Jet Propulsion Laboratory
California Institute of Technology

Benchmarking Deep Learning, Instrument Processing, and Mission Planning Applications on Edge Processors onboard the ISS

Emily Dunkel, Jason Swope, Alberto Candela, Lauren West, Faiz Mirza, Steve Chien, Zaid Towfic, and Alexandra Holloway
Jet Propulsion Laboratory, California Institute of Technology

Léonie Buckley, Juan Romero-Cañas, Jose Luis Espinosa-Aranda, and Elena Hervas-Martin
Ubotica Technologies

Mark Fernandez and Carrie Knox
Hewlett Packard Enterprise

June 21, 2023
2023 Earth Science Technology Forum

CL#23-2984

Future Missions

- Numerous current and future missions could benefit from increased onboard computation
- Rover missions could drive faster and further with greater computing power

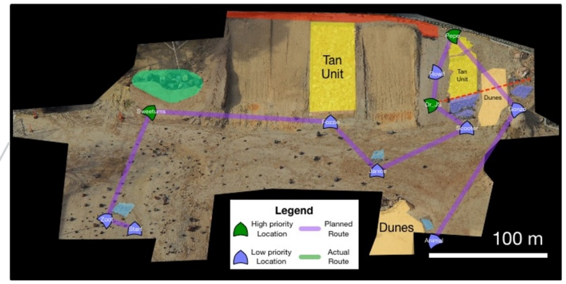

- A Europa Lander concept could use computing for robust sampling and system level autonomy

- “Smart Flyby” concepts could utilize more powerful feature detection techniques

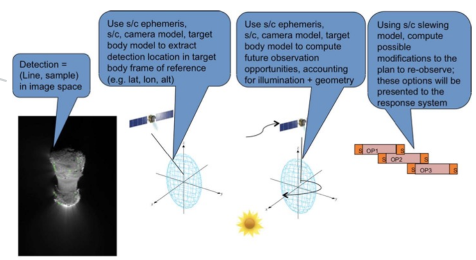

Geometric processing is required to determine valid response imaging opportunities

- Earth Science missions could use computing for “dynamic targeting”

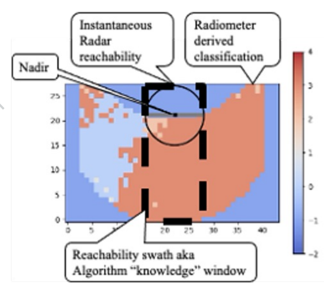

Challenge: current space processors (Rad750) have limited compute compared with modern Edge processors

Edge Processors

- Qualcomm Snapdragon 855
  - 8 core ARM CPU system
  - Adreno 640 Graphics Processing Unit (GPU)
  - Qualcomm Hexagon 690 Digital Signal Processor (DSP)
  - Neural Processing Engine
  - Running Android OS
- Intel Movidius Myriad X
  - Specialized AI/Vision Processing Unit
  - 7 SHAVE Cores
  - Direct support for CNN/Deep Learning
  - Myriad 2 (earlier version) flew on ESA’s ***Φ***-Sat (2020, 1 year)

Snapdragon 855

Myriad X

Embedded Processors

Processor	Snapdragon 855	Movidius Myriad X	Rad750	Sabertooth
Power	5W	< 1W	10+ W 5 W?	3W
Cores, Clocks	8 @ 1.7-2.8 GHz	7 SHAVE @ 700MHz	1@110-200 MHz	4@ ?
RAM, NVM	16 GB	4 GB	256 MB, 2GB	192 MB, 8 GB
Coprocessors	DSP, GPU, AIP	AIP		Motor controllers
Radiation Hardened?	No	No	Yes	No

- Rad750 is current computing for MSL and M2020

International Space Station Experiment:
Hewlett Packard Enterprise Spaceborne Computing-2

Delivered to ISS turnover: Fall 2020

Delivered to the ISS onboard Cygnus NG-15: February 20, 2021.

Powered on: May 12th, 2021.

Hewlett Packard Enterprise Spaceborne Computing-2 package:

- COTS Linux workstations from HPE
- Intel Xeon 5215 Processor (10 cores)
- 4 NVIDIA Tesla T4 GPUs
- 2 Machines aboard the ISS

2x Snapdragon 855 HDKs running Android

- Radios disabled

2x Myriad X Processors

Uplinks possible periodically to load new SW

Current status: Experiment is paused, with planned continuation

SBC-2 image credit HPE.

UAVSAR: Deep Learning for Flood Mapping

- Pixel-wise flood classification of Uninhabited Aerial Vehicle Synthetic Aperture Radar (UAVSAR) imagery
  - Model: UNET-6 fully convolutional model*
- Applications: flood mapping, alert generation
- Benchmarking results:

Processor	Full Class. Discrepancy	Binary Class. Discrepancy	Image Patches / Second
MacBook Reference	-	-	25
Snapdragon CPU	3	1	20
Snapdragon GPU	4	1	162
Snapdragon DSP	87,689 (0.74%)	44,844 (0.38%)	391
Snapdragon NPU	87,689 (0.74%)	44,844 (0.38%)	391
Myriad X	184,820 (1.56%)	82,519 (0.70%)	167

  - Model cannot be run on RAD750
  - All systems have small quantization discrepancies, and meet real-time requirement (5 patches/sec)
  - Snapdragon NPU provides 20x speedup over Snapdragon CPU
  - Models have been run on the ISS: Snapdragon 21 times, Myriad 9 times
    - No differences between ground and ISS runs

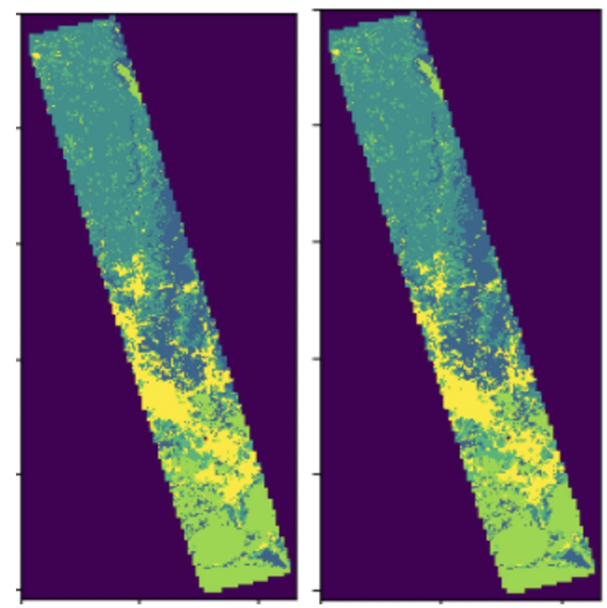

(a) Reference (b) Myriad X Results

Hurricane Harvey Classified Image

* Denbina, Towfic, Thill, Bue, Kasraee, Peacock, Lou, Flood Mapping Using UAVSAR and Convolutional Neural Networks. IEEE Xplore 2020
* Ronneberger, Fischer, Brox, 2015. U-net: Convolutional Networks for Biomedical Image Segmentation. MICCAI, 234-241.

Spectral Unmixing: Estimating components in spectral data

- Estimating abundance of coral, algae, and sand using remote sensing spectrometer data
  - Single pixel Deep Conditional Gaussian model using data from NASA CORAL mission
- Applications: Coral reef health analysis
- Benchmarking results:

Processor	Quantization Discrepancy (RMSE)	Inference Time
MacBook Reference	-	5 us
Snapdragon CPU	0	15 us
Snapdragon GPU	0	270 us
Snapdragon DSP	0.05	39 us
Snapdragon NPU	0.05	39 us
Myriad X	0.06	340 us

  - Model cannot be run on RAD750
  - All quantization discrepancies are small (given as root mean square error)
  - Snapdragon CPU is faster than NPU due to small model size and architecture
  - Model has been run 3 times on ISS and produces same results as on the ground

Spectra*

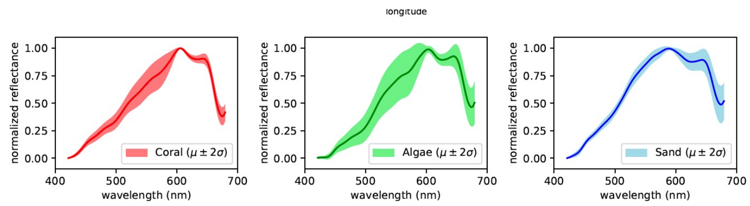

Predictions*

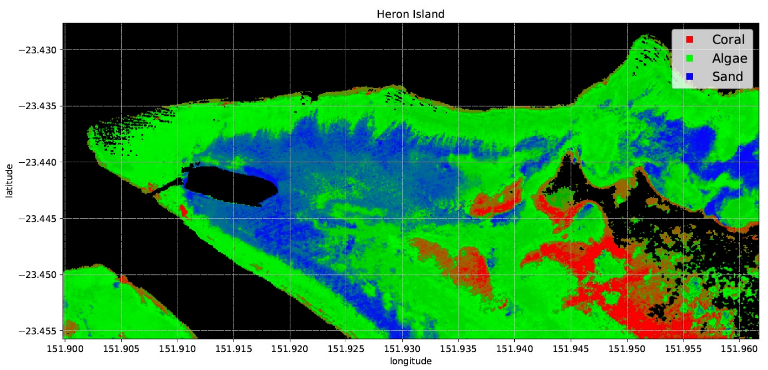

*Alberto Candela Garza. Bayesian Models for Science-Driven Robotic Exploration. PhD thesis,Carnegie Mellon University, Pittsburgh, PA, September 2021.

Spectral Algorithms

- Benchmark the match filter (MF) and spectral angle mapper (SAM) algorithms for spectral analysis
  - Both algorithms use a spectral library containing objects of interest to be searched
  - SAM is a distance function between a spectrum and an object of interest
  - MF is a linear detector that requires background statistics: the mean covariance matrix
- Used to perform mineral detection on 8 minerals on AVIRIS-NG data from Cuprite Hills, Nevada
- Implemented on the Snapdragon ARM CPU
- Dataset: 20,000 pixels with 425 bands each

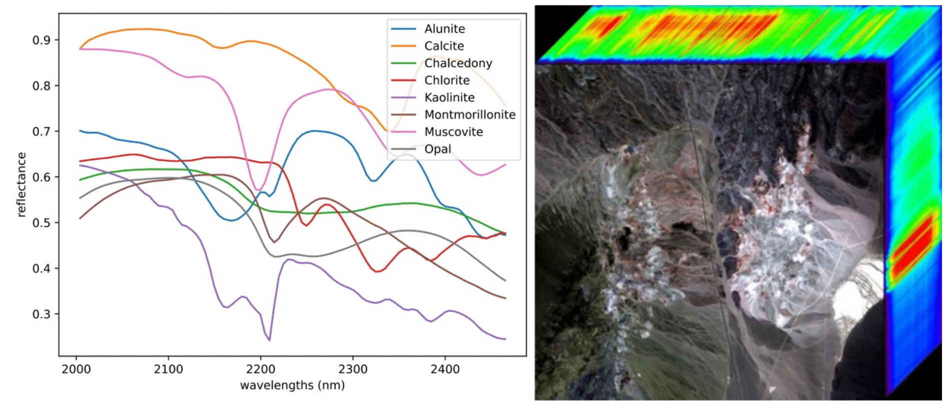

Spectral library of 8 minerals (left)
AVIRIS-NG hyperspectral image of Curpite Hills, Nevada (right)

Timing in ms per pixel	SBC2 Reference	Snapdragon CPU	Sabertooth
SAM	0.04	0.04	0.16
MF	0.04	0.06	0.15

SAM:
3.8x speedup from the Sabertooth to the Snapdragon

MF:
2.5x speedup from the Sabertooth to the Snapdragon

Dynamic Targeting Algorithms

- Dynamic targeting (DT) uses information from a lookahead sensor to identify targets for the primary, pointable sensor
- Improves science yield given energy constraints
- Scenario: maximize observations of storms clouds
- Simulations using global storm dataset GPM IMERG (Candela et al., 2022)
- We benchmark and compare 5 DT algorithms:
  - Random sampling at nadir (baseline)
  - Smart sampling only at nadir
  - Smart sampling along cross-path direction
  - Smart sampling within the primary instrument's FOV
  - Smart sampling within the primary instrument's FOV while leveraging lookahead data
- Dataset: 18,000 timesteps at 2 seconds per timestep (10 hours of simulation time, 6.3 orbits)

	SBC2 Reference	Snapdragon CPU	Sabertooth	RAD750
ms per timestep	0.39	20	1,500	2,800

75x speedup from Sabertooth to Snapdragon ARM CPU
140x speedup from the RAD750 to the Snapdragon ARM CPU

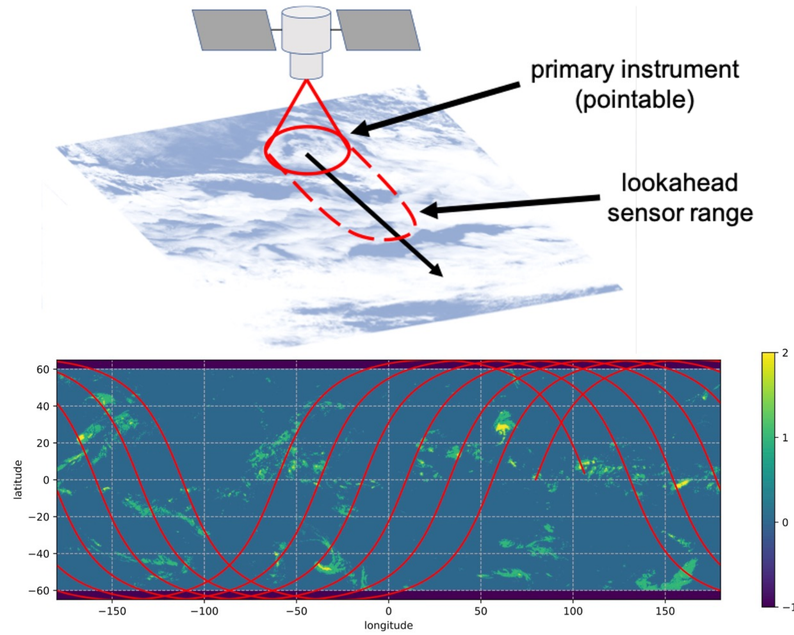

Graphic of lookahead sensor with spacecraft targeting primary sensor based on lookahead sensor data (top)
Simulated orbit path over GPM IMERG dataset (bottom)

Summary

- Edge processors provide significant speed improvement over current flying processors (RAD750), with the ability to run deep learning models

- Snapdragon DSP/NPU is optimized to run neural network inference, but speedup depends on network architecture
  - 20x speedup for UAVSAR model over CPU, but 3x slower for single pixel spectral unmixing

- Significant speedup from Snapdragon compared with RAD750
  - Snapdragon 3x-75x faster than Sabertooth for Spectral and Dynamic Targeting algorithms
  - Sabertooth 2x faster than RAD750 with Dynamic Targeting

- No discrepancies between ground and ISS runs

- We have demonstrated fast and accurate Earth science data processing with edge processors, and hope this is a step towards a new era of powerful onboard autonomy

Jet Propulsion Laboratory
California Institute of Technology

jpl.nasa.gov