

BENCHMARKING REMOTE SENSING IMAGE PROCESSING AND MISSION PLANNING APPLICATIONS ON THE SNAPDRAGON PROCESSOR ONBOARD THE INTERNATIONAL SPACE STATION

Jason Swope¹, Faiz Mirza¹, Emily Dunkel¹, Alberto Candela¹, Zaid Towfic¹, Steve Chien¹, Damon Russell¹, Joe Sauvageau¹, Doug Sheldon¹, Mark Fernandez², and Carrie Knox²

¹*Jet Propulsion Laboratory*, California Institute of Technology, USA*

²*Hewlett Packard Enterprise*

ABSTRACT

Future space missions will process and analyze imagery onboard and plan and act more autonomously placing greater demands on flight computing. Traditional flight hardware provides modest compute, even when compared to laptop and desktop computers. A new generation of commercial off the shelf (COTS) processors, such as the Qualcomm Snapdragon, deliver significant compute in a small Size Weight and Power (SWaP) and offer direct hardware acceleration in the form of Graphics Processing Units (GPU) and Digital Signal Processors (DSP). We benchmark a variety of instrument processing and mission planning software on a Qualcomm Snapdragon SoC currently hosted by HPE's Spaceborne Computer-2 (SBC-2) onboard the International Space Station.

Key words: Deep Learning, Edge Processing, Space Applications, Machine Learning, Artificial Intelligence, COTS embedded processors.

1. INTRODUCTION

Future space missions will need more powerful onboard autonomy to meet mission and science objectives by (1) handling variations in predicted execution and rapid response to science events and (2) reducing the extremely large amounts of data produced by instruments, especially hyperspectral and radar. Traditional radiation hardened flight hardware provides only modest computing for future applications. A new generation of processors, such as the Qualcomm Snapdragon 855 [1] support onboard data processing via CPU, GPU and DSP - offering the promise of more powerful edge computing. The Snapdragon 855 was chosen for this study due to its good SWaP and prior flight experience with the Snapdragon 820 on the Perseverance rover.

We benchmark remote sensing image processing and analysis algorithms on a Snapdragon processor onboard the ISS hosted by Spaceborne Computer-2 by Hewlett Packard Enterprise [2]. Advancing these ground algorithms to embedded ISS deployment is a step towards running algorithms on a free flying spacecraft or surface mission to enable onboard data analysis, targeted downloads, commanding of space assets, and onboard science interpretation.

The Qualcomm Snapdragon 855 SoC has multiple sub-systems, including a CPU cluster with 8 ARM cores, an Adreno GPU, a Compute Digital Signal Processor (cDSP), and an AI Processor (AIP). The Snapdragon Neural Processing Unit (NPU) API will use other components to optimize deep learning classification tasks. See Figure 1 for a picture of a ground testbed Snapdragon board.

The CPU of the Snapdragon 855 has been benchmarked against the other flight hardware such as the GR740, RAD750, and Jetson Nano. The 8 ARM cores on the Snapdragon produce a total DMIPS average of 138,255 compared to 1,836 on the GR740 and 500 on the RAD750. The GPU (Adreno 640) produces 950 FP32 GFLOPS compared to 472 on a Nvidia Jetson Nano. The Snapdragon 855 CPU has maximum power consumption of 6W and the GPU a maximum power consumption of 3.5W [3].

Two Snapdragon 855 handheld development boards were integrated with the HPE Spaceborne Computer-2 (SBC-2) which provides connectivity, storage, and compute support. SBC-2 was launched as part of the ISS resupply mission Cygnus NG-15 on February 20th, 2021 and the Snapdragon boards have been operational since March 2021 with scores of applications tested as of December 2021.

We have previously reported on a number of tests on the Snapdragon [4] including other instrument processing algorithms such as match filters, decision trees, and hyperspectral unmixing. Other tests including dynamic targeting and classification were also discussed.

* (c) 2022 All rights reserved. Government sponsorship acknowledged.




Figure 1. Snapdragon 855 Development Board.

We are also working on extending the benchmarking to a range of processors available in a ground testbed including the Sabertooth [5], Rad750 [6], and Nvidia Jetson Nano [7] in a ground testbed.

2. APPLICATIONS AND BENCHMARKS

The applications tested are past, present, and proposed missions that were readily available. We benchmark a range of instrument processing algorithms with applications to terrestrial, planetary, astronomy, and astrophysics. We also benchmark mission planning applications involving satellite planning and scheduling the Mars Perseverance rover.

2.1. Instrument Processing

We report on a range of instrument processing algorithms below.

2.1.1. Ocean Worlds Life Surveyor(OWLS)

OWLS [8] is a project aimed at autonomously detecting signatures of life in water at the molecular and cellular scale. The algorithms process different instrument bands from OWLS instrument by creating autonomous data science products (ADSPs) to send back the most scientifically viable subsets of information back to scientists. The application is implemented as a single threaded ARM application on the Snapdragon. It runs over 24 images of size 1024x1024. It has a total run time of 22.5s and a per-image run time of .94s as shown in Table 1.

Application	Components Used	Runtime
MEXEC	CPU (single-threaded)	1.6s
Copilot	CPU (multi-threaded)	0.63s
CLASP	CPU (single threaded)	1,400s
OWLS	CPU (single-threaded)	22.5s
NEAS Lunar Flashlight	CPU (single-threaded)	494s
Thermal Anomaly Random Decision Forest (RDF)	CPU (single-threaded)	569s
Normalized Difference Index (NDI)	CPU (single-threaded)	56s
SAM Spectral Algorithm	CPU (single-threaded)	1.7s
MT Spectral Algorithm	CPU (single-threaded)	2.4s
Dynamic Targeting	CPU (single-threaded)	357s

Table 1. Snapdragon 855 Runtime Results

2.1.2. Near Earth Asteroid Scout

This application tests co-registration and co-addition image processing applications [9] for the Near Earth Asteroid Scout (NEA Scout) [10]. The runtime for this test is shown in Table 1.

2.1.3. Thermal Emission Random Decision Forest (RDF)

The thermal emission application uses RDF machine learning to detect large thermal emissions (typically from volcanic activity but also potentially from wildfires) in Planet Skysat images. The dataset consisted of 15 Skysat images of two different volcanoes. 12 of the images were of Fagradalsfjall, Iceland volcano 2021 eruption [11] and 3 were of Kilauea, Hawaii 2021 eruption [12]. Because of the high spatial resolution of the Skysat imaging sensor (better than 1m per pixel at nadir) each image is 13 million pixels. This dataset was split into a training and test set. The training set consisted of 12 images with 10 Fagradalsfjall images and two Kilauea images. The test set consisted of two Fagradalsfjall images and one Kilauea image. Labels for the location of lava in each image were determined through band ratios analyzed by volcanologist Ashley Davies. In order to train the RDF, 100 trees were trained on each image of the training set with weights equalized by class. This resulted in a 1,200 tree

RDF with a max depth of 4. The test on the Snapdragon predicts the labels of 4,000,000 pixels from the holdout set. The dataset used in the benchmark is smaller than the entire test set due to space limitations of the Snapdragon. The runtime for the thermal emission RDF is shown in Table 1

The thermal RDF application targets onboard thermal analysis (e.g. as performed onboard ASE/EO-1 [13]) and integrated into a volcano sensorweb (e.g. [14, 15]).

2.1.4. Normalized Difference Indices (NDI)

Onboard generation of science products for low latency downlink has been proposed for space missions [16] and has been demonstrated on the Earth Observing One mission [17] as well as on the IPEX cubesat technology demonstration mission [18]. Many of the target science products are normalized difference indices (e.g. normalized difference vegetation index, normalized difference snow and ice index, etc.). Direct Readout [19] has been deployed for MODIS and VIIRS for some time but these downlink moderate resolution multispectral data. Future proposals would be for higher spatial resolution (e.g. 30m instead of 250m per pixel) and potentially dramatically greater spectral information (e.g. 220 bands instead of 10 bands). This extension necessitates selectivity (either based on location or onboard analysis) to intelligently reduce data volumes to enable use of economical ground stations.

In this test we demonstrate that the COTS embedded processors can easily process very large volumes of data analysis and/or data reduction to support such low latency direct broadcast. The runtime for the NDI test is shown in Table 1.

2.1.5. Spectral Algorithms

Earth and planetary sciences often rely upon the analysis of spectroscopic data. Measured signals are called *spectra* and contain recognizable features or patterns that can be used for composition analysis because different materials reflect, emit, or absorb energy in unique ways throughout the electromagnetic spectrum.

Herein we benchmark two common algorithms for spectral analysis [20]: 1) the spectral angle mapper (SAM) and, 2) the matched filter (MF). Both algorithms use a spectral library containing objects of interest that are to be searched. SAM is a distance function between a spectrum and an object of interest. Specifically, SAM is the angle between two n-dimensional vectors (spectra). A perfect match is the one with an angle equal to 0, and a good match is the one with an angle below a given threshold. On the other hand, MF is a linear detector that requires background statistics: the mean and covariance matrix. A perfect match is the one with a score equal to 1, and a

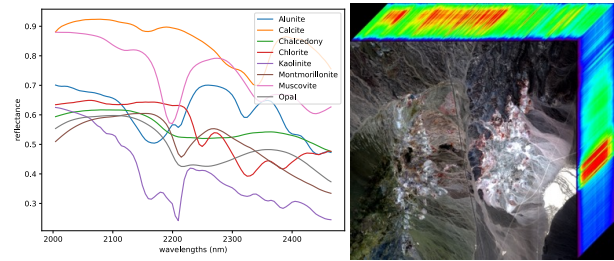


Figure 2. The dataset for benchmarking the spectral algorithms consists of a spectral library with 8 different minerals (left) and an AVIRIS-NG hyperspectral image of Cuprite Hills, Nevada.

good match is the one with a score above a given threshold.

Herein both algorithms were used to perform mineral detection (Figure 2). The spectral library consists of 8 minerals: alunite, calcite, chalcedony, chlorite, kaolinite, montmorillonite, muscovite, and opal. These minerals were searched on an Airborne Visible/Infrared Imaging Spectrometer Next Generation (AVIRIS-NG) [21] hyperspectral image covering the western portion of Cuprite Hills, Nevada; a well-studied region that is amenable to remote sensing and has a high mineralogical diversity [22]. The Snapdragon test benchmarks both algorithms over 20,000 pixels of the AVIRIS-NG dataset. Each pixel contains 425 spectral bands (350-2500 nm), of which 93 (2000-2500 nm) are used by each algorithm. The SAM algorithm finished its analysis 1.7s and the MF algorithm took slightly longer at 2.4s as shown in Table 1.

2.2. Mission Planning Applications

We also test a range of onboard autonomous response software, specifically mission planning software.

2.2.1. Multimission Executive (MEXEC)

MEXEC [23] is an application which contains a separately threaded planner and executive. The application works by taking in a "task network" and then generating conflict free plans and monitoring the execution of those plans. It also responds to deviations from the plan or external events as they arise.

Because the (re) planning portion of Mexec is the computationally demanding element, we test that capability in our benchmark. To simplify the test, we generate a single plan (where Mexec invokes the planner repeatedly when execution varies from predicted). We utilize a Europa Lander Prototype test scenario [24]. Running as a single threaded ARM application on the Snapdragon the planner generated a multi-day schedule using hierarchical planning, valid interval search, and constraint satisfaction.

Overall, one plan was able to be generated in 1.6 seconds as shown in Table 1.

2.2.2. Copilot / Mars 2020

Copilot [25] is the M2020 mission ground scheduler and is currently in use to schedule wake/sleep and preheats for the M2020 Perseverance rover operations. It uses the same scheduling algorithms as the M2020 onboard scheduler, and it works to solve wake/sleep constraints, preheat constraints, variability in execution, and complex operations handover handlings.

This test runs with 800 x 1 martian day (sol) planning problems that are generated from random variations of 7 base plans. These plans vary execution durations, incoming/outgoing energy state, and alternative action options. The test has a single threaded and multi-threaded version. During testing the single threaded version computed the 800 plans in 1200 seconds while the multi-threaded version completed the tasks in 500 seconds as shown in Table 1.

2.2.3. Compressed Large-scale Activity Scheduler and Planner (CLASP)

CLASP has been used on several missions such as NISAR [26, 27], ECOSTRESS [28], OCO-3 [29] EMIT [30], and other missions. The planner takes in scientific goals from scientists, and generates an observation schedule for a single or constellation of satellites based on scenario constraints.

This benchmark generates two years of two week schedules using ECOSTRESS data from 2018-2020. We generate a single 2 week schedule to collect the timing metric faster. This test displays the runtime of a single threaded CPU version run on the Snapdragon ARM. A version was created to investigate running portions of the application on the GPU, however, no speedup was found. CLASP runs in 1400 seconds as shown in Table 1.

2.2.4. Dynamic Targeting

This effort includes a simulation study consisting of an Earth science satellite whose mission is to analyze storm clouds. The satellite has a primary radar with a narrow swath, and a secondary radar with a wider field of view that can only be used for lookahead. The radar is able to operate with a duty cycle of 20%. General Mission Analysis Tool (GMAT) was used to simulate a realistic satellite trajectory. This simulation is run over a low Earth orbit with 65 degrees of inclination, 400km altitude, 95 minute periods, and an eccentricity of 0. The experiment consists of 18,000 timesteps at 2 seconds per time step. This equal 10 hours of simulation time. Over the simulation four different dynamic targeting algorithms are run. This

simulation takes 357 seconds to run on the Snapdragon CPU as shown in Table 1.

3. RELATED AND FUTURE WORK

This effort is in progress and therefore we are still preparing additional applications for validation on the Snapdragon 855 within the SBC-2 on the ISS. Deep learning applications are being tested on the Snapdragon 855 as well as the Intel Movidius [Dunkel 2022] and NVIDIA Jetson NANO [7]. An earlier version of the Intel Movidius chip flew on the ESA Phisat mission testing cloud detection [Giufrida et al. 2021]. Several efforts to develop advanced flight computing are ongoing such as [Goodwill et al. 2021]. Another important part of this effort is to also evaluate and compare these applications on conventional flight hardware such as the LEON4 based Sabertooth [Whitaker 2019] and RAD 750 [RAD750] as well as Linux ground-based computing.

4. CONCLUSIONS

Future space missions will use onboard autonomy to address: (1) intermittent contacts between earth and spacecraft and (2) very large amounts of data produced by instruments, especially hyperspectral and radar. A new generation of processors, such as the Qualcomm Snapdragon 855 offer the promise of more powerful edge computing via both conventional CPU as well as GPU and DSP.

We benchmark remote sensing image processing, analysis algorithms, and mission planning on a Snapdragon 855 processor onboard the ISS hosted by Spaceborne Computer-2 by Hewlett Packard Enterprise. Embedded ISS deployment is a step towards running these algorithms on spacecraft, landers, and rovers - to enable onboard data analysis, targeted downloads, commanding of space assets, and onboard science interpretation.

5. ACKNOWLEDGEMENTS

Portions of this work were performed by the Jet Propulsion Laboratory, California Institute of Technology, under contract from the National Aeronautics and Space Administration.

REFERENCES

- [1] Qualcomm. Snapdragon 855 mobile platform. <https://www.qualcomm.com/products/snapdragon-855-mobile-platform>.

- [2] HPE. Hpe spaceborne computer. <https://www.hpe.com/us/en/compute/hpc/supercomputing/spaceborne.html>.
- [3] Z. Towfic, D. Ogbe, J. Sauvageau, D. Sheldon, A. Jongeling, S. Chien, F. Mirza, E. Dunkel, J. Swope, V. Cretu, and M. Ogut. Benchmarking and testing of qualcomm snapdragon system-on-chip for jpl space applications and missions. *IEEE Aerospace Conf*, 2022.
- [4] Jason Swope, Faiz Mirza, Emily Dunkel, Zaid Towfic, Damon Russell, Joe Sauvageau, Doug Sheldon, Steve Chien, Mark Fernandez, and Carrie Knox. Benchmarking remote sensing image processing and analysis on the snapdragon processor onboard the international space station. In *International Geoscience and Remote Sensing Symposium (IGARSS 2022)*, Kuala Lumpur, Malaysia, July 2022.
- [5] William Whitaker. Sabertooth: Integrated avionics for small spacecraft missions. In *IEEE Space Computing Conference (SCC)*. IEEE, 2019. URL <https://trs.jpl.nasa.gov/bitstream/handle/2014/51550/CL%2319-4553.pdf?sequence=1&isAllowed=y>.
- [6] Rad750. <https://en.wikipedia.org/wiki/RAD750>.
- [7] NVIDIA Jetson Nano. <https://developer.nvidia.com/embedded/jetson-nano-developer-kit>.
- [8] Ocean Worlds Life Surveyor Web Site. <https://ml.jpl.nasa.gov/projects/owls/owls.html>.
- [9] Jack Lightholder, David R Thompson, Julie Castillo-Rogez, and Christophe Basset. Near earth asteroid scout cubesat science data retrieval optimization using onboard data analysis. In *2019 IEEE Aerospace Conference*, pages 1–7. IEEE, 2019.
- [10] Tiffany Russell Lockett, Julie Castillo-Rogez, Les Johnson, Joe Matus, Jack Lightholder, Anne Mariman, and Alexander Few. Near-earth asteroid scout flight mission. *IEEE Aerospace and Electronic Systems Magazine*, 35(3):20–29, 2020. doi: 10.1109/MAES.2019.2958729.
- [11] Stunning images of Fagradalsfjall volcano in Iceland, Earthsky.org. <https://earthsky.org/earth/fagradalsfjall-volcano-eruption-images-march-2021-iceland/>
- [12] Kilauea Setepmber 2021 Eruption, United States Geological Survey. <https://www.usgs.gov/volcanoes/kilauea/recent-eruption#:~:text=K%C4%ABlauea%20volcano%20began%20erupting%20on,Hawai'i%20Volcanoes%20National%20Park>.
- [13] A. Davies, S. Chien, V. Baker, T. Doggett, J. Dohm, R. Greeley, F. Ip, R. Castano, B. Cichy, R. Lee, G. Rabideau, D. Tran, and R. Sherwood. Monitoring active volcanism with the autonomous sciencecraft experiment (ase). *Remote Sensing of Environment*, 101 (4):427–446, April 2006.
- [14] Steve A. Chien, Ashley G. Davies, Joshua Doubleday, Daniel Q. Tran, David McLaren, Wayne Chi, and Adrien Maillard. Automated volcano monitoring using multiple space and ground sensors. *Journal of Aerospace Information Systems (JAIS)*, 17:4: 214–228, 2020. URL <https://doi.org/10.2514/1.I010798>.
- [15] S. Chien, J. Boerkoel, J. Mason, D. Wang, A. G. Davies, J. Mueting, V. Vittaldev, V. Shah, and I. Zuleta. Leveraging space and ground assets in a sensorweb for scientific monitoring: Early results and opportunities for the future. In *IEEE Geoscience and Remote Sensing Symposium (IGARSS)*, September 2020.
- [16] S. Chien, D. Silverman, A. Davies, and D. Mandl. Onboard science processing concepts for the hypsiri mission. *IEEE Intelligent Systems*, 24 (6), November/December 2009.
- [17] S. Chien, D. McLaren, D. Tran, A. G. Davies, J. Doubleday, and D. Mandl. Onboard product generation on earth observing one: A pathfinder for the proposed hypsiri mission intelligent payload module. *IEEE JSTARS Special Issue on the Earth Observing One (EO-1) Satellite Mission: Over a decade in space*, 2013.
- [18] S. Chien, J. Doubleday, D. R. Thompson, K. Wagstaff, J. Bellardo, C. Francis, E. Baumgarten, A. Williams, E. Yee, E. Stanton, and J. Piug-Suari. Onboard autonomy on the intelligent payload experiment (ipex) cubesat mission. *Journal of Aerospace Information Systems (JAIS)*, April 2016.
- [19] Direct Readout Web Site. <https://directreadout.sci.gsfc.nasa.gov/>.
- [20]
- [21] L. Hamlin, R. O. Green, P. Mouroulis, M. Eastwood, D. Wilson, M. Dudik, and C. Paine. Imaging spectrometer science measurements for terrestrial ecology: Aviris and new developments. In *2011 Aerospace Conference*, pages 1–7, 2011. doi: 10.1109/AERO.2011.5747395.
- [22] Gregg A. Swayze, Roger N. Clark, Alexander F.H. Goetz, K. Eric Livo, George N. Breit, Fred A. Kruse, Stephen J. Sutley, Lawrence W. Snee, Heather A. Lowers, James L. Post, Roger E. Stofregén, and Roger P. Ashley. Mapping Advanced Argillic Alteration at Cuprite, Nevada, Using Imaging Spectroscopy. *Economic Geology*, 109(5): 1179–1221, 08 2014.
- [23] M. Troesch, F. Mirza, K. Hughes, A. Rothstein-Dowden, R. Bocchino, A. Donner, M. Feather, B. Smith, L. Fesq, B. Barker, and B. Campuzano. Mexec: An onboard integrated planning and execution approach for spacecraft commanding. In *Workshop on Integrated Execution (IntEx)*

- / *Goal Reasoning (GR), International Conference on Automated Planning and Scheduling (ICAPS IntEx/GP 2020)*, October 2020. URL https://ai.jpl.nasa.gov/public/papers/IntEx_2020_MEEXEC.pdf. Also presented at International Symposium on Artificial Intelligence, Robotics, and Automation for Space (i-SAIRAS 2020) and appears as an abstract.
- [24] D. Wang, J. A. Russino, C. Basich, and S. Chien. Using flexible execution, replanning, and model parameter updates to address environmental uncertainty for a planetary lander. October 2020. URL https://ai.jpl.nasa.gov/public/papers/europa_lander_icaps2020_workshop.pdf. Also presented at Proceedings of the International Symposium on Artificial Intelligence, Robotics and Automation for Space. Also presented at Proceedings of ESA-CLAIRE Workshop on Space and AI, European Conference on Artificial Intelligence. Also presented at the Planning and Robotics (PlanRob) Workshop, ICAPS.
- [25] A. Yelamanchili, J. Agrawal, S. Chien, J. Biehl, A. Connell, U. Guduri, J. Hazelrig, I. Ip, K. Maxwell, K. Steadman, and S. Towey. Ground-based automated scheduling for operations of the mars 2020 rover mission. In *Proceedings Space Operations 2021*, May 2021. URL <https://spaceops.iafastro.directory/a/proceedings/SpaceOps-2021/SpaceOps-2021/6/manuscripts/SpaceOps-2021,6,x1385.pdf>.
- [26] Joshua R. Doubleday. Three petabytes or bust: Planning science observations for nisar. In *SPIE 9881*, New Delhi, India, May 2016. doi: 10.1117/12.2223893. URL https://ai.jpl.nasa.gov/public/papers/doubleday_spie2016_petabytes.pdf.
- [27] S. Chien, A. Yelamanchili, and J. Doubleday. Policy-based automated science coverage scheduling for earth science mission analysis and operations (nisar, ecostress, oco-3, and emit). In *Earth Science Technology Forum (ESTF 2019)*, Moffett Field, California, USA, June 2019. URL https://ai.jpl.nasa.gov/public/presentations/chien_estf2019_overview.pdf.
- [28] A. Yelamanchili, S. Chien, K. Cawse-Nicholson, and D. Freeborn. Scheduling and operations of the ecostress mission. In *Proceedings Space Operations 2021*, May 2021. URL <https://spaceops.iafastro.directory/a/proceedings/SpaceOps-2021/SpaceOps-2021/6/manuscripts/SpaceOps-2021,6,x1381.pdf>.
- [29] A. Yelamanchili, C. Wells, S. Chien, A. Eldering, R. Pavlick, C. Cheng, R. Schneider, and A. Moy. Scheduling and operations of the orbiting carbon observatory-3 mission. In *Proceedings Space Operations 2021*, May 2021. URL <https://spaceops.iafastro.directory/a/proceedings/SpaceOps-2021/SpaceOps-2021/6/manuscripts/SpaceOps-2021,6,x1382.pdf>.
- [30] A. Yelamanchili, C. Wells, S. Chien, J. Russino, R. Green, B. Oaida, and D. R. Thompson. Using automated scheduling for mission design: A case study for emit. In *Proceedings Space Operations 2021*, May 2021. URL <https://spaceops.iafastro.directory/a/proceedings/SpaceOps-2021/SpaceOps-2021/13/manuscripts/SpaceOps-2021,13,x1383.pdf>.