[SSC24-WIV-06]

# BeaverCube II: Using AI-Optimized Processors on Earth-Observing CubeSats for Autonomous Image Analysis and Intelligent Data Handling


Adam Bahlous-Boldi, Celvi Lisy, Neelambar Mondal, Brianna Ferro, Shreeyam Kacker, Mary Dahl, Madeline Anderson, Kerri Cahoy
Massachusetts Institute of Technology
70 Vassar St, Cambridge, MA
adambb@mit.edu



## ABSTRACT

The mission goal of BeaverCube II is to demonstrate autonomous on-orbit image processing and classification using the Xilinx Versal System-On-Chip (SOC). The Versal has 400 dedicated AI engines on an FPGA-based platform, enabling compute-intensive machine learning capabilities. While the Versal can consume over 100 W of power, BeaverCube II will demonstrate its on-orbit function in a low-power, size-constrained system. The system will use machine learning algorithms such as K-means clustering to autonomously assess whether features or changes of interest are present across keyframes.

BeaverCube II has a 1U imaging payload consisting of three Commercial off-the-shelf (COTS) cameras: two identical visible wavelength cameras (112.5 m GSD) and one LWIR camera (197.8 m GSD). Using this imaging payload, BC II aims to demonstrate autonomous capabilities such as cloud identification, shoreline feature recognition, and change detection. Running a high-power chip such as the Versal has demanding thermal, structural, and power requirements; the solutions to these challenges will be described in this work. BeaverCube II is a 3U CubeSat designed and built by the STAR (Space, Telecommunication, Astronomy and Radiation) Lab at the Massachusetts Institute of Technology and sponsored by the Northrop Grumman Corporation. We are targeting a Q1 2025 launch with deployment from the ISS in Q2 2025.


## INTRODUCTION

CubeSats can globally monitor various ground-based features of interest. However, bandwidth limitations can make it infeasible to downlink all the collected data in a timely manner. If the satellite is not able to downlink a sufficient amount of data in one pass, as is the case in many CubeSats, an effort should be made to downselect and prioritize data containing features of interest, while deprioritizing the rest. On-orbit feature identification using machine learning processors could be an effective mechanism to implement this on bandwidth-limited satellites.

**Figure 1: BeaverCube II 3U Cubesat.**

### *Motivation & Use-Cases*

AI can aid satellite operations, from mission operations to data processing to health management. AI is helpful for tasks dealing with non-deterministic phenomena.[1] One example is identifying features of interest in remote sensing measurements. Consider monitoring spontaneous phenomena that may be ge-


**Figure 2: Summary of BeaverCube II hardware.**

ographically and/or temporally sparse. Having an onboard AI to identify when those phenomena occur enables the satellite to swiftly respond to the situation. This can take many forms, such as the satellite pointing in a different direction,[2] prioritizing alerting assets about the identified target,[3] or choosing not to downlink data with little scientific value.[4]

This capability is valuable for various applications related to **natural disaster** identification and response. AI can be used in disaster aid for a wide variety of events including wildfires, floods, earthquakes, and landslides. Improved prediction and monitoring of these disasters allows for better deployment of resources and more informed evacuation procedures, reducing injuries and deaths.

AI is also a useful tool for **wildfire detection** and tracking. Identifying fires quickly before they grow out of control is crucial for effective wildfire management. Although fires tend to cluster in hot and dry parts of the world (California, Australia, etc.) during "fire seasons," recent changes to the Earth's climate and increases in human-caused wildfires make them more difficult to accurately measure and predict. AI-enabled orbital identification is a highly valuable tool for quick recognition and management of these fires. Real-time monitoring of wildfires allows firefighters to better deploy their resources, identify the potential path of the fire and any dangerous fuel sources or settlements nearby, and make informed decisions on evacuation routes.[5] Many wildfire identification algorithms have been designed for satellite data, including with synthetic aperture radar[6] and hyperspectral imagery.[7,8] A 6U CubeSat named KITSUNE was designed to be the first CubeSat to perform wildfire detection onboard, utilizing a Raspberry Pi running a convolutional neural network.[9,10]

Drones for Equitable Climate Change Adaptation (DECCA) is a project led by the MIT Environmental Solutions Initiative (ESI) in partnership with MIT Lincoln Laboratory. Researchers are using drone and satellite imagery of regions in Puerto Rico and Colombia to identify patterns in terrain features where **landslides** have occurred in the past. These feature maps will be used to train machine learning models in an effort to predict areas that are at high risk for future landslide events. Awareness of risk levels for any given area allows for effective evacuation planning in populated areas, and can identify the safest sites for future urbanization and development of unpopulated areas.

In areas with high mining activity, land subsidence poses a major hazard to mine workers due to the possibility of **mines collapsing**. One study uses various deep learning methods to analyze satellite imagery and identify mines in Poland with high subsidence risk.[11] This allows for mining work to be halted so that workers are protected when a site is deemed unsafe.

On-orbit AI can also be used to identify regions experiencing high **soil erosion**. This is particularly useful for agricultural applications, since eroded soil results in reduced crop yields and increases food insecurity. Researchers have analyzed soil erosion hotspots in eastern India using techniques such as neural networks and weighted regression.[12] Identifying areas in need of soil health restoration allows for targeted intervention policies, and continued monitoring of those regions will indicate the effectiveness of restorative actions so that modifications can be made if necessary.

**Flood management** is typically done with computationally intensive numerical models that may contain many errors. This can be insufficient for effectively warning people for evacuations and for rescue. Utilizing orbital assets with deep learning can help improve real-time flood warning systems[13] and disaster maps.[14]

Another application is for **ship identification** and tracking. Identifying ships that are in unauthorized locations is integral for combating illegal fishing, oil spills, and espionage. Given the vastness of the ocean and the relatively small size of these ships,

**Table 1: Overview of various AI accelerators used in space missions with comparison of maximum power and temperature.**

| Property | Units | Device | | |
|---|---|---|---|---|
| | | Xilinx Versal | NIVIDIA Jetson AGX Xavier | Intel Movidius Myriad 2 |
| Type | | FPGA + Accelerator | SBC + GPU | VPU |
| Form Factor | | Chip | Module | Chip |
| Size | mm | $40 \times 40 \times 5$ (package) | $100 \times 87 \times 9$ (module) | $9.5 \times 8 \times 2$ (package) |
| Max Power | W | 125 | 30 | 1.2 |
| Temperature Range | °C | $-40$ to 100 | $-40$ to 85 | $-40$ to 105 |
| Space Applications | | [15,16] | [17,18] | [19,20] |

especially when they are in unexpected locations, AI proves to be an invaluable tool for locating them and has been used extensively on optical satellite data.[21]

**Semantic change detection** is another valuable use of AI. Environmental and land cover changes can be used to identify known or unknown problems. This technique is often used to study problems such as deforestation, earthquakes, flooding, and urban growth. Change detection techniques have been applied to multispectral satellite images.[22] The complexity with performing change detection onboard is needing spatially and temporally matched imagery, which can be difficult for CubeSats to attain, although missions have been designed to accommodate this.[23]

In order solve these problems, we need to verify the capabilities of powerful hardware on orbit to support complex AI architectures. Many processors perform well on the ground, but are not designed for use in the space environment. Two of the most successful CubeSat technology demonstration missions are Φ-Sat-1[20] and Φ-Sat-2.[24] Φ-Sat-1 utilized a custom build of an Intel Movidius Myriad 2 EoT board with a hyperspectral camera to calculate cloud masks. Φ-Sat-2 plans to utilize the same board to continue to explore the limits of its orbital capabilities. Technology demonstration missions such as the Φ-Sat missions and BeaverCube II are integral to improving our capabilities to solve problems with increasing complexity on orbit.

### *AI Accelerators in Space*

AI accelerators have been proposed for use in space applications. Table 1 shows a summary of the most frequently considered devices, the NVIDIA Jetson AGX Xavier and the Intel Movidius Myriad 2, along with the Xilinx Versal. The key considerations for on-board AI accelerators are robustness to ionizing radiation, power consumption, and operating temperature range.

The **Intel Movidius Myriad 2** has flown on both Φ-Sat-1[20] and now on CogniSat-6.[19] It is a vision processing unit (VPU) optimized for edge applications consuming a maximum of 1.2 W and in a form factor small enough to be placed in conventional CubeSat designs. Since the device is a VPU, it is not independent and likely requires control from a separate processor. The extremely wide operating temperature range makes the device amenable to CubeSats that may not have sophisticated thermal management techniques.

The **NVIDIA Jetson AGX Xavier** is a single board computer (SBC) module with an onboard NVIDIA Tegra chip featuring both ARM cores and a GPU. The Jetson AGX Xavier module is in a form factor that is more difficult to place inside a CubeSat due to its major dimension of 100 mm. While smaller versions of the Jetson Xavier family exist such as the Jetson Nano, only the Jetson AGX Xavier comes in an Industrial version with error checking and correction (ECC) built in, which helps mitigate against ionizing radiation-induced single event upsets (SEUs).[18] Due to its design as a single board computer rather than a peripheral, the Jetson can independently perform tasks and thus support the main flight computer, or even operate in lieu of one. Additionally, the Jetson can directly connect to **MIPI** and USB devices, such as cameras, and in this case would not require another device feeding it significant amounts of data.

The **Xilinx Versal** devices operate on a compute platform that integrates FPGA fabric with hard IP for CPUs, GPUs, and AI accelerators in a single chip solution. Many space missions already utilize FPGAs for applications such as communications, which makes the Versal series a reasonable option for many space mission designs that seek to incorporate edge computing capabilities. While it consumes the most power out of all the other accelerators considered in Table 1, it is also the most configurable, with many power domains able to be shut off entirely, allowing the device to be flexible to different implementations.

### *Mission Requirements & Con-Ops*

BeaverCube II is a technology demonstration mission showing the use of the AMD Xilinx Versal SoC for on-orbit image analysis on a CubeSat.[25,26] The CubeSat is primarily made up of COTS components, significantly reducing the mission development costs.

Machine learning techniques have been previously analyzed and implemented onboard CubeSats.[1,24] Additionally, the use of the Versal SoC for space-based ML has also been discussed.[15,16] BeaverCube II will explore the intersection between these two regimes.

BeaverCube II has the following primary objectives:

- Demonstrate the successful operation of the Versal SoC on orbit on a CubeSat
- Demonstrate onboard AI-algorithms using the Versal SoC on orbit.

For the purpose of demonstrating the use of in-orbit AI on the Versal SoC, K-means and Random Forest methods will be used to identify cloud-cover on images captured by the imaging payload. These algorithms were previously tested on-board ESA's OPS-Sat Mission.[4]

As a technology demonstration, the onboard machine learning techniques applied to captured images serve as a proxy for on-orbit computation. This can inform future missions with a better equipped ADCS and imaging system for which remote sensing feature identification is the primary goal. The main objective is to demonstrate that conducting such an analysis on orbit is feasible and valuable with a Versal SoC on a CubeSat, even though the collected data for this particular mission may not necessarily be of direct scientific value on their own.

## BEAVERCUBE II DESIGN

### *Payloads*

BeaverCube II hosts two distinct payloads which work together to achieve the mission goal of demonstrating Versal Operation on orbit. These include an imaging payload consisting of 3 cameras and a compute payload consisting of the Versal SOC and Carrier board.

### *Imaging Payload*

The BeaverCube II Imaging payload consists of two Visible wavelength cameras and one LWIR camera.[27] Both cameras are COTS components which can be interfaced with through USB. All 3 cameras are connected directly to a Raspberry Pi-3 which controls their operations and stores captured images. Specifications for the cameras can be found in Table 2.

| | VIS Camera | IR Camera |
|---|---|---|
| Model | Balluff CA-IGC | Boson FLIR 640 |
| Wavelength | 390 nm to 700 nm | 8 µm to 14 µm |
| Bit Depth | 10 bit | 8 bit |
| Resolution | $752 \times 480$ | $640 \times 512$ |
| FoV | 30° | 24° |
| Focal Length | 16 mm | 18 mm |
| GSD | 112.5 m/pixel | 197.8 m/pixel |

**Table 2: Imaging Payload Camera Specifications**

### *Versal SoC & Carrier Board*

The Versal SoC is a compute platform based on an Field Programmable Gate Array (FPGA) fabric that hosts 400 AI engines, enabling it to perform complex machine-learning tasks. However, this platform is extremely resource-intensive, with the potential to consume over 125 W of power. This level of power consumption and heat dissipation is difficult to manage on a CubeSat platform. Therefore, the Versal will be limited to consuming 25 W of power, which may limit us to using only 100 of the 400 AI cores. Additionally, the reference design specified by AMD occupies a $25\,\text{cm} \times 20\,\text{cm}$ area and weighs 1 kg, which exceeds the size and mass budget constraints of our CubeSat. Consequently, the Versal

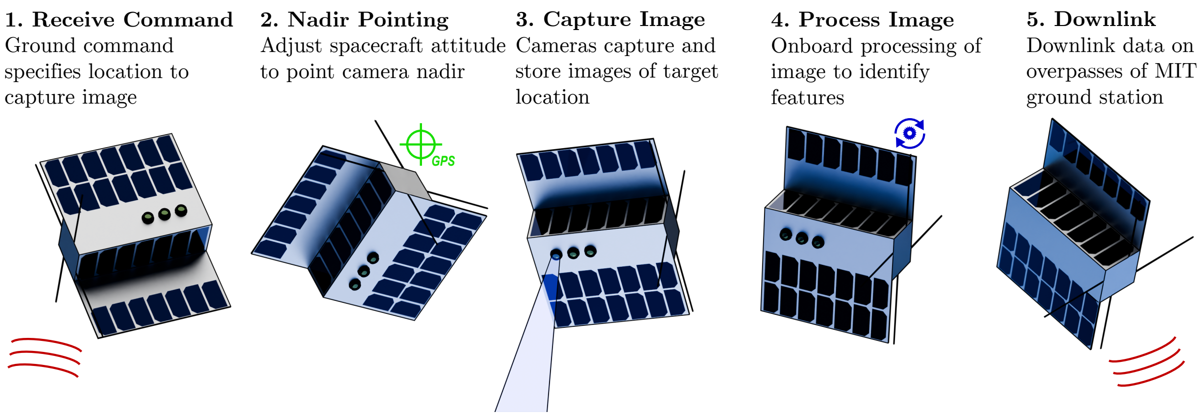


**Figure 3: Illustration of BeaverCube II ConOps**

Carrier Board had to be designed to be as minimal as possible to fit within CubeSat constraints.

Images are sent to the Versal from the Raspberry Pi via USB. The Versal will process the images using a pre-trained ML model, and send results back to the Raspberry Pi. These results, along with the original images, will be downlinked to our ground station during the next satellite pass.

### *Communications*

BeaverCube II is equipped with two radios: the Lithium Radio and the NSL Eyestar S4. The second radio adds redundancy to the system and reduces the risk of loss of contact with the satellite. MIT will utilize a UHF ground station on campus to communicate with the satellite.

| Parameter | Units | Downlink | Uplink |
|---:|:---:|---:|---:|
| Bit Rate | bps | 3500 | 3500 |
| Frequency | MHz | 401.5 | 450.1 |
| EIRP | dBW | 0.5 | 32.3 |
| Required Eb/No | dB | 13.5 | 13.5 |
| Margin | dB | 22.93 | 42.18 |

**Table 3: Radio Link Specifications**

### *Power*

The primary design driver of the BeaverCube II power system is the Versal SoC, which will be limited to consume 25 W when in use. The orbit average power consumption of BeaverCube II is 12 W. The COTS components that form the power system are the following:

- ClydeSpace Double Deploy Solar Panels: 60 minutes after deployment from the ISS, BeaverCube II's solar panels will deploy. The panels feature 7 Ultra Tripple Junction (UTJ) cells in series each, with a total of 28 cells per side. At maximum illumination, the cells can provide 24 W of power in total to the spacecraft.
- ClydeSpace EPS - the Clydespace 3G FlexU EPS interfaces with the solar panels and battery to produce 12 V, 5 V, and 3.3 V bus rails, as well as a rail at the battery voltage.
- ClydeSpace 40 Whr LiPo battery

The satellite cycles through Nominal, Eclipse, Image Capture, Image Process, Communication, and ADCS power modes throughout each orbit. Modes can also be toggled based upon either battery status or ground command, such as the low power contingency mode. The associated power consumption and duty cycles for each power mode are listed in Table 4.

### *Versal Carrier Board Design*

The Versal Carrier Board contains the Versal along with the necessary peripherals. While the reference design from AMD contains many extra devices and interfaces, size constraints for the Carrier Board limit the design to the essentials. The Carrier Board contains several power management ICs (PMICs) for powering the Versal, 1 GB of DDR4 SDRAM to store image data received from the Raspberry Pi, and 0.5 GB of flash memory for storing the boot image and Linux programs. A 30-pin connector carries power from the Electronics Power System (EPS) to the Carrier Board and transfers data between the Raspberry Pi and the Versal via USB.

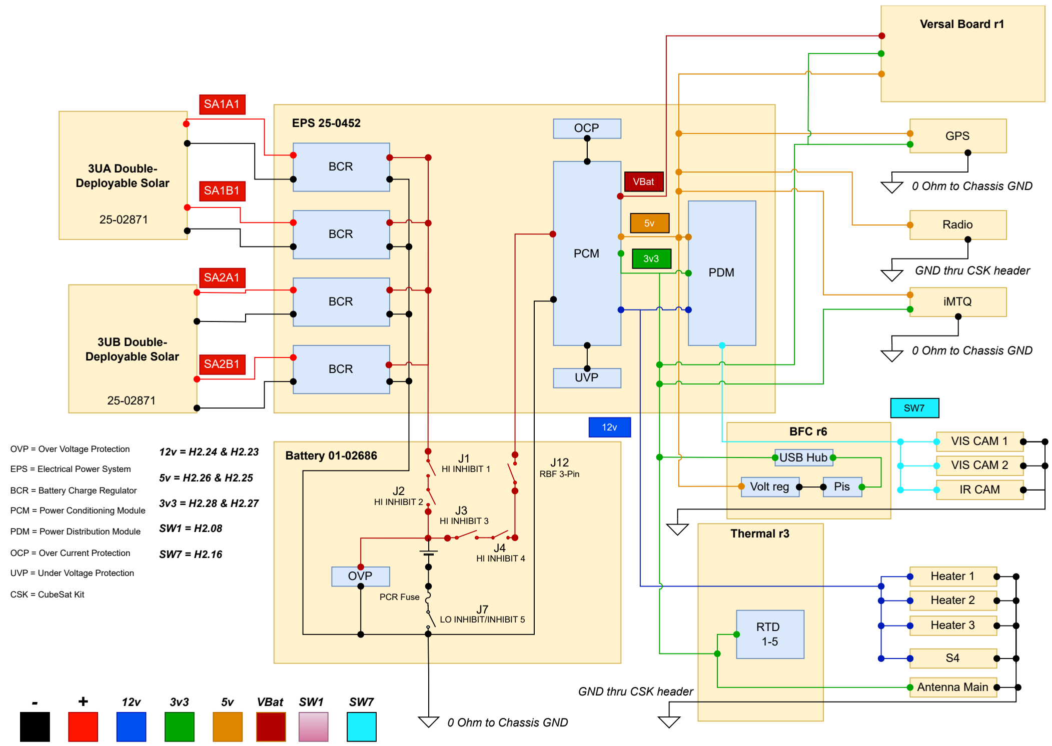


**Figure 4: CubeSat Power Diagram**

**Table 4: Satellite Power Mode Consumption and Duty Cycles**

| | Power (W) | Duty Cycle |
|---|---|---|
| Nominal Idle | 9.4 | 30% |
| Eclipse Idle | 8.2 | 50% |
| Image Capture | 17.7 | 1% |
| Image Processing | 36.2 | 11% |
| Communications | 16.6 | 3% |
| ADCS | 12.2 | 5% |
| Low Power | 5.7 | — |

For our application, the Versal requires five different voltages: 0.7 V, 0.8 V, 1.2 V, 1.5 V, and 1.8 V. The PMICs on the Carrier Board produce these voltages by stepping down the unregulated battery voltage (nominally 7.4 V). All five voltage rails cannot be turned on all at once; instead, the Versal datasheet provides a power-up and power-down sequence that must be followed. Therefore, the Versal Carrier Board also contains an MSP430 microcontroller that handles this sequencing and ensures that the PMICs are enabled and disabled in the proper order. Sufficient delay time is added in between to ensure that each voltage rail reaches its final value before the next rail is turned on.

Out of all the rails, the 0.7 V rail carries by far the most current, since it is responsible for powering the programmable logic and AI cores. Even when using only 100 of the Versal's 400 AI cores, we estimate a maximum current of 23.5 A on the 0.7 V rail, based on results from AMD's Power Design Manager tool. Carrying this large amount of current requires careful routing to minimize voltage drop across the PCB traces and to avoid overheating. The 0.7 V power trace is made as wide as possible, and is routed on two layers of the PCB using 2 oz. copper thickness. In addition, the 0.7 V PMIC has sense inputs that are Kelvin connected to the point-of-load at the center of the Versal, ensuring the voltage drop across the power traces does not affect the voltage output.

Routing the high-speed digital signals of DDR4 RAM also presents a challenge. These high-speed signals must be treated as transmission lines, so they require an adjacent ground plane for a tightly coupled return current path. These signals also have strict impedance requirements: the single-ended signals require an impedance of 50 Ω±10 %, and the differential signals require an impedance of 90 Ω±10 %. These impedances can be controlled by adjusting the trace width, trace spacing, the board dielectric ma-

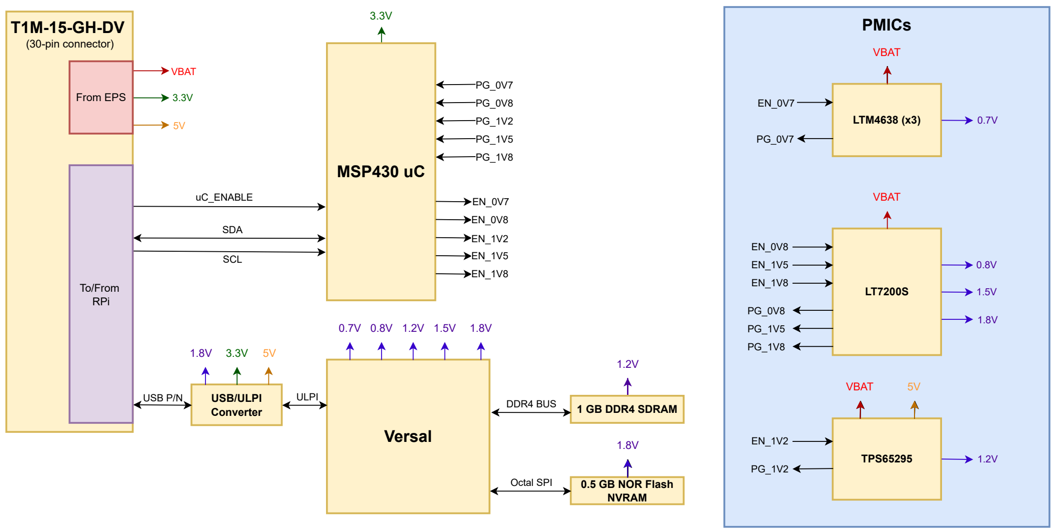


**Figure 5: Versal Carrier Board High-Level Schematic**

terial, and the distance to the adjacent ground plane. In addition, there are timing skew limits among different groups of RAM signals; for example, maximum allowed timing skew between any data line and the data strobe signal is $\pm 100$ ps. Timing skew can be minimized by tuning the length of the traces and by limiting the use of vias, which add timing delay and parasitic inductance.

### *Thermal Considerations*

While we will restrict the Versal chip's power intake to 25 W down from the full 125 W it could consume, this still poses a significant thermal challenge for a CubeSat. We implement a heatsink that interfaces through a thermal paste at the surface of the processor package. The heatsink is comprised of 0.5 kg of aluminum, acting as a thermal capacitor for the chip. It has a flat face with a high emissivity coating pointed downwards out of the CubeSat to dissipate the heat generated. The heatsink and Versal Carrier Board are mechanically isolated from the rest of the satellite board stack in order to prevent board bend from additional screw torque stress. The heatsink is also thermally connected to the spacecraft external chassis to increase heat dissipation.

The designed heatsink, while enabling the Versal to output 25 W of power, cannot dissipate that power quickly enough to allow the Versal to run for long periods of time. Thus, processes designed to be run on the Versal must be completed before the chip reaches 50 °C, which we estimate will take 10 minutes of operation at 25 W. At that point, the Versal will throttle until its operation is complete.

### *Software Con Ops*

BeaverCube II utilizes both an STM32 microcontroller and a Raspberry Pi. The STM32 software is responsible for the CubeSat's flight critical operations, while the Raspberry Pi controls the payloads. The STM32 performs the deployment of the antenna and solar panels, beaconing, detumbling, and thermal regulation. The STM32 receives ground commands from the radio and either handles the command or forwards it to the Pi, with all commands verified by checksum before execution. The STM32 is connected to most of the spacecraft peripherals including the GPS, battery, and IMUs, and it periodically sends peripheral telemetry to the Pi to be stored.

In addition to storing telemetry, the Raspberry Pi controls the cameras and the Versal. Images can be taken on command, or can be scheduled in advance. The Pi sends images to the Versal for analysis if the CubeSat state can support the use of the Versal, namely that thermal conditions are within bounds and energy is available. All data that is managed by the Pi (stored telemetry, images, Versal image analysis results, etc.) is packetized and sent to the STM32 to be downlinked upon request.

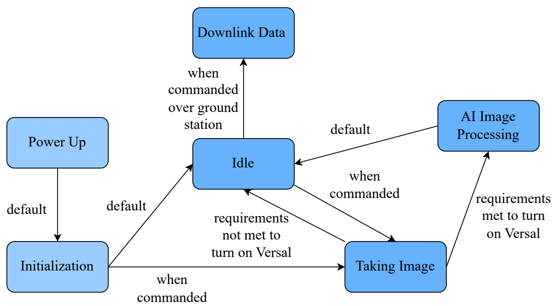


**Figure 6: Raspberry Pi software states**

## EVALUATING THE UTILITY OF AI PROCESSORS ON CUBESATS

When evaluating the practical utility of a machine learning processor on orbit, there are several mission-specific factors to consider. Here, a few of these factors will be discussed as examples of the trades that would need to be conducted to justify the use of AI.

### *Data Budget Considerations*

When using a machine learning processor to downselect data, the first consideration is what proportion of data is desirable. When searching for features with unknown spatial and temporal occurrences, a default strategy is to collect large amounts of data to ensure comprehensive coverage and increase the probability of capturing that feature. However, because most systems are bandwidth-limited, downselection is crucial. If features have predetermined locations (e.g. a region known to have various ships) but with unknown sub-feature characteristics (e.g. specific types of ships), the occurrence of these sub-features would pose a similar constraint. In either case, an autonomous machine learning downselection process would shift the data acquisition bottleneck from the link budget to the system's onboard storage, which is often less constrained and simpler to design around.

Once the feature occurrence proportion is established, the satellite's downlink budget must be considered. CubeSats are limited by the number of accessible ground stations. For a CubeSat with a single ground station, there is at most one window of opportunity per orbit to downlink data. The frequency and duration of each pass, which depends on the orbital parameters of the CubeSat, determine the amount of data that must be downlinked per pass to prevent a backlog. To plan for this, it is essential to consider both the frequency of occurrence of the feature of interest and the volume of data generated by each captured image. If the feature occurrence rate is extremely low, the false-positive rate of the feature identification algorithm may exceed the actual occurrence rate. In this case, the downlinked data will contain a large number of false positives, potentially making these false positives the driving constraint in the link budget requirement.

Even if the data collection rate (arrival rate) is less than the downlink rate (service rate), variations in either parameter can cause a queue. Instabilities in service and arrival times can significantly increase the length of this queue, thereby increasing the time between data collection and data downlinking. Using a machine learning processor to either **decrease the arrival rate through downselection** or **assign downlink priorities** to the queue based on the likelihood of containing features of interest is desirable.

### *Design Complexity Considerations*

When designing a CubeSat mission with an AI processor, one must evaluate whether the additional design complexity associated with their implementation is worth the return. This complexity stems from a number of factors. AI processors can have high power consumption, so great care must be taken with the power infrastructure design for the overall mission. Higher power requirements come with increased thermal requirements, which are not be trivial to deal with on a CubeSat.

The resources spent in developing a system to accommodate a powerful AI processor may be better spent on other aspects that could also help yield desirable results, such as increasing the radio and ground station throughput. Additionally, depending on the AI algorithm of choice, a dedicated AI processor may not be required at all, and one could substitute it for a simpler traditional processor with established space development heritage.[1]

## FUTURE WORK

One of the primary constraints that limit the utility of the Versal SoC on BeaverCube II are the power and thermal requirements. This is an area that could be improved upon without sacrificing the CubeSat form factor.

This work focuses on primarily the use of a pretrained dataset for onboard analysis. However, it is feasible to consider in-situ model training as well, either using unsupervised techniques or by leveraging a deterministic methodology to do autonomous labeling.

While we discuss the use of machine learning processors on imaging CubeSats, alternative applications will likely require their own in-depth analyses to evaluate the associated utility and design trades.

The BeaverCube II CubeSat is currently undergoing continued subsystem development and testing. It is currently scheduled to be deployed from the ISS NanoRacks deployer in Q2 of 2025.

***Acknowledgments***

This work was sponsored by the Northrop Grumman Corporation and the NASA CubeSat Launch Initiative (CSLI) as part of the BeaverCube II project.

## References

[1] Mary Dahl, Christine Page, Kerri Cahoy, and Evana Gizzi. Developing Intelligent Space Systems: A Survey and Rubric for Future Missions. *Small Satellite Conference*, August 2023.

[2] Alberto Candela, Jason Swope, and Steve A. Chien. Dynamic targeting to improve earth science missions. *Journal of Aerospace Information Systems*, 20(11):679–689, 2023.

[3] Mary Dahl, Juliana Chew, and Kerri Cahoy. Optimization of smallsat constellations and low cost hardware to utilize onboard planning. In *ASCEND 2021*.

[4] Shreeyam Kacker, Alex Meredith, Kerri Cahoy, and Georges Labrèche. Machine learning image processing algorithms onboard OPS-SAT. In *Small Satellite Conference*.

[5] Piyush Jain, Sean C.P. Coogan, Sriram Ganapathi Subramanian, Mark Crowley, Steve Taylor, and Mike D. Flannigan. A review of machine learning applications in wildfire science and management. *Environmental Reviews*, 28(4):478–505, December 2020. Publisher: NRC Research Press.

[6] Yifang Ban, Puzhao Zhang, Andrea Nascetti, Alexandre R. Bevington, and Michael A. Wulder. Near real-time wildfire progression monitoring with sentinel-1 sar time series and deep learning. *Scientific Reports*, 10(1):1322, Jan 2020.

[7] Nguyen Thanh Toan, Phan Thanh Cong, Nguyen Quoc Viet Hung, and Jun Jo. A deep learning approach for early wildfire detection from hyperspectral satellite images. In *2019 7th International Conference on Robot Intelligence Technology and Applications (RiTA)*, pages 38–45, 2019.

[8] Kathiravan Thangavel, Dario Spiller, Roberto Sabatini, Stefania Amici, Sarathchandrakumar Thottuchirayil Sasidharan, Haytham Fayek, and Pier Marzocca. Autonomous satellite wildfire detection using hyperspectral imagery and neural networks: A case study on australian wildfire. *Remote Sensing*, 15(3), 2023.

[9] Muhammad Hasif bin Azami, Necmi Cihan Orger, Victor Hugo Schulz, KITSUNE, and Mengu Cho. Demonstration of wildfire detection using image classification onboard cubesat. In *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*, pages 5413–5416, 2021.

[10] Muhammad Hasif bin Azami, Necmi Cihan Orger, Victor Hugo Schulz, Takashi Oshiro, and Mengu Cho. Earth observation mission of a 6u cubesat with a 5-meter resolution for wildfire image classification using convolution neural network approach. *Remote Sensing*, 14(8), 2022.

[11] Anna Franczyk, Justyna Bała, and Maciej Dwornik. Monitoring Subsidence Area with the Use of Satellite Radar Images and Deep Transfer Learning. *Sensors*, 22(20):7931, January 2022. Number: 20 Publisher: Multidisciplinary Digital Publishing Institute.

[12] Rabin Chakrabortty, Subodh Chandra Pal, Mehebub Sahana, Ayan Mondal, Jie Dou, Binh Thai Pham, and Ali P. Yunus. Soil erosion potential hotspot zone identification using machine learning and statistical approaches in eastern India. *Natural Hazards*, 104(2):1259–1294, November 2020.

[13] R. Bentivoglio, E. Isufi, S. N. Jonkman, and R. Taormina. Deep learning methods for flood mapping: a review of existing applications and future research directions. *Hydrology and Earth System Sciences*, 26(16):4345–4378, 2022.

[14] Bruno Adriano, Naoto Yokoya, Junshi Xia, Hiroyuki Miura, Wen Liu, Masashi Matsuoka, and Shunichi Koshimura. Learning from multimodal and multitemporal earth observation data for building damage mapping. *ISPRS*

*Journal of Photogrammetry and Remote Sensing*, 175:132–143, 2021.

[15] Michael Petry, Gabriel Wuwer, Andreas Koch, Patrick Gest, Max Ghiglione, and Martin Werner. Accelerated deep-learning inference on the versal adaptive SoC in the space domain. In *2023 European Data Handling & Data Processing Conference (EDHPC)*, pages 1–8.

[16] Noah Perryman, Christopher Wilson, and Alan George. Evaluation of xilinx versal architecture for next-gen edge computing in space. In *2023 IEEE Aerospace Conference*, pages 1–11. ISSN: 1095-323X.

[17] Man Xie, Lianguo Wang, Miao Ma, and Pengfei Zhang. Performance evaluation method for intelligent computing components for space applications. *Sensors*, 24(1):145, 2023.

[18] Ivan Rodriguez-Ferrandez, Maris Tali, Leonidas Kosmidis, Marta Rovituso, and David Steenari. Sources of single event effects in the nvidia xavier soc family under proton irradiation. In *2022 IEEE 28th International Symposium on On-Line Testing and Robust System Design (IOLTS)*, pages 1–7. IEEE, 2022.

[19] David Rijlaarsdam, Tom Hendrix, Pablo T Toledano González, Alberto Velasco-Mata, Léonie Buckley, Juan Puig Miquel, Oriol Aragon Casaled, and Aubrey Dunne. The Next Era for Earth Observation Spacecraft: An Overview of CogniSAT-6. *Authorea Preprints*, 2024.

[20] Gianluca Giuffrida, Luca Fanucci, Gabriele Meoni, Matej Batič, Léonie Buckley, Aubrey Dunne, Chris van Dijk, Marco Esposito, John Hefele, Nathan Vercruyssen, Gianluca Furano, Massimiliano Pastena, and Josef Aschbacher. The Φ-Sat-1 Mission: The First On-Board Deep Neural Network Demonstrator for Satellite Earth Observation. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–14, 2022.

[21] Urška Kanjir, Harm Greidanus, and Krištof Oštir. Vessel detection and classification from spaceborne optical images: A literature survey. *Remote Sensing of Environment*, 207:1–26, Mar 2018. PMCID: PMC5877374.

[22] Aysim Toker, Lukas Kondmann, Mark Weber, Marvin Eisenberger, Andrés Camero, Jingliang Hu, Ariadna Pregel Hoderlein, Çağlar Şenaras, Timothy Davis, Daniel Cremers, Giovanni Marchisio, Xiao Xiang Zhu, and Laura Leal-Taixé. Dynamicearthnet: Daily multi-spectral satellite dataset for semantic change segmentation, 2022.

[23] Chahira Serief, Youcef Ghelamallah, and Youcef Bentoutou. Deep-learning-based system for change detection onboard earth observation small satellites. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 16:8115–8124, 2023.

[24] Alessandro Marin, César Coelho, Florian Deconinck, Irina Babkina, Nicolas Longépé, and Massimiliano Pastena. Phi-sat-2: Onboard AI apps for earth observation. In *2021 SPACE AND ARTIFICIAL INTELLIGENCE*.

[25] Violet Felt, Shreeyam Kacker, Joe Kusters, John Pendergrast, and Kerri Cahoy. Fast ocean front detection using deep learning edge detection models. *IEEE Transactions on Geoscience and Remote Sensing*, 2023.

[26] Shreeyam Kacker, Alex Meredith, Joe Kusters, Hannah Tomio, Violet Felt, and Kerri Cahoy. On-orbit rule-based and deep learning image segmentation strategies. In *AIAA SCITECH 2022 Forum*, page 0646, 2022.

[27] Hannah Tomio, Albert Thieu, Amelia Gagnon, Sophia K Vlahakis, Shreeyam Kacker, Joe Kusters, and Kerri Cahoy. Commercially Available Imaging Payloads for CubeSat Earth Observation Missions. In *2022 IEEE Aerospace Conference (AERO)*, pages 1–19. IEEE, 2022.