



BRX-WP003: Video processing and hardware/software partitioning

White Paper

Table of Contents

1 History.....	3
2 Introduction.....	4
3 Implementation.....	5
3.1 PL subsystem.....	6
3.2 PS subsystem.....	7
3.3 Notes about SDRAM banks organization.....	8
3.4 Video latency and SDRAM arbitering.....	8
4 Future work.....	10
5 References.....	10

1 History

Version	Date	Notes
1.0.0	June 2016	First Public Release

2 Introduction

This white paper describes a video processing system - built upon Bora/BoraEVB - that satisfies specific functional and safety requirements.

From the functional point of view, it is required to (FR denotes a functional requirement):

- [FR1] acquire two independent video streams
- [FR2] mix the input stream and visualize them on a HDMI monitor
- [FR3] visualize informational and statistical data on a 7" LVDS TFT LCD.

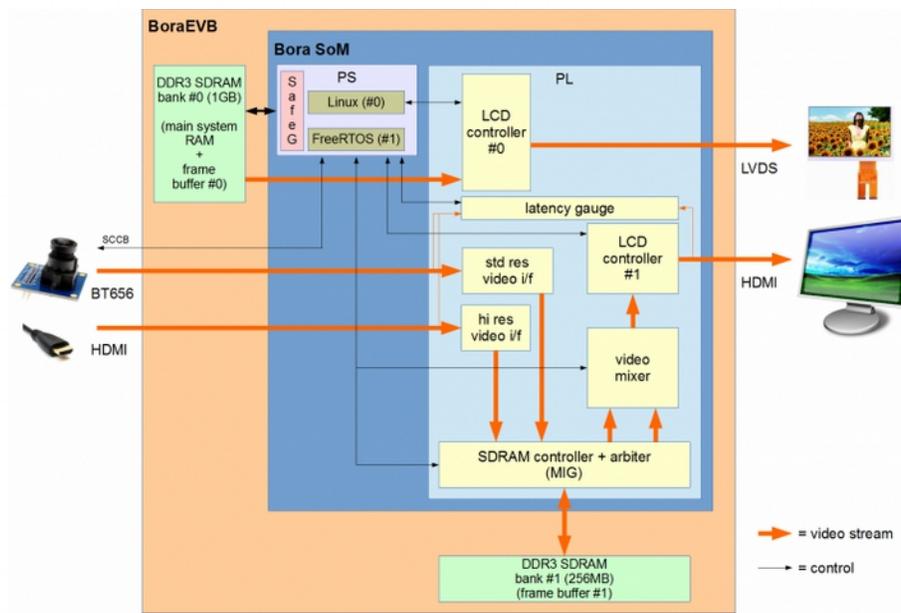
Safety requirements (SR for short) are:

- [SR1] FR1 and FR2 must be enabled as quickly as possible upon power-up
- [SR2] once started, video processing chain must be independent from the software running on PS, that is it must keep operating even in case software running on PS hangs.

The implementation combines different techniques that are available on Zynq platform, to implement a hardware/software partitioning meeting these requirements.

3 Implementation

The following picture shows a simplified block diagram of the entire system ¹.



Simplified block diagram

At top level, the natural PS/PL partitioning has been exploited: the video processing chain is entirely implemented in the PL, while PS domain is used for initializing, supervising and informational data visualization. The following section describes in more detail the actual implementation.

¹ At the time of this writing not all of the shown modules have been completed

3.1 PL subsystem

PL subsystem integrates several functional modules. The video processing chain is described first, as it represents the most important part of it.

Video processing chain supports two video sources as per requirement [FR1]:

- OV7670 camera module (640x480 @ 30fps)
- generic 1280x720 @ 60fps stream over HDMI connection.

Each video source is connected to a specific interface module:

- std res video i/f
 - converts BT656-encoded LVTTTL bus into AXI4 stream
 - performs color space conversion (YUV 4:2:2 to 24-bit RGB)
- hi res video i/f
 - converts TMDS differential signals to AXI4 stream.

Video frames encapsulated in AXI4 streams are then stored - by AXI-VDMA's not shown in the picture - in two independent buffers, both implemented on SDRAM bank #1 (this 16-bit wide bank refers to U14 component of BoraEVB board). Frames are then retrieved by AXI-VDMA's from these buffers and are forwarded to [Xilinx On-Screen Display LogiCORE IP block](#) (named video mixer in the diagram) that mixes them. The resulting stream is processed by LCD controller #1 module that converts it to TMDS signals feeding the external monitor.

The latency gauge module is used to perform video latency measurement. Specifically, it is used to measure the time a frame takes to traverse the entire video chain, from the input to the output interface.

To satisfy requirement [FR3], a second controller (LCD controller #0 in the diagram) is instantiated, feeding a 7" LVDS LCD. LVDS signals are generated by the FPGA directly.

3.2 PS subsystem

At software level, partitioning has been implemented by using the well-known AMP technique combined with SafeG monitor (for more details please see [this white paper](#)). FreeRTOS domain - associated to core #1 - is the first to come up and takes care of video processing chain initialization. Linux domain - associated to core #0 - controls LCD controller #0 to visualize informational data on 7" LCD. It also provides generic connectivity such as TCP/IP networking over Ethernet connection.

TrustZone technology - enabled with the help of SafeG monitor - allows to enforce the partitioning by preventing software running in Linux domain to access video processing chain resources.

The combination of

- short FreeRTOS domain boot time
- quick PL programming performed by FSBL

allows to minimize the video processing enabling time upon power-on (requirement SR1).

3.3 Notes about SDRAM banks organization

The block diagram shows two distinct SDRAM banks. Bank #0 is 32-bit wide and includes memory regions used by FreeRTOS and by Linux. It is also used to implement the frame buffer associated to 7" LCD screen. Bank #1 is 16-bit wide and it is used to implement buffers for video chain only. Bank #0 is accessed via Zynq's native DDR memory controller. Bank #1 is accessed via memory controller instantiated in PL and created by [Memory Interface Generator \(MIG\)](#).

SDRAM banks physical separation allows to satisfy requirement [SR2]: even if native DDR memory controller stops working ², video processing chain is unaffected because it relies on dedicated buffers.

3.4 Video latency and SDRAM arbitering

Applications that are based on such a video processing chain, often have an additional performance requirement in terms of video latency ³. For this reason the latency gauge module has been implemented. It is based on a bank of 32-bit 10ns-resolution timers that allows to measure this latency precisely. This bank - that is encapsulated in an IP - is accessible by PS via AXI-Lite bus in order to expose these measurements for offline manipulation and visualization on 7" display.

Apart from the requirement [SR2], the availability of two independent SDRAM banks provides an advantage in terms of performances too, in particular when video latency is a concern.

-
- 2 For example because a misconfiguration caused by a software bug in Linux drivers or by malicious code injected by an attacker.
 - 3 In this context, video latency is the time a video frame takes to traverse the chain, from the input interface to the visualization device.

This configuration, in fact, allows to isolate all the SDRAM traffic related to the video processing chain from all the other traffic originating either in PS or PL domain (i.e. the stream feeding 7" LCD, data read/written by Linux kernel and applications running on top of it, etc.). This solution simplifies management of concurrent SDRAM accesses dramatically, because video processing combined bandwidth requirement (there are four VDMA's accessing the memory simultaneously) is much smaller than the (theoretical) bandwidth provided by the controller/SDRAM subsystem. However, this solution comes at the price of adding a dedicated chip that is interfaced to the FPGA bank #35 ⁴.

During the development of this project, a single-bank version of the system has been tested as well. In this case all the masters that need to access the SDRAM make use of bank #0 and no SDRAM controller is instantiated in PL, resulting in a cheaper implementation. All masters compete for the unique SDRAM bank, thus it is pretty easy to create a testbed that saturates the actual available bandwidth. This may lead to clear anomalies such as flickering on display and an abnormal increase of the video latency. Although this configuration can not guarantee to meet requirement [SR2], it can nevertheless be fine tuned to overcome these performance issues. The QOS-301 module integrated in Zynq comes to help, as described in ⁵. By configuring proper priorities at Memory Interconnect level, it has been verified that, even in case of near-to-bandwidth-saturation condition, video anomalies disappear.

4 Bank #35 is optimum for such a connection, as described in detail [here](#).

5 http://wiki.dave.eu/index.php/BRX-WP003:_Video_processing_and_hardware/software_partitioning#cite_note-xapp1266-1

4 Future work

The system here described can be extended and improved in many different ways. Following is a list of works that will be possibly addressed in the future:

- Initialization of video processing chain should be moved to PL domain, in order to make it completely independent from software running on PS
- latency gauge should be completed and tested thoroughly
- additional video processing modules should be added in order to make the chain to resemble more realistic use cases
- completing SafeG integration with full TrustZone enablement.

5 References

1. Mrinal J. Sarmah, Naveen Kumar Gaddipati, XAPP1266 Using Quality of Service (QoS) Capabilities in Zynq-7000 AP SoC Devices, v1.0 September 18, 2015