## A Satellite On-Board SAR Processing Chain for Generation of Rapid Civil Alerts

Stefan Wiehle, Dominik Günzel, Björn Tings, Helko Breit, Ulrich Balss, Srikanth Mandapati German Aerospace Center (DLR), Remote Sensing Technology Institute, Stefan.Wiehle@dlr.de, Germany

## Abstract

This paper describes the concept and prototype implementation of a satellite on-board SAR processing chain designed for Maritime Situation Awareness. It aims to reduce the latency between data acquisition and product delivery to about 3-4 minutes. SAR processing is one component of a larger prototype system being developed in the frame of the H2020 EO-ALERT project. It further comprises an optical data chain, data compression/encryption, and delivery. The system employs multiple boards with Multi-Processor-System-On-Chip (MPSoC) combining FPGAs and ARM CPUs. A tailored workflow and adapted L1 and L2 processing algorithms ensure that the requirements for latency and product quality are met.

## **1** Introduction

With the growing number of Earth Observation satellites, the number of acquisitions available in a certain area per time drastically increases. However, to really benefit from these developments, users need fast access to the data. Today, the normal processing chain for satellite imagery consists of acquisition, satellite flight time to the nearest ground station supporting this mission, data downlink, image generation and processing at the ground station, and finally transfer of a product to the user via internet. Bringing the image generation, processing and delivery onto a satellite can, hence, safe the flight time to the ground station and allow a direct product delivery to any user equipped with an appropriate receiver.

For Maritime Situation Awareness this is especially important as acquisitions over the oceans may not only have a longer than average delay until transferred to a ground station, but also the retrieved information – ship positions, wind speeds and wave heights – are very time sensitive and become deprecated within minutes. A reduced product delivery delay of 3-4 minutes after the acquisition will enable this data to become a de-facto-standard for maritime operations, enhancing ship safety and security.

The following sections describe the processing concept and point out changes made to the individual algorithms in response to limited processing capabilities and temporal restrictions implied by a satellite on-board system.

## 2 System overview

In the EO-ALERT project [1], a prototype system will be built and tested supporting both optical and SAR measurements. The processing chain hardware consists of five boards each accommodating a Multi-Processor-System-On-Chip (MPSoC) device which provides programmable logic (PL) and a processing system (PS) based on ARM CPUs running a Linux-based operating system (OS). The boards will be individually configured for the different tasks such as optical processing, SAR processing, data compression and encryption, and downlink. The paper at hand focuses on the developments for the board dedicated for SAR image generation and processing. The EO-ALERT project does not involve design and development of a SAR instrument; instead the TerraSAR-X single polarization StripMap mode has been selected as a representative source of SAR data. Data rates in terms of pulse repetition frequencies, pulse lengths, and sampling rates on the one hand and available hardware resources together with an envisaged processing time limit of 210 s on the other hand drive the algorithmic layout and determine image size and coverage.

Thus, the design selected for EO-ALERT employs spectral domain SAR focusing and foresees Fast Fourier Transforms (FFTs) of 8k point in azimuth and 32k point in range direction. Range lines of lower length are filled up by zero padding at the end of the line. The raw data block with 8k x 32k size allows accommodating at least for a 12.5 km along-track and 30 km ground range image obtained from a single azimuth raw data block. Depending on the radar instrument settings a single block SAR image covers between 375 km<sup>2</sup> to 500 km<sup>2</sup>. Further development in the future may integrate the current single block processing design into a block by block processing chain which generates longer image strips.

## **3** SAR L1 processing

The task of SAR level 1 processing is generation of a focused SAR image from SAR raw data being received, demodulated, and digitized in the SAR instrument electronics. For this task, the attitude and orbit control system (AOCS) of the spacecraft is an indispensable ancillary data source providing orbit positions and velocities as well as attitude quaternions. Since instrument design and development are not covered by EO-ALERT, the SAR raw data is provided as 8bit/8bit complex integer data files derived from TerraSAR-X complex imagery by means of inverse SAR processing. This is necessary as all steps of SAR processing are to be implemented in the on-



**Figure 1**: Components of the SAR processing board and implementation of the different steps in image generation (L1), sea state and ship detection (L2).

board system, starting from raw data as provided by the SAR instrument. Orbit and attitude information is extracted from TerraSAR-X image products as well. In comparison to on-ground processing where highest possible processing accuracy is requested to open and serve wide fields of SAR applications, EO-ALERT SAR image generation is tailored to the demands of on-board ship detection and sea state analysis. Tailoring applies to the workflow as well as to the focusing algorithm and its implementation.

### 3.1 Workflow tailoring

In order to meet the latency requirement, the SAR processing workflow has been tailored accordingly:

• Instead of estimating ADC correction parameters from the SAR raw data itself, the SAR data is corrected by onground characterized values, uploaded to the spacecraft, updatable throughout the mission.

• Doppler centroid estimates are not derived from the SAR signal. Instead, on-board attitude quaternions and orbit state vectors provide the input to compute the Doppler centroid based on the viewing geometry.

• Instead of on-board chirp replica generation, replicas are modelled on-ground and uploaded to the spacecraft in advance.

• The SAR image geometry stays with the slant range acquisition geometry. The slight pixel distortion is annotated and accounted for in follow-on image processing.

• Multilooking is achieved by simple 2x2 pixel box car averaging.

#### 3.2 Hardware and workflow mapping

The target hardware for this project is a Xilinx Zynq Ultrascale+ ZU19EG MPSoC device. **Figure 1** shows an overview over the MPSoC and the processing steps. The processing system quad-core ARM Cortex-A53 is in charge of computing focusing, calibration, and annotation parameters based on auxiliary data as they are attitude, orbit, and instrument settings on input.

The programmable logic performs all SAR signal processing such as raw data correction, pulse compression, FFTs, pixel-wise complex filter multiplications, detection, and multi-looking. Since only a small amount of SAR data can be stored in the Random-Access Memory (RAM) created using Ultra RAM and Block RAM of the FPGA at a time, a 4 GiB DDR4-SDRAM is attached to the PL to provide a storage for the entire SAR raw data, for an intermediate working buffer and the final SAR image. Beside of this, auxiliary data is stored in the DDR4-SDRAM as well. **Table 1** shows the resource utilization of the SAR level 1 processing step.

Communication and data interfaces to the scheduling and storage board are realized with application-specific TCP/IP messages over Gigabit Ethernet and PCI Express transfers over QSFP+ connections, respectively.

**Table 1:** Resource usage of SAR L1 and L2 on-boardprocessing compared to available resources on theZU19EG FPGA.

	L1	L2	Total	Available	Utilization
FF	71K	28K	99K	1045K	9.5 %
LUT	61K	41K	102K	523K	19.5 %
BRAM	669	145	814	984	82.7 %
URAM	128	0	128	128	100.0 %
DSP	317	260	577	1968	29.3 %

# **3.3** SAR focusing algorithm and implementation

EO-ALERT ship detection and sea state analysis utilize 3 m resolution X-band strip map mode data. For this resolution class, imaging mode, image extent and application the focusing accuracy of the monochromatic  $\omega$ -k algorithm [2] is sufficient. This algorithm is faster and less complex than the range-Doppler algorithm (RDA), which requires interpolation, or the chirp scaling algorithm (CSA) [3].

#### 3.3.1 Parameter calculation

Once the raw and auxiliary data are in DDR4-SDRAM, the auxiliary data is transferred to the PS to compute the focusing, calibration, and annotation parameters. Because of the low computational complexity of the auxiliary data processing and need of floating-point operations, these are performed more adequately on the microprocessor and then the results are transferred to the PL. In contrast, the processing of huge amount of raw data takes full advantage of the highly parallelised processing on the FPGA.

#### 3.3.2 Fast Fourier Transform

The building blocks of signal processing are Xilinx IP cores for FFT, RAM, CORDIC, and floating-point multiplication. The Xilinx LogiCORE FFT core implements the Cooley-Tukey FFT algorithm, a computationally efficient method for calculating the Discrete Fourier Transform (DFT). The FFT IP core is runtime configurable for both forward and inverse FFT up to 32k points in a pipelined streaming I/O architecture, which reads and writes 16bit/16bit complex integer data. Internal computations

are configured as fixed point with adaptive block floatingpoint scaling to avoid clipping or wrapping of integer values.

The current design foresees 4 parallel FFT instantiations since the FFT core uses RAMs for the data and phasefactor storage; thus, availability of RAM on the deployed FPGA limits the number of parallel FFT instances in the implementation. Phases of the focusing filters are computed in double-precision floating-point format based on the intermediate results stored in PL written by PS. Here, PL utilizes the dedicated floating-point multiplication IP cores.

Prior to filtering of the integer format SAR data, floatingpoint filter phases are wrapped to baseband and the CORDIC IP core computes the complex exponential with 16bit/16bit integer output format. Integer multiplication of complex filter samples and complex SAR data is performed in the logic as well.

#### 3.3.3 Memory and Cache usage

Considering that read and write operations to the SDRAM are burst-oriented, chunks of the raw data in the SDRAM are read and written to RAM which is acting like a cache on the FPGA. Once the data is in the cache, it goes through the signal processing chain of the algorithm, which consists of pre-processing, FFT and postprocessing stages. After post-processing, the results are stored in cache for efficient corner turning of the data and transferred back to the SDRAM.

A major challenge of the implementation is efficient design of the interface between the SDRAM and cache inside the logic, which is needed for corner turning in order to enable burst mode read and write of SDRAM memory pages. Processing of an 8k x 32k raw data block involves repetitive transfers of all data between SDRAM and cache summing-up to approximately 7 GiB.

#### 3.3.4 Achieved runtime

Processing of the raw data block with four FFTs running simultaneously requires approximately 352 MCycles per FFT. This makes up the latency-critical part of the L1 signal processing chain. With the programmable logic running at a moderate clock rate of 125 MHz, the total latency for L1 processing of the scene shown in **Figure 2** (8k x 24k raw data size) including the data transfer time, FFT latency and PS processing time is 4.1 seconds. **Figure 3** compares an image detail produced by the MPSoC to the same detail in the original TerraSAR-X product.



**Figure 2:** SAR image of Napoli (12 km x 32 km) generated by the EO-ALERT MPSoC implementation.



**Figure 3:** Comparison of an image detail generated by the EO-ALERT MPSoC (left) and the same detail extracted from an original TerraSAR-X image product (right).

## 4 L2 product generation

The products generated within the EO-ALERT project either provide ship detection or extreme weather detection information. In the following, the processing chains for both types of products will be explained.

#### 4.1 Ship detection

The ship detection algorithm is adapted from [4] and [5] and consists of several steps described in the following sections and shown in **Figure 4**.



**Figure 4:** Steps of the ship detection algorithm. (a) Original image. TerraSAR-X data © DLR 2016 (b) Ship candidates after initial detection (c) Land mask (d) Filtered ship candidates.

#### 4.1.1 Initial detection

During the detection phase, ship candidates are identified on the entire image. While a fast approach of implementing this is using a constant pixel intensity threshold over the entire image, this approach is prone to errors and does not yield the desired low false alarm rate. Here, a CFAR (Constant False Alarm Rate) approach is applied, which is, however, computationally very expensive as the intensity of each pixel has to be compared to the mean intensity of its surrounding area. Connected clusters of detected pixels are merged to ship candidates.

Initial software unit testing on the ARM processing system showed that for the generated SAR image shown in Figure 2 with 8304 x 3472 pixels, this initial detection step alone would take 12.7 minutes, exceeding the target system latency of 210 seconds by far. Hence, the algorithm needs to be implemented in hardware to exploit the capabilities of parallel processing offered by applicationspecific implementation on an FPGA.

The pipelined data path presented in [6] has been customized for ship detection and adapted to the target hardware. **Table 1** shows the resource utilization of the SAR level 2 processing step. Significant speed-up has been achieved by ensuring DDR4-SDRAM address contiguity as well as optimizing memory read burst size. With the dedicated hardware core in the FPGA running at 125 MHz, the initial detection step takes just 2.7 seconds.

#### 4.1.2 Refinement

This first detection is followed by a refinement step where CFAR is applied again only around the centre of ship candidates. Previously detected pixels are ignored to find surrounding ships that were missed in the initial detection. This iterative concept was initially used for iceberg detection [7]. Computationally, this step much less expensive than the initial detection as the CFAR computation is only applied to few limited areas. Therefore, it can be executed in software without major latency impact.

#### 4.1.3 Azimuth ambiguities removal

The aliasing of the Doppler phase history during SAR processing generates echoes of bright areas in dark areas, shifted in azimuth by a distance depending on the satellite system. This can cause multiple echoes of a single real ship to appear in an image. Therefore, this step compares detected objects along the same range separated by the ambiguity distance and removes all but the brightest one. As bright objects on land can also cause echoes on the sea, detected land objects cannot be removed before this step.

#### 4.1.4 Land object removal

Whenever land is present in a scene, the detection step may find ship candidates on land, such as buildings, power poles or any similar objects with high radar backscatter. To filter these unwanted detections, a database containing polygonal coastlines is used and ship candidates on or very close to land, as indicated by this database, are removed. While on a ground station, the best available database is used, this may include data retrieved online from OpenStreetMap coastlines. Another choice is SRTM SWBD, which has data with 30 m resolution, but does not cover polar areas and has a file size of over 3 GB. For the onboard system, GSHHS (Fine) is used. It offers 40 m resolution, global coverage and a comparatively low file size of about 200 MB. A safety distance of 600 m is used to account for inaccuracies in the coastline as well as large harbour structures not included in the coastline dataset; ship candidates closer to land than this distance will be filtered out.

#### 4.1.5 Ship parameter extraction

The parameters extraction block calculates the properties of each detected ship with respect to its position, size and heading by analyzing the geometrical features of the vessel on the resulting ship-mask. Additionally, the probability of detection is estimated, using assumptions about the expected probability density function of the intensity in the surrounding of ships (in our case: in open water) as described in [8].

#### 4.1.6 Achieved runtime

With the initial detection step implemented in the PL and all other steps running in software on the ARM PS as described, processing of the generated test image shown in Figure 2 took 36.0 seconds. A 20% margin has been applied to account for image-to-image variance in computation time.

#### 4.2 Extreme weather detection

The extreme weather detection consists of detection of wind speed as shown in [9] and detection of wave height as shown in [10]. For processing, the scene is split into a grid with a cell size of about  $2x2 \text{ km}^2$ , each yielding one data point for wind and waves, respectively. Both steps are conducted similarly to the operational implementation of this algorithm running in the DLR ground station Neustrelitz.

#### 4.2.1 Land grid point removal

For the subsequent detection steps, it is necessary that land areas in the scene are filtered out since these have a substantial impact on the accuracy of detection. The same database as in 4.1.2 is used and applied to the image and all land pixels are denoted as such.

#### 4.2.2 Wind detection

The Wind Detection step uses the XMOD2 geophysical model function (GMF) described in [9], which calculates the wind speed from the measured sea surface backscatter. The function is applicable to X-band SAR data at both (HH and VV) polarizations for incidence angles between  $20^{\circ}$  and  $45^{\circ}$  and is able to give the sea surface wind speed in the range of 2–20 m/s. The algorithm requires input of the wind direction either from an external source such as wind forecast models or as a fixed input parameter valid for the entire scene, which is less accurate. Both ways will be possible in the implementation and adding spatially resolved wind direction data does not increase the runtime, as long as the data is directly accessible.

#### 4.2.3 Wave detection

The Extreme Ocean Wave Detector uses the empirical XWAVE algorithm presented in [10]. After calculation of

the image spectrum via FFT, spectral analyses of the grid cells are performed. Due to the influence of wind on sea state, the previously created wind-field is given as additional input. XWAVE derives the wave height directly from the SAR image spectrum rather than the wave spectrum and is tuned for TerraSAR-X input products.

#### 4.2.4 Achieved runtime

All steps of the extreme weather detection will be implemented in software exclusively. Outputs only used for further analysis and their calculations were removed from the existing code and the processing time was optimized in several calculation steps without changing the results. In particular, the FFTs mentioned in the previous section have been sped-up by utilizing the NEON SIMD instructions of the ARMv8 processor. Execution with the test image in Figure 2 took 38.8 seconds, again including a 20% margin.

#### 4.3 Alerts

The outcomes of both detections are so-called alerts. These are small packets of data designed to contain only the required information for a warning: coordinates, time, properties of the detection (e.g. wind speed, ship dimensions) and error information, as well as scene and alert identification. The ship alert also contains a 100x100 pixel thumbnail image of the ship. These alerts have a low data volume and can be directly transferred to users on the ground if satellite and user are equipped with the appropriate transceivers.

## 5 Conclusions

Transferring the SAR image generation and further processing from ground to a satellite is a challenging task. With the capabilities of an FPGA increasing the performance of suitable parts of the algorithm, the computation can become fast enough for a significant reduction in total delivery time, allowing for improved Near Real Time service. With SAR, this information can be provided regardless of illumination or cloud coverage. Compromises being made w.r.t quality result only in minimal effects on the final application products, which are still of high quality and provide valuable low-latency information to the end user.

While these investigations were carried out using TerraSAR-X StripMap data as baseline, the results are applicable to all types of SAR data. The limits of SAR image size depend on the specifically used FPGA and CPU for resource usage and latency requirements. When planning a new satellite mission with on-board processing, these need to match the goals of the satellite mission regarding area coverage, resolution, acquisition duration and latency.

## 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 DEI-MOS Space. More information on the project at eo-alerth2020.eu.

## 6 Literature

- Kerr, M., Tonetti, S., Cornara, S., Bravo, J.-I., Hinz, R., Latorre, A., Membibre, F., Solimini, C., Wiehle, S., Breit, H., Tings., B., Koudelka, O., Teschl, F., Magli, E., Bianchi, T., Migliorati, A., Motto Ros, P., Caon, M., Freddi, R., Benetti, M., Milani, F., Curci, G., Fraile, S., Garcia, L., Marcos, C., Fiengo, A. (2020) EO-ALERT: A Novel Architecture for the Next Generation of Earth Observation Satellites Supporting Rapid Civil Alert, 71st International Astronautical Congress (IAC 2020), 12.-14. Oct. 2020
- [2] Raney, R. K., Vachon, P. W. (1989) A phase preserving SAR processor, Proc. IGARSS '89, pp. 2588-2591
- [3] Bamler, R (1992) A Comparison of Range-Doppler and Wavenumber Domain SAR Focusing Algorithms, IEEE Trans. On Geoscience and Remote Sensing, 30 (4), 706-713
- [4] Brusch, S., Lehner, S., Fritz, T., Soccorsi, M., Soloviev, A., Van Schie, B. (2011) Ship Surveillance with TerraSAR-X. IEEE Trans Geosci Remote Sens 49, 1092–1103. doi:10.1109/TGRS.2010.2071879
- [5] Tings, B., Bentes da Silva, C. A., Lehner, S. (2016) Dynamically adapted ship parameter estimation using TerraSAR-X images, Int J Remote Sens, 37:9, 1990-2015, DOI: 10.1080/01431161.2015.1071898
- [6] Gregorek, D., Günzel, D., Rust, J., Paul, S., Velotto, D., Imber, J., Tings, B., Frost, A. (2019) FPGA Prototyping of a High-Resolution TerraSAR-X Image Processor for Iceberg Detection, Asilomar Conference on Signals, Systems, and Computers, 53, 1832-1835, DOI:10.1109/IEEECONF44664.2019.9048985
- [7] Frost, A. Ressel, R., Lehner, S. (2016) Automated iceberg detection using high-resolution X-band SAR images, Can J Remote Sens 42.4, 354-366
- [8] Frost, A., Ressel, R., Lehner, S. (2015) Iceberg detection over northern latitudes using high resolution TerrSAR-X images, 36 Can Symp Rem Sens Abstr
- [9] Li, X.-M., Lehner, S. (2014) Algorithm for sea surface wind retrieval from TerraSAR-X and TanDEM-X data, IEEE Trans Geosci Remote Sens., 52, 5, 2930-2941
- [10] Pleskachevsky, A.L., Rosenthal, W., Lehner, S. (2016) Meteo-Marine Parameters for Highly Variable Environment in Coastal Regions from Satellite Radar Images. ISPRS J. Photogr. Remote Sens., 119, 464–484