GR740 Development Status

Presenter: Magnus Hjorth, Cobham Gaisler

Agenda

- Component status update
  - Overview
  - Qualification schedule
- Deep dive into architecture
  - Architecture overview
  - Cache system
  - Options for mitigating timing interference

Component status update

Quad-core radiation tolerant processor on 65 nm CMOS

- Cobham Gaisler digital IP (LEON4FT and IO peripherals)
- STMicroelectronics standard cell library and analog IP (C65SPACE technology platform)
- Development done under ESA contract

European supply chain:

- ASIC manufacturing and test: ST (FR)
- Ceramic package assembly and space qualification: ST (FR)
- Column attach: Micross (UK)
- Commercialization, support: Cobham Gaisler (SE)

GR740 project timeline

- 2016: First silicon (rev 0)
  - Issues with fault-tolerance and ESD lead to re-spin
- 2018: Second silicon (rev 1)
- Complete test program working at +125C and -55C.
- Passed validation:
  - Functional validation
  - ESD validation
  - Heavy-ion and proton SEE testing
  - Burn-in trials
- TID test campaign scheduled for end of 2019
  - Confirm results already obtained on cell level from test vehicles
- Production start of qualification batch in December 2019
- Qualification effort finished end of Q3 2020
- Qualification approved estimated end of 2020 (DLA pending)

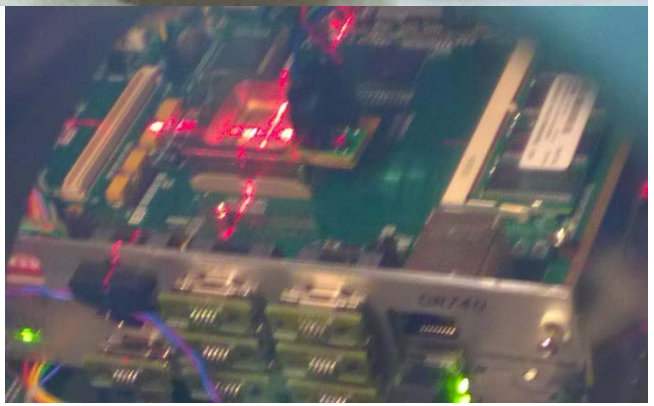

GR740 schedule

Product table

Code	Description	Availability
GR740-CP-LG625 GR740-CP-CG625	Prototype, room temp	Available
GR740-MP-LG625 GR740-MP-CG625	Prototype, full temp range	Available
GR740-MSEV-LG625 GR740-MSEV-CG625	QML equivalent (-MSEV)	Planned Q2 2020 Planned Q3 2020
GR740-MSV-LG625 GR740-MSV-CG625	QML qualified (-MSV)	After DLA approval
GR740-DD-LG625 GR740-DD-CG625	Daisy dummy packages	Available
GR740-DC-LG625 GR740-DC-CG625	Daisy chain packages	Available
GR-CPCI-GR740	Evaluation board	Available

Parts can be provided either with columns attached (-CG) or without columns exposing ceramic package lands (-LG).

All prototypes sold now are based on rev 1 (latest) silicon revision, functionally equivalent with flight parts.

Qualification batch planned to start production in December 2019.

Deep dive into architecture

Architecture overview

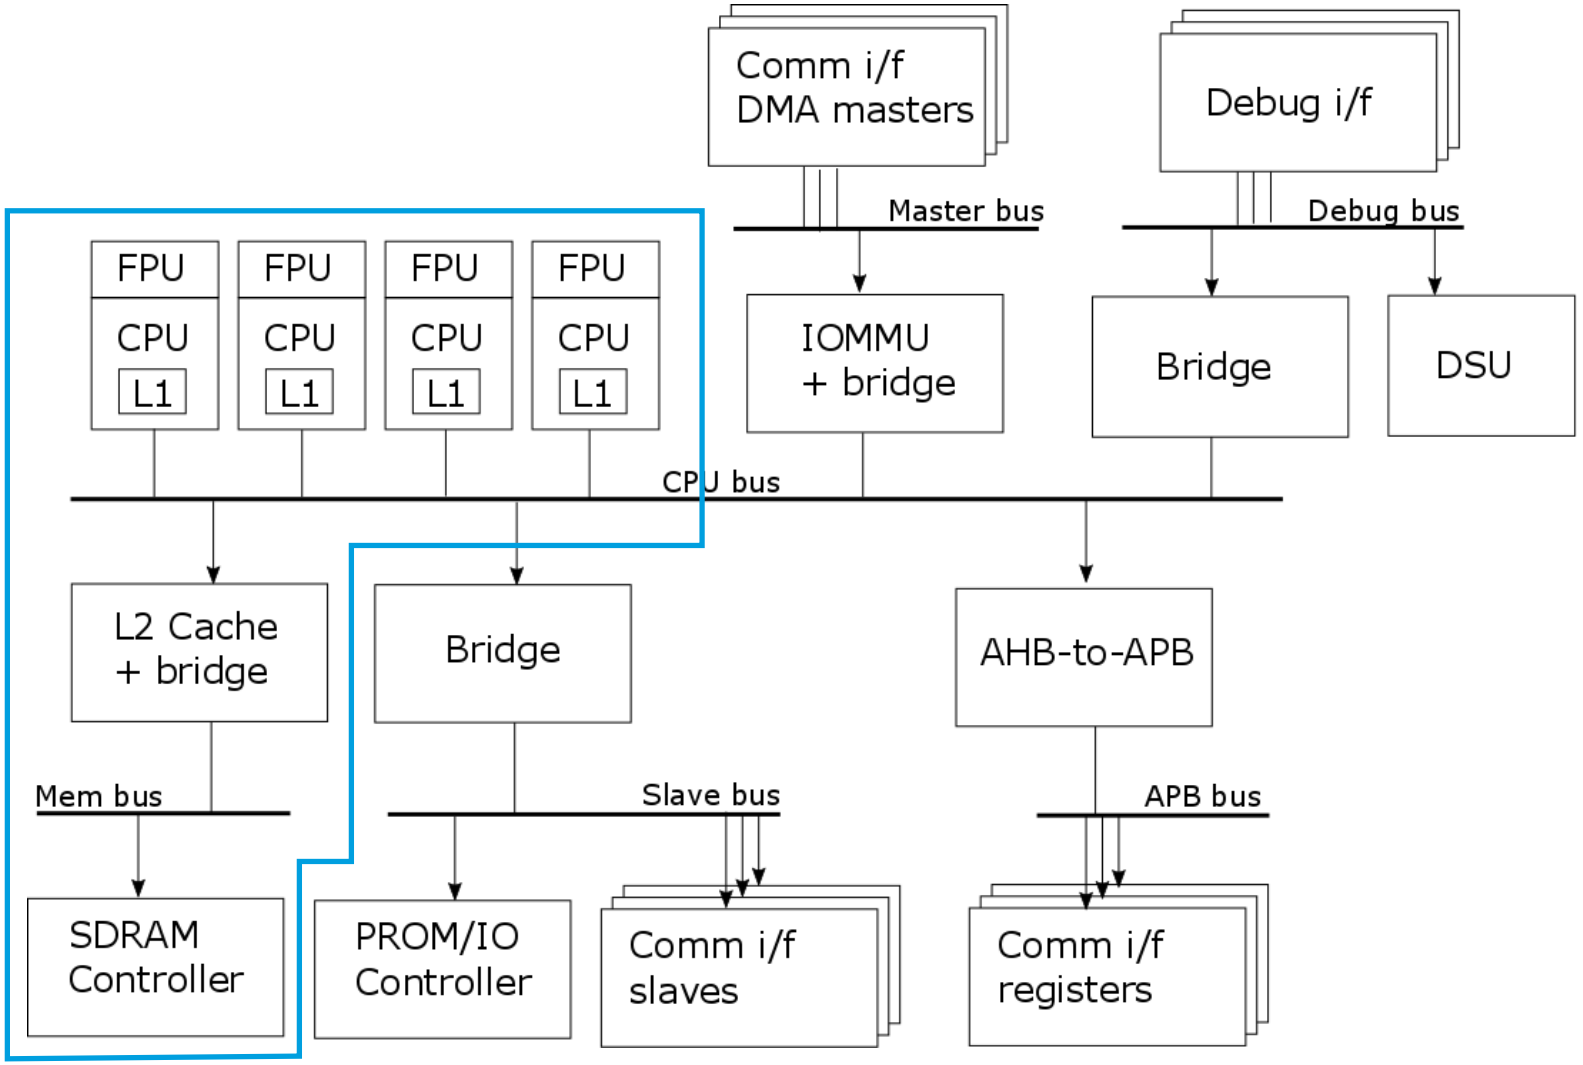

System based on multiple AHB buses.

Processor cores sharing a level-2 cache via the CPU bus

Behind Level-2 cache is the external SDRAM memory controller

Peripherals connected via bridges.

Cache system
Level-1 cache

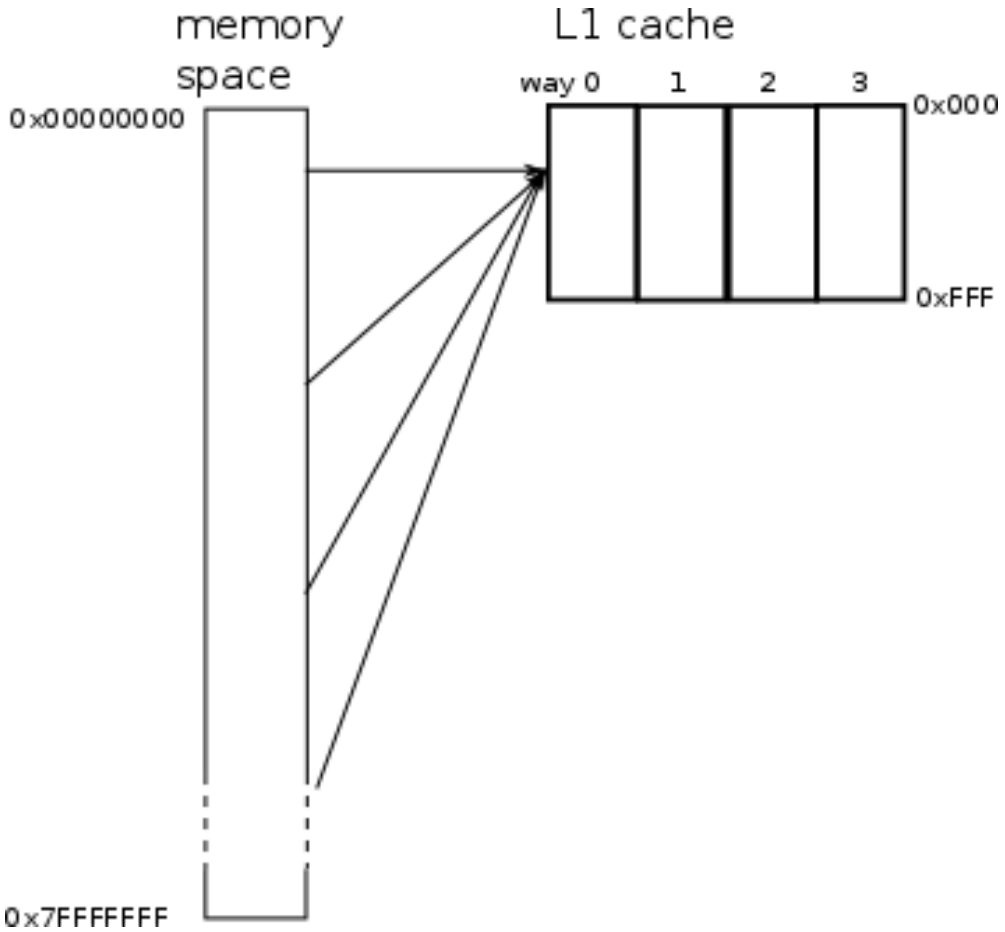

- Level-1 cache separate for each CPU
- Separate for instruction and data
- 4 KiB/way x 4 ways = 16 KiB
- 32 bytes per cache line
- Bits 11:5 of address determines cache set
  - Address N, N+4KiB, N+8KiB, …, in same set
  - At most 4 cache lines of same set can be held in cache at any given time
- LRU replacement by default
- Data cache is write through, allocation on load only.

Cache system
Level-2 cache

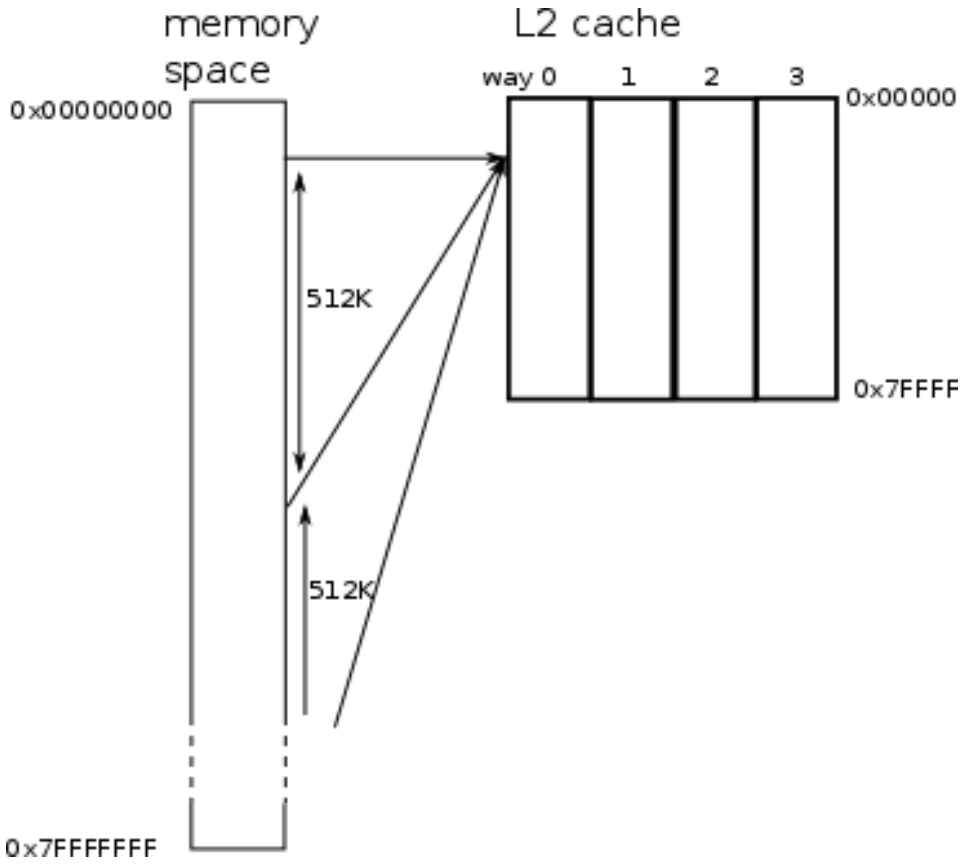

- L2 cache, shared between processors (and also DMA units)
- 512 KiB/way, 4 ways
- 32 bytes per cache line (same line size as L1)
- Bit 18:5 of address determines L2 cache set
  - Address N, N+512KiB, N+1MiB, …, in same set
  - At most 4 cache lines of same set held in L2 cache at one given time
- Copy-back policy (by default), allocate on both read and write

Cache system

Interaction between level-1 and level-2 caches

- L1 caches and L2 cache are not in general synchronized
  - Lines that map to same set in L1 may map to different sets in L2
  - Store that misses in L1 does not allocate line in L1 but allocates in L2
  - L2 is affected by other masters (CPU and DMA)
- We have various different cases to consider:

Access	L1	L2	L1 line replaced	CPU bus access to L2 cache	L2 line replaced	L2 access to SDRAM
Load	Hit	-	No	No	No	No
Load	Miss	Hit	Yes	Yes	No	No
Load	Miss	Miss	Yes	Yes	Yes	Yes
Store	-	Hit	No	Yes	No	No
Store	-	Miss	No	Yes	Yes	Yes

Level-2 cache timing interference
Introduction

- Multiple cores running in parallel may be competing for space in the L2 cache. This complicates execution time analysis.
- The architecture of the GR740 provides some options to manage this for real time applications:
  - One way per master mode
  - Locked ways
  - MMU translation based approach

Level-2 cache timing interference

One way per master mode

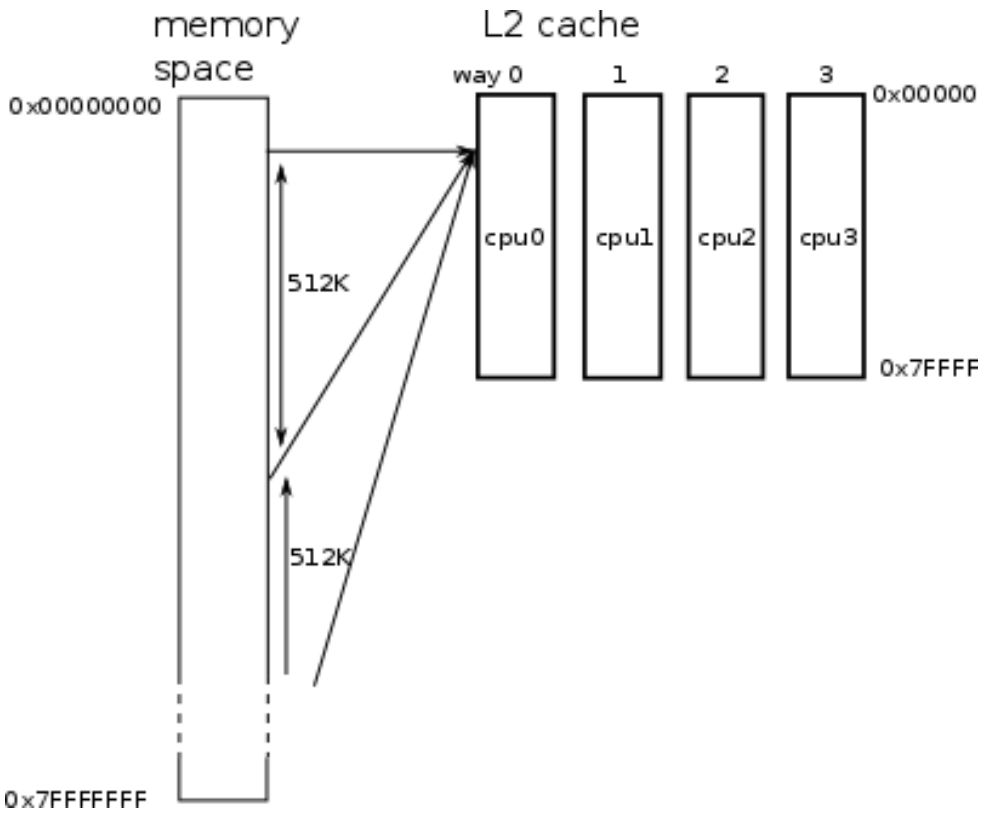

- One way per master mode option in L2
- Each CPU can only evict cache lines from a specific way (CPU 0 way 0, CPU1 way 1, etc). CPUs can therefore not evict each other's cache lines.
- L2 cache acts almost as a direct mapped 512 KiB cache for each processor.
- Only affects replacements, the masters can still read and write to all cache lines cached inside the L2 cache regardless of which set.
- Pros: Functonally transparent to software
- Cons: Less efficient caching on average, migration of tasks between CPUs can be an issue.

Level-2 cache timing interference

Locked ways in Level-2 cache

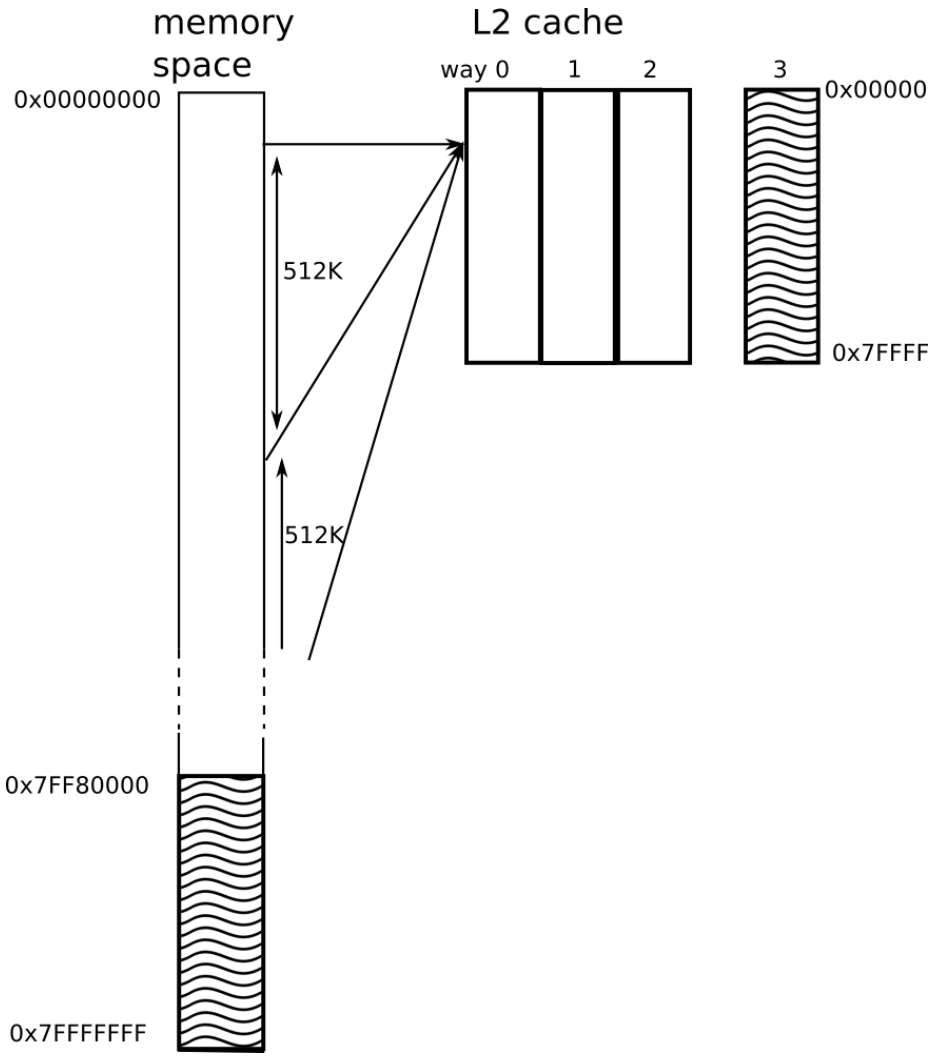

- Locked ways are another approach
- 1-4 ways can be locked down and set to always map to a specific address range.
- Those addresses are then guaranteed to always stay in the L2 cache, can be used to hold all data needed for real time critical tasks. Does not have to be backed by RAM.
- Remaining ways are shared between all tasks with LRU replacement as before.
- Pros: Locked address range can be accessed from any CPU.
- Cons: Quite coarse grained (512 KiB steps)

Level-2 cache timing interference

Using MMU translation to partition L2 cache

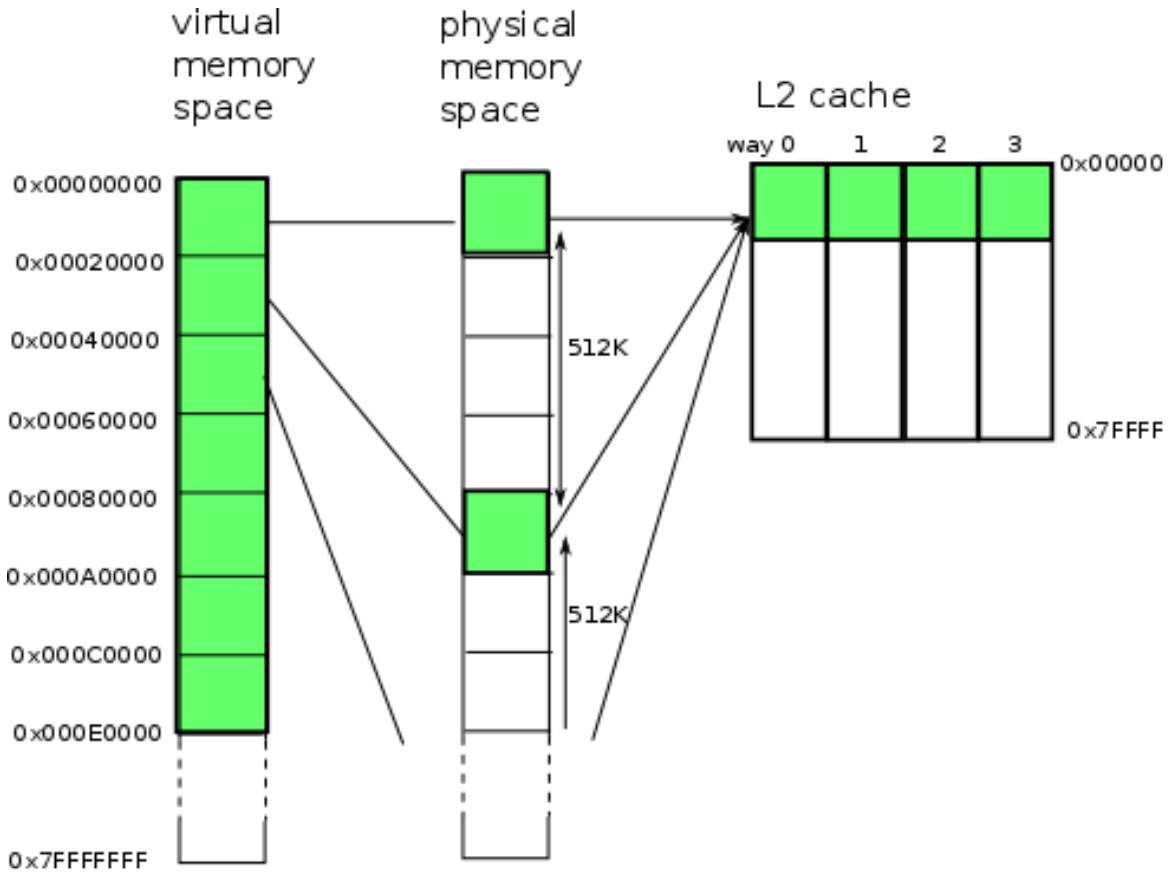

- Idea for software based approach to avoid L2 cache interference
- Using MMU translation tables, one could arrange so the memory for a task can only end up in specific sets in the L2 cache. As MMU translation works on 4 KiB pages, bits 18:11 of the L2 cache tag are under software (operating system) control.
- This allows a very flexible partitioning of the L2 cache, where some parts may be used by only one (critical) task, where other parts are shared between several tasks.
- The downside of this approach is the need to set up level-3 page tables for all addresses used.

Level-2 cache timing interference

Simultaneous level-2 cache misses

COBHAM

- Another source of timing interference is simultaneous L2 cache misses.
- As the L2 cache can only handle one L2 miss at a time, other processors have to wait for the other misses to complete.
- The L2 cache has hit-under-miss support, so while one CPU is waiting for a L2 cache miss to complete, the other CPUs can access the data in the L2 cache.
- One way to optimize for this is through task scheduling, tasks expected to create many L2 cache misses may be better to run serially on one CPU than in parallel, if other tasks can run on the other CPUs that produce few L2 cache misses.

Impact of SDRAM frequency on performance

- On the GR740, SDRAM only needs to be accessed when the L2Cache misses
  - Writes from the processors can be ”absorbed” by the L2 cache as long as they hit.
- Impact of SDRAM frequency on performance is therefore highly application dependent. Example from running SPEC2000 (250 MHz system clock):

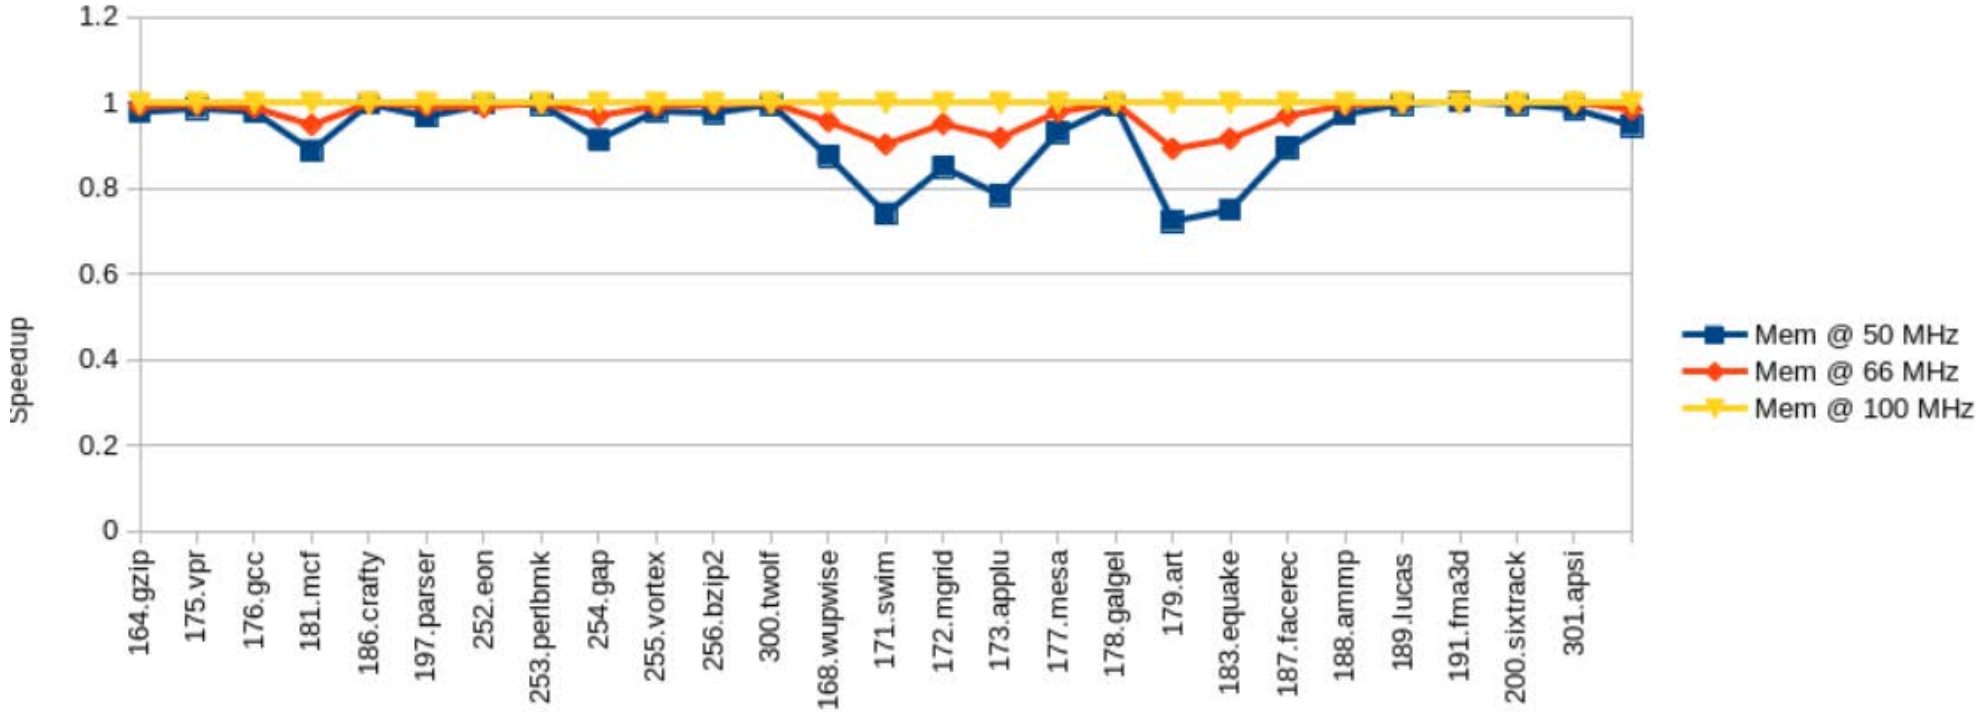

Benchmarking support

- The GR740 provides a number of features for benchmarking performance:
  - Statistics counter module with counters for many event types (# of cache misses, CPU hold cycles)
  - Global cycle counter that can be accessed from CPU application specific register.
  - Timer module that can latch interrupt times (for interrupt response time measurements)
- These may also be combined with the scripting capabilities of GRMON to perform more advanced instrumentation tasks.

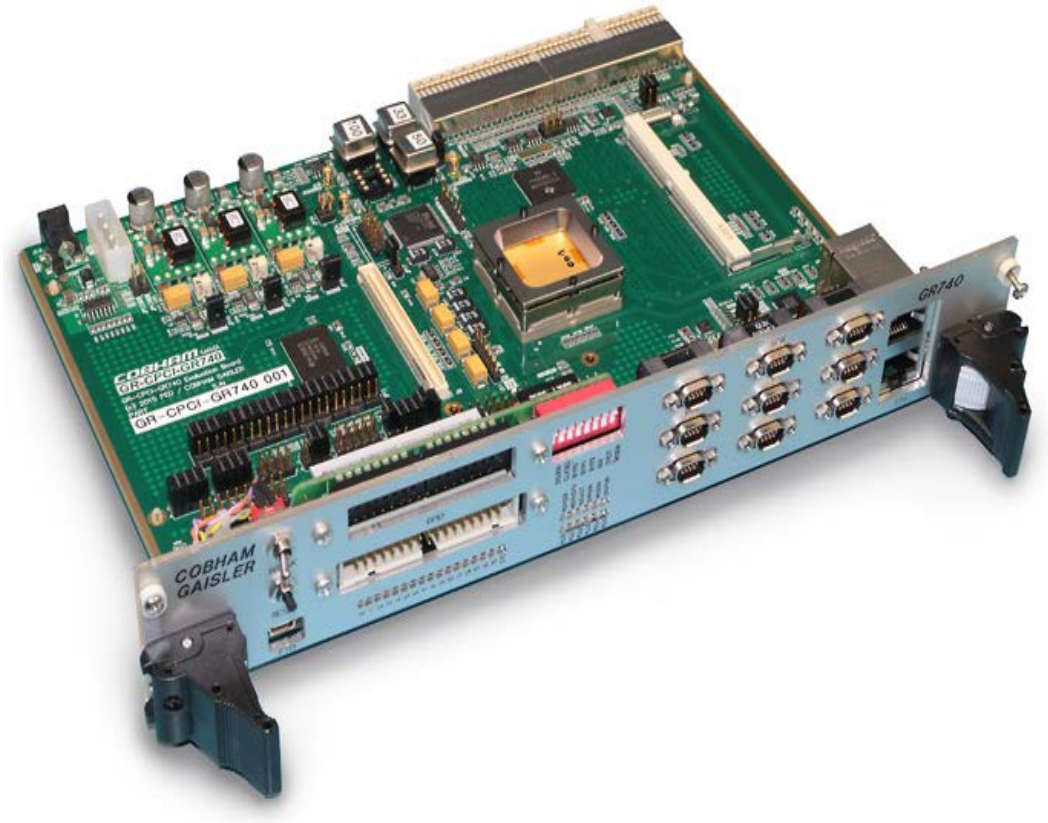

In summary

GR740 has been successfully developed to prototype stage, gone through pre-qualification preparations and is now entering qualification.

With four processor cores at 250 MHz and a large Level-2 cache, the GR740 provides the best performance relative to SwAP on the space market today, enabling more advanced space missions.

The architecture provides several features to support mixed criticality workloads, and to allow analysis, debug and instrumentation.

Thank you for your attention!