White Paper

HMI Memory Demystified

Renesas Electronics America Inc.

July 2015

1 Introduction

In the past it was sufficient for developers of industrial and consumer applications to integrate monochrome LCDs or low resolution-TFTs into their HMI-designs. For these designs they were quite satisfied with the performance, available interfaces and memory size of conventional microcontrollers.

Now more and more designers of HMI applications are facing the fact that they have to integrate TFT-Displays with increasing pixel resolutions into their designs.

Additionally there are requirements for animated graphics that demand double frame buffering and multiple layers for alpha blending. These additional demands lead to a substantial increase of the required RAM size. For example a WVGA-TFT-Display with a resolution of 800 x 480 and 24-bit color depth already needs 1.1 MByte of RAM for a frame buffer with one layer. With double buffering the required RAM size is 2.2 MByte.

The internal RAM of conventional microcontrollers is too small to store the frame buffer for this kind of TFT display. Therefore the designers have to search for solutions to increase the available RAM for the application.

This paper studies the options available to them and will be of interest for all designers of HMI systems, whatever their background and previous experience.

2 RAM Solutions

2.1 Introduction

There are basically two ways to add more RAM to a HMI application. The designer can choose from the two memory technologies SRAM and SDRAM.

2.2 DDR-SDRAM

Meanwhile DDR-SDRAM devices are used as external RAMs. DDR-SDRAM is the abbreviation for DoubleDataRate- Synchronous-Dynamic RAM. Because of the DRAM architecture the memory cells need to be refreshed periodically, causing a higher power consumption of the DRAM compared to a Static RAM (SRAM). SDRAM devices have a special memory interface and therefore cannot be connected to a standard processor bus. The used processor needs to contain a DDR Memory Controller for I/O interactions with the DDR-SDRAM device.

Up to now 4 generations of DDR-SDRAM devices exist: DDR, DDR2, DDR3 and DDR4. The newest generation of DDR4 is for future use.

2.2.1 DDR-SDRAM design issues

There are a lot of design issues which need to be taken into account when designing a reliable DDR-SDRAM memory 2015.01 HMI Memory Demystified interface. Routing a layout for the DDR-SDRAM interface is quite a challenging task, especially if this is a new topic for the designer. The designer has to observe power integrity and signal integrity characteristics because of their low voltages and high switching frequencies.

Power-Integrity: Because of the low supply voltage, it is necessary to have a stable supply for the DDR-SDRAM devices. Especially the impedance of the power supply at the switching frequencies has to be guaranteed. The placement of the decoupling capacitors and their properties (ESR, ESL) needs to be checked by simulation. To minimize the inductance of the supply, additional power planes are mandatory, which leads to rising PCB costs.

Signal Integrity: Because of the short rise times of the DDR-SDRAM signals the PCB traces of a memory interface cannot be treated as simple wires anymore. Instead these traces need to be treated as RF-transmission lines and effects like reflections and crosstalk have to be observed. Simulation tools need to be used to verify the signal integrity of the design. The memory manufacturer provides IBIS models of the DDR-SDRAM-devices for simulating. To provide good signal integrity for the memory interface a lot of design requirements need to be met by the designer. Here is a summary of the most important SI-design considerations:

- DDR-SDRAM signals must not cross gaps or slots in their related reference plane.
- The signals for address, control and clock traces have to be organized in a symmetrical tree structure with minimal length differences to minimize reflections.
- Signals of the data bus are arranged in byte groups. Signals within the same byte group should have the same length, should be routed on the same layer and use the same number of vias.
- Sufficient spacing to adjacent signals needs to be provided. Minimum 2.5 x distance to reference plane.
- Single ended signals need to be routed with 50 Ω impedance.
- Differential signals need to be routed with 100 Ω impedance.
- Proper termination of address, command and control signals according to their trace impedance.

To route the memory interface in the symmetrical tree structure for a great number of DDR-SDRAM devices is quite a complex task. As a new feature DDR3-SDRAM devices offer an alternative configuration where the address, command and control busses are routed in a chain from chip to chip. The so called fly-by routing eases the layout but it can only by applied if it is supported by the memory controller of the used CPU. Additional termination of these signals is necessary.

Before the DDR-SDRAM is ready to be use by the application code, the memory-controller and the DDR-SDRAM devices need to be properly initialized, which is not a trivial task.

If DDR-SDRAMs are not working correctly it is quite difficult to find the reason for the malfunction. There is no easy way to measure signals of a DDR-SDRAM interface, because the layout constraints for high-speed signals do not allow placing test-points for measurement to the signal traces. Simulating and analysis of the layout is usually the only way to check the layout.

A big advantage of DDR-SDRAM devices are their low costs per bit. They are available for memory sizes > 128 Mbyte.

2.3 SRAM

A second approach to provide sufficient RAM for HMI applications is to add static RAM. They offer some advantages compared to DDR-SDRAM devices.

SRAMs are usually available with a parallel interface which can be directly connected to the processor bus of the CPU without the need for a special memory controller. They don't need to be refreshed periodically like DDR-SDRAMs. Therefore they have a much smaller power consumption. SRAMs have a lower clock frequency and switching times than DDR-SDRAMs. That is why they don't have special constraints for power- and signal integrity. This eases the layout for the memory interface and eventually reduces the number of required PCB layers. SRAMs don't require special supply voltages which reduces costs. SRAM memory cells are more complex compared to DDR-SDRAM, which leads to their greatest disadvantages— they are very expensive and the maximum memory size for SRAMs is in the range of 4 Mbyte.

3 Different flash types for HMI applications

3.1 Introduction

Having more sophisticated TFT-displays with more graphical features like animations will not only lead to an increase of needed RAM but also to an increase of needed Flash memory. Basically there are two dominating memory technologies for non-volatile memory available: NAND- and NOR-Flash.

3.2 NAND-Flash

NAND-flash devices split up into Single-Level-Cells (SLCs) and Multi-Level-Cells (MLC.) SLCs are able to store one single bit per memory cell. They allow more erase and write cycles and have a better availability for the extended temperature range than MLCs.

Typical numbers of erase/ write cycles for SLCs are 50.000 – 100.000. For these reasons SLC-Flash devices are preferred for industrial HMI applications despite of their higher costs compared to MLCs.

For cost sensitive applications the usage of Multi-Level- Cells (MLCs) might be a solution. The costs for MLCs are by far lower than for SLCs of the same size. The MLC technology is capable to store two bits per memory cell, causing a higher error rate than SLCs. Because of this, MLCs have a smaller amount of erase/write cycles compared to SLCs. Typical numbers of MLC erase/write cycles are 3.000 – 10.000. But for consumer applications with a low number of write cycles to the flash, MLCs might be a good option.

The major advantage of the NAND-Flash technology is its high memory density. This leads to relatively low costs per bit for NAND-Flash devices and makes this kind of memory devices to a good choice for cost sensitive applications.

A disadvantage of the NAND-Flash devices is their need for an Error Correction Code (ECC). ECC is an algorithm to detect and correct corrupted bits inside of a memory block by using special control bits. Because of its complexity the ECC algorithm is usually built in hardware. If a block cannot be corrected by the ECC anymore, it is marked as "bad block" in a bad block table and its content is moved to a new block in the spare region of the NAND flash. The ECC is part of the NAND-flash controller and the management of the bad block table has to be controlled by the application software.

All NAND Flash devices are organized in blocks. Read and write access to the NAND-Flash is only possible for whole blocks. They don't support single byte access and therefore execute in place is not possible. A memory block with application code has to be copied from the NAND-Flash to the RAM before it can be executed.

NAND-Flash devices are available with a variety of different interfaces, which are listed in the following.

3.2.1 RAW-NAND

For interactions with the RAW-NAND the processor needs to have a Flash controller which puts the correct commands, addresses and data onto the bus. Because RAW-NANDs don't have an internal ECC this has to be done by the flash controller of the processor. It is up to the application software to control the bad blocks of the RAW-NAND. Because there is no standardized format for the bad block table, the application software needs to be updated if the RAW-NAND device in the design is changed.

3.2.2 eMMC Managed NAND Flash devices

eMMC devices internally consist of a RAW-NAND (MLC) and a flash memory controller. All ECC and bad block table tasks are performed by the internal memory controller. The eMMC device is connected to the CPU via a standard MMC interface.

eMMC are also called managed NANDs because they provide ECC functionality and bad block management. These tasks don't have to be controlled by the application software anymore. By using a standard interface it is not necessary to update the software if the eMMC device is changed.

3.2.3 eSD Managed NAND Flash devices

eSD memory devices are very similar to eMMC devices, except for the interface to CPU. The eSD device is connected to the CPU via a standard SD interface.

3.3 NOR-Flash

The second major technology for non-volatile memory is the NOR-flash technology. The memory density of NOR-flash devices is considerable smaller than the density of NANDflash devices. This leads to significant higher costs for NOR flash devices compared to NAND flash of the same memory size. An advantage of NOR-flash devices is their 5 – 10 times higher data retention and their higher number of erase/write cycles compared to NAND-flash devices. An ECC is not necessary.

Because of their higher cost and their small memory density, NOR-devices are often used to store smaller application firmware or boot-loaders. Another advantage of NOR-Flash is its much higher data rate during a read access compared to NAND-Flash devices. This characteristic of the NOR-Flash helps to keep the start-up time of an HMI-System low.

NOR-Flash devices are organized in blocks of 128 Kbyte. Write access to NOR-Flash memory is only possible by writing complete blocks whereas read access is possible as single byte read. NOR flash devices are available with different interfaces, which are listed in the following.

3.3.1 NOR-Flash with parallel interface

These NOR Flash memory devices can be connected directly to the processor bus of the CPU. Read access to the NOR- Flash is performed by referencing the address of the NOR-Flash connected to the bus space. A major advantage of the SRAM like interface and its single byte reads is the ability to execute code directly out of the flash memory (execute in place) without the need to copy if to the RAM first.

3.3.2 Serial NOR-Flash with SPI interface

NOR-Flash devices are connected by a serial SPI interface. The advantage of serial Flash devices is that they have a smaller pin count than devices with a parallel interface which helps to save space on the PCB and eases the layout. Serial Flash is quite cost efficient.

But serial Flash devices have lower data rates compared to parallel NOR-Flash and execute in place is not possible with them. They are often used to store boot-loaders.

4 Special solution from Renesas for HMI applications

4.1 Introduction

Renesas developed the new RZ/A processor family for HMI applications with TFT panels which range from medium scale to large scale. Renesas' approach for the RZ/A processor is to combine the advantages of SRAM and serial NOR-Flash and improve them for these kind of applications.

4.2 Use of internal SRAM

The RZ/A CPUs from Renesas are equipped with a large sized internal SRAM. Currently the CPU with the largest internal RAM available is the RZ/A1H with a total RAM size of 10 MByte. CPUs with 3 MByte and 5 MByte are also available. The usage of internal SRAM this adds further benefits to the design:

- The memory interface of the RAM is completely encapsulated inside of the CPU. This improves the EMC behavior and noise immunity characteristics of the design.
- Using internal RAM saves space on the PCB and reduces layout effort.
- Internal SRAM has fastest access time.
- The design is not affected by discontinuation of RAM memory devices anymore.
- Reducing the number of components lowers the MTBF rate in the design.

4.3 Use of serial NOR-Flash with Quad-SPI interface

Renesas RZ/A processors are equipped with a serial Quad-SPI interface (also referred to as Multi-I/O SPI). This allows connecting a serial Quad SPI NOR-Flash device to the processor which offers additional benefits compared to other serial Flash devices.

Quad-SPI Flash devices have four data lines and it is possible to connect two Quad-SPI devices in a dual mode to the RZ/A processor which doubles the available data lines to eight.

The data rate of these devices is comparable to devices with a parallel interface. The RZ/A processor supports the execute-in-place functionality for Quad-SPI memory devices. This allows flash memory to be accessed by directly reading from Quad-SPI bus space.

Quad-SPI flash devices combine the advantage of flash devices with processor bus together with the reduced pin count of serial flash devices. Similar to serial Flash devices they are cost efficient, assumed the required flash size is not in the range above 128MByte. This makes Quad-SPI flash devices to a good choice for HMI-applications.

5 About the costs

When it comes to the costs of the different design approaches, a lot of parameter needs to be taken into account: quantities per year, packages, temperature ranges etc.. This makes it quite difficult to give an exact evaluation of the cost.

But in general the RZ/A-processor with its internal SRAM and a serial Quad-SPI-Flash device is competitive compared to others processors with similar CPU performance which use DDR-SDRAM and NAND-Flash devices.

6 About lifetime and availability

First and foremost, the longevity of memory products is depending on the obsolescence policy of the particular manufacturer. But in general, the longevity of DDR-SDRAM and Flash devices is in the range of approximately five years. Of course the manufacturers will inform their customers about

obsolete components. The period of time between the product change notification and the last order can reach from several months to a year. This enables the customers to purchase sufficient components until the end of life cycle of the application is reached, or until a redesign with a new memory device is made.

The Replacement of memory devices in a current design can be a complex task, especially for NAND-Flash devices. Because the internal structure of NAND-Flash device, like page size, block size and plane size, the software of the application needs to be updated too. For DDR-SDRAM memory it may be necessary to update the timing setting for the memory controller which also will require an update of the software.

This is certainly a positive benefit of the RZ/A processor with its internal SRAM, which is of course not affected by discontinued DDR-SDRAM or NAND-Flash devices.

7 Conclusions

DDR3-SDRAM and NAND-Flash memory devices offer low costs per bit, but they are only available with memory densities > 128 MByte. But a lot of HMI applications don't need RAM and Flash memory above 128 MByte.

For HMI applications with the demand for more RAM and Flash memory, but not as much as 128MByte, processors form the RZ/A family are a good choice. They deliver sufficient internal SRAM (3 – 10 MByte) and support cost efficient Quad-SPI Flash with comparable data rates without the disadvantage of a complex layout.