Demonstrating a SAR Satellite Onboard Processing Chain

7th Workshop on RF and Microwave Systems, Instruments & Sub-systems +
5th Ka-band Workshop
10-12 May 2022

Helko Breit[(1)], Stefan Wiehle[(2)], Srikanth Mandapati[(1)], Dominik Günzel[(2)], Ulrich Balss[(1)]

[(1)]German Aerospace Center
Remote Sensing Technology Institute
Münchener Str. 20, 82234 Weßling, Germany
Email: {helko.breit, srikantha.mandapati, ulrich.balss}@dlr.de

[(2)]German Aerospace Center,
Maritime Safety and Security Lab Bremen
Am Fallturm 9, 28359 Bremen, Germany
Email: {stefan.wiehle, dominik.guenzel}@dlr.de

INTRODUCTION

With the rising number of Earth Observation (EO) satellites, especially driven by missions consisting of many small EO satellites designed for high revisit frequencies of several times per day, spaceborne EO gains more and more attention by time critical applications. The needs of such applications push the EO data chain towards faster and faster delivery of products to their users. The major factors limiting the utilization of today's spaceborne SAR systems are onboard storage capacity, ground station contacts and transmission network speed summing up to comparatively high latencies of product delivery. Onboard data processing and information extraction is one of the solutions to overcome such limitations. This is especially advantageous for event-driven processing and rapid provision of information products such as ship detections or extreme weather and sea state parameters. The EO-ALERT project [1], being implemented by a European consortium in the frame of the European Union's Horizon 2020 programme, has developed, built, and tested a prototype demonstrator chain comprising onboard raw EO data (SAR and optical) Level 1 processing, Level 2 image processing and transfer of final EO alert products to the end user. The goal of the project is to demonstrate the feasibility of the concept aiming at overall latencies below 5 minutes even down to 1 minute.

Besides fast onboard processing, the EO-ALERT concept includes immediate low rate downlink options in the absence of ground station contacts. In order to minimize the transfer latency, the output results of the processing chains are specified as small packets of data, so-called alerts, containing only the demanded information. Hence, in the proposed communication scenarios the user does not directly obtain the entire L1 image product due to data volume limitations. Instead, the raw data is optionally stored onboard and may be downlinked during the next ground station contact and fed into currently existing EO services with higher latency.

The task of DLR within the project was to implement a low latency SAR onboard Level 1 and Level 2 processing chain. SAR processing is well known for its high demand for computational resources. The targeted latencies together with low power and low mass constraints require the use of Field-Programmable Gate Array (FPGA) technology. Multi-Processor System on a Chip (MPSoC) devices were chosen for all onboard data chain developments. SAR image formation (IF) and SAR image processing (IP) are integrated in a single bitstream file and loaded into a single MPSoC device.

The project did not aim for implementation of functional input interfaces to a specific optical or SAR payload or other onboard electronics such as an Attitude and Orbit Control System (AOCS). In the context of the project, these interfaces are substituted by data files providing the required inputs for processing. In the case of SAR, the TerraSAR-X instrument has been selected as a source of representative input data and the algorithmic layout of the onboard SAR processing chain is tailored to the Stripmap (SM) acquisition mode. For interfacing the EO-ALERT chain, the sensor specific and complicated instrument source packet (ISP) structure of TerraSAR-X data takes have been converted into simply structured data files containing 8bit/8bit complex integer raw data samples supplemented by a set of parameters annotating the timing and data take specific instrument settings.

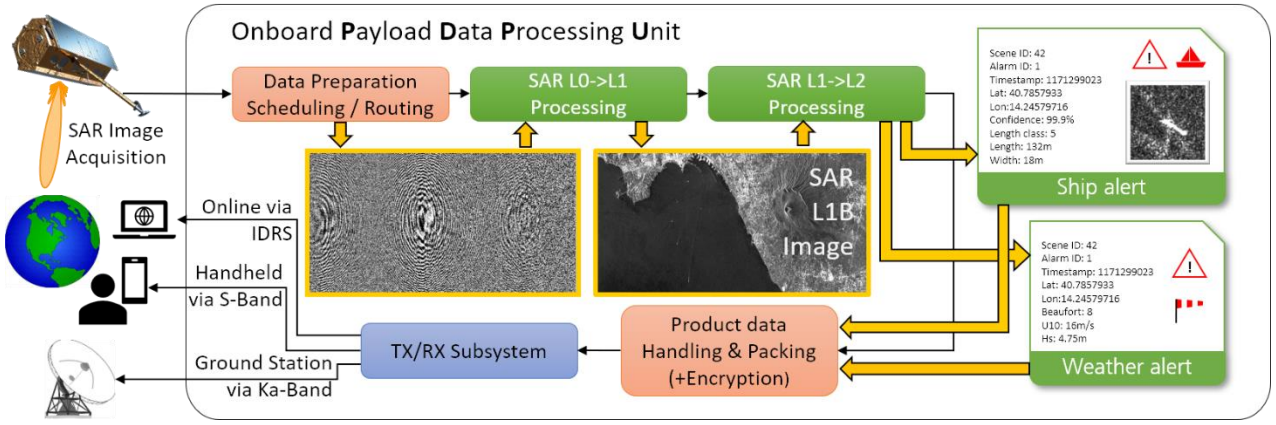

Fig. 1: Workflow of the demonstrator. After SAR data acquisition, the raw data arrives at the Compression, Encryption and Data Handling (CEDH) board (red), and is then transferred to the SAR processing board (green). After processing, the CEDH board handles packing and encryption and forwards the generated alerts and optionally the raw data to the communications board (blue) for downlink. Three downlink modes were designed or implemented in the project: IDRS, S-Band and Ka-Band.

The following sections provide an overview of the demonstrator chain built in the project, the SAR processing chain implementation and the achieved test results.

THE SATELLITE ONBOARD PROCESSING DEMONSTRATOR

The demonstrator chain built in the EO-ALERT project consists of a test bench and additional hardware for communications. The test bench consists of four MPSoC boards. One board is used for data handling [2] and three boards are used for the EO data processing. The SAR processing chain uses only one board and the optical processing chain uses up to three boards for its data processing. After processing, the encrypted results (raw data, processed data, and alerts) are transferred from the test bench to an external communications board, which takes care of transmitting the results in three possible ways [3]. First, a simulated low rate direct S-band downlink to a hand-held user terminal, which was built in the project, simulating global availability of the alerts for end users equipped with an appropriate receiver. Second, a direct communication link between the EO-ALERT test bench and the Inmarsat Inter Satellite Data Relay System (IDRS) was set up. For that purpose, a space-representative broadband global area network (BGAN) transceiver hardware was attached to the test bench, which transmitted the processed results to the IDRS making the data finally accessible on a web server via an IP connection. As a third option, Ka-band downlink to a ground station has been emulated. The full workflow of the processing chain for a SAR data acquisition is sketched in Fig. 1.

HARDWARE OVERVIEW

The SAR processing chain is implemented and runs on a prototyping board equipped with a Zynq UltraScale+ ZU19EG MPSoC from Xilinx. The Xilinx Zynq UltraScale+ series MPSoC integrates a 64-bit quad-core ARM Cortex-A53 and dual-core ARM Cortex-R5F based processing system (PS) and UltraScale+ architecture FPGA programmable logic (PL) in a single chip. The ARM Cortex R5F is not employed in the design of the SAR processing chain. The components of the PS are connected together and to the PL with a multi-layered ARM AMBA Advanced eXtensible Interface (AXI) non-blocking interconnect that enables multiple master-slave transactions simultaneously. Furthermore, 4 Gigabytes (GB) DDR4 SDRAM memory are connected to the PS of the MPSoC and 5 GB DDR4 SDRAM multiport external memory are connected to the PL. Several peripheral connectivity interfaces, such as Gigabit Ethernet and PCI Express extension boards, are connected for high speed data transfers. For flashing the board and storing static ancillary data, an 8 GB micro-SD card is connected to the PS. A purpose-built embedded Linux operation system (OS) deployed on the PS runs the control program and parts of the algorithms in software and offers customisation of device drivers, applications, libraries, and boot parameters.

SAR PROCESSING IMPLEMENTATION

Since the complete SAR processing chain is designed to run on a single MPSoC and share its resources, it is important to assess the algorithms of Level 1 and Level 2 processing with respect to processing speed, computational complexity and resource availability. Especially the workload and suitability of the algorithms for implementation in the PL were investigated. As pointed out in [4][5], image formation is based on the monochromatic ω-k algorithm [6]. The main

constituents of the algorithm are 4 types of Fast Fourier Transforms (FFTs), which are applied in forward and backward direction in the range and azimuth dimensions, respectively. As depicted in Fig. 3, in-between the FFTs the focusing filters are computed and multiplied with the SAR signal data. These signal processing steps are characterized by high computational complexity and a high demand for parallelization, thus, favorably implemented and executed in pipelined streaming mode in the PL. On the other hand, less computing time demanding parameter calculations such as iterative geometric calculations based on orbit and attitude data are far more efficiently developed in software (SW) running on the PS OS. The PS also hosts the main control of execution and the communications daemon interfacing as a client program with a server program running on the CEDH board [2].

The entire algorithmic work and data flow is depicted in Fig. 3. SAR image formation is started with the availability of the input data: i.e., an annotated SAR raw data file together with a chirp pulse replica and an orbit state vectors and attitude quaternions file. Level 2 processing takes over as soon as a focused and annotated SAR image has been formed and becomes available in the PL-attached DDR4 memory. After Level 2 processing, the number of generated alerts is signaled via the communications daemon to the CEDH board and the alert data is picked up from the PL DDR4.

Details on the implementation of SAR image formation and SAR image processing are presented in the following subsections as well as in [4][5][7].

Level 1 SAR Image Formation

SAR IF necessitates image focussing through consecutive signal processing of raw data. Limited by the available PL-attached DDR4 memory resources and due to an FPGA FFT IP core imposed 2^n length constraint, the SAR processor developed in the project is designed to process raw data blocks containing exactly 8192 azimuth lines and up to 32768 raw data pixels per line. Each raw data pixel is an 8-bit/8-bit complex integer value. The granularity of the block size ensures a spatial coverage of 375 km^2 to 500 km^2 depending on the acquisition geometry (incidence angle), the pulse repetition frequency and the pulse range bandwidth.

Along with 8-bit/8-bit complex integer SAR raw data, timing data and instrument settings as well as Attitude and Orbit Control System (AOCS) data are required as inputs for SAR IF. A geoid model with 10 km lateral resolution is used in place of a high-resolution DEM to provide local sea surface heights w.r.t. WGS84 ellipsoid. SAR IF either generates a Single Look Complex (SLC) image or a Multi look, Slant range Detected (MSD) image and allows to flexibly choose the desired output type. By default, an MSD image is generated as input for subsequent SAR IP since the ship detection or wind and sea state determination algorithms are based on calibrated pixel intensities and do not exploit the Doppler spectrum of an SLC image. The focused image is σ_0-calibrated with a resolution between 6 m to 8 m, offering approximately 4 radiometric looks, and it is accompanied with a geo-reference grid which maps image radar time coordinates to geographic coordinates utilising the geoid model.

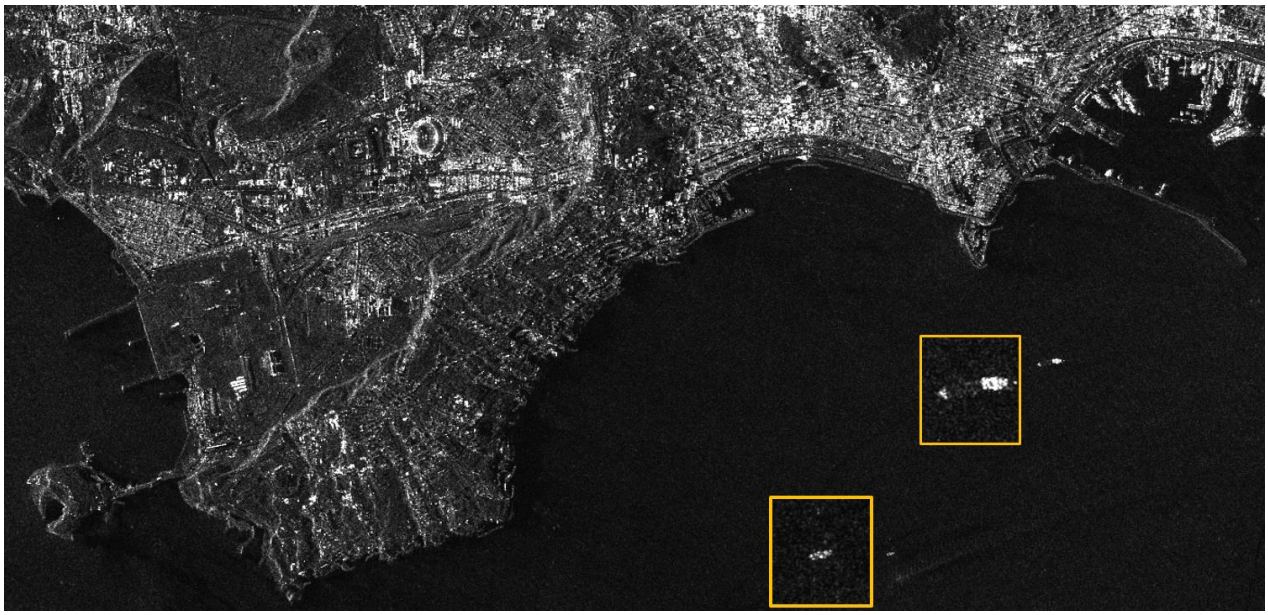

Fig. 2: Hardware-processed TerraSAR-X image detail depicting parts of the city of Napoli and highlighting two detected ships. The 10 km × 5 km detail is extracted from a focused azimuth block covering 32 km in range and 13 km in azimuth direction.

Signal processing of the sensor's SAR data is implemented in the PL. The four different FFT and filter kernel multiplication steps, see Fig. 2, are handled by a single generalized data path design which is configurable at runtime for an individual step of the monochromatic ω-k algorithm. The circuitry consists of a pre-processing section connected to an FFT IP core, followed by a post-processing and corner turning section. Pre-processing is performed on the digitized sensor data to correct artefacts introduced in Analog-to-Digital-Converters (ADCs) due to unavoidable manufacturing tolerances. The FFT IP core is configured either for 8k azimuth FFTs or 32k range FFTs in forward or backward direction. The post-processing circuitry, which is connected to the output of the FFT IP core, comprises pipelined pixel-wise spectral filter computation and complex multiplication with the SAR signal or detection and multilooking. Corner turning is a key challenge for efficient SAR IF and it is required after almost each focusing step since each of the different FFT steps accesses the data in alternating direction, azimuth or range. The data path design is instantiated in parallel four times on the device and it is connected to an on-chip input and output cache. This cache is build up from UltraRAMs and allows for 16 × 32 k range lines or 64 × 8 k azimuth lines to be cached at a time. The software running on the PS is responsible for generation of filter coefficients required for subsequent signal processing in the PL. Computation of the points of the geo-reference grid is based on an iterative numerical geo-location algorithm running in the PS as well. The grid serves as geometric annotation of the focused image.

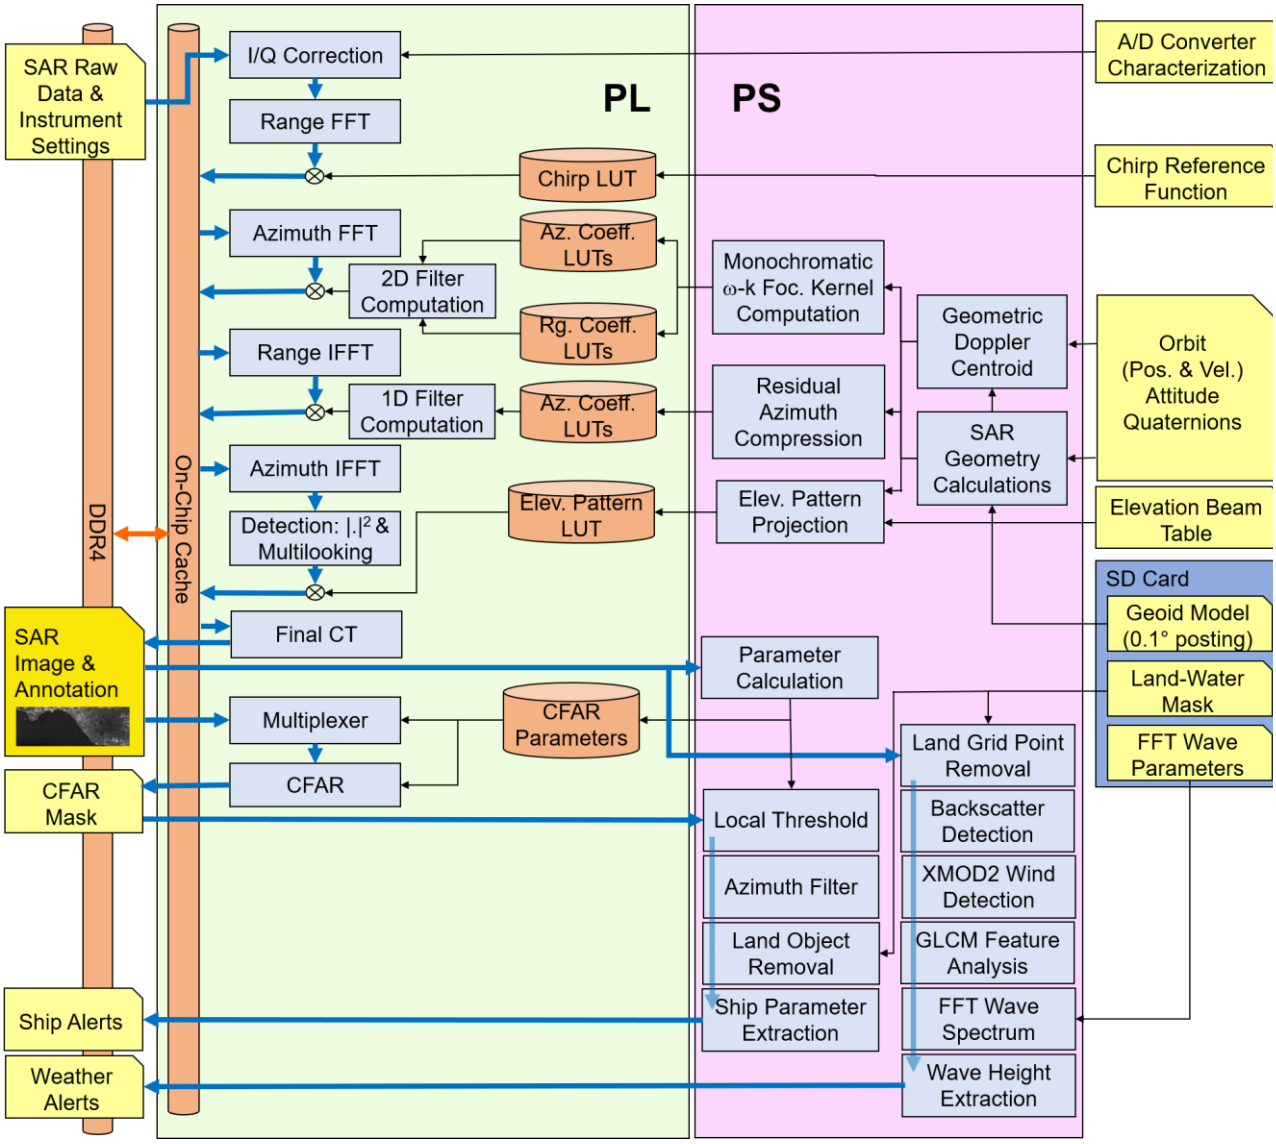

Fig. 3: Data workflow and PL/PS partitioning of the algorithmic components of integrated Level 1 SAR image formation and Level 2 information extraction and alert generation. The data workflow starts with availability of SAR raw data (top left side). Prior to PL side SAR signal processing, the PS side fills the look up tables (LUT) within the PL with all required coefficients. Level 1 processing ends with the final corner turning step completing SAR image formation. Level 2 processing takes over and ends with the generation of alerts (bottom left side).

Level 2 SAR Image Processing

In the context of the EO-ALERT project Level 2 processing aims at ship detection and extreme weather indicators such as sea surface wind speed and wave heights. Both algorithms are running fully operationally in the ground station today and have been adopted for onboard processing in the project instead of developing new processors from scratch.

The ship detector employs a constant false alarm rate (CFAR) algorithm to discriminate ships from background, followed by several filtering steps [8]. This CFAR step was the only part of SAR IP which had to be implemented in the PL as it is computationally too extensive to be performed in the PS in reasonable time. A processing pipeline consisting of 29 stages supplied with pixel data by an efficient streaming cache in the form of an indexed queue was implemented in the PL, significantly speeding up the processing while still allowing for flexibility with respect to window size. This is required as the pixel spacing of the MSD product is not fixed but the guard and background windows used are determined by ship sizes and should therefore not depend on the image resolution. Other steps of the ship detector are run in the PS and were adopted from the operational ground processor.

The extreme weather processor is also based on its operational ground counterpart, but was recoded in several aspects to improve the latency. Its main constituents are the wind detector based on the XMOD2 geophysical model function (GMF) [9], and the sea state detector [10], which combines a spectral analysis via FFT and a gray-level co-occurrence matrix (GLCM) analysis. The processor divides the scene into a grid with cell size of 2 km by 2 km and one value for wind and sea state is calculated per cell, indicating the average value in the respective area.

RESULTS

SAR processing was successfully implemented and executed in the targeted Xilinx Zynq UltraScale+ ZU19EG MPSoC. Table 1 lists the resource utilization of the L1 and L2 circuits. The limiting factor was the available UltraRAM of which all available 128 tiles were used. All other resources (slice registers, slice look-up tables, Block RAM tiles and DSP slices) were only partially used. The available resources allowed the implementation of 4 parallel data paths for SAR IF. The latency scales almost reciprocally proportional with the number of data paths running. The design runs at a clock rate of 125 MHz.

Table 1: Resource utilization of the SAR L1 and L2 onboard processors integrated into the Zynq UltraScale+ ZU19EG FPGA

Site type	L1	L2	Total	Available	Utilization
Slice registers	71K	28K	99K	1045K	9.5 %
Slice LUTs	61K	41K	102K	523K	19.5 %
Block RAM tiles	669	145	814	984	82.7 %
UltraRAM tiles	128	0	128	128	100.0 %
DSP slices	317	260	577	1968	29.3 %

Comparing the image quality of the onboard product to the underlying TerraSAR-X ground processor at corner reflectors, peak position and side lobe parameters are very similar with differences in the order of a few centimetres and less than 1 dB, respectively. The complex data products are within 0.2 dB m² of average signal energy, indicating correct radiometric calibration of the onboard processor.

The onboard ship detector has detected 18 of 21 ships across all test scenes used in the project, whereby the undetected ships had a size of 21 m or below; the next larger vessel was above 60 m and was detected, as were all larger ships. In additional scenes acquired during an experimental campaign, the ship chartered by the project had a length of 29 m and was detected in all 3 acquisitions. The detection of smaller vessels is challenging with the reduced pixel spacing of the onboard product, but similar to ground processing if similar pixel spacing is present.

The wind speed and sea state processors were compared to their ground counterparts running on the original TerraSAR-X products. For wind, the ground processor retrieves the wind direction from modelling results prior to processing, which is not considered to be available on a satellite, so the onboard processor uses a fixed wind direction. A root-mean-square deviation (RMSD) of 0.7 m/s was achieved. For wave heights, calibrating the processor for a new product with different resolution is an extensive task requiring hundreds of scenes, which were not available in the project. Hence, the unaltered TerraSAR-X calibration was used. The measured RMSD is 0.2 m, indicating that both products are still very similar despite the different resolutions available.

The SAR processing chain is part of the full onboard data chain successfully built and demonstrated in the EO-ALERT project, also containing optical processing, data handling and encryption, and transmission to the end user. Generated

alerts were transmitted from the test bench via a BGAN transceiver and a ground antenna to the IDRS which makes the data finally available via TCP/IP access to a remote server. This is a representative globally available dissemination scenario for highly condensed information products generated on spaceborne platforms.

The prime goal and the ambitious top-level requirement of the project is rapid delivery of relevant information to the end user with a total system latency below 5 minutes. This includes all onboard data handling, onboard processing as well as alert delivery to the end user. This requirement leaves up to 210 s for the full SAR processing. Latency tests were performed for 9 TerraSAR-X test acquisitions. On average, L1 image formation took 3.7 s. Minor variations occur depending on different image dimensions. L2 processing is divided into ship detection and wind and sea state detection. For ship detection, average latency was 16.0 s and wind and wave height detection took 30.8 s. In both scenarios, latencies vary substantially due to varying land coverage in the processed scenes. Combining the SAR L1 and L2 processing times, even in the worst case the 210 s requirement was easily fulfilled, facilitating the project goal of rapid alert delivery to the end user.

OUTLOOK

In the frame of the demonstrator project EO-ALERT the design of SAR IF is tailored to the Stripmap mode of TerraSAR-X. Nevertheless, the building blocks such as FFTs, focusing filter computations, complex multiplications, and corner turning may be rearranged to form other focusing algorithms such as Spectral Analysis (SPECAN) for ScanSAR imaging modes or Baseband Azimuth Scaling (BAS) for the Terrain Observation Mode SAR (TOPS) employed by the Copernicus Sentinel-1 satellites. Comparable to the Stripmap mode, where the granularity for Level 1 and Level 2 processing is an 8k azimuth block, the burst mode raw data will be focussed to individual burst-level images for subsequent Level 2 processing. Stitching together small burst images to a product of larger coverage, which is a memory intensive task, prior to level 2 processing is not required and should be omitted onboard.

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. This work reflects only the author's view and the EC is not responsible for any use that may be made from the information contained in this work. The EO-ALERT project was coordinated by DEIMOS Space. More information on the project at eo-alert-h2020.eu. TerraSAR-X data were provided by DLR under AO proposal OCE3625.

REFERENCES

[1] M. Kerr et al. “Advanced Data Chain Technologies for the Next Generation of Earth Observation Satellites Supporting On-Board Processing for Rapid Civil Alerts,” *IAC 2020, 71th International Astronautical Congress*, 2020.

[2] M. Caon, P. Motto Ros, T. Bianchi, M. Martina, E. Magli, “Low Latency On-Board Data Handling for Earth Observation Satellites Using Off-The-Shelf Components,” 2nd *European Workshop on On-Board Data Processing (OBDP2021)*, 14-17 June 2021, doi: l10.5281/zenodo.5570469.

[3] O. Koudelka, F Teschl, W. Gappmair, M. Kerr, S. Tonetti, S. Cornara, “Advanced Communications Solutions for the Next Generation of Earth Observation Satellites,” *IAC 2020, 71th International Astronautical Congress*, 2020.

[4] S. Wiehle, D. Guenzel, B. Tings, H. Breit, U. Balss, and S. Mandapati, "A Satellite On-Board SAR Processing Chain for Generation of Rapid Civil Alerts," *EUSAR 2021; 13th European Conference on Synthetic Aperture Radar*, Mar 29-Apr 01, 2021

[5] H. Breit, S. Mandapati, and U. Balss, "An FPGA/MPSoC Based Low Latency Onboard SAR Processor," *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*, 2021, pp. 5159-5162, doi: 10.1109/IGARSS47720.2021.9553539.

[6] R. Bamler, “A comparison of range-doppler and wavenumber domain SAR focusing algorithms,” *IEEE T Geosci Remote Sensing*, vol. 30, no. 4, pp. 706–713, 1992.

[7] S. Wiehle, D. Günzel, and B. Tings, "SAR Satellite On-Board Ship, Wind, and Sea State Detection," *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*, 2021, pp. 8289-8292, doi: 10.1109/IGARSS47720.2021.9554096.

[8] B. Tings, C. A. B. da Silva, and S. Lehner, “Dynamically adapted ship parameter estimation using TerraSAR-X images,” *Int. J. Remote Sens*., vol. 37, no. 9, pp. 1990–2015, 2016.

[9] X.-M. Li and S. Lehner, “Algorithm for Sea Surface Wind Retrieval From TerraSAR-X and TanDEM-X Data,” *IEEE T. Geosci. Remote*, vol. 52, pp. 2928–2939, May 2014

[10] A. Pleskachevsky, B. Tings, and S. Jacobsen, “Sea State from Sentinel-1 SAR Wave Mode Imagery for Maritime Situation Awareness,” *EUSAR 2021; 13th European Conference on Synthetic Aperture Radar*, Mar 29-Apr 01, 2021.