

EO-ALERT: A Novel Architecture for the Next Generation of Earth Observation Satellites Supporting Rapid Civil Alerts

Original

EO-ALERT: A Novel Architecture for the Next Generation of Earth Observation Satellites Supporting Rapid Civil Alerts / 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, Andrea; MOTTO ROS, Paolo; Caon, Michele; Freddi, R.; Benetti, M.; Milani, F.; Curci, G.; Fraile, S.; Garcia, L.; Marcos, C.; Fiengo, A.. - (2020). (71st International Astronautical Congress (IAC) Virtual event 12-14 October 2020).

Availability:

This version is available at: 11583/2858693 since: 2020-12-22T16:26:05Z

Publisher:

IAF

Published

DOI:

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

EO-ALERT: A Novel Architecture for the Next Generation of Earth Observation Satellites Supporting Rapid Civil Alerts

M. Kerr^{a*}, S. Tonetti^a, S. Cornara^a, J. I. Bravo^a, R. Hinz^a, A. Latorre^a, F. Membibre^a, C. Solimini^a, S. Wiehle^b, H. Breit^b, B. Tings^b, O. Koudelka^c, F. Teschl^c, E. Magli^d, T. Bianchi^d, Andrea Migliorati^d, Paolo Motto Ros^d, Michele Caon^d, R. Freddi^e, M. Benetti^e, F. Milani^e, G. Curci^e, S. Fraile^f, L. Garcia^f, C. Marcos^g, A. Fiengo^h

^a DEIMOS Space S.L.U., Tres Cantos – Madrid, Spain, Email:

[/murray.kerr](mailto:murray.kerr), stefania.tonetti, stefania.cornara, juan-ignacio.bravo, robert.hinz, antonio.latorre, francisco.membibre, chiara.solimini / [@deimos-space.com](mailto:deimos-space.com)

^b Deutsches Zentrum für Luft- und Raumfahrt e.V., Germany, Email:

[/stefan.wiehle](mailto:stefan.wiehle), helko.breit, bjoern.tings / dlr.de

^c Technische Universität Graz, Graz, Austria, Email: koudelka@tugraz.at, franz.teschl@tugraz.at

^d Politecnico di Torino, Turin, Italy, Email:

[/enrico.magli](mailto:enrico.magli), tiziano.bianchi, andrea.migliorati, paolo.mottoros, michele.caon / [@polito.it](mailto:polito.it)

^e OHB Italia Spa, Milan, Italy, Email:

[/rffreddi](mailto:rffreddi), mabenetti.ext, fmilani.ext, gcurci.ext / [@cgspace.it](mailto:cgspace.it)

^f Deimos Imaging S.L.U., Spain, Email: silvia.fraile@deimos-imaging.com, lucia.garcia@deimos-imaging.com

^g Agencia Estatal de Meteorología, Spain, Email: cmarcosm@aemet.es

^h Formerly at Space S.L.U., Tres Cantos – Madrid, Spain

* Corresponding Author

Abstract

Satellite Earth Observation (EO) data is ubiquitously used in many applications, providing basic services to society, such as environment monitoring, emergency management and civilian security. Due to the increasing request of EO products by the market, the classical EO data chain generates a severe bottleneck problem, further exacerbated in constellations. A huge amount of EO raw data generated on-board the satellite must be transferred to ground, slowing down the EO product availability, increasing latency, and hampering the growth of applications in accordance with the increased user demand.

This paper provides an overview of the results achieved by the EO-ALERT project (<http://eo-alert-h2020.eu/>), an H2020 European Union research activity led by DEIMOS Space. EO-ALERT proposes the definition and development of the next-generation EO data processing chain, based on a novel flight segment architecture that moves optimised key EO data processing elements from the ground segment to on-board the satellite, with the aim of delivering the EO products to the end user with very low latency (quasi-real-time). EO-ALERT achieves, globally, **latencies below five minutes** for EO products delivery, reaching **latencies below 1 minute** in some scenarios.

The proposed architecture solves the above challenges through a combination of innovations in the on-board elements of the data chain and the communications. Namely, the architecture introduces innovative technological solutions, including on-board reconfigurable data handling, on-board image generation and processing for the generation of alerts (EO products) using Artificial Intelligence (AI), on-board data compression and encryption using AI, high-speed on-board avionics, and reconfigurable high data rate communication links to ground, including a separate chain for alerts with minimum latency and global coverage.

The paper presents the proposed architecture, its performance and hardware, considering two different user scenarios; ship detection and extreme weather observation/nowcasting. The results show that, when implemented using COTS components and available communication links, the proposed architecture can deliver alerts to ground with **latency lower than five minutes, for both SAR and Optical missions**, demonstrating the **viability of the EO-ALERT concept and architecture**. The paper also discusses the implementation on an avionics test bench for testing the architecture with real EO data, with the aim of demonstrating that it can meet the requirements of the considered scenarios in terms of detection performance and provides technologies at a high TRL (4-5). When proven, this will open unprecedented opportunities for the exploitation of civil EO products, especially in latency sensitive scenarios, such as disaster management.

Keywords: Satellite Architecture, Earth Observation, On-Board Processing, AI, Real-Time, Low Latency

1. Introduction

Earth observation (EO) data delivered by remote sensing satellites provide a basic service to society, with great benefits to the civilian. The data is nowadays ubiquitously used throughout society for a range of diverse applications, such as environment and resource monitoring, emergency management and civilian security [1].

Over the past 50 years, the EO data chain that has been mastered involves the acquisition of sensor data on-board the satellite, its compression and storage on-board, and its transfer to ground by a variety of communication means, for later processing on ground and the generation of the downstream EO image products.

While the market is growing, the classical EO data chain generates a severe bottleneck problem, given the very large amount of EO raw data generated on-board the satellite that must be transferred to ground, slowing down the EO product availability, increasing latency, and hampering applications to grow in accordance with the increased user demand for EO products.

This paper provides an overview of the EO-ALERT project (<http://eo-alert-h2020.eu/>), an H2020 European Union research activity that addresses the challenge of a “high speed data chain” and the need for increased EO data chain throughput. EO-ALERT proposes the definition and development of the next-generation EO data and processing chain, based on a novel flight segment architecture that moves optimised key EO data processing elements from the ground segment to on-board the satellite. The objective is to deliver the EO products to the end user with very low latency for increased throughput [22].

Achieving this goal poses great challenges on the flight system, to be addressed through a combination of innovations in the on-board elements of the data chain and the communications link. As such, this goal necessitates **innovation in several critical technological areas**; namely on-board reconfigurable data handling, on-board image generation, on-board image processing, high-speed on-board avionics, on-board data compression and reconfigurable high data rate communication links to ground. Such innovations will also provide capabilities for the optimisation of the classical EO data chain towards a data chain with greatly improved data throughput.

The paper presents the project results on the above technological areas, with particular attention to the development of a flexible and reconfigurable data handling architecture integrating different on-board technologies for both Synthetic Aperture Radar (SAR) and optical sensor data, including image generation, image processing for rapid alerts, joint compression and encryption algorithms. The paper also discusses possible applications of the envisioned high-speed data

chain and introduces open issues regarding the implementation of the proposed architecture, from the use of dedicated hardware platforms to commercial off-the-shelf (COTS)-based implementations.

The aim is to raise interest in this ambitious project from platform developers, commercial EO mission operators, governmental agencies, and the on-board data compression and processing community. This should lead to the sharing of needs and views between the different stakeholders, on the technological challenges and solutions, and foster fruitful discussions towards new ideas and the expatiation of the proposed solutions.

2. EO-ALERT Overview

Data latency has become a key requirement in the EO market, since end-users require that data is available in a very short time interval. In case of polar platform satellites, what it is currently referred to as near real-time (NRT) consists of the provision of image products in the range of 1 to 3 hours; e.g. Sentinel-1 makes ocean products available within 1 hour of observation over NRT areas with a subscription [2].

Current market trends are moving beyond NRT applications, to applications with latencies in the order of 30 minutes to 15 minutes. Based on this, the latency performance concept behind EO-ALERT is to achieve latencies well below 15 minutes for the EO products. To be precise, the definition of latency can be taken to be: the time from the collection of the last photons through to the time that the data is converted to a specified EO product and delivered to the user portal. Based on this performance concept, and latency definition, EO-ALERT has a **goal latency of less than one minute** and requires a maximum latency of less than 5 minutes, for both SAR and optical image products.

This in turn flows down to requirements on the technology building blocks, which drive the EO-ALERT solution at all technology levels.

Satellites on-board data handling has traditionally been managed in a simple fashion, where raw data is collected, compressed and scheduled for transmission to the ground segment, where data processing takes place. This entails a significant delay in the generation of knowledge that can reduce the “value” of the acquired imagery. However, the recent availability of more powerful computational systems for onboard data handling allows for a redesign of the onboard processing chain. With modern systems, it is conceivable to generate data products directly on-board, reducing the time needed to generate products, from hours to a few minutes (Fig. 1). Achieving this ambitious objective requires the development of new concepts and novel technologies in order to 1) define a set of building blocks implementing the main processing functionalities, and 2) use these building blocks sensibly, defining an innovative reconfigurable

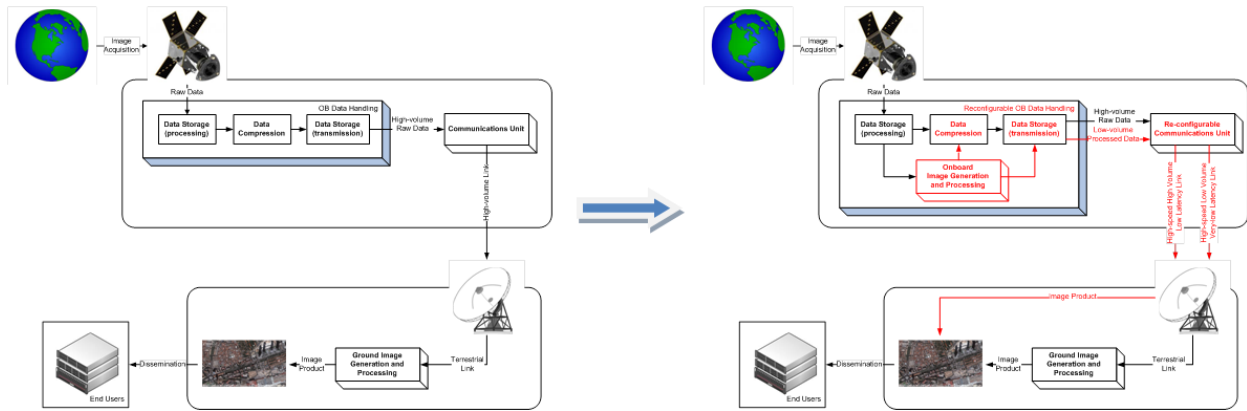


Fig. 1. EO-ALERT data chain moving ground processing and data management to the flight segment for very low latency EO image products: (left) classical data chain based on raw data compression and transfer, (right) innovative data chain showing its key elements and new data flows (in red)

onboard data handling architecture, that is able to achieve the desired goals, with reliable performance levels.

The proposed satellite processing chain has implications on several technological areas, including high speed avionics, Flight Segment/Ground Segment (FS/GS) communications, on-board compression and data handling and on-board image generation and processing. The project develops both the technologies and the data handling architecture, with an approach aimed at optimizing the use of onboard resources, as well as making the most out of the available image data through direct processing on the flight segment. As far as compression is concerned, recently developed image compression concepts aimed at improving image quality for a given compression ratio, while maximizing throughput, will be adopted for the first time in space applications.

Both the technologies and the integrated technologies chain will be **verified and experimentally evaluated** during the project on an avionics test-bench, using first relevant EO historical sensor data. They will then be demonstrated through relevant End-to-End (E2E) tests, in which EO data acquired specifically for EO-ALERT will be injected in the system and results will be obtained, evaluated and benchmarked.

Ground truth data for evaluation has been obtained through an experimentation campaign on a real and representative test-field. Namely, EO-ALERT identified two main scenarios for testing the potential of the proposed high- speed data chain: ship detection and extreme weather monitoring, that are described below.

2.1 Ship Detection Scenario

A deep analysis of the most common needs of maritime users has been performed in order to detect the specific ones [20]. The most important aspects are to have a ground resolution as small as possible, at least

1m, and low revisit time. But these requirements are not the only ones, in maritime there are some important details to be consider as key needs:

- *Flexibility*: in order to mitigate risks and choose the image acquisition target date.
- *Responsiveness*: in order to make last minute changes to the image request, one hour is the target.
- *Revisit*: less than 6 hours of revisit time between consecutive observations.

To increase the results, a combination of satellite technology has been selected. On one side the SAR satellites are the most suitable for ship detection, whereas very high-resolution optical imagery reduces false positives and allows ship identification. For the EO-ALERT project, Deimos-2 and TerraSAR-X have been selected for the optical imagery and radar data, respectively, since their characteristics are perfectly adjusted to the project requirements.

Bearing in mind the user needs requirements and with the main intention to develop the most suitable technology, the following tests are proposed in the Mediterranean Sea:

- *Ship Monitoring*: the objective is focused on illegal fishing and cargo monitoring between 10-20 meters. TerraSAR-X will be used for ship location and detection, and Deimos-2 for ship identification to confirm the SAR detection.
- *Coastal Monitoring*: the objective is focused on maritime safety and illegal migration. Deimos-2 will be forced to detect a small boat on an open sea area.

The main intention is to be able to send an alert, similar to the EMSA Vessel Detection Service (VDS), to an end user including the following parameters with a very short latency (below 5 minutes):

- Position Information.
- Movement Information.

- Ship details

And in a second step, with a more relaxed latency (below 30 minutes): ship identification and clipping ship.

2.2 Extreme Weather Scenario

For the extreme weather scenario, two types of detections have been considered: convective storms and wind speed.

2.2.1 Storm Detection

One of the main tasks of meteorological forecasters consists in issuing very early alerts of hazardous phenomena, like convective storms, to reduce their negative impact on the population and property. To do so, they use satellite products, mainly derived from Geostationary (GEO) platforms, to detect those phenomena. One of the key points for the use of those products in Nowcasting tasks is their availability with a very low latency. The EO-ALERT technology can be very helpful and beneficial for the generation of meteorological products for which high quality and low latency is required.

For these reasons, the EO-ALERT project has included an Extreme Weather Scenario where those issues are being handled [21].

Convective storms are going to be detected in all the convection stages using GEO satellite data, paying special attention to their monitoring, providing specific information on individual storm location, trajectory and characteristics. A four step algorithm that consists of candidate cell extraction, tracking, convective stage discrimination and alert generation, is being developed. OPERA [11] radar network data is being used as ground truth. Qualitative comparison of the first version of the EO-ALERT algorithm with the operational NWCSAF [12] RDT-CW (Rapidly Developing Thunderstorms – Convection Warning) shows compatible results for the candidate cell extraction and convective discrimination. A full quantitative comparison to NWCSAF RDT-CW and OPERA based convection detection algorithms will follow. Information on overshooting top detection will also be provided by the EO-ALERT algorithm.

2.2.2 Wind Speed

Wind speed and wave height constitute crucial information for vessels on sea, as underestimation of the weather situation endangers the entire ship and crew. The trail and development of a storm field on the open sea are often hard to accurately model in advance: e.g., the information available to ships might show a storm warning, but the indicated location of the storm centre might be several kilometres off its real location. This error can be prevented using satellite data to derive the desired information directly from the sea surface. With the developments in the EO-ALERT project, this

satellite-derived information can be directly provided to ship officers who can then find safer routes to circumvent storm.

3. EO-ALERT Architecture

3.1 Functional Architecture

To achieve the target latency in EO products delivery while maintaining the system suitable for several mission scenarios, the proposed architecture is designed to be modular, scalable and reconfigurable. The entire data-chain is divided into several functional blocks, each one implemented on dedicated software and/or hardware computing resources. Each function can be configured or changed with no or little impact on the others, and the available processing power can be assigned to each function based on the mission requirements. With this approach, the system can process different data types (e.g. optical and SAR data) from several sensors over a wide range of dataset sizes. The top-level architecture of the EO-ALERT system can be divided into three main physical blocks: the Sensor Board, the Processing Board and the Transmitter/Receiver (TX/RX) Subsystem.

The Sensor Board acquires raw image data (optical and/or SAR) from the sensor(s) and transfers it to the Processing Board using an acquisition memory buffer.

The Processing Board consists of several Multiprocessor System-on-Chip (MPSoC) devices that integrate a multi-core Processing System (PS) and a Programmable Logic (PL) part. Each device has access to a dedicated memory block (Random Access Memory) that is accessible to both the PS and the PL, and to an additional memory block dedicated to the PL. A storage device (e.g. a Solid-State Drive) is connected to the PS of one of the MPSoC devices. Bidirectional, point-to-point, high data rate links are used to transfer input and output data among the MPSoC devices, and a shared low data rate communication channel is used to exchange commands and telemetry information.

The Processing Board fetches the raw data from the acquisition buffer and implements all the data processing tasks required to produce the EO products (e.g. alerts and generated images). These tasks can be divided further into the following sub-modules:

- Optical image generation and processing.
- SAR image generation and processing.
- Central Processing Unit (CPU) Scheduling, Compression/Encryption and Data Handling (CS-CEDH).

Computationally expensive tasks, such as image processing or image compression/encryption, are partially or completely implemented in hardware. Depending on the workload, one or more MPSoC devices may be assigned to the same task, in a master-slave configuration. Less demanding operations

(execution flow control, data routing, simple data processing) are performed by the software running on the PS CPU cores of each device.

The TX/RX Subsystem receives data (EO products and alerts, raw data, ancillary data etc.) from the Processing Board, and forwards it to the ground station using either a very-low latency channel for critical data (i.e. alerts) or a high data rate channel for large data (e.g. generated images).

The following sections describe in detail the three sub-modules of the Processing Board (sections 3.3, 3.4 and 3.5) and the TX/RX subsystem (section 3.6).

3.2 Physical Architecture

The physical design of the Avionics is implemented as a hybrid solution, i.e. a solution that uses both Commercial off-the-shelf (COTS) and space qualified components [8] [9]. COTS are used in conjunction with mitigation techniques to increase robustness of the design against radiation effects, whereas space qualified components are used for the critical functions. This choice allows keeping weight, volume and cost of the Payload Data Processing Unit (PDPU) low with respect to an all space-grade design and it takes advantage of the state-of-the-art technology and processing power of the latest COTS components. This last point has proven to be fundamental to reach the programme’s latency goals.

Reliability in the Low Earth Orbit (LEO) environment is addressed through the employment of radiation mitigation techniques, redundancy, extensive telemetry collection and fault detection, isolation and recovery (FDIR) actions so that the solution is more robust than an all-COTS design.

The PDPU should weigh about 8Kg in a 10U volume and absorb 200W peak power during processing, although much smaller Size Weight and Power (SWaP) constraints can be met if needed. The architecture is based on the Compact Peripheral Component Interconnect (cPCI) Serial Space standard (Fig. 2) which guarantees a modular and scalable approach to on-board processing [10].




Fig. 2. CPCI Serial Space Architecture

This standard provides a backplane with two System Slots and up to seven peripheral boards. The whole

system is supervised by a shelf controller that can check the status of the boards and control their power supply.

The System slots are at the centre of a star connection with point-to-point high-speed links to the peripheral boards, which in turn are interconnected with a mesh of high-speed links.

In the EO-ALERT configuration (Fig. 3) System Slots are used for scheduling, compression, encryption and data handling tasks, while five peripheral slots are dedicated to processing functions. Each System Slot is connected directly, using the star links, to two “Master” processing boards, one for Optical (OPT) data processing and the other for Synthetic aperture radar (SAR) data processing. Master boards can offload computation to slave boards using mesh links.

Each peripheral board can be reconfigured from System Boards to handle OPT or SAR data so that the system can dynamically adapt to workloads and recover from failures; the two System Slots implement an intrinsic redundancy without the need of duplicating the entire PDPU: if the Main System Board or one of the Master processing boards fail the system can switch to the Redundant System Board which will reconfigure processing boards accordingly.

The shelf controller design features all space-grade components and contains Latch-up Current Limiting (LCL) circuitry for each board and a supervising microcontroller.

All boards are based on the powerful Xilinx Zynq US+ ZU19EG Multi-processor System-on-Chip (MPSoC) featuring a quad core ARM processor and a large Field-programmable gate array (FPGA) built onto the same die so that processing algorithms can benefit from hardware acceleration.

Due to the large amount of data that has to be moved across the system, high-performance protocols have been selected: PCI Express Gen3 x2 is used for the high-speed links on the backplane (both star and mesh connections) ensuring transfer speeds above 1GB/s, while Time-triggered Ethernet (TTE) with 1Gb/s data rate is used as interface between the System slot and payload sensors (SAR/OPT) as well as the Transmitter/Receiver (TX/RX) subsystem.

All boards are also connected to a redundant Controller Area Network (CAN 2.0) bus to exchange telemetry and telecommands (TM/TC) among themselves and with the shelf controller, while universal asynchronous receiver-transmitter (UART) links are implemented in System Boards to interface with other ancillary data sources from the host Spacecraft (GPS, AOCS etc.)

Lastly, each System Board uses a solid-state drive (SSD) for storage.




Fig. 3. Avionics Architecture

3.3 Optical Image Generation & Processing

The optical image processing will be performed in a Zynq® UltraScale+™ MPSoC, using all its computing resources (fundamentally multi-core and FPGA). Depending on the type of algorithm implemented in the Image Generation and Image Processing, it is important to carry out a previous co-design to assign the different processing tasks to specific resources and therefore, obtain the best performance of the system.

3.3.1 Optical Image Generation

The objective of the on-board Optical Image Generation is twofold: generate a high-resolution image product, for scenarios where latency is a key driver, and enable the Image Processing stage by generating a denoised and artifact free image that can be processed by the subsequent stages. In the framework of the EO-ALERT project, the Optical Image Generation has been only considered for the ship scenario, and thus it is the one presented here.

The on-board L1 product consists of:

- 1) High resolution calibrated and denoised image
- 2) Geolocation information
- 3) High resolution sea-land binary mask (3m per pixel)

First, the raw data obtained from the payload is calibrated to remove pixel inconsistencies and to convert the pixel digital counts to radiances. The calibrated image is processed with an edge-aware denoising algorithm based on optimized convolutional operations.

The position provided by the satellite GPS and the attitude provided by the Attitude and Orbit Control System (AOCS) are used to geolocate the image corners. The geolocation algorithm is prepared to integrate Earth Orientation Parameters (EOP) that can be sent to the spacecraft. The geolocation information is used to retrieve the sea-land pixel binary mask that is stored on-board using a Run-Length Encoding compression algorithm. To avoid decompressing the whole land-mask on-board, which could lead to memory issues, the minimum chunk of information required to generate the sea-land mask for the image is extracted and decompressed on the on-board memory. Using projective transformations, the sea-land mask is projected onto the image coordinates to assign sea-land information to each of the image pixels.

3.3.2 Optical Image Processing

The **ship detection algorithm** is designed to work on high-resolution panchromatic images and consists of a three-step approach:

- 1) Candidate ship extraction
- 2) Ship discrimination
- 3) Fusion

The first step, candidate ship extraction, is designed to obtain salient regions of the image that have ship-like shapes and intensities. Otsu thresholding [17] combined with intensity and shape metrics are used to create a binary image containing the areas likely to contain a ship. This information is fused with the land-sea mask to remove salient areas detected over land regions. This step, opposed to sliding-window techniques that classify each region of the image, reduces the search space leading to a more efficient solution suitable for on-board execution at the cost of not providing constant processing times.

In the ship discrimination step, each image region from the binary image is labelled independently using the algorithm from [18]. The regions from the image, corresponding to the labels from the binary image, are classified using a combination of Gabor features and Support Vector Machine classifier to assess the presence of a ship.

The last step, fusion, is designed to remove ships detected on overlapping areas of the image (in case the image is divided over different boards) or to suppress different detections of the same object (non-maximum suppression)

For the **extreme weather scenario**, convective storms are detected from GEO satellite images using a three-step algorithm comparable to the approach of the NWCSAF [12] RDT product, consisting of:

- 1) Candidate cell extraction
- 2) Candidate cell tracking
- 3) Convective cell discrimination

In the first step, candidate convective cells are detected as local temperature minima in brightness temperature images derived from GEO satellite data. The position and velocity of these cells is then tracked over time by matching candidate cells between images of subsequent acquisition times based on spatial overlap in the corresponding candidate maps. Ambiguities in the assignment introduced by cell merging or splitting are filtered in order to produce unique tracks. Finally, a machine learning classifier trained on ground truth data generated from composites of Meteosat-SEVIRI images and OPERA weather radar network data is used to perform convective cell discrimination. Each cell is characterized by its corresponding brightness temperatures in 5 infrared channels in SEVIRI imagery and their respective evolution over time. Features for the classification are derived taking historical data from up to four previous and the present acquisition times by calculating statistics on brightness temperatures (e.g., minimum, mean, maximum) and their respective intra-channel differences between acquisition times and inter-channel differences for same acquisition times. A

Gradient Boosted Decision Trees (GBDT) ensemble technique is used for classification of convective/non-convective cells.

3.4 SAR Image Generation & Processing

The Processing Board includes a dedicated Multiprocessor System-on-Chip (MPSoC) for SAR processing. In detail, the tasks to be computed are the SAR image generation (alternatively known as SAR focusing) and SAR image processing. These two steps are consecutive processes necessary to generate image products (ship and weather alerts) from the raw data collected by the Sensor Board.

3.4.1 SAR Image Generation

The generation of a precise high resolution SAR image from satellite raw data is in general a complex and computationally expensive task, which is tackled with dedicated powerful computers at ground receiving stations. Therefore, taking into consideration TerraSAR-X as space segment, a full adaptation of the existing focusing processor, which uses the Chirp Scaling algorithm [7], is not foreseen. SAR image formation algorithms employed for spaceborne system with resolution capabilities less than 3 meters are the Range-Doppler Algorithm (RDA), the Chirp Scaling Algorithm (CSA), and the wave number domain Omega-K (ω KA) [3]. After a careful review and considering the objectives posed in EO-ALERT, the latter is the algorithm selected. In particular, the approximate version, named monochromatic ω KA, seems to fit well with the resolution class and scene size aimed at in the context of EO-ALERT.

The objectives of SAR image generation are signal processing of the sensor data and computation of processing and annotation parameters. Signal processing of the sensor data such as raw data correction, Fast Fourier Transform/Inverse Fast Fourier Transform (FFT/IFFT), antenna pattern correction, detection and multilooking demands high computation power, resources and input/output throughput so this part of algorithm is implemented in the programmable logic (PL) of the Xilinx Zynq US+ ZU19EG with parallelization. Computation of processing and annotation parameters such as geometric doppler centroid determination, SAR focusing parameter and geolocation are calculated on the processing system (PS)'s ARM cores of the Xilinx Zynq US+ ZU19EG as a software.

3.4.2 SAR Image Processing

The selection of the SAR image processing algorithms to be implemented is strictly related to the identified scenarios described in 2.1 and 2.2.2. The applied **ship detection algorithm** is based on [4] and

involves the implementation of the following three image processing steps:

- 1) Initial detection
- 2) Refinement
- 3) Filtering

A Constant False Alarm Rate (CFAR) algorithm (REF) is used in the first step. As each pixel's intensity has to be compared to the mean intensity of its surrounding area, this step is computationally very expensive. Hence, it is implemented in hardware. Neighbouring detections are merged to ship candidates. The refinement step investigates the surrounding of all ship candidates to find additional ships close-by which were missed in the initial detection. The filtering consists of two parts: first, azimuth ambiguities created during the image processing are removed; second, land masking is applied to remove all detections on land. Detections left after this step are considered detected ships and their properties (e.g. location, dimensions, heading) are gathered.

For the **extreme weather scenario**, the ocean surface wind speed and wave height are derived from the SAR image. The image processing steps involved are basically the following:

- 1) Image tiling
- 2) Wind speed detection
- 3) Sea state detection

The first step subdivides the image into a grid of sub-scenes of configurable size, which determines the spacing of the final products. Wind and sea state detection are then performed on each individual sub-scene. For wind speed detection, the project will employ the Geophysical Model Function XMOD2 [5] previously tuned on TerraSAR-X archived images, calculating the wind speed from sea surface backscatter. For sea state detection, the empirical XWAVE algorithm [6] is used. The resulting product includes wind speed and wave height on a per sub-scene basis.

3.5 CPU Scheduling, Compression, Encryption & Data Handling (CS-CEDH)

Besides performing general flow-control and coordination tasks during the entire EO products generation process, the CS-CEDH module is responsible for several steps in the EO-ALERT data chain:

- Acquire the raw data from the Sensor Board and the ancillary data from the on-board ancillary data source.
- Forward the raw, ancillary and configuration data to the target optical or SAR Image Processing modules.
- Acquire the EO products (e.g. alerts, generated images, etc.) from the Image Processing modules once they are ready.

- Compress and/or encrypt the raw data, ancillary data, and EO products.
- Store the compressed/encrypted data on the on-board storage device and remove it once it is no longer required (e.g. after it has been transmitted).
- Send the compressed/encrypted data to the TX/RX Subsystem for transmission to the ground segment.

All the above operations are performed by the software running on the CPU in the PS of one of the MPSoC devices, except for image compression and encryption, which is implemented in hardware on the PL to increase the compression and encryption throughput and offload intensive operations from the CPU.

3.5.1 Compression and Encryption

Compression and encryption functionalities are available for all the data types consumed and produced by the Processing Board: raw sensor data, configuration data, ancillary data and EO products (e.g. generated images, alerts).

The compression of optical and SAR raw sensor data and generated images is achieved using a hardware implementation of a method based on the CCSDS 123.0-B-2 standard [13], modified in order to employ an out-of-loop prequantizer [14]. The compression algorithm has been extended to also embed image encryption by sign-randomization of the prediction residuals [15]. The resulting algorithm has been entirely implemented in hardware, by means of a specific custom-developed IP core in the PL part of the MPSoC. It is connected to the CPU through a High-Performance AXI4 interface to load/store the input/output data (directly from/to the main CPU memory), and an AXI4-Lite interface for commands and parameters; the use of a hardware interrupt allows for a non-blocking management of the compression/decryption operations in software. The IP has been developed with a High-Level Synthesis (HLS) approach, leveraging the ability of state-of-the-art Electronic Design Automation (EDA) tools to synthesize Register Transfer Level (RTL) code starting from a (subset of) ad-hoc function written in the common C programming language. With this approach, and by properly refactor the C model, it has been possible to design a flexible and reconfigurable digital core not tied to a specific FPGA family, eventually targeting different timing constraints, resource utilization, and performance. In particular, all the main parameters used in the envisioned scenarios (e.g., image size, number of bands, dynamic range, compression quantization, encryption parameters) can be set at runtime on a per-image basis, while others can be optionally set at synthesis time so to better exploit runtime performance. The core has been optimized, without losing in generality, for the most common used set of parameter values. The result is a hardware core

deployed on a Xilinx Zynq US+ ZU19EG board, running at 200 MHz and occupying about 50% of the PL resources. Performance tests on the target platform, on real-case test images and with timing measurements that include all the involved memory transfers (i.e., in-memory end-to-end compression and encryption, from the software perspective), showed promising results. Median values, considering both optical and SAR multi-band images, range from 97ns/pixel to 124ns/pixel, with a median value of 100ns/pixel (or 19MiB/s for images with 16-bit pixels). That means about 10s to compress and encrypt an optical 10000x10000 image.

The software running on the PS splits each image into multiple sub-sets of along-track lines (tiling), and each sub-set is compressed and encrypted individually by the hardware IP core. This allows most of the image to be successfully reconstructed even when some data is lost while delivering the image to the ground segment. The cost in terms of compression efficiency is very low: using EO-ALERT compliant optical and SAR images, the overall compressed image size with 128-line tiling is at most 3% bigger than the one of the same image compressed at once (without tiling), mainly due to the overhead of the compression headers. The degradation of the compression throughput (mostly due to the additional memory transfers) is negligible. Tiling also enables the possibility to implement parallel image compression and encryption (on multiple IP cores) if the target platform features enough hardware resources.

The other non-image data types are compressed by the software running on the MPSoC PS CPU using common generic compression algorithms (e.g. zlib's deflate method). The compressed data is encrypted using an assembly-optimized software implementation of the Keccak Pseudo Random Generator (PRG) [16], initialized with a shared encryption key and a random or pseudo-random initialization vector fetched from an on-board source. Smaller data types, like alerts, are only encrypted using the same method to reduce the processing latency, as their size does not represent a bottleneck in the transmission process. The single-core software compression throughput on the target Xilinx Zynq Ultrascale+ MPSoC ZU19EG 64-bit Cortex A53 processor (unloaded) ranges from 5MiB/s to 30MiB/s (depending on the data content), when using test data from EO-ALERT compliant optical and SAR datasets. The encryption throughput on the same platform is about 28MiB/s.

3.5.2 Storage Management

To promote an efficient use of the PS memory on the CS-CEDH MPSoC, compressed and encrypted data is immediately stored on the on-board mass-storage device. Depending on the mission requirements, the stored data may be transferred to the TX/RX Subsystem for transmission to the ground segment. The software

running on the PS CPU keeps track of the available compressed and encrypted data and its transmission status, allowing for an efficient storage-management policy. Data that has already been transmitted (or that shall not be transmitted at all) is removed.

3.5.3 Data Transmission Scheduling

The software running on the CS-CEDH MPSoC PS CPU also schedules data transfers to the TX/RX Subsystem, so that data can be transmitted to the ground segment. The transmission policy is based on a two-priority transmission queue, where alerts are given a "high" priority, and other data types (raw sensor data, ancillary data, etc.) are scheduled with "low" priority. Low priority data is only transmitted when there is no high priority data available for transmission. To further reduce the latency in alert delivery, alert data is forwarded to the TX/RX subsystem directly from the PS memory, as soon as it is encrypted. Low priority data is fetched from the mass-storage device just before being transferred instead. The smallest transmission unit is an individually compressed/encrypted file (or sub-file from a tiled image). Each transmission unit is split into a stream of CCSDS packets, sent to the TX/RX Subsystem over a high speed, unidirectional communication channel, that on the target platform has a throughput of 1Gbit/s. Taking into account the overhead of the packetization process and the storage read performance, the largest transmission unit in EO-ALERT is transferred in under 100ms. This is also the maximum delay from the moment when an alert encryption completes and the one when its transfer to the TX/RX subsystem begins.

3.6 Communications

The communications system has to deliver EO products, bulk (raw and generated image) data and alerts (processed and generated on board) locally and globally at high speed with minimum latency. A variety of technical possibilities have been investigated. The latency requirement is 5 minutes for alerts for both local and global delivery and 30 minutes maximum for raw data/images for global delivery. Local delivery means direct-to-ground communications to existing ground stations, e.g. for the Mediterranean area the stations in Matera and Puertollano (augmented by the station in Svalbard) can be considered.

For high-speed delivery of data, different solutions have been investigated and were traded against the requirements.

- 1) For microwave communications systems:
 - direct to ground communications in a local scenario
 - via a global ground station network in a global scenario
- 2) For optical communications systems:
 - direct to ground communications in a local scenario

- via a geostationary relay satellite in a global scenario

Optical communications have the advantage of very high bandwidth and license-free spectrum, the disadvantage is the high costs of the optical terminals.

Emerging big LEO constellations as a relay service (LEO to LEO and to the end-user via the constellation network) were investigated, but had to be discarded due to technical and regulatory reasons. The main problem is that these systems have been optimised for terrestrial broadband Internet access. The antenna beams are narrow and tracking antennas for an EO LEO spacecraft are at present not provided. But such a solution may become available in the future.

A solution for global delivery of alerts meeting the requirements was found using the INMARSAT satellite network and was adopted for the short-term solution. A compact product for the on-board transceiver is available by the company ADDVALUE [19] (iDRS service).

A potential delivery service for alerts can also be provided by store-and-forward transponders on board of a constellation of LEO spacecraft in a ring topology, transmitting the alerts from one satellite to next one until the alert reaches a spacecraft which is in coverage of a ground station.

In addition, an S-band communications system to deliver alerts to hand-held terminals was included in the system design.

3.6.2 Local Solution

Fig. 4 shows the block diagram of the on-board communications system. The first payload for high-speed data operates in Ka-band. This band has been chosen because the frequency band of 25.5 – 27 GHz is assigned for Earth Observation and provides significantly more bandwidth than X-band. The Ka-band transmitter consists of a modem/codec (supporting QPSK, 8PSK, 16APSK and 64APSK modulation schemes with serial concatenated convolutional coding) and an upconverter unit. A 10 W Ka-band solid-state power amplifier (SSPA) amplifies the signal and delivers it to a horn antenna with 25 dBi gain. The system is fully redundant. One chain operates in right-hand circular polarisation (RHCP), the other one in left-hand circular polarisation (LHCP). Data rates up to 2.6 Gbit/s can be supported. The Ka-band system is used for both bulk data and alerts to local ground stations.

The second payload consists of an S-band transmitter for alerts only, supporting data rates up to 1 Mbit/s to small and inexpensive hand-held terminals for rescue teams that do not have Internet connectivity.

Data (bulk data and alerts) are delivered by the avionics system to the respective communications payloads via the multiplexer/interface based on the data type indicated in the packet header.

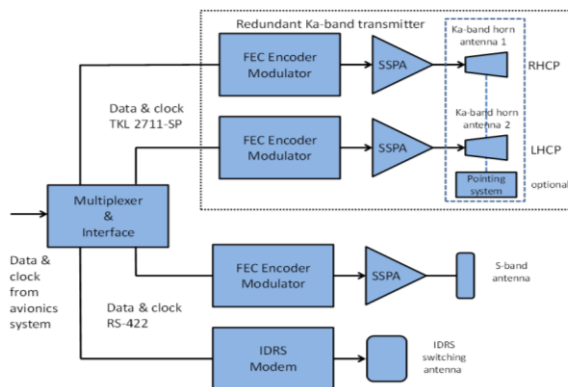


Fig. 4. Block Diagram of the Communications System

3.6.1 Global Solution

Bulk data can be delivered globally within 30 minutes via Ka-band, if a ground station network of at least 13 stations around the globe is employed (Fig. 5):

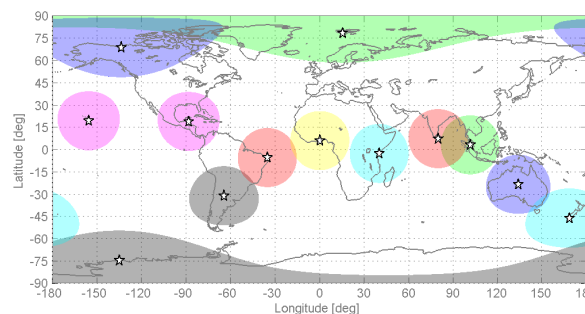


Fig. 5. Required Ground Station Network for Global Delivery in Less than 30 min

An alternative global broadband solution is the employment of an optical terminal on board of the EO spacecraft and a data relay service such as the European Data Relay Satellite System (EDRS).

For the global delivery of alerts the iDRS service is envisaged for the short-term solution which provides a data rate of 250 kbit/s. About 100 alerts can be transferred in 37 seconds.

4. High-Speed Avionics Test-Bench

The Avionics Test-Bench consists of a scaled-down version of the Avionics Subsystem composed by commercial hardware (not cPCI Serial Space compliant) and offering four boards instead of seven.

In order to resemble the PDPU as much as possible, one board is dedicated to CPU Scheduling and Compression, Encryption & Data Handling functions (CS-CEDH) while the other three boards are dedicated to processing. As for the Avionics, CS-CEDH board communicates directly to a Master Processing Board which can offload part of the calculations to two Slave Processing boards. Moreover, the same per-board

processing power is provided by using the same MPSoC (Zynq US+ ZU19EG).

Test-Bench's boards are also interconnected with PCIe links, as foreseen for the Avionics, so that the complete data chain can be reproduced and a realistic latency measurement can be obtained. A standard Gigabit Ethernet connection (not Time-Triggered Ethernet, but same data rate) connects the Test-Bench to a TX/RX subsystem emulator which is needed to test the transmission to Ground Segment (GS).

The emulator implements two different data paths: the first is a FPGA simulation of the transmission channel which sends product data over Ethernet to a personal computer (PC) for decryption and display; the second transmits products over S-band to a Raspberry PI module acting like a hand-held device to which EO-ALERT products could be delivered in the real application. The Raspberry PI runs the decryption algorithm and displays alerts on a small screen.

The Test-Bench can be configured manually to process OPT or SAR data using an external PC and dedicated Ethernet links to each board. In order to inject real OPT and SAR data in the Test-Bench, raw images are loaded into a SSD connected to the CS-CEDH board, providing the inputs that should come from sensors in the final system.

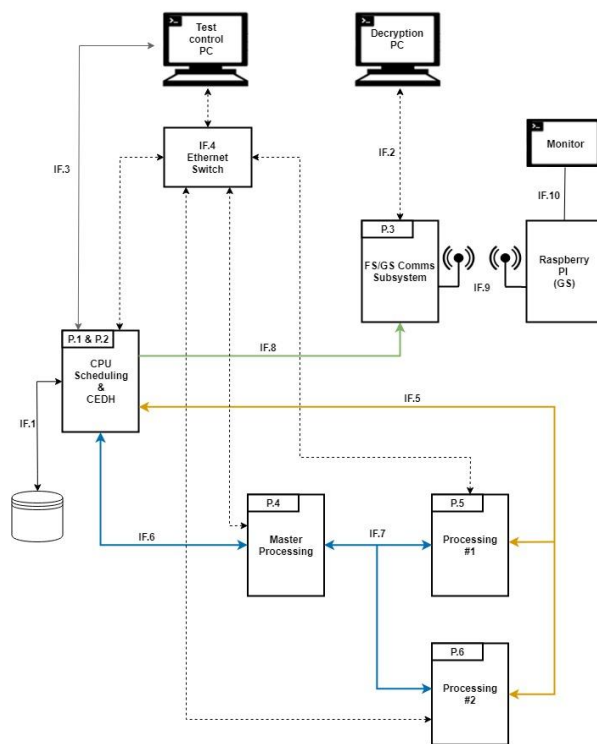


Fig. 6. Test-Bench Architecture

5. Results

This section presents the results of the project in terms of latency and processed area for optical and SAR data processing chains.

5.1 Optical Data Processing Chain

For the ship scenario the optical processing is performed in a configurable multi-board scheme; 1 to 3 boards can be added to decrease latency or increase target area size. Each board will process 2 raw data blocks of about 8000 x 6000 points and covering about 9 km x 6 km (54 km²) each, therefore about 100 km² per board. To process this area, the **entire processing chain takes 35 s running on a single board**. In a multi-board scheme where it is possible to process 300 km² in parallel, 900 km² can be processed in about 110 s and 1200 km² in about 150 s, assuming transfer delays and management tasks.

For the weather scenario data is processed in a dual-board scheme with only one processing board. The processed images have about 1200pixels times 650pixels with 3 km/pixel of GSD, covering almost 7·10⁶ km². Assuming transfer delays and management task, it is possible to have ready the products to be sent in about 15 s.

5.2 SAR Data Processing Chain

The SAR processing is run on a single Xilinx Zynq US+ ZU19EG board. The sensor input consists of 1-3 raw data blocks of 8192 x 32768 points each. Each block covers a surface area of 12.5 km x 30 km. The image generation for one such block currently takes 4 s with FPGA input clock frequency of 125MHz and image processing takes 30 s for the ship scenario and 35 s for the extreme weather scenario using a raster size of 2 km x 2 km. Hence, the **total runtime of SAR processing for a single block is below 40 s**. Multiple blocks are calculated successively, hence, the total runtime for three blocks covering an area of 1125 km² is below 120 s.

6. Experimental Campaign

To validate the performances of the high-speed avionics test bench in a representative environment, a testing campaign using real EO data from the chosen scenarios has been executed at the beginning of July 2020.

For the ship detection scenario, the CRV LEONARDO ship of the Centre for Maritime Research and Experimentation (CMRE) has been deployed off the coast of La Spezia (Italy) during five days. The locations of the ship have been selected in order to fall at the centre of Deimos-2 and TerraSAR-X swaths, in order for the space assets to take images of the ship. A feasibility analysis on the observation opportunities for both Deimos-2 and TerraSAR-X, taking into account

orbital dynamics and payloads constraints, has been performed in order to identify the exact date and time of observations of pre-defined locations in the Tyrrhenian Sea. The raw data acquired by the space assets will be used as input in the avionics test bench of the EO-ALERT data processing chain and the in situ measurements provided by the ship will be used to validate the results of the generated EO products.

For the meteorological scenario, since it not possible to foresee and control an extreme event, archived data from SEVIRI and TerraSAR-X will be used as input raw data, while historical in situ measurements will be extracted from OPERA network and sea buoys to validate the results of the EO-ALERT data processing chain.

7. Conclusions & Way Forward

In this paper, we have presented an overview of the EO-ALERT project, its objectives, structure and innovations and a description of the results obtained individually by each partner on their breadboards. Preliminary results show that the demanding objective of providing EO products with latency below 5 min can be achieved. Furthermore, the results show **global EO product latencies below 1 min in realistic scenarios**.

The architecture is targeted to the generation of alerts in two example scenarios, ship detection and extreme weather monitoring, however it remains quite general and can be easily adapted to alternative scenarios. The proposed architecture can be efficiently implemented relying on a hybrid solution combining space qualified components and high performance COTS components.

The EO-ALERT project will run until mid 2021 and the roadmap for the remaining activities includes the full data chain integration and the verification into the avionics test bench.

The components of the full data chain will be integrated into the avionics test bench in OHB Italy premises, collecting all the breadboards individually developed and tested by the partners along the EO-ALERT project.

Overall results of the entire project, including the experimental campaign and validation of the avionics test bench, will be presented to end-users and possible commercial partners at the end of the project during the final EO-ALERT workshop.

Acknowledgements

The research leading to this publication has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 776311.

DEIMOS would like to acknowledge the Centre for Maritime Research and Experimentation (CMRE) for

their role in and support to the experiential testing for the ship detection scenario.

References

- [1] Business Innovation Observatory. Applications related to Earth Observation. Study Case 63. Ares (2016) 2153857 - 06/05/2016
- [2] ESA webpage: <https://sentinel.esa.int/web/sentinel/missions/sentinel-1/data-distribution-schedule>, last retrieved 21/09/2018
- [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] 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
- [5] 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
- [6] 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
- [7] Raney, R.K., Runge, H., Bamler, R., Cumming, I.G., Wong, F.H.(1994) "Precision SAR Processing Using Chirp Scaling," IEEE Transactions on Geoscience and Remote Sensing (TGRS), Vol. 32, pp. 786-799
- [8] D. Sinclair, J. Dyer, "Radiation effects and COTS parts in SmallSats", in Conference on Small Satellites, SSC13-IV-3K, 2013
- [9] A. LaBel, "Commercial Off The Shelf (COTS): Radiation Effects Considerations and Approaches", in NEPP Electronic Technology Workshop, 2012
- [10] D. Gleeson, M. Melicher, "The smart backplane – lowering the cost of spacecraft avionics by improving the radiation tolerance of COTS electronic systems", in 33rd Space Symposium, 2017
- [11] <http://eumetnet.eu/activities/observations-programme/current-activities/opera/>, last retrieved 18/09/2020
- [12] <http://www.nwcsaf.org/>, last retrieved 18/09/2020
- [13] Consultative Committee for Space Data Systems (CCSDS), "Low-Complexity Lossless and Near-Lossless Multispectral and Hyperspectral Image Compression," Blue Book, no. 1, February 2019
- [14] D. Valsesia and E. Magli, "High-Throughput Onboard Hyperspectral Image Compression With Ground-Based CNN Reconstruction," in IEEE Transactions on Geoscience and Remote Sensing, vol. 57, no. 12, pp. 9544-9553, Dec. 2019

- [15] A. Migliorati, T. Bianchi, E. Magli, “Selective encryption in the CCSDS standard for lossless and near-lossless multispectral and hyperspectral image compression”, Proceedings Volume 11533, Image and Signal Processing for Remote Sensing XXVI; 1153312 (2020)
- [16] Bertoni G., Daemen J., Peeters M., Van Assche G. (2010) Sponge-Based Pseudo-Random Number Generators. In: Mangard S., Standaert FX. (eds) Cryptographic Hardware and Embedded Systems, CHES 2010. CHES 2010. Lecture Notes in Computer Science, vol 6225. Springer, Berlin, Heidelberg
- [17] Otsu, N. (1979). A threshold selection method from gray-level histograms. IEEE transactions on systems, man, and cybernetics, 9(1), 62-66
- [18] Wu, K., Otoo, E., & Suzuki, K. (2009). Optimizing two-pass connected-component labeling algorithms. Pattern Analysis and Applications, 12(2), 117-135
- [19] <https://www.addvaluetech.com/inter-satellite-data-relay-system-idrs/>
- [20] S. Tonetti, S. Cornara, G. Vicario de Miguel, L. Carzana, M. Kerr, R. Fabrizi, S. Fraile, C. Marcos Martín, D. Velotto, “EO-ALERT: Next Generation Satellite Processing Chain for Security-Driven Early Warning Capacity in Maritime Surveillance and Extreme Weather Events”, 2019 Living Planet Symposium, 13-17 May 2019, MiCo - Milano Congressi, Milan, Italy
- [21] A. Fiengo, C. Marcos, J. I. Bravo, M. Kerr, “EO-ALERT – Extreme Weather Scenario: Towards Convective Storm Nowcasting via On-Board Satellite Processing”, 3rd European Nowcasting Conference, 24-26 April 2019, Madrid, Spain
- [22] Kerr M., Cornara S., Latorre A., Tonetti S., Fiengo A., Gomez de Agüero S., Bravo J. I., Velotto D., Breit H., Balss U., Koudelka O., Teschl F., Magli E., Bianchi T., Freddi R., Benetti M., Fabrizi R., Fraile S., Marcos C., “EO-ALERT: A Novel Flight Segment Architecture for EO Satellites Providing Very Low Latency Data Products”, ESA Phi-week 2019, 9-13 September, ESA-ESRIN, Frascati, Rome