Application Report
SPNA238–April 2018

MPU and Cache Settings in TMS570LC43x/RM57x Devices

ABSTRACT

This application report provides an overview of various features available in CPU MPU.

Contents

List of Figures

Trademarks

1 Introduction

TMS70LC43x/RM57x are dual-core lockstep devices built around Arm® Cortex®-R5F CPU. Cortex-R5 has an in-built Memory Protection Unit (MPU) module that helps configure the memory types and attributes as defined in a processor's memory ordering mode. The MPU is specific to each core in the system and can only modify the memory ordering model of the CPU to which it is attached.

Apart form the ARM MPU module available in the core, this device also includes an additional module Enhanced Memory Protection Unit (NMPU) that enables to control the memory access rights of bus masters in the system other than the host CPU. The programmer's model for the NMPU is similar to but a subset of the host CPU's own MPU.

This application report provides an overview of various features available in CPU MPU.

2 Device Memory Architecture

The TMS570LC43x/RM57x microcontrollers are based on the TMS570 Platform architecture, which defines the interconnect between the bus masters and the bus slaves. Figure 1 shows a high-level architectural block diagram for the superset microcontroller.

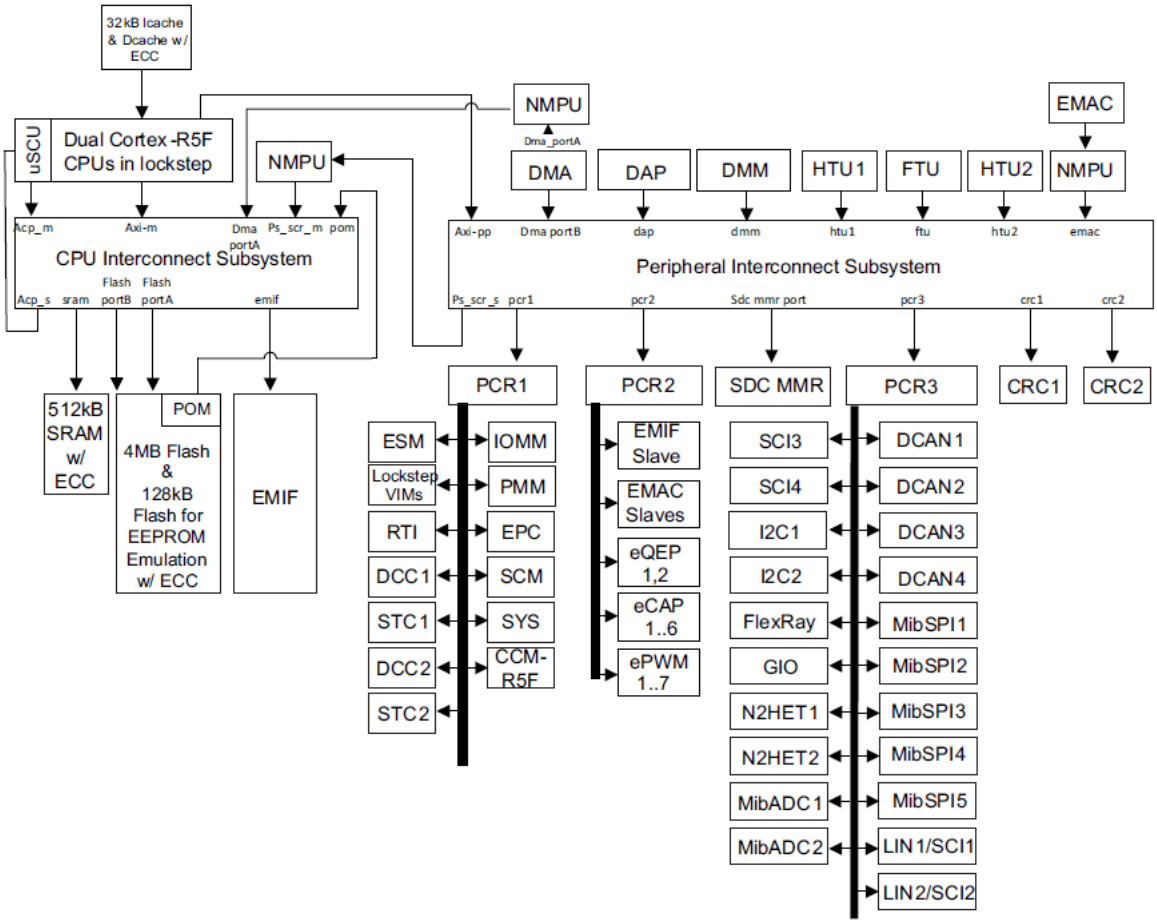

Figure 1. Functional Block Diagram

ARM Cortex-R5F CPU present in this device follows Harvard Level one (L1) memory system with:

- 32kB of instruction cache and 32kB of data cache implemented
- ARMv7-R architecture MPU with 16 regions

The CPU uses a 32-bit address bus, giving it access to a memory space of 4GB. This space is divided into several regions as shown in Figure 2.

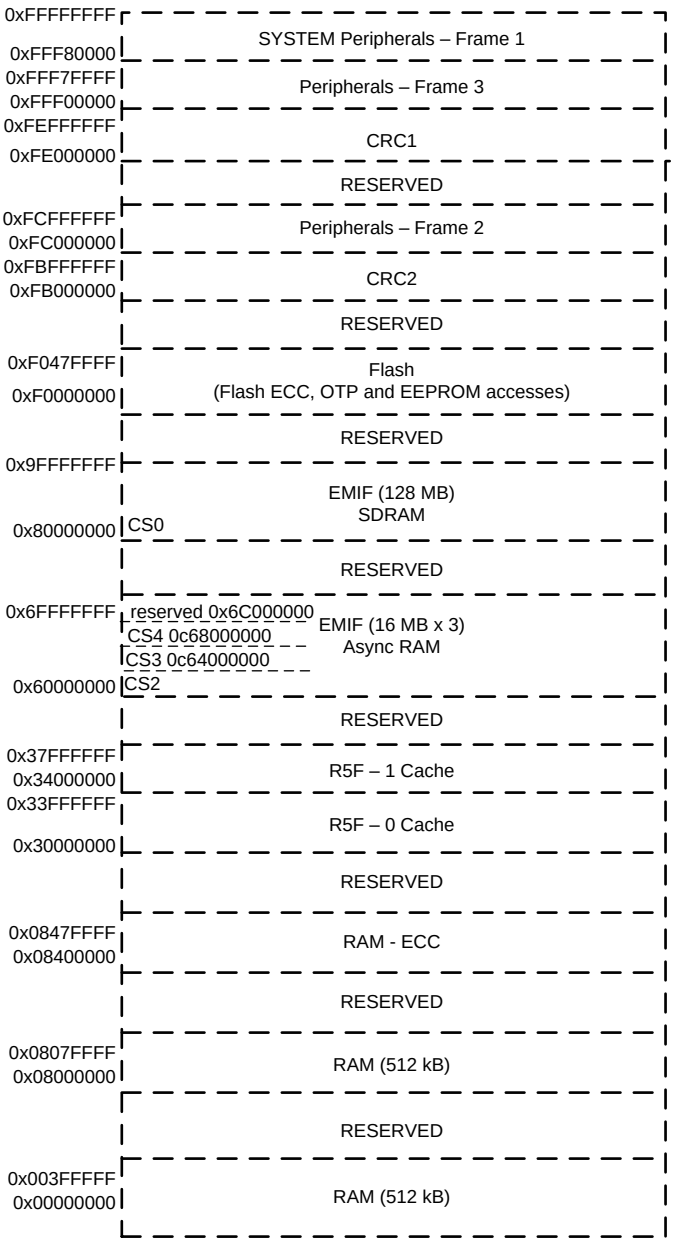

Figure 2. Memory Map

It is to be noted that the SRAM and Flash memories in this device are not tightly coupled memories (TCM) as opposed to other devices in the Hercules family. They are L2 memories in this device. This implies that the accesses to these memories are not as fast as compared to other Hercules devices which has TCM. For this reason, it is strongly recommended to enable cache for these regions (in case the memory is only accessed by the CPU). Further details on cache are mentioned in the following sections.

3 Bus Masters

Apart form the CPU, the following are the various bus master available in this device:

- Direct Memory Access (DMA)
- Ethernet Media Access Controller (EMAC)
- High-end timer Transfer Unit (HTU)
- FlexRay Transfer Unit (FTU)
- Debug Access Port (DAP)
- Data Modification Module (DMM)
- Parameter Overlay Module (POM)

The DMA has a memory protection unit internally built-in that allows you to configure the memories accessible by DMA. For other bus masters, the NMPU module can be used for setting the access attributes. DMA MPU and NMPU provides configuration only for read and write permissions for bus masters in the device.

4 ARM Cortex-R5 MPU Settings

4.1 Memory Types

- Strongly Ordered
  - All memory accesses to Strongly Ordered memory occur in the program order.
  - An access to memory marked as Strongly Ordered acts as a memory barrier to all other explicit accesses from that processor, until the point at which the access is complete
  - All Strongly Ordered accesses are assumed to be shared
  - It is recommended to configure the external peripheral or FIFO logic (accessible via EMIF) as strongly ordered
- Device
  - Defined for memory locations where an access to the location can cause side effects
  - The load or store instruction to or from a Device memory always generates AXI transactions of the same size as implied by the instruction
  - Can be shared or non-shared
  - It is recommended to configure the peripheral register spaces as device type
- Normal
  - Defined for memories that store information without side-effects. For example, RAM, Flash
  - Can be shared or non-shared
  - Can be cached or non-cached

In case there are multiple accesses to a normal memory, the CPU might optimize them leading to a different set of accesses of different size or number. The order of accesses may also be altered by the CPU. CPU makes an assumption that the order or number of accesses to a normal memory is not significant. For example, two 16-bit accesses to consecutive normal memories may be combined to a single 32-bit access. Whereas, in case of device and strongly-ordered memories, the CPU always performs the accesses in the order specified by the instructions. CPU does not alter the order, size or number of accesses to these memories. Device accesses are only ordered with respect to other device accesses, while strongly ordered memory accesses are ordered with respect to all other explicit accesses. It is to be noted that strongly-ordered memory leads to a larger performance penalty.

4.2 Shared and Non-Shared Memories

Shared memory attribute permits normal memory access by multiple processors or other system masters whereas non-shared memories can only be accessed by the host CPU.

The processor's L1 cache does not cache shared normal regions. This means that a region marked as shared is always a non-cached region (this device does not support L2 cache).

4.3 Cache Settings

This device only supports L1 cache. Cache property is only applicable for normal memories. Due to the unavailability of L2 cache, cache is applicable only for normal non-shared memories.

The following are various configurations available for cache:

- WTNOWA - Write-Through, No Write-Allocate
- WBNOWA - Write-Back, No Write-Allocate
- WBWA - Write-Back, Write-Allocate

If an access to a cached, non-shared normal memory is performed, cache controller does a lookup in the cache table. If the location is already present in the cache, that is a cache hit, the data is read from or written to the cache. If the location is not present, that is a cache miss, it allocates a cache line for the memory location. That means, the cache is always Read-Allocate (RA). In addition, data cache can allocate on a write-access, if the memory is marked as Write-Allocate (WA). Write accesses that are cache-hit, are always written to the cache locations. If the memory is marked as Write-Through (WT), the write is performed in the actual memory as well. If the memory is marked as Write-Back (WB), the cache line is marked as dirty, and the write is only performed on the actual memory when the line is evicted.

5 MPU Regions

The MPU present in this device supports up to 16 regions. Each region has 8 sub-regions. For more details on sub-regions, see Usage of MPU Subregions on TI Hercules ARM Safety MCUs. The memories accessed by the CPU can be partitioned up-to 16 regions (with region 0 having the lowest priority and region 15 having the highest). Each can be configured to a specific memory type and assign required permissions.

When the CPU performs a memory access, the MPU compares the memory address with the programmed memory regions. If a matching memory region is found, it checks whether the required permissions are set. If not, it signals a Permissions Fault memory abort. If the matching memory region is not found, the access is mapped onto a background region. If background region is not enabled, it signals a Background Fault memory abort.

6 How to Configure MPU Settings in a Multi-Master Application

MPU settings can be easily done using the HALCoGen MPU Configuration Tab.

Recommended MPU settings for the memories:

- Flash
  - Accessed only by the CPU
  - Can be split into Privileged and non-privileged regions. Typically, in an RTOS context, the tasks are executed in the user mode and the kernel code is executed in the privileged mode. The tasks can be placed in the non-privileged flash section and kernel code can be placed in privileged section to ensure that no user task accidentally executes the kernel functions.
  - Privileged Flash Region
    - Type : Normal, Non-Shareable, Cacheable
    - Permission : Privileged Read Only, Executable

    Non-privileged Flash Region
    - Type : Normal, Non-Shareable, Cacheable
    - Permission : Privileged/User Read Only, Executable
- RAM
  - Can be accessed by CPU and other bus masters
  - Can be split into shared and non-shared regions
    - Shared RAM (accessed by other bus master master like DMA or EMAC)
      - Type : Normal, Shareable, Non-Cacheable
    - Non-Shared RAM (accessed only by the CPU)
      - Type : Normal, Non-Shareable, Cacheable
      - Use WriteBack mode for faster accesses
      - Use WriteThrough mode if any of the other masters does a read to this memory
  - Can be split into privileged and non-privileged regions. Privileged RAM are typically used to store data accessible from a privileged code
    - Privileged RAM
      - Permission : Privileged Read-Wriite, Non-Executable
    - Non-Privileged RAM
      - Permission : Privileged/User Read-Wriite, Non-Executable
  - RAM can also be used to store the code to which CPU can branch and execute
    - Executable RAM
      - Permission : Privileged (or Privileged/User) Read-Write, Executable
- Peripherals
  - Accessed only by CPU
  - Type : Device, Non-Shareable
  - Permission : Privileged/User Read-Write, Non-Executable
- External memories (accessed via EMIF module)
  - External SDRAM (accessed only by CPU)
    - Type : Normal, Cacheble, Non-Executable/Executable
  - External peripherals or FIFO memories
    - Type: Strongly-Ordered

IMPORTANT NOTICE FOR TI DESIGN INFORMATION AND RESOURCES

Texas Instruments Incorporated ('TI") technical, application or other design advice, services or information, including, but not limited to, reference designs and materials relating to evaluation modules, (collectively, "TI Resources") are intended to assist designers who are developing applications that incorporate TI products; by downloading, accessing or using any particular TI Resource in any way, you (individually or, if you are acting on behalf of a company, your company) agree to use it solely for this purpose and subject to the terms of this Notice.

TI's provision of TI Resources does not expand or otherwise alter TI's applicable published warranties or warranty disclaimers for TI products, and no additional obligations or liabilities arise from TI providing such TI Resources. TI reserves the right to make corrections, enhancements, improvements and other changes to its TI Resources.

You understand and agree that you remain responsible for using your independent analysis, evaluation and judgment in designing your applications and that you have full and exclusive responsibility to assure the safety of your applications and compliance of your applications (and of all TI products used in or for your applications) with all applicable regulations, laws and other applicable requirements. You represent that, with respect to your applications, you have all the necessary expertise to create and implement safeguards that (1) anticipate dangerous consequences of failures, (2) monitor failures and their consequences, and (3) lessen the likelihood of failures that might cause harm and take appropriate actions. You agree that prior to using or distributing any applications that include TI products, you will thoroughly test such applications and the functionality of such TI products as used in such applications. TI has not conducted any testing other than that specifically described in the published documentation for a particular TI Resource.

You are authorized to use, copy and modify any individual TI Resource only in connection with the development of applications that include the TI product(s) identified in such TI Resource. NO OTHER LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE TO ANY OTHER TI INTELLECTUAL PROPERTY RIGHT, AND NO LICENSE TO ANY TECHNOLOGY OR INTELLECTUAL PROPERTY RIGHT OF TI OR ANY THIRD PARTY IS GRANTED HEREIN, including but not limited to any patent right, copyright, mask work right, or other intellectual property right relating to any combination, machine, or process in which TI products or services are used. Information regarding or referencing third-party products or services does not constitute a license to use such products or services, or a warranty or endorsement thereof. Use of TI Resources may require a license from a third party under the patents or other intellectual property of the third party, or a license from TI under the patents or other intellectual property of TI.

TI RESOURCES ARE PROVIDED "AS IS" AND WITH ALL FAULTS. TI DISCLAIMS ALL OTHER WARRANTIES OR REPRESENTATIONS, EXPRESS OR IMPLIED, REGARDING TI RESOURCES OR USE THEREOF, INCLUDING BUT NOT LIMITED TO ACCURACY OR COMPLETENESS, TITLE, ANY EPIDEMIC FAILURE WARRANTY AND ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT OF ANY THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

TI SHALL NOT BE LIABLE FOR AND SHALL NOT DEFEND OR INDEMNIFY YOU AGAINST ANY CLAIM, INCLUDING BUT NOT LIMITED TO ANY INFRINGEMENT CLAIM THAT RELATES TO OR IS BASED ON ANY COMBINATION OF PRODUCTS EVEN IF DESCRIBED IN TI RESOURCES OR OTHERWISE. IN NO EVENT SHALL TI BE LIABLE FOR ANY ACTUAL, DIRECT, SPECIAL, COLLATERAL, INDIRECT, PUNITIVE, INCIDENTAL, CONSEQUENTIAL OR EXEMPLARY DAMAGES IN CONNECTION WITH OR ARISING OUT OF TI RESOURCES OR USE THEREOF, AND REGARDLESS OF WHETHER TI HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

You agree to fully indemnify TI and its representatives against any damages, costs, losses, and/or liabilities arising out of your non-compliance with the terms and provisions of this Notice.

This Notice applies to TI Resources. Additional terms apply to the use and purchase of certain types of materials, TI products and services. These include; without limitation, TI's standard terms for semiconductor products http://www.ti.com/sc/docs/stdterms.htm), evaluation modules, and samples (http://www.ti.com/sc/docs/sampterms.htm).

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265