

# Satellite on-board processing of synthetic aperture radar data for rapid delivery of latency sensitive maritime information products

Dominik Günzel<sup>\*a</sup> and Srikanth Mandapati<sup>b</sup>

<sup>a</sup>Team SAR Oceanography, Remote Sensing Technology Institute, German Aerospace Center (DLR), Am Fallturm 9, 28359 Bremen, Germany

<sup>b</sup>Team SAR Processing, Remote Sensing Technology Institute, German Aerospace Center (DLR), Münchener Straße 20, 82234 Weßling, Germany

## ABSTRACT

Remote sensing satellites generate a large amount of data that is conventionally down-linked to ground stations for further processing, but the time from acquisition until the satellite makes contact to the nearest ground station incurs additional latency, ultimately increasing the delay before the desired processing results are available to stakeholders. For certain applications though, latency is critical as the contained information decays rapidly. For example, the location of a moving ship on satellite imagery becomes outdated after a short time. Moving data processing from the ground to on-board of the satellite has the potential to significantly speed up product delivery. This contribution introduces developments regarding satellite on-board processing of synthetic aperture radar (SAR) data, addressing products concerning maritime safety and security in particular. Focusing of raw SAR data into interpretable images is implemented on specialized hardware to keep within the constraints on-board of satellites. The focused SAR images are processed further using conventional algorithms and artificial intelligence to extract desirable information, such as ship location, ocean wave height, and sea ice classification. Compared to the raw SAR echoes or focused image, the data size of derived products is so small that low data-rate transfers via satellite data relays become a possible path of delivering the information directly to the users, further facilitating rapid delivery. The on-ground prototype developed in this work intends to serve as a first step towards high-resolution satellite on-board data processing in near-realtime by using hardware representative for use in low Earth orbit.

**Keywords:** synthetic aperture radar, SAR, image processing, on-board processing, maritime surveillance, ship detection, sea ice, significant wave height

## 1. INTRODUCTION

The positions of moving targets, such as ships, on satellite imagery deprecate after a short time. Furthermore, the state of the sea surface itself may change rapidly depending on meteorological conditions, currents and other influences. In order to enable novel time sensitive applications of remote sensing, frequent observations of target areas are therefore desired. For this reason, the number of Earth observation satellites in orbit continuously increases, with a particular focus on constellations of comparatively smaller satellites.<sup>1</sup> This includes satellites carrying synthetic aperture radar (SAR) instruments, which is a technology commonly employed in remote sensing due to its ability to penetrate rain clouds and, as an active system, independence of natural illumination.

Frequent acquisitions alone are, however, not sufficient for approaching remote surveillance in near-realtime as users also need fast access to the gathered information. Today, in the typical SAR satellite remote sensing data chain, the SAR echoes received during acquisition are digitized and stored on-board of the satellite until next contact to the nearest ground station supporting the mission along the flight path. Only then the large amount of data is downlinked to Earth using high-bandwidth links. This incurs latency between time of acquisition and availability of the data on the ground, especially for acquisitions over the oceans where the nearest ground station

---

\*Email: [dominik.guenzel@dlr.de](mailto:dominik.guenzel@dlr.de)

may be further away. After the so-called Level-0 raw data has been downlinked, sophisticated algorithms are applied to generate a focused Level-1 SAR image. The SAR images can be manually examined, but are usually processed further to automatically extract derived information delivered as Level-2 products to stakeholders. Overall, the latency from time of acquisition to delivery of the Level-2 product to the end user ranges from multiple minutes to several hours depending on the actual system design and level of access.

To enable timely delivery of latency sensitive products approaches outside of the traditional ground station paradigm may be explored. By moving the data processing on-board the satellite processing could start right after acquisition, and the data volume of derived Level-2 products is generally significantly smaller compared to the Level-0 data, so that alternative downlinking concepts become an option. For example, the position of a detected ship can be stored in data packet so small that it could be transferred even via low-bandwidth inter-satellite data relays directly to the user on ground.

In the following, the concept of on-board SAR processing for rapid delivery of latency sensitive information products in the context of maritime situation awareness will be explored. Section 2 discusses the type of hardware that could be employed in such systems and introduces the SAR Level-1 and Level-2 on-board processing prototypes that have been developed in the frame of the H2020 projects EO-ALERT<sup>2</sup> and dAIEDGE\*. Afterwards, the implementation and results of L1 image focusing and derived L2 products are detailed in Sec. 3 and 4, respectively.

## 2. ON-BOARD HARDWARE

State-of-the-art algorithms for SAR Level-1 image focusing and Level-2 product generation have high computational complexity, so that servers with general-purpose x86-architecture many-core central processing units (CPUs) are usually employed. This type of hardware is versatile as it supports a wide range of instructions and is therefore suitable for different applications. With the popularity of artificial intelligence (AI) in recent years, graphics processing units (GPUs) are now widespread as well due to their advantages in workloads such as training and inference of convolutional neural networks (CNNs). On satellites however, such powerful hardware is unusable due to energy, cooling and further constraints, so that the approach of downlinking the acquired Level-0 raw data before further processing outlined in the previous section has been well-established.

Because of the susceptibility of electronic components to damage from ionizing radiation even in low Earth orbit (LEO), where SAR satellites such as TerraSAR-X or Sentinel-1 reside, on-board processing systems must be hardened to withstand the environmental conditions for the duration of the mission. The development and validation of radiation hardened on-board processors is a complex and time-intensive task, so that the use of already existing designs is encouraged. Combined with the long planning and development cycle of satellites, processing capabilities on-board of satellites can be summarized as substantially more limited compared to what is available on ground. Nevertheless, by utilizing specialized hardware with higher efficiency and hardware-accelerated algorithms an acceptable tradeoff between hardware requirements, processing latency and product quality may be achieved.

Within the EO-ALERT project, a prototype on-board SAR processing chain was developed, including image focusing, ship detection, wind speed and wave height calculation.<sup>3</sup> The prototype is implemented on a base board equipped with an AMD Zynq UltraScale+ ZU19EG multi-processor system on a ship (MPSoC), which is a device combining a field programmable gate array (FPGA) with an ARM Cortex-A53 quad-core CPU. An FPGA is a type of integrated circuit that can be repeatedly configured to implement digital functions. It offers potentially high performance for workloads that benefit from a large degree of algorithm parallelization at an energy-efficiency that is usually better than that of a general-purpose CPU but worse than an application-specific integrated circuit (ASIC). The FPGA and CPU are interconnected on the MPSoC so that the FPGA can be used to accelerate algorithms with particularly high computational complexity while the CPU handles less demanding tasks. This facilitates faster implementation as developing a hardware-accelerated function on the FPGA is more labor-intensive compared to the CPU, which uses standard software development techniques. As the project did not include building an actual satellite, commercial off-the-shelf (COTS) hardware was used and a requirement placed on transferability of the design to hardware actually qualified for use in LEO. For example, the design

---

\*<https://daiedge.eu/>



Figure 1. TerraSAR-X image acquired over the Saronic Gulf focused on MPSoC including coastal areas and bright scatterers on sea. The scene covers 33.5 km in range and 13.1 km in azimuth direction. TerraSAR-X data © DLR 2016.

could be ported to the radiation hardened ThalesAlenia multiMIND<sup>4</sup> or KP Labs Leopard<sup>5</sup> on-board processing boards for LEO satellites, which feature MPSoCs from the same Zynq UltraScale+ family.

The SAR on-board system in EO-ALERT uses conventional algorithms, but many applications in Earth observation nowadays favor AI to extract desired information, with a focus on convolutional neural networks (CNNs) for image segmentation. Therefore, in the frame of the dAIEDGE project a CNN-based sea ice classifier was developed for the same MPSoC to show the feasibility of creating ice charts directly on-board of a SAR satellite.

### 3. IMAGE FOCUSING

The focusing of SAR images requires reconstructing the complex reflectivity of a ground surface acquired by a sensor. A prevalent method for achieving this is through the matched filter approach, where the raw data is convolved with the complex conjugate and time-inverted range-variant point scatter response. Most SAR focusing algorithms efficiently execute this convolution in the spectral domain. Notable spectral-domain SAR focusing techniques include the range-Doppler algorithm (RDA), chirp scaling algorithm (CSA), and omega-K algorithm ( $\omega$ KA).

#### 3.1 TerraSAR-X StripMap

In the EO-ALERT project, monochromatic  $\omega$ KA described in 6 is selected due to its effective balance of sufficient image quality for the target application and level of computational complexity that is feasible for on-board implementation on the MPSoC introduced in Sec. 2. A more detailed breakdown of the reasoning for algorithm selection is published in 3. The selected imaging mode is StripMap (SM), which offers a good tradeoff between coverage and resolution for the intended derived L2 products, i.e. ship detection and wind speed & wave height (see Sec. 4.1 and 4.2). As no new SAR instrument is developed in this project, TerraSAR-X single polarization SM acquisition data acts as a representative source of L0 raw data. The TerraSAR-X SM mode features a swath width of 30 km and a typical resolution of about 3 m.

The block size for on-board SM processing is configured to handle exactly 8,192 azimuth lines and a maximum of 32,768 range pixels. Depending on the acquisition parameters, each block covers an area between 375 km<sup>2</sup> and 500 km<sup>2</sup>. The output of the SAR processing yields images with a resolution of approximately 6 to 8 meters, presented as 2 by 2 multi-looked, slant range detected (MSD)  $\sigma_0$ -calibrated images. Thus, the radiometric resolution is enhanced at the expense of spatial resolution. The detection and multi-looked processes decrease the data size of the focused image by a factor of 8 which is an advantage for onboard storage and faster downlinking. If preferred, these steps can be skipped by adjusting the processor, allowing for generation of single look complex (SLC) images.

Figure 1 shows an image focused on the MPSoC from data acquired by TerraSAR-X over the Saronic Gulf including the Greek coast, islands and bright scatterers on sea. The generation of MSD images like the one shown takes just 3.7 seconds on the MPSoC, and 4.3 seconds for the SLC mode. This is achieved by implementing a dedicated hardware accelerator for the SAR signal processing on the FPGA, utilizing a pipelined data path with a high degree of parallelism.<sup>3</sup> Before the focusing begins, necessary geometric parameters and SAR signal filter coefficients are first calculated on the ARM-based CPU and transferred to the FPGA, where then the comparatively much more computationally demanding SAR signal processing of the raw data is executed. The data pipeline consists of five major steps, which are raw data in-phase and quadrature-phase (I/Q) correction, forward and backward fast Fourier transforms (FFTs) in both range and azimuth, pixel-wise complex multiplications of focusing filter kernels computed from lookup tables, detection, and, finally, multi-looking. The partitioning of the algorithm into demanding steps on the FPGA and lighter calculations on the CPU is an important task at the beginning of the design phase to ensure matching performance goals while minimizing development time.

To assess the geo-location accuracy of the focused images, positions of a selection of ships, i.e. bright scatterers, have been compared between the on-board generated images and the original TerraSAR-X L1 products. On average, the difference between the measured coordinates from the TerraSAR-X and EO-ALERT images is in the decimeter range concerning both azimuth and range, and the standard deviation is less than two meters. Therefore, the localization accuracy exceeds the requirement of ten meters set in the project.

### 3.2 Sentinel-1 TOPSAR

For the dAIEDGE project, the Sentinel-1 sensor was selected primarily because both raw L0 data and L1 images are freely available. This has led to developments of many different applications that benefit from having large datasets, such as sea ice classification using AI (see Sec. 4.3). The Sentinel-1 L0 data products are compressed using flexible dynamic block adaptive quantization (FDBAQ) to reduce the instrument downlink data rate. For the data to be usable, it needs to be decompressed first, which is done using software as part of the preparatory work.

Considering the intended application of sea ice classification, the terrain observation with progressive scans SAR (TOPSAR) mode is preferred due to its larger coverage area compared to StripMap. The algorithm selected for processing the TOPSAR mode is chirp scaling due to its superior geometric accuracy compared to monochromatic  $\omega$ KA.<sup>7</sup> While full  $\omega$ KA would also offer good accuracy, it necessitates Stolt interpolation, which can be challenging to implement in hardware. The processor, which is currently in development, is able to focus both interferometric wide (IW) and extra wide (EW) swath modes with swath widths of 250 km and 400 km, respectively. This is achieved by merging three sub-swaths in case of IW and five sub-swaths for EW.

The hardware processor, which again targets the MPSoC introduced in Sec. 2, is configured to handle block sizes of exactly 8,192 azimuth lines for either mode and a maximum of 32,768 and 8,192 range pixels for IW and EW, respectively. Since the same hardware as in the EO-ALERT project is being used for this project as well, the overall design strategy remains unchanged. Again the computation of parameters is implemented in software executed on the ARM processor. SAR signal processing is carried out on the FPGA component of the MPSoC to make use of the high degree of parallelization offered. As development is ongoing the implementation results as well as performance and latency metrics are still pending, but based on the experiences from the EO-ALERT implementation and theoretical calculations it is expected that the image focusing will achieve processing times of 8 seconds per sub-swath in IW mode and 2 seconds for each EW sub-swath with an uncertainty of  $\pm 20\%$ .

## 4. PRODUCT GENERATION

### 4.1 Ship Detection

The on-board ship detector developed in the scope of the EO-ALERT project is based on software that has been in operation in the Neustrelitz ground station of the German Aerospace Center for several years, where SAR images covering areas of interest are automatically processed. It has been validated on images from the TerraSAR-X and Sentinel-1 SAR missions. The basis of the detector is described in 8 and 9. It consists of several steps to achieve a high probability of detection and low false alarm rate. As shown in Fig. 2a ships and other human-made objects appear as clusters of bright pixels on SAR images due to higher reflectivity compared




Figure 2. Section of a SAR image acquired over the Saronic Gulf including islands and bright scatterers on sea. (a) Image generated by the EO-ALERT on-board prototype from synthetic TerraSAR-X raw data. (b) Ships detected by the on-board processor are framed in red. TerraSAR-X data © DLR 2016.

to the surroundings. Therefore, potential targets are identified by a cell-averaging constant false alarm rate (CFAR) algorithm. After initial screening of the whole image, the algorithm is applied again but only around the centers of previously found ship candidates. In this iteration already detected pixels are masked so that some ships that were missed previously are found. Next, so-called azimuth ambiguities are filtered out, which are echoes of bright scatterers sometimes appearing on an image a certain distance in flight direction away from the actual scatterer.<sup>10</sup> After the filtering a land mask is applied to keep only objects that are located on open water. Finally, the location of each ship is saved and parameters including length, width, and heading are derived from geometric features. Figure 2b shows ships detected by the on-board detector.

The development of the on-board ship detector has been explained in detail already in 3; therefore, only the results will be summarized here. When porting the ship detector to the MPSoC, it was identified that the initial application of the CFAR algorithm is by far the most computationally demanding step as the backscatter statistics of the surroundings have to be calculated for each pixel of the image, so that this step alone would take around 16 minutes for a sample image on the ARM-based CPU. In contrast, all other steps outlined above finish within seconds on the CPU. Accelerating the CFAR algorithm in hardware on the FPGA allowed for a massive speed-up, reducing the latency of this step to just 2.7 seconds for the same image. With the initial CFAR step on the FPGA and all other steps running on the ARM CPU, total processing time for one on-board SAR image is between 8.2 and 28.8 seconds (average 16.0 seconds) depending on image size and land coverage, proving the feasibility of performing ship detection directly on-board of a satellite by utilizing specialized hardware.

## 4.2 Wind Speed and Wave Height

The second scenario explored in the EO-ALERT project is computation of ocean wind speed and wave height to detect extreme weather events. One method to derive wind speed from the sea surface backscatter on SAR images is application of a geophysical model function (GMF). In this case, the XMOD2 GMF described in 11 is used, which was developed to retrieve wind fields from TerraSAR-X data. Wave height is derived using the linear regression method CWAVE\_EX published in 12.

For the on-board processing, the L1 SAR image is tiled into a grid with 2 km spacing in range and azimuth direction to achieve a good compromise between processing time and resolution of the product. Grid tiles containing more than 10% land are immediately discarded, and to the remaining tiles XMOD2 is applied to obtain the mean U10 wind speed, i.e. the wind speed at 10 m height above water. The wind speed is also an input for wave height calculation, so the steps are executed sequentially. Significant wave height ( $H_s$ ) is computed on the same processing grid. The visualization in Fig. 3 shows the output of the processing for a scene acquired at the coast of Naples, Italy.

During algorithm partitioning it was found that both algorithms can be executed on the ARM CPU of the MPSoC sufficiently fast after several code optimizations, so this scenario serves as an example on-board

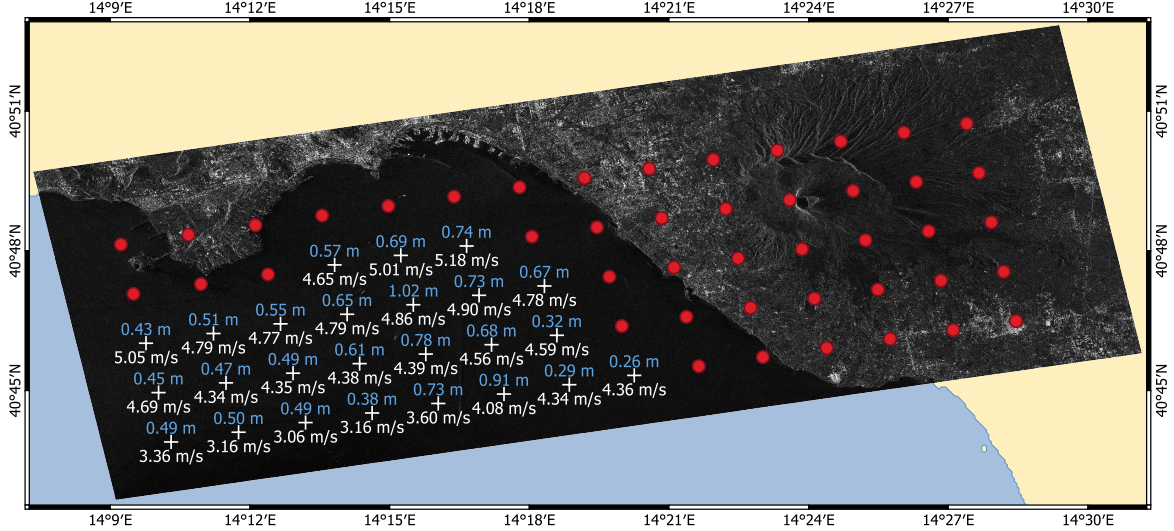


Figure 3. Wind and sea state processing for a scene around Naples, Italy. Red points show invalid grid cells on land, for which no calculation is performed. White crosses mark valid grid cells with the respective results for wind speed  $U_{10}$  (white) and mean wave height  $H_s$  (blue). TerraSAR-X data © DLR 2016. Source: 3 licensed under CC-BY-4.0

application that can be executed even on very light-weight hardware. The combined processing time for  $U_{10}$  and  $H_s$  is between 16.4 and 38.5 seconds with an average of 30.8 seconds for the scenes in the test set. As only grid cells with less than 10% land are processed, the latency scales linearly with the presence of land on the image. A detailed analysis of the results is again provided in 3, including a comparison of the accuracy of on-board  $U_{10}$  and  $H_s$  compared to using operational TerraSAR-X L1 data. In summary, it was found that the used algorithms are susceptible to even minor radiometric changes in the input data, so that they should be re-trained on a larger set of images focused with the on-board processor in a real deployment scenario.

### 4.3 Sea Ice Classification

To demonstrate the feasibility of on-board applications involving AI, sea ice classification has been chosen as it is a common use case of AI for SAR image processing. In the scope of the dAIEDGE project, a sea ice classifier based on the UNet CNN is first trained on ground to derive ice classes from Sentinel-1 L1 products. Inference of the trained model is implemented on the MPSoC described in 2 to enable rapid delivery of the generated ice charts. As the work on the classifier is an ongoing development this section discusses the current implementation and results, which are already in an advanced state but still subject to change.

#### 4.3.1 Dataset

In the AutoICE challenge, participants were tasked with developing deep learning models for retrieving sea ice parameters from a combination of SAR and other data sources.<sup>13</sup> For this the AI4Arctic sea ice challenge dataset has been assembled, which contains input data for training and test but also reference ice charts as training labels. The reference ice charts are manually labelled by analysts of the Canadian Ice Service (CIS) and the Danish Meteorological Institute (DMI). Therefore, they contain some uncertainties and are not ground truth in the literal sense. A comprehensive description of the dataset is provided in 14. The input data includes Sentinel-1 EW dual-polarization imagery, corresponding data from the advanced microwave scanning radiometer 2 (AMSR2) satellite sensor, and European Centre for Medium-Range Weather Forecasts (ECMWF) reanalysis v5 (ERA5) weather data.

In a satellite on-board context, the combined availability of these sources is less conceivable, so that only the Sentinel-1 data is considered in this work. Furthermore, the dAIEDGE project only foresees retrieval of ice stage of development (SOD), whereas the AI4Arctic dataset contains charts for sea ice concentration (SIC) and floe size (FLOE), too. In total, the dataset includes 533 scenes, 513 for training and 20 for testing the model. For now, this work uses the pre-processed *ready-to-train* version of the dataset to facilitate focusing on model

development. Thus, the Sentinel-1 L1 data has been noise-corrected and down-sampled from the original 40 m to 80 m pixel spacing, which results in image dimensions of around 5,000 by 5,000 pixels. The reference ice charts have been converted from polygons to pixel rasters matching the SAR geometry.

For model development, the dataset is split randomly by scenes into 80 % for training and 20 % for validation. Raster data channels of each scene are sliced into square 512 by 512 pixel patches to match the input and output size of the model, which will be specified in Sec. 4.3.2. Patches containing more than 10 % of pixels marked as invalid are discarded. Pixels are marked as invalid in the dataset if no reference ice chart data is available, for example because the area is on land. Discarding some patches after splitting scenes into training and validation dataset results in slight deviation from the intended 80/20 ratio, but this method is preferred over splitting after patch creation to avoid patches from the same scene being in both datasets, which could negatively impact model performance by leaking scene-specific weather and SAR effects. In total, 26,296 training and 6,785 validation patches are available.

### 4.3.2 Model Specification

Following the winning approach of the AutoICE challenge, a U-Net CNN for image segmentation is constructed using the PyTorch framework, although the inputs and outputs are reduced as explained in the previous section.<sup>13</sup> Furthermore, no additional down-scaling is applied to the data in order to preserve the resolution of the generated product. The network has three input channels; Sentinel-1 SAR HH polarization, Sentinel-1 SAR HV polarization, and the incidence angle of the SAR sensor for each pixel. Invalid pixels are marked with a special value so that they can be ignored when calculating metrics. The single output has seven channels, one for each SOD ice class plus one to denote invalid pixels. Finally, ArgMax is applied to find the dominant class for each pixel. The differentiated classes for SOD are open water, new ice, young ice, thin first-year ice (FYI), thick FYI, and old ice. A five stages deep VGG13 decoder with pre-trained ImageNet weights is used, which might change in future to optimize the balance between model complexity and performance for the target on-board application. The model in its current state has around 18.4 million trainable parameters.

### 4.3.3 Training and Validation

The model is trained using the 512 by 512 pixel size patches described in Sec. 4.3.1, with batch size 64 and over 100 epochs. Training loss is evaluated with a weighted cross-entropy loss function to account for the class imbalance in the labelled dataset. Likewise, the validation score is computed as a weighted F1 score. Training was carried out on a server equipped with an NVIDIA A100 Graphics Processing Unit (GPU) and converged after around 50 epochs with loss approaching zero and validation score of 87.3 %.

### 4.3.4 Test

To test the model, the scenes of the test dataset are sliced into 512 by 512 pixel size patches in memory at different offsets as explained in 15, so that for each scene multiple sets of patches are created. As a result, each pixel on the image is contained in multiple patches. The purpose of this process is to run inference on each scene multiple times using patches cut at different positions. Then, by merging the predicted classes for each pixel from different patches, applying weighting to pixels depending on their distance to the respective patch center, edge effects are mitigated so that the generated ice chart is smoothed. With this method, using two sets of patches per image, the trained model achieves F1 score of 80.1 % across the whole test set, which fulfills the project goal of 80.0 %. Compared to the results of the AutoICE challenge, the model would have been placed third with regard to SOD F1 score, which is noteworthy considering that only the SAR data and no other input data channels are used.<sup>13</sup>

### 4.3.5 On-board Integration

Integration of CNNs such as the developed U-Net on the AMD MPSoC is possible by utilizing the deep learning processor unit (DPU) for Zynq UltraScale+ MPSoC in combination with the Vitis AI development tools. The DPU is a soft core with its own instruction set that can be integrated in the FPGA and is controlled by software code running on the ARM CPU. To port the trained PyTorch model to the DPU, the model is first inspected to ensure that all operations in the model layers can be mapped to the DPU. As models for GPUs generally use floating-point arithmetics but the DPU supports only integers, the model has to be quantized. It is expected

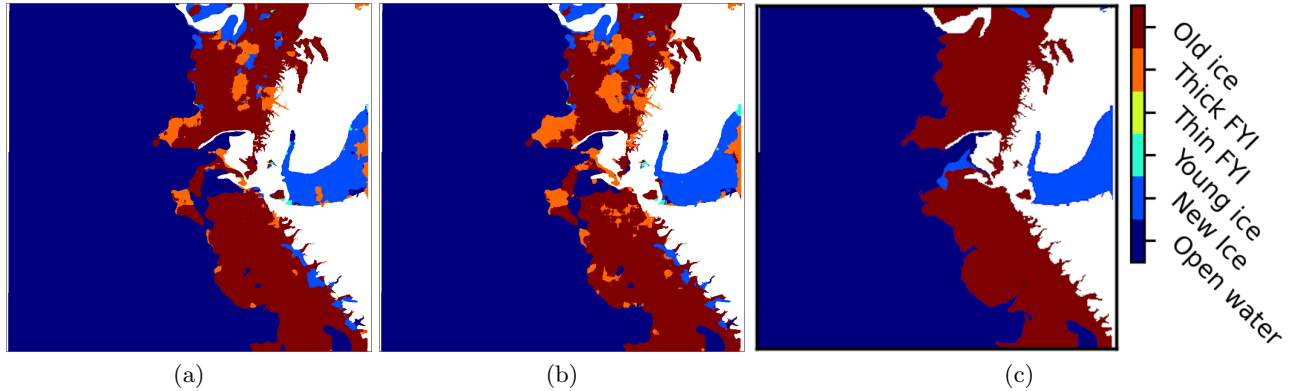


Figure 4. Sea ice SOD charts generated by the developed classifier for a Sentinel-1 EW scene acquired over the east coast of Greenland on 2020-10-13 at 08:04 UTC. (a) Chart generated by the on-ground floating point model on GPU. (b) Chart generated by the quantized model on MPSoC for on-board processing. (c) Reference chart based on manual labelling. White pixels are masked areas that cover land.

that quantization negatively impacts the accuracy of the model, but the degree of degradation depends on a multitude of factors and can be negligible in some cases. Indeed, in this case the validation F1 score of the quantized model incurs minor decrease by just 0.17 percentage points. Finally, the quantized model is compiled for the target DPU.

The software running on the ARM CPU is functionally similar to the code for testing the floating-point model except that inference is executed on the DPU. The processing steps include reading the input data, slicing into patches as explained in Sec. 4.3.4, programming the DPU and executing inference on it, merging the output patches again, applying ArgMax to find the dominant class, and saving the ice chart. Processing on the MPSoC takes between 26.6 and 56.8 seconds per scene (average 49.9 seconds) depending on image size, which is within the project target of one minute and demonstrates the feasibility of generating even high-resolution ice charts on-board of a satellite.

Despite quantization F1 score across the whole test set actually improves slightly from 80.1 % to 80.5 %, which may be explained by the image slicing and merging process quantitatively influencing the results. Figure 4 shows SOD ice charts generated for the same scene on GPU using the floating-point model (Fig. 4a) and on MPSoC using the quantized model (Fig. 4b). The reference chart based on manual labelling by DMI is also included. A detailed analysis of the results has not been carried out yet and will be published at a later date. In general however, both models are able to distinguish ice and open water well, but have trouble differentiating younger ice classes in particular.

## 5. CONCLUSIONS

This contribution demonstrated that data processing on-board of satellites to derive latency sensitive information from SAR is possible, which was illustrated based on three different applications. The developments target COTS hardware that is representative of currently available radiation-hardened devices. By utilizing specialized hardware, in this case an MPSoC with an FPGA, SAR images could be focused in less than ten seconds with minimal compromises to quality and resolution under the constraints on-board a satellite in LEO. For example, focusing an image covering 439 km<sup>2</sup> from TerraSAR-X raw data took just 3.7 seconds. This was followed by ship or extreme weather detection, which took on average 16.0 or 30.8 seconds, respectively. Even in the longest case, the combined processing time of L1 SAR focusing and L2 product generation was below one minute in these two scenarios.

Another common use case, sea ice classification on Sentinel-1 imagery, is currently being developed and showing promising results already. It is expected that an EW sub-swath will be focused on the same hardware in around 2 seconds. A standard U-Net was trained on the AI4Arctic sea ice challenge dataset, achieving F1 score of 80.5 % on the test dataset that is competitive with the highest-ranking submissions of the AutoICE challenge

despite quantization for execution on the MPSoC. In future, performance with regard to score but also latency may be improved further by evaluating different encoders, including those intended for edge deployment on less powerful but more efficient hardware.

The prototypes developed in this work are a first step towards achieving near-realtime delivery of latency sensitive EO products from satellites. Users could obtain the desired information within minutes after the acquisition by breaking with the existing paradigm of down-linking first and then processing. Down-linking the full acquired data to process on ground using powerful hardware will certainly stay relevant for achieving maximum quality and long-term storage. Nevertheless, processing directly on-board may improve existing systems for rapid response and enable novel applications in the future.

## ACKNOWLEDGMENTS

This work has received funding from the European Union Horizon 2020 research and innovation programme, partly by the EO-ALERT project under grant agreement No 776311 and in part by the dAIEDGE Network of Excellence (NoE) under grant agreement No 101120726. The EO-ALERT project is coordinated by DEIMOS Space and more information is available at [www.eo-alert-h2020.eu](http://www.eo-alert-h2020.eu). The dAIEDGE NoE is coordinated by Deutsches Forschungszentrum für Künstliche Intelligenz GmbH (DFKI) and more information is available at [www.daiedge.eu](http://www.daiedge.eu). This work reflects only the views of the authors and the European Commission (EC) is not responsible for any use that may be made from the information contained in this work. TerraSAR-X data were provided by DLR under Announcement of Opportunity (AO) proposal OCE3625.

## REFERENCES

- [1] Cratere, A., Gagliardi, L., Sanca, G. A., Golmar, F., and Dell’Olio, F., “On-board computer for cubesats: State-of-the-art and future trends,” *IEEE Access* **12**, 99537–99569 (2024).
- [2] Kerr, M., Tonetti, S., Carnara, S., Bravo, J. I., Hinz, R., Latorre, A., Membibre, F., Ramos, A., Wiehle, S., Koudelka, O., Magli, E., Freddi, R., Fraile, S., and Marcos, C., “EO-Alert: A satellite architecture for detection and monitoring of extreme events in real time,” in *[2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS]*, 168–171 (2021).
- [3] Wiehle, S., Mandapati, S., Günzel, D., Breit, H., and Balss, U., “Synthetic aperture radar image formation and processing on an MPSoC,” *IEEE Transactions on Geoscience and Remote Sensing* **60**, 1–14 (2022).
- [4] Pawlitzki, A. and Steinmetz, F., “multiMIND – high performance processing system for robust newspace payloads,” in *[2nd European Workshop on On-Board Data Processing (OBDP2021)]*, (2021).
- [5] Nalepa, J., Kuligowski, P., Gumieła, M., Drobik, M., and Nowak, M., “Leopard: A new chapter in on-board deep learning-powered analysis of hyperspectral imagery,” in *[Proc. IAC-CyberSpace Ed.]*, 1–4 (2020).
- [6] Bamler, R., “A comparison of range-Doppler and wavenumber domain SAR focusing algorithms,” *IEEE Transactions on Geoscience and Remote Sensing* **30**(4), 706–713 (1992).
- [7] Raney, R., Runge, H., Bamler, R., Cumming, I., and Wong, F., “Precision SAR processing using chirp scaling,” *IEEE Transactions on Geoscience and Remote Sensing* **32**(4), 786–799 (1994).
- [8] Brusch, S., Lehner, S., Fritz, T., Soccorsi, M., Soloviev, A., and van Schie, B., “Ship surveillance with TerraSAR-X,” *IEEE Transactions on Geoscience and Remote Sensing* **49**(3), 1092–1103 (2011).
- [9] Tings, B., da Silva, C. A. B., and Lehner, S., “Dynamically adapted ship parameter estimation using TerraSAR-X images,” *International Journal of Remote Sensing* **37**(9), 1990–2015 (2016).
- [10] Changcheng Wang, Y. W. and Liao, M., “Removal of azimuth ambiguities and detection of a ship: using polarimetric airborne C-band SAR images,” *International Journal of Remote Sensing* **33**(10), 3197–3210 (2012).
- [11] Li, X.-M. and Lehner, S., “Algorithm for sea surface wind retrieval from TerraSAR-X and TanDEM-X data,” *IEEE Transactions on Geoscience and Remote Sensing* **52**(5), 2928–2939 (2014).
- [12] Pleskachevsky, A., Tings, B., Wiehle, S., Imber, J., and Jacobsen, S., “Multiparametric sea state fields from synthetic aperture radar for maritime situational awareness,” *Remote Sensing of Environment* **280**, 113200 (2022).

- [13] Stokholm, A., Buus-Hinkler, J., Wulf, T., Korosov, A., Saldo, R., Pedersen, L. T., Arthurs, D., Dragan, I., Modica, I., Pedro, J., Debien, A., Chen, X., Patel, M., Cantu, F. J. P., Turnes, J. N., Park, J., Xu, L., Scott, K. A., Clausi, D. A., Fang, Y., Jiang, M., Taleghanidoozdoozan, S., Brubacher, N. C., Soleymani, A., Gousseau, Z., Smaczny, M., Kowalski, P., Komorowski, J., Rijlaarsdam, D., van Rijn, J. N., Jakobsen, J., Rogers, M. S. J., Hughes, N., Zagon, T., Solberg, R., Longép e, N., and Kreiner, M. B., “The AutoICE challenge,” *The Cryosphere* **18**(8), 3471–3494 (2024).
- [14] Buus-Hinkler, J., Wulf, T., Stokholm, A. R., Korosov, A., Saldo, R., Pedersen, L. T., Arthurs, D., Solberg, R., Longép e, N., and Brandt Kreiner, M., “AI4Arctic sea ice challenge dataset.” <https://doi.org/10.11583/DTU.c.6244065.v2> (Nov 2022).
- [15] Murashkin, D. and Frost, A., “Arctic sea ice mapping using sentinel-1 sar scenes with a convolutional neural network,” in [*2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*], 5660–5663 (2021).