

Next generation FPGAs and SOCs – How embedded systems can profit

Felix Eberli

Supercomputing Systems AG

Technoparkstrasse 1, CH-8005 Zürich, Switzerland

felix.eberli@scs.ch

Abstract

New SOC like the Xilinx Zynq 7045 allow researchers and developers to combine the advantages of writing software for control functionality and having accelerators in the FPGA logic for the number crunching. The dual core Cortex-A9 ARM processor runs with up to 1 GHz and the FPGA has up to 900 DSP slices allowing a performance of up to 1,334 GMACs. SCS is porting a lot of algorithms like SGM stereo [1], Stixel clustering or an optical flow [2] to such devices allowing new cars to see their environment and react appropriately. The new developed SCS Zynq 7045 module will allow accelerated development using this technology.

1. Introduction

In 2003 Xilinx announced their Spartan 3A FPGAs in 90nm technology. It (XCS1400A) had up to 25k logic cells, 576Kb Block RAM bits and 32 dedicated multipliers. As predicted by Moore's law, the technology-shrink went to 45nm with Spartan 6 and in 2011 to 28nm with the Artix family. There are already plans to shrink the next generation FPGAs to 20nm. What is the advantage of shrinking the size of a transistor? It takes much less area on the wafer and therefore these next generation devices

give another 50% price-performance-per-watt improvement as well as twice the memory bandwidth and also next generation transceivers. This will make it possible to implement functionality that we were dreaming of for a long time. For example the new Xilinx Zynq 7045 has 350k logic cells, 19,620 Block RAM bits and 900 DSP slices computing 1,334 GMACs [3].

2. Next generation SOC

The new Xilinx Zynq 7000 is the first device combining FPGA fabric with a dual core Cortex-A9 ARM processor with Neon engine that runs with up to 1 GHz. As most algorithmic problems can be divided into control code (80% code, 20% time) and computation (20% code, 80% time) these devices will accelerate development a lot. It will also allow software programmers to improve their algorithms running on the ARM and accelerate the number crunching with FPGA blocks that are connected to the ARM Cortex A9 over the AMBA AXI bus. A few examples on such accelerator blocks that were implemented by SCS will follow in the next chapters.

3. SCS 10 Bit JPEG Decoder IP

There are cameras with chipsets that produce a 10 Bit JPEG compressed stream over Ethernet. To decode such a stream it would be possible to do this on the ARM core. As this is quite repetitive work it suits much better on the

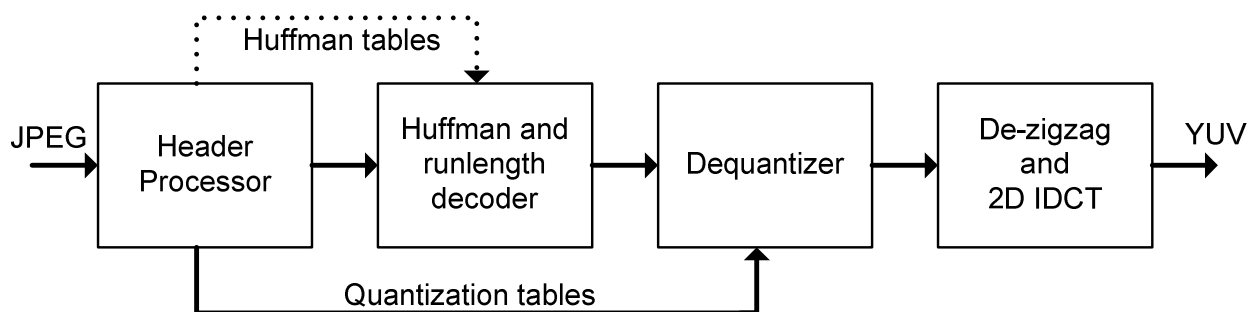


Figure 1: Overview of JPEG decoder blocks.

FPGA fabric and keeps the processors free for more complex high-level tasks. JPEG decoding consists mainly on the building blocks that are shown in Figure 1. A more detailed description can be found on Wikipedia [4].

We implemented the IP in two versions. One version has a static Huffman table and the other gets the Huffman table dynamic from the JPEG header. As you can read in Table 1, using a static Huffman table uses much less resources and will also reach higher speed.

Resources	Huffman-static	Huffman-dynamic
Max. clock	135 MHz	133 MHz
FF	1,700	2,298
LUT	2,000	2,427
DSP48	5	5
RAMB18	3	3
RAMB36	1	1

Table 1: Resource utilization of the 10 Bit JPEG decoder in a Zynq 7020-1 device.

main blocks are shown in Figure 2 and the used FPGA resources can be found in Table 2.

The cache allocation and the fitting of the linear approximation is optimized offline by software. Like this, the algorithm is capable to rectify a 2048x1536 image with 16fps.

Resources	
Max. clock	130 MHz
FF	2,000
LUT	2,000
DSP48	4
RAMB18	12

Table 2: Resource utilization of the rectification IP.

4. SCS Image Rectification IP

To compute stereo disparity maps, many algorithms require rectified stereo input images. Depending on the amount of the camera lens distortion and mounting errors the size of the internal cache has to be adapted to reduce the external memory bandwidth. Our lookup table based algorithm produces these images by bilinear interpolation of the required output pixels from the input image. The

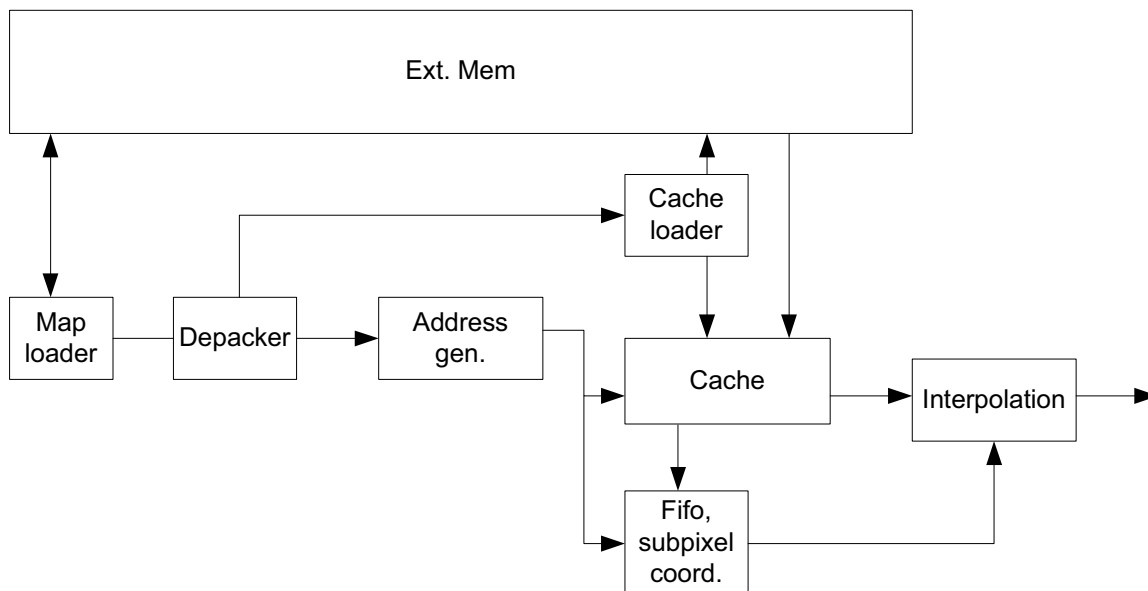


Figure 2: Overview of the image rectification block.

5. Semi Global Matching Stereo

The Semi Global Matching algorithm and its optimizations by Daimler are described 2009 in “A real-time low-power stereo vision engine using semi-global matching” [1]. The computation on a PC took 2s before we started to implement the FPGA solution. The FPGA accelerator was implemented for Spartan 6 and tested with the SCS Accelerator box. The needed resources are listed in Table 3. It is able to compute a disparity map as shown in Figure 3 from 1024x400 pixel input images with 128 disparities range and 25 fps. The combination of SGM and an optical flow is also known as 6D-Vision [5] that has won the Beckurts price in 2012 and was in the final of the German Future Prize in 2011.



Figure 3: The SGM disparity map computed on the prototyping system. red=near, green=far.

Resources	
Max. clock	130 MHz
FF	12,000
LUT	12,000
DSP48	0
RAMB18	120

Table 3: Resource utilization of the SGM IP.

6. Stixel

Based on the excellent stereo computed by SGM, researchers from Daimler developed the Stixel representation [6]. Disparities are clustered together to rectangular bars called Stixel as you can see in Figure 4. The image shows the representation of an urban scene, the colors encode the distance (red=near, green=far). A big advantage is that subsequent steps such as obstacle detection, freespace computation or attention control don't have to analyze half a million 3D points but only about 500-100 Stixels per image.



Figure 4: Based on Stixel representation, it is much less computation effort to analyze the scene.



Figure 5: The SCS Accelerator Box with direct camera input connectors helps to prototype new systems.

7. SCS Zynq 7045 module

To reduce the effort to build a system, SCS developed a module that allows using most of the functionality that is supported by the Zynq 7045 SOC. As shown in Figure 6, the module uses two comexpress connectors and allows using more than 140 FPGA GPIO's (differential pairs, single-ended or analog) and 8 transceivers. There is a DDR3 memory (x32) connected to the ARM Cortex processing system (PS) and two separate DDR3 memories (x16) connected to the FPGA processing logic (PL). The module is implemented on an 18 layer PCB as shown in Figure 7. By using such a module and a customized

module-baseboard, the developer can focus on the development of his algorithms and enjoy all the amenities that the Linux environment provides.

8. Conclusion

New SOCs that combine ARM processors and FPGA fabric will help to build complex and performing systems. As shown with a few example FPGA accelerators in this paper, it is possible to offload the processors from the number crunching tasks. Together with the Linux these systems will solve problems embedded, that were not feasible just a few years before.




Figure 6: Block diagram of the SCS Zynq 7045 module.




Figure 7: Layout of the SCS Zynq 7045 module. The actual module size is 60mm x120mm.

References

- [1] S. Gehrig, F. Eberli, and T. Meyer. A real-time low-power stereo vision engine using semi-global matching. In International Conference on Computer Vision, Liege, Belgium, October 2009.
- [2] F. Stein. Efficient computation of the optical flow using the census transform. In 26th DAGM Symposium on Pattern Recognition, Tuebingen, August 2004. Springer-Verlag Berlin Heidelberg.
- [3] Xilinx. <http://www.xilinx.com>, 2013.
- [4] Wikipedia. <http://en.wikipedia.org/wiki/JPEG>, 2013.
- [5] <http://www.6d-vision.com>, 2013.
- [6] D. Pfeiffer, U.Franke. Towards a Global Optimal Multi-Layer Stixel Representation of Dense 3D Data, British Machine Vision Conference, August 2011.