AN FPGA/MPSOC BASED LOW LATENCY ONBOARD SAR PROCESSOR

Helko Breit, Srikanth Mandapati, Ulrich Balss

German Aerospace Center (DLR)
Remote Sensing Technology Institute
82234 Weßling

ABSTRACT

This paper describes the concept and prototype implementation of a low latency spaceborne onboard Synthetic Aperture Radar (SAR) processor runing on a Multi-Processor-System-On-Chip (MPSoC) computing device combining an ARM processor and a Field-Programmable-Gate-Array (FPGA). The SAR processor is designed to generate SAR imagery from TerraSAR-X stripmap data for subsequent ship detection and sea state determination. Low latency data processing is a key development goal. Currently, a raw data block of 8k x 32k samples, covering 375 km^2 to 500 km^2, is focused on the hardware within 4 s. Together with an attached level-2 ship detection, wind, and sea state processor, running on the same device, a SAR data processing chain for generation of maritime alerts is formed. This chain is part of a larger prototype system being developed in the frame of the H2020 EO-ALERT project which further comprises an optical data chain, data compression/encryption, and scheduling on multiple reconfigurable MPSoC boards.

Index Terms— SAR, on-board processing, TerraSAR-X, FPGA, MPSoC, EO-ALERT

1. INTRODUCTION

Spaceborne SAR is an indispensable data source for a manifold of operational monitoring services aiming at near real time maritime information, such as sea ice monitoring, ship detection, wind field and wave height estimation, or oil spill detection. Current SAR missions rely on a sparse network of dedicated ground stations to receive vast amounts of acquired SAR raw data on the order of 20 to 40 Giga bytes (GB) per contact. For many applications, the information being finally retrieved from that data comprises only a few kilo bytes as is the case for ship coordinates or wind vector fields. In these scenarios information latency could be drastically reduced by bringing SAR data processing on board and immediate broadcasting of retrieved information via low bit-rate data communication channels. Another promising scenario for onboard SAR processing is autonomous onboard decision making such as instrument tasking based on most recent information being extracted from onboard processed SAR images.

In contrast to on-ground processing where multi-core servers in combination with massive parallel GPGPUs ensure high throughput, while power consumption is of secondary interest, the use of FPGAs enable sufficient low latencies in space at a much lower power requirement. Modern MPSoC devices allow for straight forward implementation and to run non-time critical calculations on a CPU, while demanding signal processing is designed as an efficient pipelined streaming mode FPGA circuitry.

One of the main goals of the EO-ALERT project [1] is to demonstrate by prototype development that low latency on-board data processing starting from level 0 SAR raw data up to level 2 application specific products is feasible. In order to setup a realistic development and demonstration scenario the payloads of operational satellite missions have been selected as data source, TerraSAR-X and Deimos-2, respectively.

Compared to a SAR processor being developed to run on general-purpose computer hardware, possible use scenarios for a FPGA based processor are less flexible due to limited resources and fixed size caches and hardwired data paths. In principle the onboard SAR processor being presented here is able to process stripmap SAR raw data from different sensors, but the design has been optimized for the parameter space of the TerraSAR-X stripmap mode. The processor supports a wide range of pulse repetition frequencies (PRF), different pulse lengths and bandwidths as well as different elevation beams.

2. OVERVIEW OF THE HARDWARE

The on-ground breadboard hardware featured for the SAR processing module in the EO-Alert project is a proFPGA motherboard hosting a Xilinx Zynq UltraScale+ ZU19EG device. The Zynq UltraScale+ integrates a 64-bit quad-core ARM Cortex-A53 and dual-core ARM Cortex-R5 based processing system (PS) and programmable logic (PL) in a single device. PS has onboard interfaces like USB UART for debugging and SDIO. Further interfaces like PCIe, Gigabit Ethernet, and 5 GB Double Data Rate Synchronous Dynamic Random-Access Memory (DDR4 SDRAM) at PL are added as extensions to the Zynq UltraScale+ board. In addition, it offers 4 GB DDR4 SDRAM memory at PS side and a single

Quad Serial Peripheral Interface (QSPI) flash memory. On the ARM Cortex-A53, a Linux operating system (OS) is built and deployed. The entire level-1 SAR and subsequent level-2 processing are performed on a single board hosting the MPSoC. Details on level-2 processing are presented in [2].

3. ONBOARD PROCESSING CONCEPT

In the framework of this prototyping project the interfaces to the satellite payload, such as the SAR instrument, Attitude and Orbit Control System (AOCS), and orbit propagator have been abstracted by bulk transfers of binary datasets to the PL attached DDR4 SDRAM: orbit data, attitude quaternions, 8bit/8bit non-quantized SAR raw data, and acquisition specific instrument settings. Output of level-1 processing is a SAR image together with a geographical coordinates grid. The prototype scenario specific data interfaces are schematically depicted in Fig. 1.

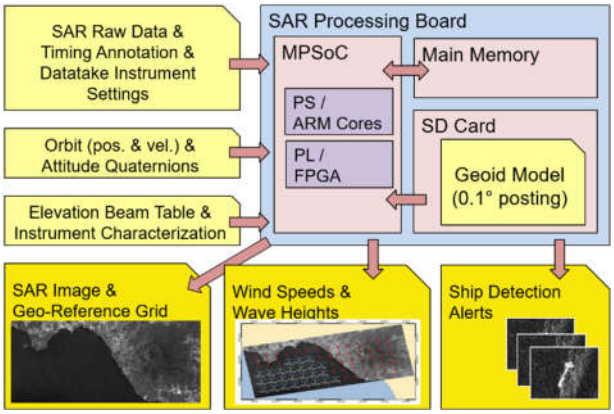

Fig. 1: Prototype onboard SAR processing data interfaces

In contrast to ground segment-based processing which serves a manifold of applications with different image product variants, onboard SAR processing in the context of EO-ALERT is designed to meet the requirements of ship detection and sea state determination. Output of SAR processing is an approximately 6 m to 8 m resolution, 2 by 2 **M**ulti-look, **S**lant range **D**etected (MSD) σ_0-calibrated image supplemented by a geo reference grid which maps the image's radar time coordinates onto a low-resolution Geoid model. Compared to a high-resolution DEM data base the 10 km resolution Geoid model requires only 13 MB onboard storage and still allows for high accuracy SAR focusing parameter calculations as well as high accuracy geolocation over ocean surfaces. Furthermore, SAR raw data screening is skipped. In fact, it is assumed that calibration information and instrument characterization parameters are uploaded to the satellite in advance and stored onboard. Instead of estimating Analog Digital Converter (ADC) correction parameters and Doppler centroid values from the SAR raw data itself, pre-characterized ADC corrections are applied and the Doppler centroid is geometrically derived from onboard AOCS data. Rather limited main memory resources of 5 GB at PL side restrict image generation to a single raw data block of 8k pulses in azimuth containing up to 32k range samples. Thus, the full TerraSAR-X 30 km stripmap swath is covered in all elevation beam scenarios, but the along track extent of one image is limited to 11 to 15 km. Larger continuous azimuth coverage is obtained by repeated image generation of overlapping subsequent 8k azimuth raw data blocks.

4. FOCUSING ALGORITHM AND PROCESSING SCHEME

In the EO-ALERT project, preference is given to a spectral domain focusing algorithm instead of time domain back projection. This decision is based on long term inhouse development experiences and on reports of successful FPGA implementations by other teams [3,4,5]. It also reduces development risks. Compared to the chirp scaling algorithm being adopted by the TerraSAR-X ground segment processor, the finally selected monochromatic ω-k algorithm is less complex and does not require a range pre-compression step which is required for higher order chirp replica focusing and which implies two additional range FFTs. In contrast to the original ω-k algorithm the monochromatic variant omits the cumbersome Stolt mapping. However, this is at the cost of uncompensated residual range migration Δr. According to [6], it amounts to

$$\Delta r(k_x, r) = \frac{c^2}{8} \cdot (r - r_0) \cdot \frac{k_x^2}{\omega_0^2} \text{ with } k_x = \frac{2\pi f_D}{v_{eff}}.$$

For the envisaged maritime onboard applications, the processed Doppler frequencies do not exceed $f_D = \pm 1000$ Hz which corresponds to a single look azimuth resolution in the order of 4 m. TerraSAR-X's effective velocity is in the order of $v_{eff} = 7350$ m/s and the circular frequency $\omega_0 = 2\pi \cdot 9.65 \cdot 10^9$ s^{-1}. For a maximum swath of $r - r_0 = \pm 15$ km the maximum residual range migration amounts to $\Delta r = \pm 0.034$ m. This is an acceptable blurring of less than 1% compared to the envisaged 4 m single look resolution.

A schematic overview of the SAR data workflow and its interfaces is shown in Fig. 2. While the entire signal processing is implemented in the PL, mostly employing integer arithmetic, all geometric calculations and preparatory SAR signal filter coefficient calculations are implemented and executed in the PS in floating-point precision: orbit interpolation, geometric Doppler centroid determination, geolocation, projection of the elevation antenna pattern into slant range timing coordinates, calibration and annotation parameter calculations, computation of filter phase coefficients to be loaded to PL Random-Access-Memories (RAM) which act as Look-Up-Tables (LUTs) for subsequent pipelined computation of the 2D monochromatic ω-k bulk focusing kernel and the range variant residual azimuth compression filters.

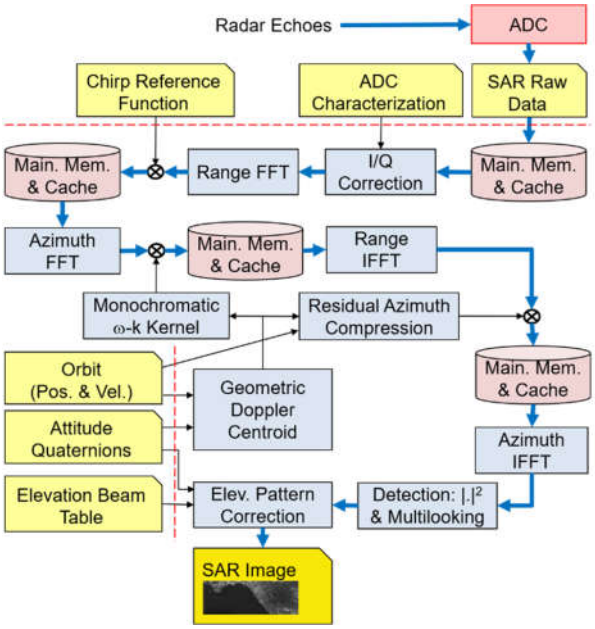

Fig. 2: Processing scheme based on the monochromatic ω-k algorithm, SAR data workflow, and data interfaces.

5. DESIGN AND IMPLEMENTATION

5.1. Data transfer and storage

Processing of a raw data block of 8k x 32k of 8bit/8bit complex integer data size requires at least 2.5 GB of memory for storage of SAR data between the different FFT steps as depicted in Fig. 2. FPGA on chip storage capacity is insufficient to store the entire SAR data during processing. As a solution, an external DDR4 SDRAM is attached to the PL and the Xilinx Memory Interface Generator (MIG) IP core is used. The MIG IP is used to interface with the physical layer of the DDR4. Along with the raw data, attitude, orbit, and instrument settings are needed for SAR focusing. These auxiliary data sets are stored in PL DDR4 as well. In the integrated EO-ALERT system the data is written via a PCIe interface to the DDR 4. Once the filter phases and elevation pattern corrections are computed on PS side, the results are transferred to the LUTs in the PL RAM. Then PL takes control and performs all signal processing such as ADC corrections, pulse compression, FFTs, pixel-wise complex filter multiplications, detection, and multi-looking.

5.2. Cache and matrix transpose

SAR signal processing is performed on small subblocks of data which fit to the storage and processing capacity of the FPGA. First, a subblock is buffered from PL DDR4 to a first group of URAMs forming the input cache. Next, the data go through the bit width conversion, and get streamed to the signal processing data path. Finally, the output of the data path is written to a second group of URAMs forming the output cache, see Fig. 3. URAM buffering of data facilitates matrix transposition and enables burst-mode I/O operations to the DDR4 SDRAM. The number of available URAMs allow for buffering of 16 x 32k range lines or 64 x 8k azimuth lines. Thus, for each FFT type, see Fig. 2, the cache is filled and emptied 512 times.

5.3. Signal data path

The signal data path is the major component in the PL design of SAR focusing. The data path contains the Xilinx IP cores for FFT, RAM, COordinate Rotation DIgital Computer (CORDIC), data width converters (DWC) and floating-point multiplication. The Xilinx FFT is implemented with the Cooley-Tukey FFT algorithm which is a computationally efficient method for calculating the Discrete Fourier Transform (DFT). The main advantage of the Xilinx FFT IP core is its configurability at runtime for both, forward and inverse FFT of up to 32768 points each of 16bit/16bit complex integer data. The FFT IP core supports pipelined streaming I/O and is configured as fixed point with adaptive block floating-point scaling to evade the clipping or wrapping of integer values.

Available FPGA resources allow to instantiate 4 parallel FFT pipelines. Each FFT instance processes 4 x 32k range lines or 16 x 8k azimuth lines stored in the URAM cache. In total the FFT instances process 2 x 8192 x 32k ranges lines and 2 x 32768 x 8k azimuth lines. Running 4 parallel data paths significantly decrease the overall latency. Real-valued focusing filter phases are calculated in double-precision floating-point format based on the coefficients stored in PL LUTs by using dedicated floating-point multiplication IP cores. Floating-point filter phases are wrapped to baseband and the CORDIC IP core computes the complex exponential with 16bit/16bit integer output format and finally complex filter samples and complex SAR data are multiplied. The IP cores, data type and bit width changes in the data path are depicted in Fig. 3. As the last step, the complex data is converted into amplitude using the CORDIC IP core and multi-look summation is applied on pixel amplitudes.

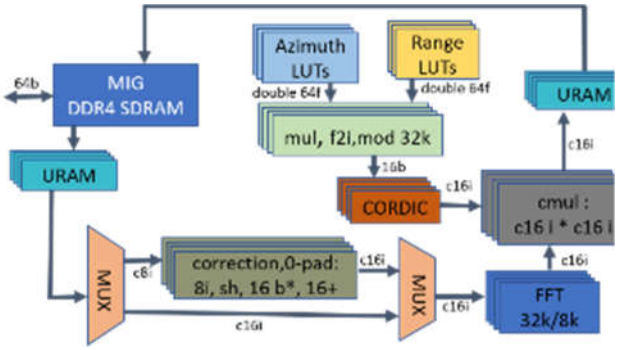

Fig. 3: Concept and building blocks of the SAR signal data path.

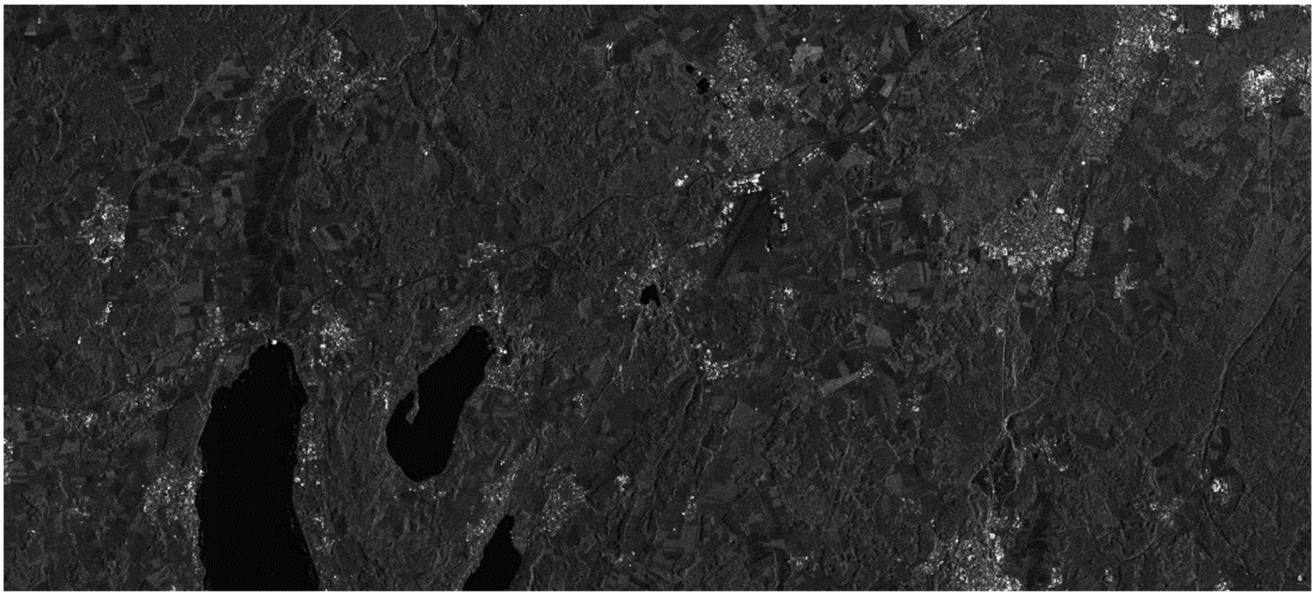

Fig. 4: MPSoC/FPGA processed 12.5 km x 31 km SAR image of Oberpfaffenhofen based on 8192 TerraSAR-X SAR raw data lines. In order to provide adequate terrain heights significantly above sea-level, the Geoid model on the processing board's SD card has been replaced for this image by a subsampled version of the GLOBE DEM, with 13 MB data size and 10 km resolution.

5.4. Resource utilization and latencies

The number of available BRAMs and URAMs is the limiting factor of the SAR focusing implementation. BRAMs are used in the FFT IP cores and as LUTs for storage of filter coefficients. All available URAMs are used as SAR data caches in the FPGA. Table 1 summarizes resource utilization of the SAR level-1 processing. Remaining resources suffice implementation of the ship detection CFAR core [2].

Table 1: Resource utilization of the implemented SAR focusing algorithm on Zynq UltraScale+ ZU19EG MPSoC.

Resource type	Available	Used	Utilization
Flip flop	1045K	71K	6.8 %
LUTs	523K	61K	11.7 %
Block RAM	984	669	68 %
UltraRAM	128	128	100 %
DSP Slices	1968	317	16.1 %

Totalling the processing times of PS, signal processing, and I/O throughput of the DDR4 SDRAM MIG gives the total latency of level-1 SAR processing on the Xilinx Zynq US+ ZU19EG. Running the PL design at 125 MHz clock, the latency is measured as 4.1 s. PS processing contributes 0.3 s, the SAR signal data path 2.9 s, and I/O 0.9 s.

6. ACKNOWLEDGEMENTS

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 776311. The EO-ALERT project is coordinated by DEIMOS Space. More information on the project at eo-alert-h2020.eu.

7. REFERENCES

[1] M. Kerr, S. Tonetti, S. Cornara, et al., "EO-ALERT: A Novel Architecture for the Next Generation of Earth Observation Satellites Supporting Rapid Civil Alerts, " 71st International Astronautical Congress (IAC 2020), 12.-14. Okt. 2020.

[2] S. Wiehle, D. Günzel, B. Tings, "SAR Satellite On-board Ship, Wind, and Sea State Detection," Proc. IGARSS 2021

[3] M. Pfitzner, F. Cholewa, P. Pirsch and H. Blume, "FPGA based architecture for real-time SAR processing with integrated motion compensation," Proceed. of 2013 Asia-Pacific Conference on Synthetic Aperture Radar (APSAR), Tsukuba, 2013, pp. 521-524.

[4] C. Yang, et al., "A Spaceborne Synthetic Aperture Radar Partial Fixed-Point Imaging System Using a Field- Programmable Gate Array—Application-Specific Integrated Circuit Hybrid Heterogeneous Parallel Acceleration Technique,". Sensors (Basel, Switzerland). 2017, 17, 1493; doi: 10.3390/s17071493.

[5] Y. Sugimoto, S. Ozawa, N, Inaba, "Spaceborne synthetic aperture radar signal processing using field-programmable gate arrays, " J. Appl. Remote Sens. 12(3), 035007, 2018, doi: 10.1117/1.JRS.12.035007

[6] R. Bamler, "A Comparison of Range-Doppler and Wavenumber Domain SAR Focusing Algorithms," IEEE Trans. On Geoscience and Remote Sensing, 30 (4), pp. 706-713, July 1992.