→ DISCOVERY

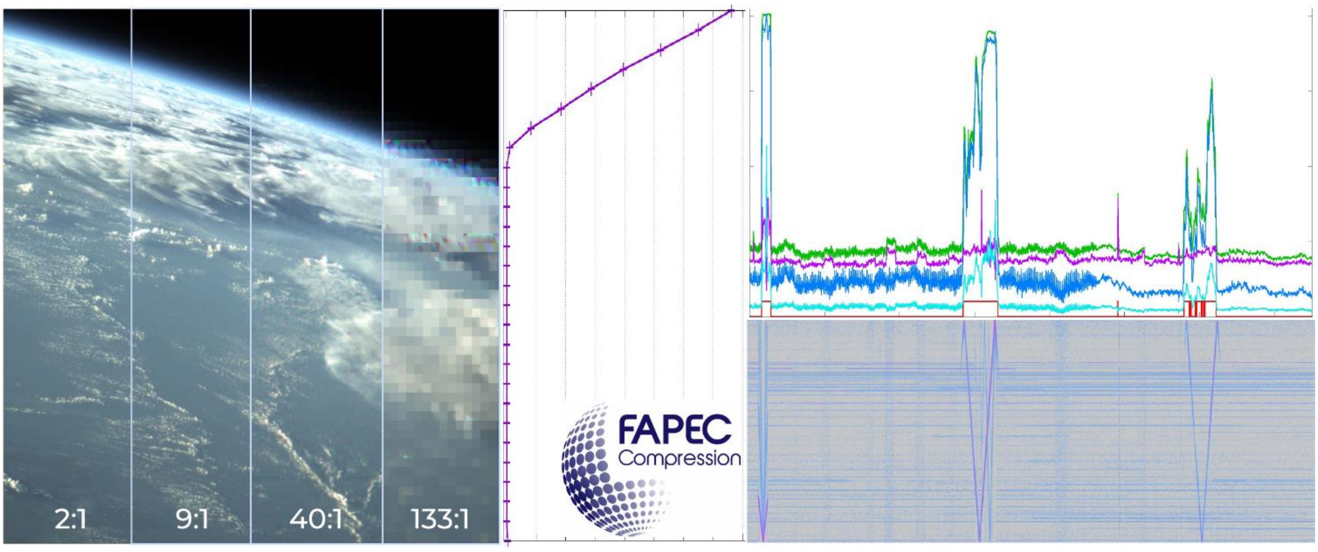

Versatile data compression software for sustained high-throughput in-orbit data acquisition

Executive summary

Activity type: Study

OPS-SAT OSIP Campaign

Prepared by: Jordi Portell
Affiliation: DAPCOM Data Services S.L. (Prime)

Activity summary:
We have developed a high-performance data compression software with embedded data analysis features for radio-frequency (RF), multi-band images and video. It offers excellent compression ratios and quality levels with very small computing requirements which even allow for real-time operation on modest processors. Tests on OPS-SAT data reveal image ratios over 1:10 with excellent visual quality, and RF ratios over 1:3 with nearly identical spectrograms thanks to a smart lossy approach. Demanding and innovative applications are now possible in cubesats, including continuous optical or radio-frequency monitoring, maximizing the amount of useful information downlinked.

→ THE EUROPEAN SPACE AGENCY

Publishing Date: 25 November 2022
Contract Number: 4000137290
Implemented as ESA Initial Support for Innovation

ESA Discovery & Preparation
From breakthrough ideas to mission feasibility. Discovery & Preparation is laying the groundwork for the future of space in Europe
Learn more on www.esa.int/discovery
Deliverables published on https://nebula.esa.int

1. INTRODUCTION AND SCOPE

This is the Executive Summary Report of the OPS-SAT OSIP Campaign Study for ESA, “Versatile data compression software for sustained high-throughput in-orbit data acquisition”, under ESA Contract No. 4000137290 with DAPCOM Data Services S.L. The activity is abbreviated RICSDAC (Radio-frequency and Image data Compression Software for Demanding Applications in Cubesats). The kick-off took place by teleconference on 3-March-2022. DAPCOM is a spin-off company of the Technical University of Catalonia (UPC) and the University of Barcelona (UB), founded in 2013.

The objective of this activity is to develop new data compression algorithms for some of the most important and sizeable data generated by nanosatellite payloads, focusing on the case of OPS-SAT. Specifically, we aim at multi-band images (including Bayer Colour Filter Arrays) and digital radio-frequency data. It is a quite ambitious and challenging study, given the complexity of the problem, the limited time and resources, and the many solutions already available – especially for image and video compression. Some difficulties have been found, which have been partly compensated by developing an added value. Namely, we have implemented basic data analysis techniques embedded in the same data compression software, allowing to generate tiny overviews of the data contents, enabling for on-board (autonomous) or on-ground (manual) decisions on which data files that are worth being downloaded. This, combined with the enhanced data compression results achieved, allows to maximize the amount of useful information retrieved through the limited downlink of nanosatellites such as OPS-SAT, furthermore with a small computational cost.

2. BACKGROUND

Data compression is a strongly recommended technique in most satellites, especially in cubesats, to maximize the net downlink capacity (information retrieved from remote sensing) and get the most from their payloads. Many efforts have been put on image data compression, given the continuously increasing amount of data, generated by sensors with increasing spatial resolution and spectral bands. There are several recommendations and standards, but they have some limitations. Radio-Frequency (RF) data compression has received much less attention despite of its high interest given the New Space paradigm, tightly associated with cubesats and, in many cases, with the use of Software Defined Radio (SDR) systems onboard. Fortunately, this kind of data have similarities with digital stereo audio data, for which many more solutions exist – although not all of them are applicable to RF data, and others have limitations as well.

The New Space paradigm is a new philosophy of space missions often based on the cubesat standard. One of these is OPS-SAT, a technology demonstrator of the European Space Agency (ESA). It features a high-resolution camera with a Bayer Colour Filter Array and a capability to acquire bursts of up to five frames per second. It also includes a Software Defined Radio (SDR) front-end and a patch antenna, able to acquire radio-frequency (RF) signals with 12-bit in-phase and quadrature samples (I/Q). OPS-SAT also has a powerful dual-core ARM microprocessor at a clock speed of 800 MHz and with 1 GB of RAM.

FAPEC is a data compression software provided by DAPCOM. It features a variety of pre-processing or decorrelation algorithms, followed by a high-performance adaptive entropy coding core, altogether supported by a multi-threaded framework providing several command-line interface (CLI) and application process interface (API) options to the user. It is already running in OPS-SAT since late 2020. One of the algorithms available is the so-called “Wave” decorrelator based on the LPC method. FAPEC provides image decorrelators as well, including CILLIC which provides lossless and lossy multi-band image compression. This activity intends to improve both algorithms and extend them with new functionalities.

3. CILLIC V2: MULTI-SPECTRAL IMAGE AND VIDEO COMPRESSION

We consider four main pillars for the improvement of FAPEC image compression: better spatial decorrelation, better spectral decorrelation, new motion decorrelation to support video compression, and better fixed-rate lossy option. The spatial decorrelator gives the name to the algorithm (Context Interpolation Lossless and Lossy Image Compressor), as it progressively builds the pixel references which are later used to predict a pixel from the interpolation of its neighbours. The spectral decorrelator essentially predicts a pixel from the previous band plus a common inter-band offset. The mixed decorrelator combines spatial and spectral prediction, estimating each pixel from the previous band plus the average inter-band offsets from the west and north pixels. In its near-lossless mode, residuals are quantized and the pixels are then reconstructed during compression to avoid error propagation.

We have increased the block size of CILLIC to 17×17 pixels, which lead to some benefits in the software implementation and in the prediction accuracy. We have also revised the interpolation for most of the pixels to avoid the accumulation of errors in high loss (high compression) modes, leading to higher quality images. All these changes should lead to better ratios, qualities and speeds, not only on OPS-SAT images but also on other multispectral or hyperspectral instruments. We have added a new feature to calculate image thumbnails (embedded in the compressed output file), which are simply the average pixel value per CILLIC block. With the objective of better compressing raw video files, we have implemented a motion estimator which can then be applied to a simplistic inter-frame decorrelator (based on the same inter-band algorithm). Inter-band decorrelation is done by predicting each pixel from the same pixel location in the previous band, plus an inter-band offset determined from the north and east pixels of both bands. It provides a fast operation and reasonably good compression performance. Given the interest in compressing video files, compression speed is of utmost importance. For this reason we have included the option to use FASEC, a fast and simple entropy coder, as the coding stage of CILLIC v2. With this, we can get significantly higher compression throughputs, yet at the cost of a modest reduction in the compression ratios.

Finally, we have revised the data analysis capabilities embedded in the FAPEC framework, generating a summary or statistics file with information on the image. Specifically, CILLIC determines the number of blocks in each chunk, the fraction of these handled by each of the decorrelator types (spatial or mixed), and also the fraction of blocks found to be *flat* – that is, with a smaller variability than the quantization step (in case of lossy compression). The latter is remarkably interesting to easily identify images with no relevant content, as well as images with large areas of clouds, sea or space.

The ESA OPS-SAT team has kindly provided us a rich set of test images, on which we have executed several tests: compression ratios in lossless, near-lossless and lossy mode; quality of near-lossless and lossy compressed images; data compression speed (evaluated on a Raspberry Pi 400 for now); and results of the embedded data analysis.

Typical lossless compression ratios are around 2.0 depending on the OPS-SAT scene, but these can be easily doubled without noticing any quality reduction. By selecting the adequate FAPEC settings, ratios above 10 can be reached while still keeping a good image quality. Remarkably, this algorithm is quick, achieving typical throughputs above 30 MB/s on the Pi 400 in single-thread mode. This figure can be boosted to almost 200 MB/s if we switch to the simplistic FASEC alternative, although the ratios and quality levels achieved are slightly worse.

In the following figure we show a small selection of images covering different situations. The small sizes used to embed them in this document makes it difficult to see the artefacts, so we just focus on quite extreme configurations:

CILLIC v2, lossless (ratios 3.07, 1.99, 1.98, 2.07 and 1.98, 45 MB/s single-thread on Raspberry Pi 400):

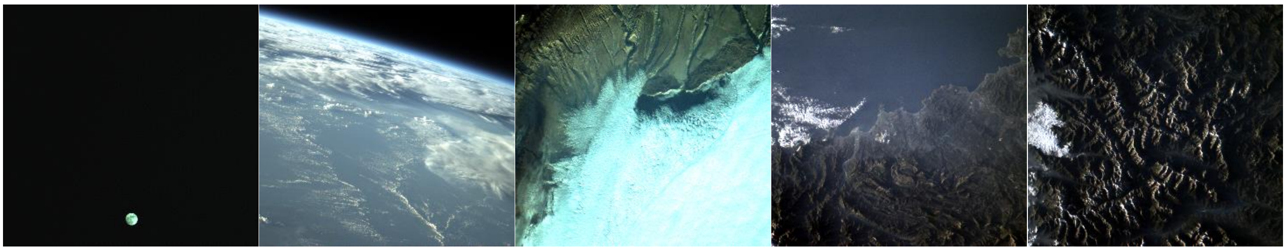

CILLIC v2, near-lossless level 14 (ratios 185.5, 40.0, 53.5, 68.4 and 38.9, 58 MB/s):

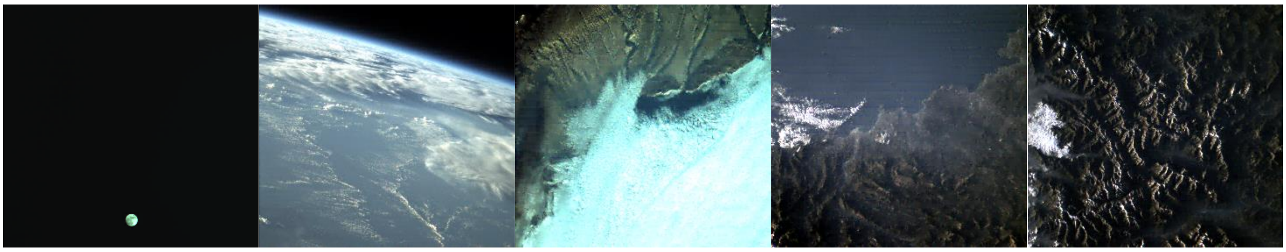

CILLIC v2, near-lossless level 20 (ratios 193.1, 133.4, 150.2, 148.8 and 125.9, 80 MB/s):

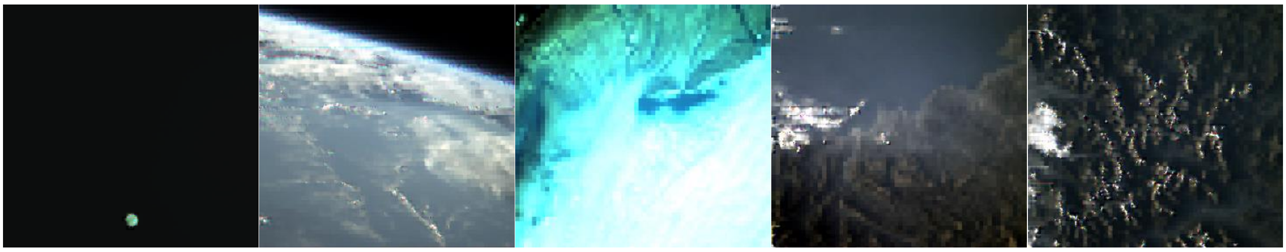

FASEC, near-lossless level 12 (ratios 38.5, 9.8, 9.6, 12.6 and 9.1, 180 MB/s):

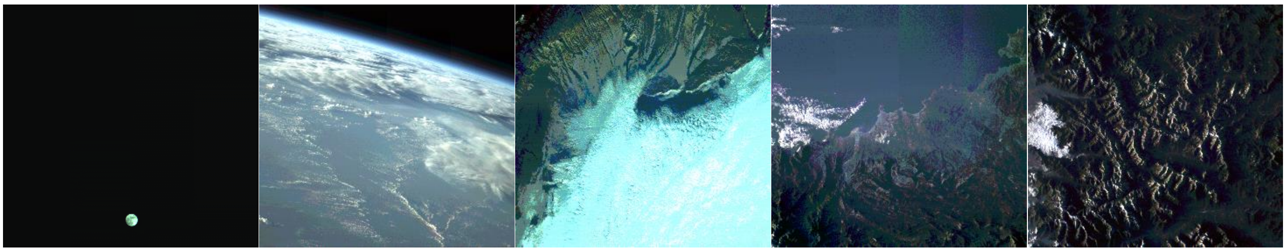

Regarding the data analysis capabilities, the following figure shows its outcome. Each of the panels show, in the horizontal axis, the percentage of flat blocks detected for each of the image sections. The vertical axis corresponds to the image strips, so these plots can be put side by side with their associated image above. Thus, a panel with the plot completely to the left means that no flat blocks were detected at all, whereas the points towards the right mean an increasing amount of flat blocks detected. For example, the first panel clearly indicates a "useless" image (100% flat blocks in nearly all chunks), except for a slight decrease towards the bottom of the image: it corresponds to the Moon, so it clearly demonstrates the correct operation and usefulness of this method. The second panel reveals the variation due to the Earth limb. Note that each text file created by FAPEC and used to generate each of these plots is just 3.3 KB (0.6 KB with *gzip*), compared to the 8 MB of the raw image data:

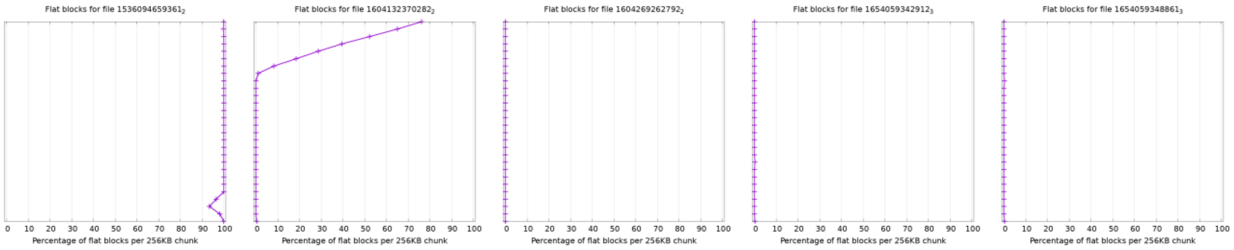

4. WAVE V2: RADIO-FREQUENCY DATA COMPRESSION

RF data compression is very challenging due to the huge variety in signal types, modulations, bandwidths, Doppler ranges, or channel multiplexity, just to put some examples. For our work, we have opted to focus just on OPS-SAT data files which provide a rich enough dataset covering several of these cases, even including spread spectrum signals (such as those from GNSS). Given this variety of scenarios, we will not aim at any demodulation-based approach, since it would restrict too much the applicability of our solution. Instead, we will consider observational signals in general, and thus will look for generic approaches to their compression.

We have analysed the data using different methods: inspection of their time series, autocorrelation, MSE (or MSD) at different lags, spectrum (FFT), and spectrograms. We have done this through Python libraries, simple C programs plus *Gnuplot*, and even with the *Audacity* audio program. The latter has appeared as an extremely useful tool, specially for its capability to easily generate spectrograms.

We have identified some interesting methods to decorrelate or transform IQ data (mainly RF but also observational). We have evaluated approaches like DWT-1D (already available in FAPEC), DCT and FFT, methods and software such as FLAC, etc. None of them have revealed a significant improvement compared to the Wave method of FAPEC. Thus, Linear Predictive Coding (LPC) seems to be the most convenient option considering our limitations, namely, a generic approach suitable for observational signals with a modest computational cost. We have made some updates to the algorithm, such as an adaptive LPC order or a larger FAPEC adaptive block, leading to a tiny increase in the ratios.

Additional methods have been investigated, such as signal and noise estimation, aiming at their use in a "smart lossy" option – allowing for an automatic adaptation of the quantization level depending on the noise level found. That is, basically, an adaptive near-lossless compression. In cubesats, SDR receivers can often generate data files mostly containing noise. In these cases, we considered that it would be useful to automatically detect the periods that actually contain some sort of signal, keeping a very low level of losses there (or even use the lossless mode), and significantly increasing the losses for the noise-only periods. However, signal detection is a challenging problem, especially considering the vast range of use cases of a satellite SDR receiver, including GNSS signals which are based on a spread spectrum approach and are hidden below the noise ground. We have followed and tested two approaches. First, a very simplistic one based on the prediction error already calculated by the Levinson-Durbin recursion and the accumulated energy from the autocorrelation and the LPC coefficients. This approach is further refined by providing user options to reduce the quantization aggressivity in periods with higher or lower SNR, as well as the SNR threshold to determine which periods are dominated by signal and which ones by noise. With this, different scenarios can be better handled, such as monitoring of parasitic signals, decoding of faint signals while masking louder parasitic signals, etc. The second approach is inspired in classical spectrum sensing problems and it involves a higher computational load, although it is probably acceptable given the capacity of the OPS-SAT SEPP. This second approach has been implemented as a library, “Spectra”, and published in GitHub. We have taken the results from this library as the reference to calibrate the first method (simpler and faster).

Similarly to the CILLIC case, these Wave improvements make possible a basic on-the-fly data analysis of SDR data files, allowing to generate short summaries of their contents which, for example, can aid in the decision to download each of the files generated on-board. Specifically, we generate small text files indicating the estimated level of signal or noise.

The ESA OPS-SAT team has also provided a rich set of RF data files generated by the onboard SDR from different frequency bands (433 MHz, 1.575 GHz and 1.602 GHz), acquired at 3 Msamples/s, a low-pass filter of 0.75 MHz, and a receiver gain of 60 dB or 66 dB depending on the case. Each file typically contains just 0.5 seconds of acquisition, except for some files which cover 2 seconds. RF data is more difficult to illustrate than images. We try to give a hint of their content by means of spectrograms (generated with the Audacity program): as usual, the horizontal axis corresponds to the time, and the vertical axis to the frequency. We have carried out tests on lossless, near-lossless and smart-lossy compression ratios, also checking the resulting quality by means of visual spectrogram inspection and PSNR values. Compression speed is completely satisfactory, reaching 20 MB/s on the Pi 400.

Tests reveal lossless ratios around 1.4-1.8 for the UHF files (which contain quite some strong signals) and around 2.0-3.7 for the GNSS signals. Near-lossless ratios can be significantly higher, with GNSS ratios of 3.9 to 8.9 depending on the case, and UHF ratios around 2.1 to 3.3. When using smart lossy, we can reach quite spectacular ratios of almost 50 in some cases, although the correct recovery of the signals therein should be carefully evaluated. Using a reasonably balanced configuration, we reach ratios of 3.3-4.6 in UHF files, and 2.8-5.2 in GNSS files. As can be seen, the benefit for UHF files is quite remarkable.

The following figure shows the spectrogram of one of these UHF files which presents very interesting features: first, there seem to be several horizontal lines, which probably correspond to different narrow-band communication channels. Then, we see several vertical lines, which may just correspond to interferences. Finally, we see very loud signals with a strong Doppler effect. We suspect that these correspond to communication signals from other satellites with an extremely high relative speed with respect to OPS-SAT (maybe from polar orbits, for example). Below this, we show the spectrogram for the same file after smart lossy compression, which increases the ratio from 1.5 (lossless) to 3.3. As can be seen, both spectrograms look identical, without any remarkable artifacts, so we think that this kind of lossy compression should not pose a problem when decoding it:

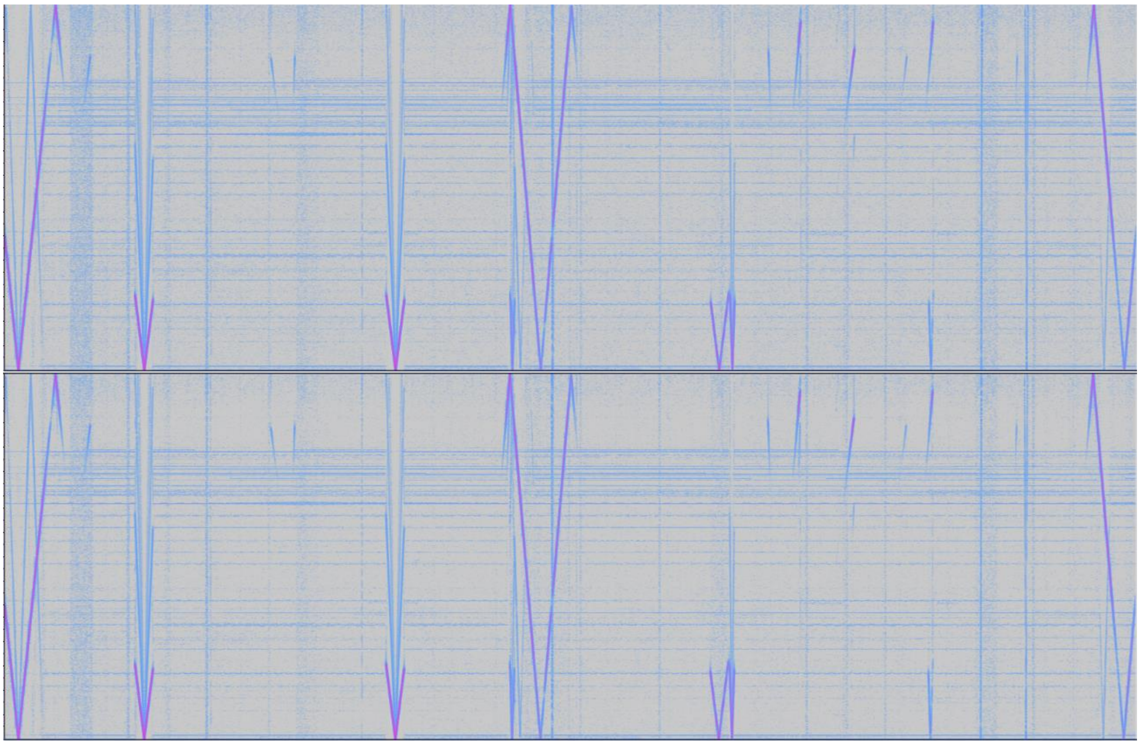

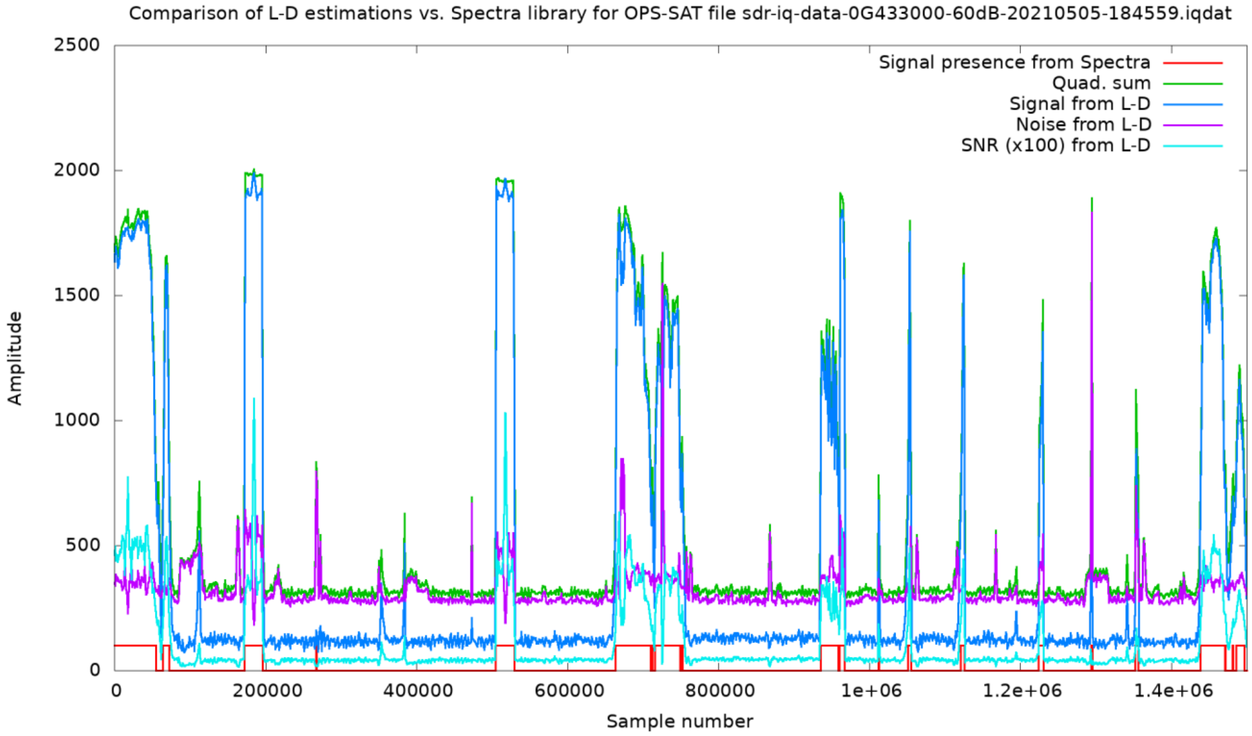

The plot below the spectrograms shows the estimators from the simplistic algorithm (that is, signal and noise levels), as well as their ratio (SNR). Additionally, it shows the signal presence result from the information-theoretic method (our Spectra library). As shown, both the SNR of the simplistic method and the Spectra signal presence are strongly correlated with the clear signal structures seen in the spectrogram. Perhaps the most interesting result is the strong correlation between the Spectra outputs and the SNR peaks: it means that the simplistic method should be quite reliable, thus offering a fast method to detect not only the presence of a signal, but also its level above the noise. Finally, it is worth mentioning that the text file created by FAPEC to generate these plots (with the signal and noise estimations) are just 11.1 KB for a 6 MB data file and 1024-samples periods, which can be further reduced (with *gzip*) to just 3.9 KB.

5. CONCLUSIONS AND FORTHCOMING WORK

In this work we have developed and implemented several improvements to the FAPEC data compression software for image and radio-frequency data. The algorithms are called CILLIC and Wave, respectively. They were already available before the start of this activity, but here we have carefully revised, updated and tested them. For the tests, we have used real in-orbit data from OPS-SAT, which has allowed us to perform an exhaustive analysis under a large variety of scenarios and compression configurations. The initial conclusion is that lossless ratios are very difficult to improve, both on images and RF data, so near-lossless and lossy methods have been carefully revised to provide the best possible compromise between ratios and quality.

On CILLIC, we have put many efforts on improving the spatial decorrelator due to its interest in Bayer cameras (such as OPS-SAT), where spatial-only decorrelation seems to perform better, as well as in certain cameras with artefacts or misregistration, and raw YUV-like videos which are essentially greyscale images. For video compression, the motion vectors are correctly determined, but its application does not significantly improve the ratios, perhaps due to the decorrelator being too simplistic. On this regard, we have revised the inter-band decorrelator, but it is clear that it requires significant updates to reach higher ratios. An approach based on mean square differences to determine the weights of neighbour and previous-band pixels for the prediction is very promising. The main improvements achieved in CILLIC are the capability to reach higher compression ratios while still keeping a good image quality. It is now also able to generate thumbnails, as well as tiny text files with the outcome of a basic data analysis (focused on the fraction of flat blocks detected). A very fast option for image and video compression has also been developed, offering lower ratios, a limited near-lossless quality, and no fixed-rate lossy option (at least for now). Pending onboard tests, it should be able to compress OPS-SAT videos in real-time.

Regarding the Wave algorithm, we have slightly improved its lossless performance, and especially, we have implemented a smart lossy algorithm able to detect the presence (or absence) of signals in a data file. With this, we can significantly improve the ratios in RF files, and also generate tiny summary files with the signals detection outcome. Spectrograms and PSNR figures reveal good quality levels when applying this approach.

As an extremely interesting and useful side product of this activity, we have created the *Spectra* library and published it on GitHub. It allows for a quite efficient spectral analysis of RF data files, including noise power estimation and signal detection. We have used it to adjust some of the parameters of the smart lossy options in Wave for the tests shown here.

The embedded data analysis features can be used to optimize the download of image and RF files, prioritizing those for which FAPEC has detected relevant contents. Thus, it can be seen as a “smart downlink” enabling technology, even allowing to automatize the file downlink decision onboard using simple scripts.

In general, FAPEC confirms to be an excellent option for cubesats, owing to its versatility, ease of integration, basic data analysis capabilities, excellent performance and good ratios under nearly any situation.

The following are some cases that may be enabled by the outcome of this activity:

- Continuous monitoring of RF signals, generating short-duration files (around one second) compressed in real-time by FAPEC using smart lossy, and checking the summary files to know which are worth being downloaded. We may just download those which show a clean signal without interferences, or even the contrary (in case we are just interested in studying the interferences and parasitic signals).
- Optical monitoring of potential sources of parasitic signals (e.g., observation of nearby satellites): depending on the orientation of the antenna and the camera, we may continuously monitor RF signals, and when the summary files reveal a strong signal, promptly trigger the acquisition of a picture (or several pictures), eventually analyzing their summary files and just download those with enough non-flat blocks.
- Continuous optical monitoring of space debris: with a camera orientation above the Earth limb (thus expecting just the dark sky), acquire pictures and compress them generating summary files. Then, we may detect those with some non-flat blocks, which may potentially have some trace of space debris.
- Long video for outreach or monitoring: Continuous video acquisition with real-time FASEC-based compression (if downlink allows), or just high-cadence pictures acquisition with high-ratio CILLIC compression (but still acceptable quality).