BENCHMARKING REMOTE SENSING IMAGE PROCESSING AND ANALYSIS ON THE SNAPDRAGON PROCESSOR ONBOARD THE INTERNATIONAL SPACE STATION

Jason Swope[1], Faiz Mirza[1], Emily Dunkel[1], Zaid Towfic[1], Steve Chien[1], Damon Russell[1], Joe Sauvageau[1], Doug Sheldon[1], Mark Fernandez[2], Carrie Knox[2]

[1]Jet Propulsion Laboratory, California Institute of Technology, USA
[2]Hewlett Packard Enterprise

ABSTRACT

Future space missions will process and analyze imagery onboard 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 Qualcomm Snapdragon, deliver significant compute in 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 analysis software (including machine learned classifiers) on a Qualcomm Snapdragon SoC currently hosted by HPE's Spaceborne Computer-2 (SBC-2) onboard the International Space Station.

Index Terms— *Edge processing, Space Applications, Machine Learning, Artificial Intelligence*

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 [Qualcomm 2021] 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 [HPE 2021]. Advancing these ground algorithms to embedded ISS deployment is a step towards running algorithms on a satellite or Mars Rover, to enable onboard data analysis, targeted downloads, commanding of space assets, and onboard science interpretation.

The Qualcomm Snapdragon 855 SoC has multiple subsystems, 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.

The CPU of the Snapdragon 855 has been benchmarked against the other flight hardware such as the GR740 and RAD750. 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 energy consumption of the Snapdragon 855 CPU has a max of 6W and the GPU takes a max of 3.5W [Towfic et al. 2022].

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.

2. APPLICATIONS AND BENCHMARKS

The applications chosen were readily available past, present, and proposed missions. We benchmark a range of instrument processing algorithms with applications to terrestrial, planetary, and astronomy and astrophysics. We also benchmark instrument targeting algorithms for the "Dynamic Targeting" concept in which a lookahead sensor is used to identify targets (e.g. convective storms) [Swope et al. 2021] or avoidances (e.g. clouds [Hasnain et al. 2021] to inform targeting and configuration of a primary sensor.

3.1. Instrument Processing

We report on a range of instrument processing algorithms below.

Copyright © 2022, California Institute of Technology. U.S. Government sponsorship acknowledged.

Decision Trees. Manual decision tree of modest size (~10 nodes) for per pixel classification of thermal anomalies (volcanic or wildfire) [Davies et al. 2006] or cryosphere (snow water, ice, land) [Doggett et al. 2005]. The benchmark runs one image through both the thermal and cryosphere decision trees. CPU and GPU implementation runtimes are reported.

Application	Components Used	Runtime
Decision Trees (Thermal + Cryosphere)	CPU, GPU	CPU: 21s GPU: 13s
Synthetic Aperture Radar (SAR) Image Formation	CPU+GPU	217s
Match Filters (Cuprite)	CPU	850s
Match Filters (lunar)	CPU	108.4s
Hyperspectral Compression	CPU, GPU, DSP	See Table 2
Hyperspectral unmixing (SMACC)	CPU	16.9s
SMICES Classification	CPU	See Table 3
Salience Detector	CPU	23s
Landing Vision System	CPU	COARSE: 2.46s FINE: 2s
HOWFS	CPU	2.2h, 1.8h
Europa Lander Stereo Vision	CPU	19s - 15.6 min

Table 1: Image Processing Applications

Synthetic Aperture Radar (SAR) Image Formation. Pipeline of 3 CPU and 2 GPU applications adapted from the Uninhabited Aerial Vehicle SAR (UAVSAR) project [Hawkins and Tung 2019]. Primarily a row wise and column wise 2D FFT with filters. Image size is 27916 x 26880 and takes 217s where performance goal is 240s which corresponds to keeping up with instrument acquisition rate. Performance likely could be further improved as the GPU utilization is approximately 60%.

Match Filters. CPU application of signature detection for remote sensing imaging spectroscopy [Thompson et al. 2015]. First test is mineral detection at Cuprite, NV site using AVIRIS 2014 data. Most of the runtime is I/O. For an image size of 670x2512x425(spectra), 8 images, one mineral signature is 850s on the Snapdragon CPU. For a second lunar dataset using M3 data searching for several forms of water (OH, molecular H_2O, and H_2O ice) with an image size of 304 x 1000 x 301 produces a runtime of 108.4s.

Hyperspectral Compression [Hernandez-Cabronero et al. 2021]. This test performs lossless compression on a set of test images for the Earth Surface Mineral Dust Source Investigation (EMIT) mission [EMIT 2021]. Each image is 64 lines x 640 samples per line with 481 spectral bands per pixel. CPU, DSP, and GPU ports are reported. MSamples/sec = lines x samples per line x bands / runtime and the EMIT Target is 23.1 MSamples/sec (instrument throughput). See runtimes in table below.

	CPU	GPU	DSP	Virtex-5	GTX 580
CCSDS Standard	123.0.B-2	123.0.B-1	123.0.B-2	123.0.B-1	123.0.B-1
Compression	Lossless				
Runtime (ns)	14.12	6.5	184.5	25	16.18
Sample Rate (MSamples/sec)	70.82	153.85	5.42	40.00	61.80
Power (W)	6.1	3.5	1.9	2	>100

Table 2: Performance of Hyperspectral Compression on various Hardware Configurations

Hyperspectral. This implements Sequential Maximum Angle Convex Cone (SMACC) spectral endmember extraction as previously flown on the Earth Observing One spacecraft [Thompson et al. 2012]. This is a CPU implementation that extracts the top 5 endmembers from the AVIRIS-NG Data. The image size is 638 x 679 x 425 spectra and the runtime is 16.9 seconds (compared to 6h onboard EO-1).

SMICES Classification. SMICES is an instrument concept for a “smart” deep ice convective storm hunting radar [NASA 2019, Swope et al. 2021, Bosch-Lluis et al. 2021]. In the SMICES concept a lookahead radiometer acquires data to detect deep convective ice storms and a radar is used to study detected storms in greater detail. The SMICES machine learning classification application [Chien et al. 2021] classifies simulated radiometer data into five separate cloud types (clear, thin cirrus, cirrus, rainy anvil, and convection core) to identify the location of the deep convective storms. The application runs a random decision forest (RDF), multi-layer perceptron (MLP), simple vector machine (SVM), and naïve Bayes Gaussian classifiers over 198,016 pixels with 8 bands of radiance. Each classifier is run on the Snapdragon CPU in a single threaded python application. The runtimes for each classifier are listed below.

Classifier	RDF	MLP	SVM	Bayes
Runtime (s)	0.5	0.55	1316.7	.27

Table 3: Snapdragon 855 CPU runtime for SMICES Classifiers

Salience Detector. CPU application for generating salience maps of large image swaths. This detector can be used to reduce data volume from imagery obtained by the High

Resolution Imaging Experiment (HiRISE) instrument onboard the Mars Reconnaissance Orbiter. The salience detector uses computer vision techniques, with parameters optimized by a genetic algorithm, to detect landmarks of potential interest to planetary scientists [Wagstaff et al. 2021]. Runtime shown is for a single grayscale image of size 2048 x 4032 pixels.

Landing Vision System (Astrotipping). This is a descent landing system that estimates cross velocity based on descent imagery and uses this to both localize and avoid landing hazards. The baseline system uses a hybrid FPGA+ CPU system. The problem is divided into coarse landmark matching and fine landmark matching phases. In the COARSE phase a 1024 x 1024 image is warped and matched to a template via FFT, taking 2.46s on the Snapdragon CPU. In the FINE phase a 1024 x 1024 image is processed using normalized cross correlation, taking 2s. Work is ongoing to improve this application using pre-optimized functions provided by the FastCV library [Qualcomm, 2021].

Higher Order Wavefront Sensing (HOWFS). Proposed onboard image processing for the Roman Space Telescope Coronagraph Instrument [Krist et al. 2018]. Currently single threaded Python port to CPU processor. Double precision calculation takes 2.2 hours, single precision 1.8 hours. Moving to the GPU, multi-threading, and porting from python to C++ would all improve performance significantly. In context, a similar comparison for the CGI instrument could not be accomplished with existing compute platforms, so computation necessitated a ground-in-the-loop implementation with a 100-hour requirement (due to the telecommunication logistics involved).

Europa Lander Stereo Vision. Benchmark for stereo depth extraction from a pair of images, targeting the Europa Lander Mission Concept [Europa 2021]. This application runs on 24 image pairs, with each image natively 5120x3840 resolution. The benchmark is run at three different image resolutions (runtimes in parentheses): 5120x3840 (15.6 minutes), 2560x1920 (2.1 minutes), and 1280x960 (19 s).

3.1. Targeting Remote Sensing Instruments

Another onboard autonomy application is “Dynamic Targeting” in which a lookahead sensor is used to provide information to an algorithm that dynamically targets the primary sensor and also may turn on/off or configure the primary sensor to optimize science. We have tested two such applications on the Snapdragon CPU and are working on more such applications.

SMICES. SMICES is an instrument concept for a “smart” deep ice convective storm hunting radar (described above). The SMICES pointing planning application plans out the targeting (pointing) and operations of the radar considering the detected storms and observing according to the provided science policy while managing energy constraints. In this test the instrument planner operates over 15,232 timesteps where 1 timestep = 2 seconds for a total of approximately ~8.5 hours CPU time. This takes 53.6 seconds to run on the Snapdragon CPU in a single threaded python application.

Cloud Avoidance. This is an operations concept in which a lookahead sensor is used to detect clouds and imaging is directed to target around clouds improving data quality. This technique is already being used in the TANSO-FTS instrument [L3Harris 2020]. We benchmark an application [Hasnain et al. 2021] re-implemented in Rust for the Snapdragon CPU that simulates such tasking based on historical MODIS cloud masks. The application runs over 50 images with each size 1354 x 2030 and the range of search methods runtimes (mean milliseconds per image) are: Greedy search 49.7 ms, Adaptive Grid Graph Search 199.2 ms, Mixed Grids Graph Search 13.7 ms, and Fixed Grid Graph Search 9.7 ms.

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. 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) time delay 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 and analysis algorithms 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.

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

Bosch-Lluis, X.; Chien, S.; Q.Yue; Swope, J.; P., P. T.; Ogut, M.; Ramos, I.; Kangaslahti, P.; Deal, W.; and Cooke, C. Developing Radiometer and Radar synergies using Machine Learning. In Intl Geoscience and Remote Sensing Symposium, July 2021.

Chien, S.; Swope, J.; Yue, Q.; Lluis-Bosch, J.; and Deal, W. Using a Digital Twin Weather Research and Forecasting (WRF) Model for Machine Learning of Deep Convective Ice Storms. Proc Fall Meeting of the American Geophysical Union, New Orleans, LA, USA, 2021. American Geophysical Union

Davies, A.; Chien, S.; Baker, V.; Doggett, T.; Dohm, J.; Greeley, R.; Ip, F.; Castano, R.; Cichy, B.; Lee, R.; Rabideau, G.; Tran, D.; and Sherwood, R. Monitoring Active Volcanism with the Autonomous Sciencecraft Experiment (ASE). Remote Sensing of Environment, 101 (4): 427-446. April 2006.

Doggett, T.; Greeley, R.; Davies, A.; Chien, S.; Cichy, B.; Castano, R.; Williams, K.; Baker, V.; Dohm, J.; and Ip, F. Autonomous On-Board Detection of Cryospheric Change. Remote Sensing of Environment, 101 (4): 447-462. 2005.

E. Dunkel, J. Swope, Z. Towfic, S. Chien, D. Russell, J. Sauvageau, D. Sheldon, J. Romero-Canas, J. L. Espinosa-Aranda, L. Buckley, E. Hervas-Martin, M. Fernandez, C. Knox, "Benchmarking Deep Learning Inference of Remote Sensing Imagery on The Qualcomm Snapdragon and Intel Movidius Myriad X Processors Onboard the International Space Station, Intl Geoscience and Remote Sensing Symposium, 2022, under review.

Europa Lander Mission Concept
https://www.jpl.nasa.gov/missions/europa-lander

Giuffrida, G, et al. The Φ-Sat-1 mission: the first on-board deep neural network demonstrator for satellite earth observation. IEEE Transactions on Geoscience and Remote Sensing, 2021.

J. Goodwill, G. Crum, J. MacKinnon, C. Brewer, M. Monaghan, T. Wise, C. Wilson, "NASA SpaceCube Edge TPU SmallSat Card for Autonomous Operations and Onboard Science-Data Analysis," In Proceedings of the Small Satellite Conference 2021 Aug 7 (No. SSC21-VII-08). AIAA.

Hasnain, Z.; Mason, J.; Swope, J.; Vander Hook, J.; and Chien, S. Agile Spacecraft Imaging Algorithm Comparison for Earth Science. Intl Wkshp Planning & Sched Space (IWPSS), July 2021.

Hawkins BP, Tung W. UAVSAR Real-Time Embedded GPU Processor. IEEE International Geoscience and Remote Sensing Symposium (IGARSS) 2019 Jul 28 (pp. 545-547). IEEE.

Hernandez-Cabronero M, Kiely AB, Klimesh M, Blanes I, Ligo J, Magli E, Serra-Sagrista J. The CCSDS 123.0-B-2 Low-Complexity Lossless and Near-Lossless Multispectral and Hyperspectral Image Compression Standard: A comprehensive review. IEEE Geoscience and Remote Sensing Magazine. 2021 Feb 10.

HPE, "HPE Spaceborne Computer," 2021, https://www.hpe.com/us/en/compute/hpc/supercomputing/spaceborne.html

J. Krist, R. Effinger, B. Kern, M. Mandic, J. McGuire, D. Moody, P. Morrissey, I. Poberezhskiy, A. J. Riggs, N. Saini, E. Sidick, H. Tang, and J. Trauger, "WFIRST coronagraph flight performance modeling," in Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave, M. Lystrup, H. A. MacEwen, G. G. Fazio, N. Batalha, N. Siegler, and E. C. Tong, Eds., vol. 10698, International Society for Optics and Photonics. SPIE, 2018, pp. 788 – 810.

L3Harris, TANSO-FTS Data Sheet, retrieved https://www.l3harris.com/sites/default/files/2020-08/sas-ss%26g-tanso-sell-sheet.pdf

NASA, Advanced Information Systems Technology 2019 awards. Retrieved 20 December 2021. https://esto.nasa.gov/files/solicitations/IIP_19/ROSES2019_IIP_A49_awards.html#deal

Qualcomm, "Snapdragon 855 mobile platform," 2020, https://www.qualcomm.com/products/snapdragon-855-mobile-platform

RAD 750 Wikipedia Page, retrieved 20 December 2021. https://en.wikipedia.org/wiki/RAD750

Swope, J.; Chien, S.; Bosch-Lluis, X.; Yue, Q.; Tavallali, P.; Ogut, M.; Ramos, I.; Kangaslahti, P.; Deal, W.; and Cooke, C. Using Intelligent Targeting to increase the science return of a Smart Ice Storm Hunting Radar. In International Workshop on Planning & Scheduling for Space (IWPSS), July 2021.

Thompson DR, Bornstein BJ, Chien SA, Schaffer S, Tran D, Bue BD, Castaño R, Gleeson DF, Noell A. Autonomous spectral discovery and mapping onboard the EO-1 spacecraft. IEEE Transactions on Geoscience and Remote Sensing. 2012 Dec 24;51(6):3567-79.

Thompson DR, Leifer I, Bovensmann H, Eastwood M, Fladeland M, Frankenberg C, Gerilowski K, Green RO, Kratwurst S, Krings T, Luna B. Real-time remote detection and measurement for airborne imaging spectroscopy: a case study with methane. Atmospheric Measurement Tech. 2015 Oct 19;8(10):4383-97.

Towfic Z. Ogbe D, Sauvageau S, Sheldon D, Jongeling A, Chien S, Mirza F, Dunkel E, Swope J, Ogut M, Cretu V, "Benchmarking and Testing of Qualcomm Snapdragon System-on-Chip for JPL Space Applications and Missions," IEEE Aerospace Conference, 2022.

K. Wagstaff, S. Lu, E. Dunkel, K. Grimes, B. Zhao, J. Cai, S.B. Cole, G. Doran, R. Francis, J. Lee, L. Mandrake, "Mars Image Content Classification: Three Years of NASA Deployment and Recent Advances," The Thirty-fifth AAAI Conf on Artificial Intelligence, pp. 15204-15213, 2021.

W. Whitaker, "Sabertooth: Integrated Avionics for Small Spacecraft Missions," 2019 Space Computing Conference. https://trs.jpl.nasa.gov/bitstream/handle/2014/51550/CL%2319-4553.pdf