Evaluation of Hybrid Satellite On-Board Computer Architectures for Edge-Compute and Machine-Learning Focussed Applications

J. Langille [1], A. Amellal [1]*,

[1] Langille & Amellal, Dalhousie University, Halifax, NS - (jack.langille, aaa)@dal.ca

KEY WORDS: Edge-AI, Climate Change, CubeSats, Computer Architecture, Multispectral Earth-Observation.

ABSTRACT

The Dalhousie University Space Systems Laboratory (DSS) is developing and deploying a cube satellite (CubeSat) to detect the presence of harmful algal blooms (HAB) in coastal regions and monitor the rate-of-change of phenological states in forested regions world-wide. These observations uniquely monitor the impact of climate change. Additionally, HABs are unpredictable episodic events that require timely warnings to avoid the area and to enact measures that mitigate the environmental damage. These are the benefits of remote sensing from Low Earth Orbit (LEO). To detect / monitor both phenomena, the satellite uses multi-spectral camera payloads to capture imagery and onboard processors to analyze the imagery in situ. The results from such analyses are transmitted back to Earth-based stations. Given multi-spectral imagery, considerable analysis effort is required. The challenge is to perform such analysis on the CubeSat with a machine learning (ML) based detection. This is the contribution of this paper and will be briefly discussed next.

A thorough evaluation of DSS' proposed solution for on-board computing and edge AI is presented. Edge AI, or the deployment of ML on hardware situated close to the data source, affords numerous technical benefits. It means raw data (multi-spectral imagery), need not be transmitted from the edge to a centralized processing centre and thus eliminates the requirement for high-bandwidth connections. Data security also improves, as raw source data is not transmitted over a network. With such advantages comes several design challenges. For instance, inferencing, often performed with graphical processing units (GPUs), is energy intensive. A balance is needed between acceptable inference times and detection accuracy, and hardware constraints like the physical GPU size and power draw. This trade-off is the key design driver in the ML framework. This balance is achieved through model optimization (weight quantization), pre-processing (identifying redundancies in the feature vector), and memory management towards efficient software. Additional pre-processing with segmentation and clustering strategies reduce a raw satellite image of a region to a smaller localized region of interest, e.g., segmenting a treed region within an image from non-forest features like rocks or farmland. Also presented are mission environment challenges like the discrepancies associated with collecting and inferencing images during periods of sun-coverage, compared to dark imagery, or disturbance of image quality from cloud coverage. Lastly, the training strategy which ranges from sourcing and curating labeled datasets to optimizing the models for compute hardware are discussed.

Conventionally, a single On-Board Computer (OBC) is the centre of a satellite's Command & Data Handling (C&DH) subsystem. When integrating OBC with hardware suitable for edge-computing applications, novel concerns emerge. Overall system reliability, power consumption, power sensitivity and low-level hardware abstractions motivate a split architecture where the edge-compute and critical C&DH capabilities run on separate systems-on-a-chip (SoC). This paper will elaborate on the predicted effect of this split architecture on C&DH reliability, i.e., reduce single points of failure and complexity while augmenting system uptime.
ML capabilities are inherently power-intensive. It is therefore expected that an ML-capable OBC consumes more baseline and maximum power than a non-ML one. In practice, the difference is an order-of-magnitude or more. The attainment of critical power modes and the induced power supply fluctuations therein can lead to undefined behaviour in the OBC (shutdowns, power-cycles and data loss). Therefore, the use of a monolithic OBC architecture that is ML-capable is thought to undesirably inflate the baseline power requirements for system C&DH and negatively impact the threshold for power-related complications as opposed to the hybrid model. Further, relegating system-critical functions and portions of C&DH to specialized lower-power hardware enables an additional power mode option where the ML-capable OBC is offline during system critical functions. Given the nominal difference in baseline power

* Corresponding (presenting) author

consumption can vary by an order-of-magnitude (or more), an entirely new class of operational contingency modes become possible when it would have otherwise been: (1) not possible in the case of monolithic systems or (2) require specialized implementations at the subsystem level.

The outcome from this work contributes to Edge AI implementations within the CubeSat space which is relevant for in situ analysis of real-time sensor data.