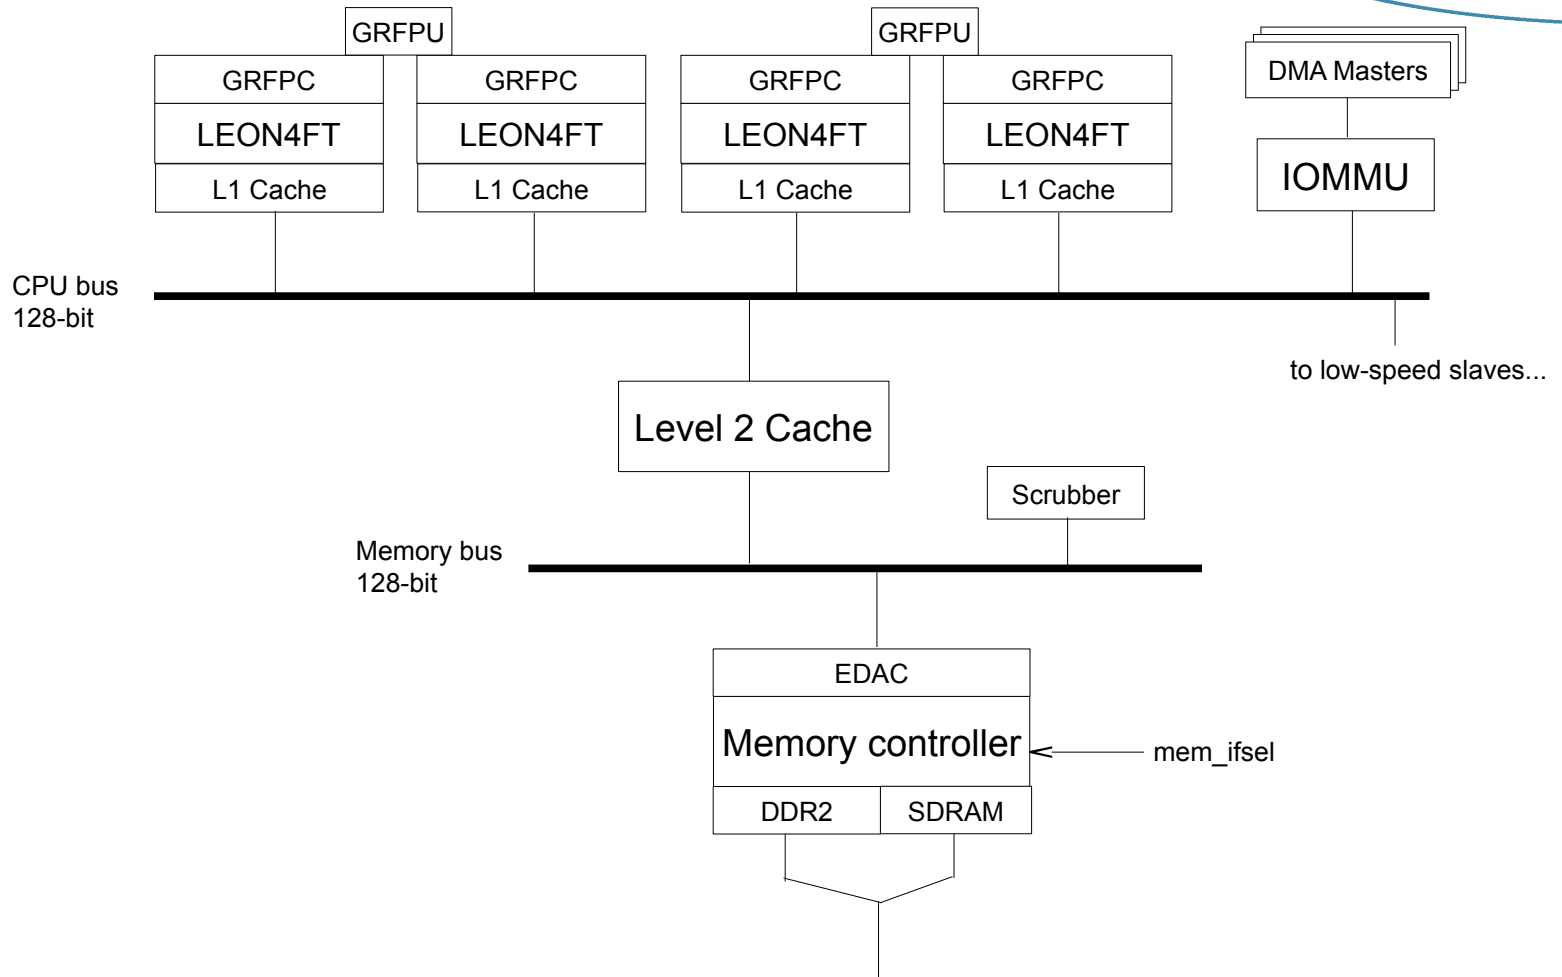


NGMP Project Overview



- NGMP is an ESA activity developing a multi-processor system with higher performance than earlier generations of European Space processors
- Part of the ESA roadmap for standard microprocessor components
- Aeroflex Gaisler's assignment consists of specification, the architectural (VHDL) design, and verification by simulation and on FPGA. The goal of this work is to produce a verified gate-level netlist for a suitable technology.
- As an additional step in the development of the NGMP, a functional prototype ASIC “NGFP” has been developed, also under ESA contract.

NGMP Architecture Overview (1/2)





to either DDR2 or PC133 SDRAM

Presentation does not contain US Export controlled information (aka ITAR)


NGMP Architecture Overview (2/2)

- Quad-core Leon4
- GRFPU, pairwise shared



CPU bus
128-bit

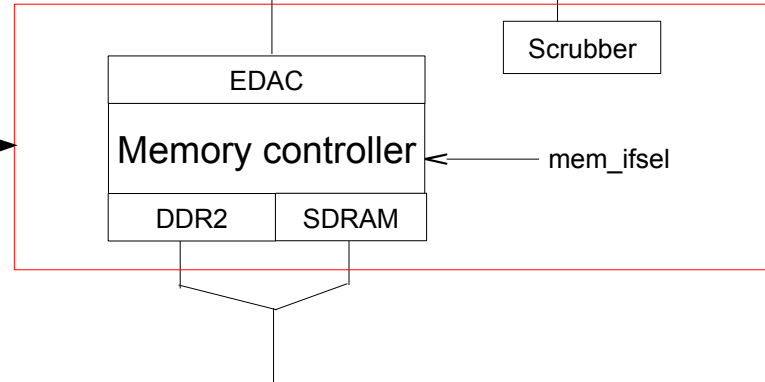
- L1 and L2 Caches



to low-speed slaves...

Memory bus
128-bit

- Memory controller



to either DDR2 or PC133 SDRAM

Presentation does not contain US Export controlled information (aka ITAR)

NGMP Overview - Features



- 4x LEON4FT with 128-bit AHB bus interface
 - High-performance GRFPU, one FPU shared between two CPUs
 - 16 KiB I-cache, 16 KiB D-cache, write-through
- Level-2 cache, bridge in the bus topology
 - Configurable, copy-back operation, can be used as OC RAM
- External memory: DDR2 or SDRAM, selected with bootstrap signal
 - Powerful interleaved 16/32+8-bit ECC giving 32 or 16 checkbits (SW selected, can be switched on the fly)
- Memory error handling (memory controller, scrubber, CPU together)
 - Hardware memory scrubber for initialization, background scrub, error reporting and statistics
 - Rapid regeneration of contents after SEFI
 - Graceful degradation of failed byte lane, regaining SEU tolerance
 - Example code for RTEMS available

NGMP Overview – I/O Interfaces



- Large number of I/O interfaces:
 - 8-port SpaceWire router
 - 32-bit 33/66 MHz PCI Master/Target with DMA
 - 10/100/1000 Mbit Ethernet
 - MIL-STD-1553B
 - 2xUART, SPI master/slave, 16 GPIO
- Debug interfaces:
 - Ethernet
 - USB
 - Spacewire (RMAP)
 - JTAG

NGMP Overview – Improvements



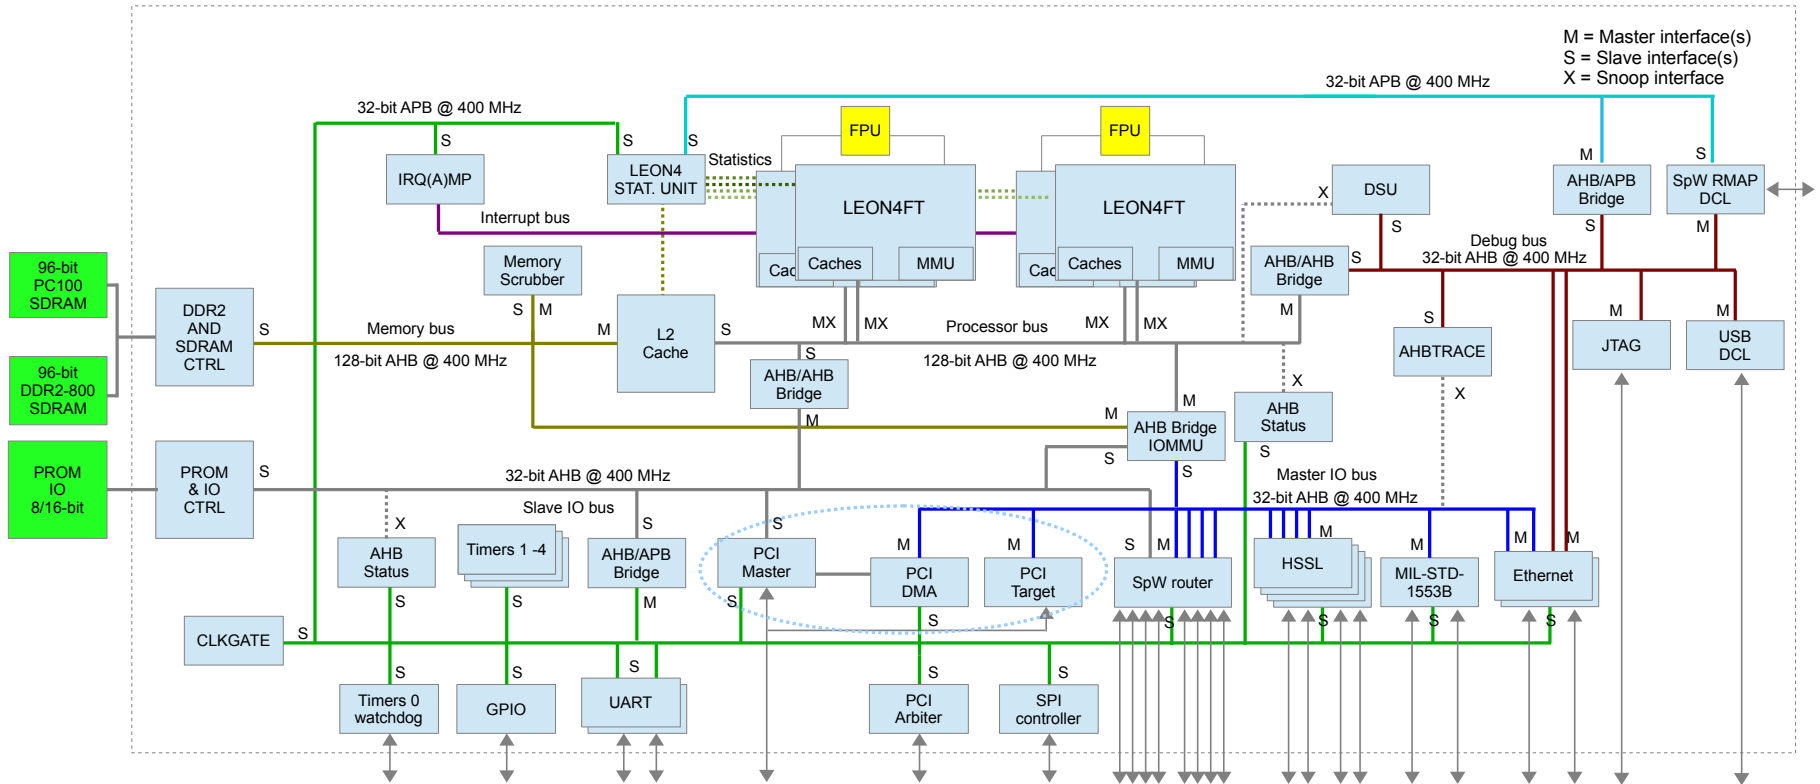
- Resource partitioning
 - The architecture has been designed to support both SMP, AMP and mixtures (examples: 3 CPU:s running Linux SMP and one running RTEMS, 4 separate RTEMS instances, 2x Linux/1x RTEMS/1x VxWorks, etc.)
 - The L2 cache can be set to 1 way/CPU mode
 - Each CPU can get one dedicated interrupt controller and timer unit, or share with other CPU
 - Peripheral register interfaces are located at separate 4K pages to allow restricting (via MMU) user-level software from accessing the wrong IP in case of software malfunction.
 - IOMMU
- Improved debugging
 - Dedicated debug bus allows for non-intrusive debugging
 - Performance counters, AHB and instruction trace buffers with filtering, interrupt time stamping

NGMP - PROM-less / SpW applications



- Extended support for PROM-less boot
- PROM-less booting possible via SpaceWire
 - Connect via RMAP
 - Configure main memory controller
 - Use HW memory scrubber to initialize memory
 - Enable L2 cache (optional)
 - Upload software
 - Assign processor start address(es)
 - Start processor(s)
- SpaceWire router, with eight external ports, is fully functional without processor intervention
- Device can also act as a software/processor-free bridge between SpW and PCI/SPI/1553 etc.
 - IOMMU can be used to restrict RMAP access

NGMP Overview – Block Diagram




NGMP Overview – FT and target tech



- Fault-tolerance
 - External DDR2/PC100 SDRAM: Reed-Solomon. PROM: BCH
 - LEON4FT
 - 4-bit parity on L1 cache
 - Protected register files (both CPU and FPU)
 - L2 Cache
 - BCH protected memories, Built-in Scrubber
 - General
 - Block RAM contents in IP cores protected by ECC
 - Rad-hard flip-flops and logic by process, library or TMR on netlist
- Baseline target technology: ST Space DSM (65 nm)
 - Alternate target technologies also currently under evaluation

NGFP Evaluation Board

Evaluation board providing
interfaces of the NGFP device
6U CPCI form factor



Demo in Aeroflex booth here at
FSW!

NGMP Summary

- Quad core LEON4FT architecture with shared L2 cache.
- Baseline memory interfaces: DDR2-800 SDRAM, PC100 SDRAM
- Non-intrusive debugging, performance counters, trace buffer filters
- I/O interfaces; PCI 2.3, SpaceWire, 1553, Gigabit Ethernet, SPI, UART, ...
- Debug interfaces: Ethernet, JTAG, USB, SpaceWire

- Part of ESA roadmap for standard microprocessor components

- Functional prototype device on evaluation board on display in Aeroflex booth here at FSW

- NGMP website: <http://microelectronics.esa.int/ngmp/>

LEON SoC Architecture Summary



- Current systems use single AMBA AHB bus with one or several LEON3FT or LEON4 processor cores.
- Trend is toward multicore with more or less complicated bus topologies. New developments still started with traditional single core, single bus.
- NGMP architecture introduces several new features for European space processors in terms of separation, FT Level-2 cache, error recovery and improved debugging support.
- Next incremental improvement: Mitigating effects of shared memory

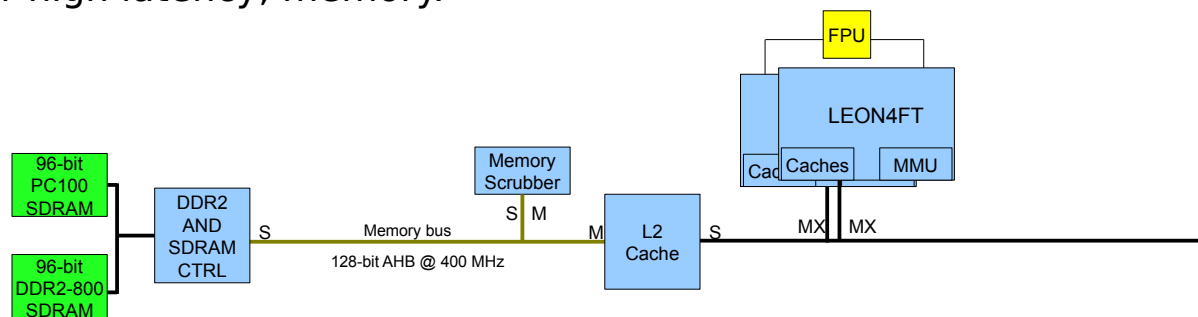
Thank you for listening!

Questions?

EXTRA SLIDES

NGMP Architecture - Level-2 Cache

- 256 KiB baseline, 4-way, 256-bit internal cache line
- Replacement policy configurable between LRU, Pseudo-Random, master based.
- BCH ECC and internal scrubber
- Copy-back and write-through operation
- 0-waitstate pipelined write, 5-waitstates read hit (FT enabled)
- Support for locking one more more ways
- Support for separating cache so a processor cannot replace lines allocated by another processor
- Fence registers for backup software protection
- Essential for SMP performance scaling, reduces effects of slow, or high latency, memory.




NGMP Architecture - Memory scrubber

- Can access external DDR2/SDRAM and on-chip SDRAM
- Performs the following operations:
 - Initialization
 - Scrubbing
 - Memory re-generation
- Configurable by software
- Counts correctable errors with option to alert CPU
- Supports, together with DDR2 SDRAM and SDRAM memory controller, on-line code switch in case of permanent device failure



NGMP Architecture - IOMMU

- Uni-directional AHB bridge with protection functionality
- Connects all DMA capable I/O master through one interface onto the Processor bus or Memory bus (configurable per master)
- Performs pre-fetching and read/write combining (connects 32-bit masters to 128-bit buses)
- Provides address translation and access restriction via page tables
- Provides access restriction via bit vector
- Master can be placed in groups where each group can have its own set of protection data structures



Benchmarks - Overview




- Benchmarks
 - I/O traffic
 - FPU sharing
 - Scaling
 - Comparison with AT697, UT699, GR712RC

Benchmarks – I/O traffic routing

- Traffic simulations* on system at target frequency
- Tests benefited from L2 cache, however, decent transfer rates can be achieved without causing traffic on the Processor bus

Configuration	1x Eth	2x Eth		SpW		Combined			
		Eth0	Eth1	Per port	Total	Eth0	Eth1	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
Bypassing L2 cache using IOMMU connection directly to Memory bus	1.4 Gb/s	1.2 Gb/s	1.2 Gb/s	697 Mb/s	2.8 Gb/s	746 Mb/s	850 Mb/s	338 Mb/s	1.4 Gb/s



* Exact figures no longer fully applicable due to updates of L2 cache and memory controller

Benchmarks – FPU sharing (0)

- Quad instance runs of single/double precision Whetstone on quad CPU system show no measurable difference between having 1x FPU, 2xFPU or 4xFPU (one dedicated FPU per CPU)
- Runs of some of the benchmarks included in SPEC CPU2000 on dual CPU system with 1x and 2x FPU:

Test	ML510-A (1x FPU) exec. time (s)	ML510-B (2x FPU) exec. time (s)	Difference (A-B) (s)	Dedicated FPU speed-up (A/B)
168.wupwise	1467	1451	16	1.01
171.swim	540	534	6	1.01
172.mgrid	1331	1311	20	1.02
173.applu	637	623	14	1.02
177.mesa	1632	1629	3	1.00
178.galgel	2094	2067	27	1.01
179.art	450	448	2	1.00
183.quake	1298	1271	27	1.02
188.ammmp	2210	2165	45	1.02
191.fma3d	12476	12348	128	1.01
200.sixtrack	3265	3024	241	1.08
301.apsi	494	478	16	1.03

Benchmarks – FPU sharing (1)


- FPU sharing, for this particular system, more noticeable with applications using division (FDIV, FSQRT):

```
int main(void)
{
    int i;
    double a = 3, b = 0.1, c;
    volatile double d;


    for (i = 0; i < 10000000; i++) {
        c = a/b;
        b += 0.1;
    }

    d = c;

    return 0;
}
```




```
1042c: 82 00 60 01    inc %g1
10430: 95 a3 89 c8    fdivd %f14, %f8, %f10
10434: 80 a0 40 02    cmp %g1, %g2
10438: 12 bf ff fd    bne 1042c <main+0x28>
1043c: 91 a2 08 4c    fadd %f8, %f12, %f8
```




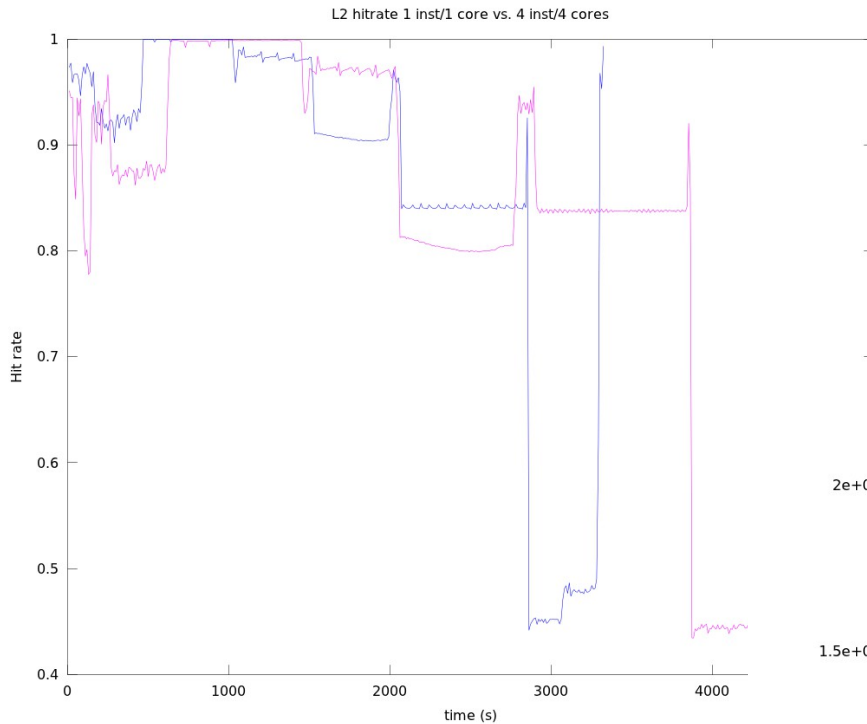
Test	ML510-A (1x FPU) exec. time (s)	ML510-B (2x FPU) exec. time (s)	Difference (A-B) (s)	Dedicated FPU speed-up (A/B)
One instance	3.53	3.53	0	1.00
Two instances	4.68	3.53	1.15	1.33
Three instances	4.73	3.55	1.18	1.33
Four instances	4.82	3.58	1.24	1.35

Benchmarks – Scaling (0)


- Running *gcc*, *parser*, *eon*, *twolf*, *aplu* and *art* CPU2000 benchmarks first sequentially and then in parallel on one to four cores:




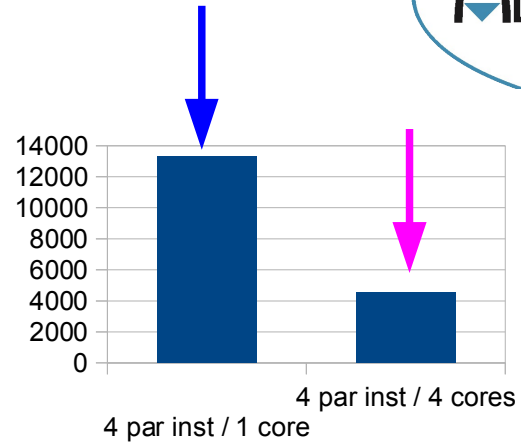

Benchmarks – Scaling (1)



Four parallel instances on four cores
versus one instance on one core



Benchmarks – Scaling (2)



Four parallel instances on four cores
versus four parallel instances on one core

Benchmarks - Comparison

- Comparison with existing processors targeted for space
- Benchmark scores relative to AT697:

Benchmark	AT697	UT699	GR712RC	NGMP
164.gzip	1	0.94 (0.66)	1.1 (1.1)	1.31 (5.24)
176.gcc	1	0.79 (0.55)	0.97 (0.97)	1.3 (5.2)
256.bzip2	1	0.93 (0.65)	1.06 (1.06)	1.33 (5.32)
AOCS	1	1.2 (0.84)	1.52 (1.52)	1.79 (7.16)
Basicmath	1	1.3 (0.91)	1.46 (1.46)	1.62 (6.48)
Coremark, 1 thread	1	0.89 (0.62)	1.09 (1.09)	1.21 (4.84)
Coremark, 4 threads	1	0.89 (0.62)	2.05 (2.05)	4.59 (18.36)
Dhrystone	1	0.94 (0.66)	1.05 (1.05)	1.39 (5.56)
Dhrystone, 4 instances	1	0.94 (0.66)	1.61 (1.61)	4.81 (19.24)
Linpack	1	1.2 (0.84)	1.26 (1.26)	1.71 (6.84)
Whetstone	1	1.94 (1.36)	2 (2)	2.22 (8.88)
Whetstone, 4 instances	1	1.94 (1.36)	3.7 (3.7)	8.68 (34.72)

All benchmarks were compiled with GCC-4.3.2 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)
Note that the maximum operating frequency of the devices differ, here all tests are run at 50 MHz
Values in parentheses are scaled for maximum frequency.