OPTIMIZING THE PATH FROM PHOTONS ON-ORBIT TO MBITS ON THE GROUND

Thys Cronje(1), Gregor Dreijer(1), Jacu Vos(1), Søren Pedersen(2), Mathias Persson(2), Henrik Magnusson(2), Fredrik Bruhn(2), Andy Phipps(3)), Daniel Åkerman(3), Lucy Donnell(4), Murray Ireland(4), Francis Doumet(5)

(1) Simera Sense, Old Paardevlei Road, Somerset West, 7130, South Africa, +27(0)218526450, thys.cronje@simera.com
(2) Unibap AB, Kungsängsgatan 12, 4 FL, SE-753 22 Uppsala, Sweden, +46 18-320 330, soren.pedersen@unibap.com
(3) Spacemetrics, PO Box 1575, Guildford, GU1 9BW, UK, +44(0)7535431542, ap@spacemetric.com
(4) Craft Prospect, Suite 12, 1048 Govan Road, Glasgow G51 4XS, UK, +44(0)7938566172, lucy@craftprospect.com
(5) Metaspectral, 400-319 W Hastings St., Vancouver, BC V6B1H6, Canada, 778-814-6553, francis@metaspectral.com

The amount of data generated by Earth Observation instruments used on nanosatellites has increased significantly over the last couple of years. This tendency motivates the use of advanced onboard image processing capabilities, mainly due to current restrictions in terms of existing hardware and downlink bandwidth. This development in the earth observation industry requires a new data processing and management approach to enable new data products from all types of earth observation instruments.

This work presents a novel onboard processing pipeline to increase the data efficiency of a hyperspectral instrument onboard a nanosatellite developed for the King Abdullah University of Science and Technology (KAUST). The proposed optical instrument is a 32 band VNIR hyperspectral imager with a GSD of 30 m, at 500 km orbital height and a swath of 120km.

The novel processing solution presented as SpaceCloud with its heterogeneous computing solution enables a hybrid approach to evolving processing tasks on orbit more dynamically. This approach allows processing tasks and other moving parts of the data processing and filtering pipeline to be transferred directly onto the spacecraft.

As part of the image processing pipeline being implemented on the spacecraft, the RAW image data is processed from level 0 to level 1A data, and the processing unit performs radiometric corrections in the imager, geometric corrections, band co-alignment, and geo-location of the raw image data being performed by Spacemetric Keystone on SpaceCloud. For data reduction purposes, area of interest identification and cloud masking is performed by Craft Prospect on SpaceCloud. In contrast, data is lossless compressed down to as little as 30% of the original size in real-time using an implementation of the CCSDS 123.0-B-2 algorithm developed by MetaSpectral. Furthermore, the implemented architecture allows for uploading user-specific applications on the fly, turning raw image data into actionable information.

This solution is currently being implemented using the HyperScape50 from Simera Sense and the iX5-100 SpaceCloud computing solution from Unibap. This solution represents a significant step forward regarding what is now being done for Earth Observation data processing onboard a 6U Cubesat. This sensor capability is enabled through the SpaceCloud pre-processing framework that allows for fast integration of any instrument and using the

Spacecloud Framework generic sensor interface (gSIF) to read this data from user-specific applications. It now enables a wide range of functionalities and applications to be brought onto the spacecraft for data preparation, filtering purposes, or generation of actionable data points like identifying critical events on the ground, e.g., flooding or forest fire detection.

This opens a new approach to providing EO data to users faster on the ground either because it is tagged on the edge with different valuable information for further ground processing and prioritized when moving data to the ground. Secondly, the data can be delivered in the form of metadata in text messages, possibly combined with small selected data packets over real-time links directly to users on-ground.

1 INTRODUCTION

The number of smaller satellites carrying Earth Observation payloads is exploding [1][2]. Since Earth Observation data became available for civilian use in the early 1970s with the Landsat program, it created the backbone for land surface monitoring and characterization. For the first four decades after the launch of Landsat 1, the commercial market was mainly dominated by assets owned and operated by governments or sustained by government and public procurement. However, major shifts have been seen within the EO landscape over the last decade. These shifts are driven by the so-called new space movement where the traditional space industry meets the information technology world and is funded by private investments [2][3].

This transition from traditional or classical space to the new paradigm is a result of the exponential growth in the demand for data, the lower cost of small satellites and the associated launch costs, advancements in sensor technology, better connectivity to store and distribute data, and new ways of processing large volumes of data [4]. However, the commercial EO market is still evolving, and smaller satellites play a more prominent role in this sector.

1.1 SPATIAL RESOLUTION

Ground Sampling Distance (GSD) is still king, especially in the defense market where very high-resolution data (< 30cm GSD) combined with high accuracy is a crucial requirement. On the other hand, high accuracy spatial, radiometric, and spectral data requires expensive solutions and large platforms. Therefore the data within this segment tends to be costly and without the reach of many commercial applications [5].

1.2 TEMPORAL RESOLUTION

By lowering the EO system's revisit time to increase the frequency of data collected and services delivered, new demands in the commercial sector are addressed. As a result, data freshness is becoming a dominant driver within these markets, focusing on change detection (precision agriculture, business intelligence, and location-based services) [3][5]. The data freshness challenge can be solved by either capturing every point or any point on earth at regular intervals, for example, daily. Most traditional service providers use the latter option, while the new space actors prefer the first option.

In both cases mentioned above, significant amounts of pixels are collected, increasing the data handling complexity. However, regularly covering every point on the earth's land surface is the most challenging.

For a complete global land surface coverage at 30m GSD, about 165 gigapixels are captured. The above numbers must be multiplied by the number of spectral bands when using a multi- or hyperspectral instrument. Therefore, the system's ability to efficiently handle large volumes of data is an essential factor to consider when designing an EO small satellite.

1.3 OVERVIEW OF ONBOARD PROCESSING

As the need for more spectral data in the visible and infrared spectrum increases, the data output from especially multi- and hyperspectral sensors continues to grow, combined with the increasing output from the development of COTS EO sensors, demand an increase in the ability to filter data close to the sensor. This has been an approach applied on the ground for a long time, and that trend is moving into space for a wide variety of EO sensors to avoid data congestion and provide actionable information with low latency to users. Several architectures are available when designing a processing system for satellites, including CPUs, FPGAs, and GPUs. They all come with different advantages, and when applied, the CPU offers flexibility e.g. the ARM Cortex-series offers a popular processing choice, the FPGA offers speed, and high throughput with the Xilinx SoC 7000-series has in the last decade been the popular choice for complex, and demanding computing tasks on-orbit. Furthermore, the GPU brings the latest addition to space-based processing, i.e., hardware acceleration, with multiple offerings currently being tested on orbit, including Nvidia Jetson, AMD, and the Myriad2/X from Intel characterized on-orbit [6][7].

1.4 CURRENT DRIVERS FOR ONBOARD PROCESSING

Onboard processing for EO data is set on a track where parts of the processing pipeline are moved onto the spacecraft to achieve two major objectives, supporting autonomous operations on-orbit and minimizing the cost of bringing data to the ground. The two main objectives are based on new processing architectures being brought to space and deployed to solve some of the processing challenges related to the goals.

Tip-and-Cue for more prominent constellations is a significant driver for on-orbit processing, where autonomous tasking in space highly depends on the ability to process, filter, and relay information. This is achieved by combining different known terrestrial tools and bringing them to orbit. This is done in the software arena, where large terrestrial data providers integrate their terrestrial networks with the space-based infrastructure. Commercial examples are Amazon Web Service and the Kuiper project, Microsoft is launching Azure Orbital, and SpaceX is offering real-time connectivity with StarLink. Here all the prominent actors are trying to move the sensor closer to the data distribution infrastructure on the ground with well-known goals like lower latency, ease of access, and flexibility to maintain and upgrade algorithms running on the edge [6][7].

1.5 MAJOR CHALLENGES FOR ONBOARD PROCESSING

A typical earth observation workload runs in the cloud and on desktop computers on the ground. Using tools and computing environments familiar to the developers implementing these applications reduces the step of running the application in orbit. Things that are problematic in this regard are [6][7]:

- Custom compute architectures require expensive adaption of algorithms that are otherwise available off the shelf.

- The standard file-based interfaces found on the ground using already pre-processed data are not available off-the-shelf on the satellite. Instead, applications may need to pre-process and/or adapt the interfaces to the application.
- Specific operational requirements for a satellite limits other requirements on when and how computation can happen. Power budgets for a small satellite may require only a limited compute window during the day, while a typical ground-based application can use significant amounts of power for a long time. Finding the correct abstraction for application developers is key to solving this
- Ultimately having sufficient performance to perform the calculations is a roadblock and may initially require careful selection of algorithms and or partitioning applications to run partly on the ground and partly in orbit. Devising the correct abstraction for how to handle data communication between the application on the satellite and the application on the ground is critical

2 KAUST MISSION OVERVIEW

The King Abdullah University of Science and Technology (KAUST) contracted Spire Global, Simera Sense, and Unibap to develop a collaborative nanosatellite mission to collect high-quality data across global terrestrial, coastal, and ocean ecosystems intended for an extensive range of applications [8].

The payload onboard the 6U Cubesat consists of three main components:

- A GNSS-R instrument to provide timing and positioning information from reflected signals for weather and ground moisture readings at a 100% duty cycle;
- A hyperspectral instrument to capture up to 32 user-configurable bands across the visible and near-IR spectral range with the ability to capture at 10 minutes of continuous strips per orbit or at least 10 scenes per day, and;
- An onboard processing unit, with the associated software environment, to process and handle large amounts of data onboard the satellite.

Figure 1: An illustration of the HyperScape50 optical instrument.

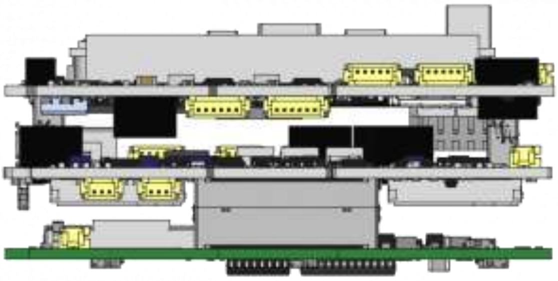

Figure 2: An illustration of the iX5 computing platform.

2.1 THE HYPERSPECTRAL INSTRUMENT

The HyperScape50 instrument was selected for the Kaust mission. This instrument utilizes a continuous variable filter aligned in the focal plane on a 4096 x 3072 pixel CMOS global shutter detector. A maximum of 32 spectral bands from 442 band options, each with a Full-Width Half Maximum of 3.5% of the selected central wavelength, can be captured at a time. The 32 spectral bands can be selected on the fly for each planned imaging session.

Furthermore, the Hyperscape50 was designed to fit within a 1U form factor, weigh less than 0.5 kg and operate over a wide thermal range, -10°C to +50°C. A high-speed LVDS data interface is used between the HyperScape50 and iX5 hardware. In addition, a redundant data path is provided via the spacecraft bus.

Table 1: The HyperScape50 specifications.

OPTICS	
Focal Length	93.9 mm ±1 mm
Aperture Diameter	11.75 mm
Full Field of View	13.68° (across-track); 9.2° (along-track)
IMAGING	
Configuration	Line-scan (push broom)
Sensor Technology	CMOS, Global Shutter
Cross Track Resolution	4096
Pixel Size	5.5 µm
Swath @ 500 km	122 km
GSD at 500 km	30 m
Spectral Bands	1 Panchromatic Band from 500 nm – 750 nm 442 Hyperspectral Bands from 442 – 884 nm spaced 1 nm apart
Hyperspectral FWHM	3.5% of the central wavelength

Hyperspectral Accuracy	± 1 nm in High Accuracy Mode ± 4 nm in High Band Count Mode
# of dTDI stages	Up to 32 stages per band
Storage Capacity	128 Gigabyte ECC NAND Flash
Power Consumption	3.5 W when idle or during image data readout 7.0 W during imaging
ENVIRONMENTAL AND MECHANICAL	
Mass	0.44 kg ±5%
Dimensions	95.9 x 90.2 x 117.5 mm
Operating Temperature	-10°C to +50°C

2.2 THE iX5-100 AND SPACECLOUD SOLUTION

Spacecloud is the name of the whole family of technologies from hardware to a framework that makes up a useable system for in-orbit data processing.

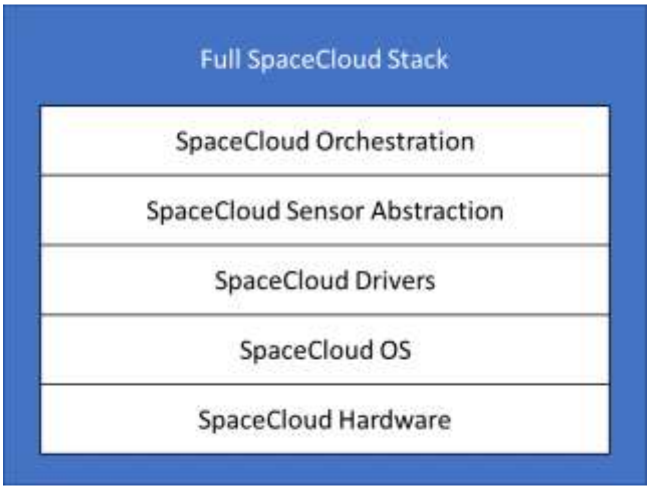

Figure 3: An overview of the SpaceCloud architecture.

Spacecloud hardware, in this case, consists of the iX5 with a specific extension board.

Spacecloud OS is an Ubuntu packaged for use in space. It contains the following core features

- Drivers that integrate the HW sensors of Spacecloud (e.g. standard sensors and robustness systems like the Safetychip watchdog)
- An optimized driver for the graphics stack found in the Spacecloud Hardware
- BIOS tweaks

On top of the drivers sits a sensor abstraction that contains an image pipeline that transforms the sensor data to a useable product by the downlink pipe or Spacecloud Framework

On top of this sits Spacecloud Framework with an abstraction layer for serving the image data to applications and orchestration for triggering applications and tracking their lifetime.

3 THE IMAGE PROCESSING PIPELINE

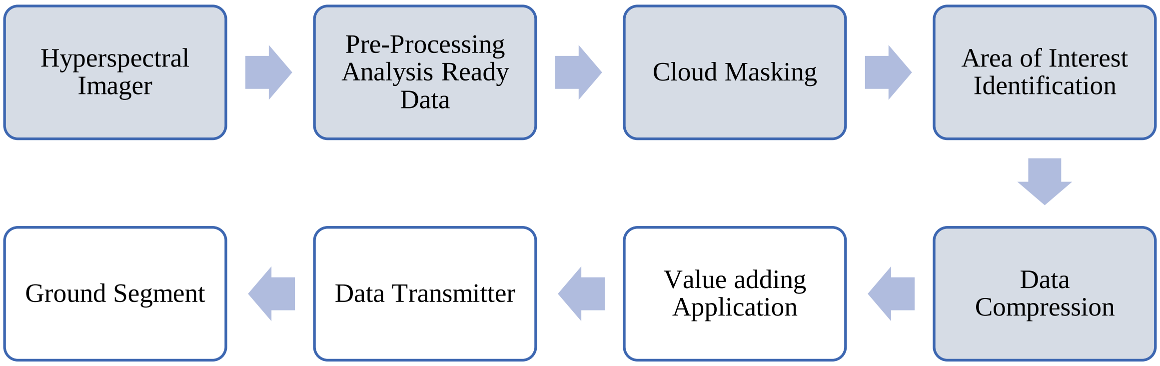

Figure 4: A simplified image processing and data handling pipeline.

The image pipeline exists as its application right under the OS and can serve use cases in Spacecloud Framework or custom functionality. In the current solution, the downlink capability is implemented as a second consumer of the pre-processed data to streamline the solution when no in-orbit processing is required.

The imaging pipeline is built around a traditional pipelined DAG where all individual components are scheduled to run parallel to maximize throughput and utilize the platform to the max. In practice, this means that cloud masking is started as soon as a tile is finished and can run on the Movidius deep learning accelerator while the data preparation IP core, creating analysis-ready data, is run. At the end of the pipeline is the data compression core which runs on the FPGA. The pipeline leaves the GPU available, which can be used in SCFW for additional parallel algorithms. Data is stored in a datastore/database for prioritized downlinking based on a policy decision from the ground (prioritizing between region, cloud coverage or age).

The data stored are 1024x1024 cloud-optimized TIFs. The size is optimized for the processing elements and sufficient granularity to avoid downlinking cloud-covered areas. In this implementation, the data is shared to Spacecloud Framework via a request interface to the database. Similar to the downlink interface, an application can request a particular set of data, the data is served via gRPC/http interface to applications running as docker containers.

3.1 DATA CAPTURE

With a scene resolution of 4096x4096 pixels and 32 spectral bands, each pixel sampled at 12 bits, the KAUST payload generates a significant data size, making data handling extremely challenging. Furthermore, the data stream needs to be as efficient as possible to maximize the return on investment. Table 2 the available storage space, the data generated per scene, and the satellite's download capacity.

Table 2: The payload's data capturing and handling capabilities.

Data Parameters	
GSD at 500 km	30 m

Scene Resolution	4096 x 4096 pixels
Number of bands	32 Bands
Quantization	12 bits
Time to capture a scene	17 seconds
Storage Space	128 GByte (1 Tbit)
Scene Size	
1 Band	16.8 Mpixels 192 Mbits
32 Bands	537 Mpixels 6 Gbits
Satellite Download Capacity	7 Gbits/day

3.2 ANALYSIS-READY DATA

The software module responsible for creating analysis-ready data is designed to process hyperspectral imagery from L0 to L1 products – defined as radiometrically sensor corrected and band registered imagery - and store these on board. This expertise has been extensively utilized on-ground by using the company's Keystone image processing suite.

The key design challenges in orbit demonstration evolve around reducing the 1 million lines of code in the ground processing solution to a manageable number compatible with resources available on a spaceborne processor. Consequently, the architecture has been substantially streamlined. In addition, it is desirable to ensure the module can process image data as it arrives in the system to prevent a data bottleneck. Therefore, a module that can operate in a more memory-constrained environment is required. Furthermore, consideration was given to developing computationally efficient data processing to limit the number of cycles the CPU and GPU must undertake; this minimizes overall power consumption. Lastly, to allow other algorithms to run on other execution hosts in parallel, the module is designed to do image processing with CPU core only.

A series of steps are required to accomplish this activity. First, raw data from the hyperspectral sensor is ingested into the module. Secondly, the parameters required to reduce radiometric effects like static noise and objective light fall-off were estimated and applied to the data. Next, the geometric differences between the 32 bands are measured, and the time-dependent attitude changes of the imager are calculated. Then, the module creates an output file where 31 bands are resampled to fit the master band. All 32 bands are now registered – a process that means that all pixels in the 32 wavelength bands are co-aligned. Finally, this image is outputted in a format that subsequent stages can utilize.

3.3 CLOUD DETECTION

Cloud detection is performed autonomously onboard for data reduction, identifying data regions that contain cloud cover and, therefore, obscure valuable information. The detection is carried out by a compact machine learning (ML) model, which performs semantic segmentation, indicating cloud presence at the pixel level. The model returns a cloud detection result in the form of a binary mask with each pixel belonging to one of two classes: clear or cloudy.

The model has been trained on a dataset of multispectral (MS) RGB images captured by Sentinel-2. The dataset provides truth labels in the form of cloud pixel masks and other labels such as land type, cloud type, and percentage cloud cover, which are valuable for assurance. For assurance of the ML component, datasets used during development must be relevant and balanced. The model is tested and verified on independent datasets to understand limitations and identify where retraining may be needed.

A transformation is performed on the hyperspectral (HS) data received from the payload, where multiple HS bands are aggregated into three wider RGB bands, which the trained ML model then recognizes.

Figure 5 is an RGB image containing Sentinel-2 data input to the cloud detection model. The image was captured over agricultural land and featured clear and cloudy areas. The cloud detection mask generated by the model is overlayed onto the image in a transparent blue color.

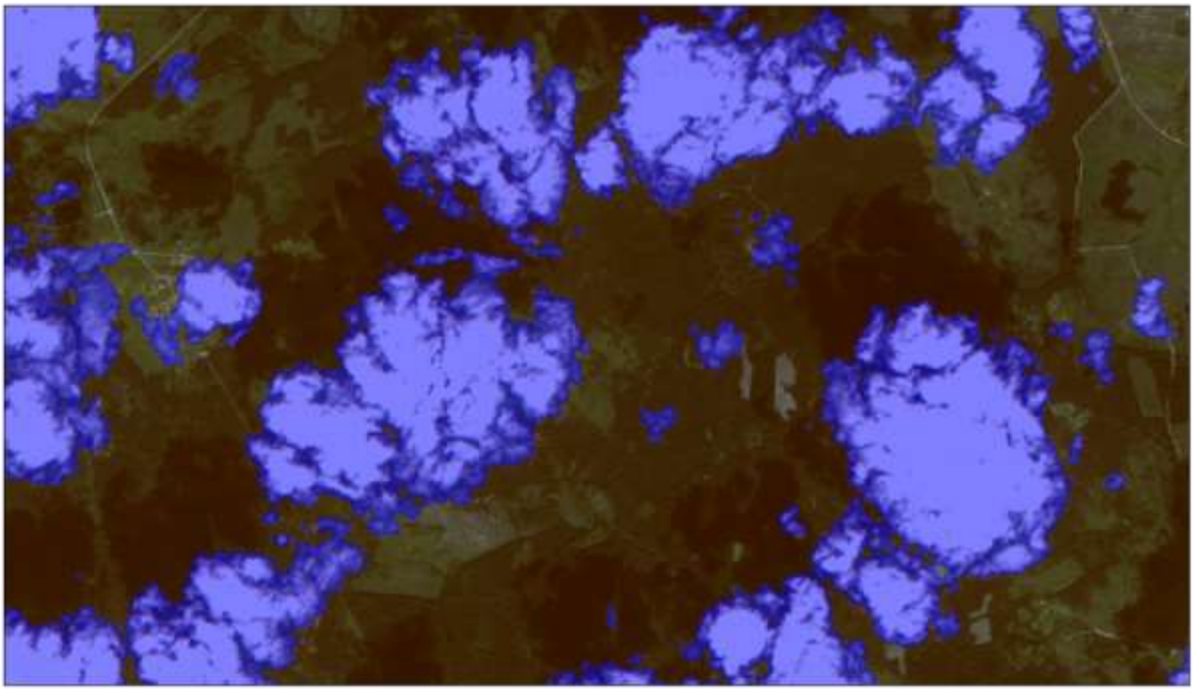

Figure 5: A typical image with clear and cloudy areas.

3.4 DATA COMPRESSION

Novel data compression algorithms solve the vast amount of data volume generated by hyperspectral imagers. Specifically, the recently-finalized 123.0-b-2 standard was purposefully developed for lossless and near-lossless compression of hyperspectral data by the Consultative Committee on Space Data Systems (CCSDS), a consortium of 11 major space agencies and over 150 observers and industrial participants worldwide.

The CCSDS 123.0-b-2 algorithm implements a two-step process for data compression. The first is a residual calculation from neighboring pixels, and the second is encoding.

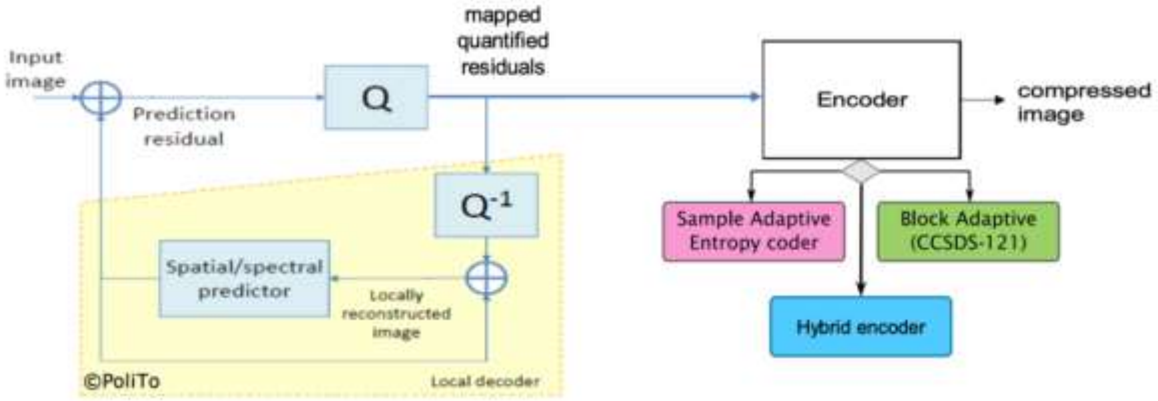

Figure 6: CCSDS-123.0-B-2 algorithm. Credit: Roberto Camarero – Lucana Santos | OBDP 2019

Operations of the CCSDS 123.0-b-2 standard were parallelized and implemented on a Size, Weight, and Power (SWaP) optimized hardware Field Programmable Gate Array (FPGA) onboard the spacecraft to achieve very-high throughput (up to 125 Msamples/sec) and at a very low power (< 1W) consumption. The CCSDS 123.0-b-2 algorithm consumes approximately 3.4 times less logic resources than the often-used alternative, JPEG-LS. Table 3 summarizes the resource consumption of the compression algorithm on the Microsemi SmartFusion2 System-on-Chip (SoC) located on the SpaceCloud hardware.

Table 3: A summary of the FPGA resource allocation.

	Logic Resources	**Memory Resources**	**Frequency (Hz)**	**Throughput (Msamples/sec)**
CCSDS123.0-B-2	6,174	1K uSRAMs: 16 18K LSRAMs: 34	125	125
JPEG-LS(*)	21,014	1K uSRAMs: 71 18K LSRAMs: 22	80	100

(*) Estimate. Based on an 8-bit-per-sample configuration by CAST, Inc.

Sample 31-band data from the Moderate Resolution Imaging Spectroradiometer (MODIS) imager was losslessly compressed using the CCSDS 123.0-b-2 algorithm. The average data compression is 86% and exceeded JPEG-LS compression by 9%. Detailed results are shown in Table 4.

Table 4: Data compression results.

	Raw Size (bytes)	**CCSDS 123.0-b-2 Compressed Size (bytes)**	**JPEG-LS Compressed Size (bytes)**
MODIS AQUA (EV) L1A	180,544,000	23,978,544	26,070,893

MODIS TERRA (EV) L1A	180,544,000	24,906,096	27,259,112

A slightly better compression ratio is expected with CCSDS 123.0-b-2 on this system. In addition, since more bands are captured (32 vs. 31), the CCSDS 123.0-b-2 algorithm could locate more correlated data, which increases the compression ratio accordingly.

4 CURRENT STATUS

While in orbit processing is still in its infancy, the current implementation does address some of the roadblocks to in-orbit processing [6][7]:

- Similar to many ground applications, a cloud-optimized GeoTIFF interface is used to serve data to applications data to the
- The platform runs common application frameworks like Docker and can run all the usual software that practitioners are using on the ground like TensorFlow, PyTorch, OpenCV, SNAP, and GDAL for processing
- The Spacecloud Framework implements an abstraction for data processing that simplifies scheduling for the developer.

5 CONCLUSION

This implementation brings a flexible platform for efficient data downlink and in-orbit data processing from a high-performance imager. Furthermore, given the ability to update applications, the functionality can continue to evolve during the satellite's lifetime introducing things like area of interest selection and other capabilities based on the end-users requirements.

6 REFERENCES

[1] Smallsats by numbers 2022, A Bryce Tech publication, February 7, 2022, https://brycetech.com/reports

[2] Allied Analytics, LLC, Report "CubeSat Market by Size, Application, End-User and Subsystem: Global Opportunity Analysis and Industry Forecast, 2020-2027", January 2021.

[3] Denis, Gil & Claverie, Alain & Pasco, Xavier & Darnis, Jean-Pierre & Maupeou, Benoît & Lafaye, Murielle & Morel, Eric. (2017). Towards disruptions in Earth observation? New Earth Observation systems and markets evolution: Possible scenarios and impacts. Acta Astronautica. 137. 10.1016/j.actaastro.2017.04.034.

[4] van Genderen, John & Goodchild, Michael & Guo, Huadong & Yang, Chaowei & Nativi, Stefano & Wang, Lizhe & Wang, Cuizhen. (2019). Digital Earth Challenges and Future Trends. 10.1007/978-981-32-9915-3_26.

[5] Earth Observation data and services market, an extract, A 360° view of the demand & supply related to the Earth observation data & services market, October 6, 2021, 14th Edition.

[6] Bruhn, F.C., Tsog, N., Kunkel, F. et al. Enabling radiation-tolerant heterogeneous GPU-based onboard data processing in space. CEAS Space J12, 551–564 (2020). https://doi.org/10.1007/s12567-020-00321-9

[7] Oskar, F., et al. Spacecloud cloud computing and in-orbit demonstration, European Workshop On On-Board Data Processing (Obdp2021), 14-17 June 2021

[8] Kaust, This Saudi university team has big plans for a tiny space-bound satellite, Wired, Technology, October 31, 2021, https://wired.me/science/space/this-saudi-uni-team-has-big-plans-for-a-tiny-space-bound-satellite/