SSC23-X-03

BrainStack: A Platform for Artificial Intelligence & Machine Learning Collaborative Experiments on a Nano-Satellite

Marcus S. Murbach, Eric Barszcz, L. Seth Schisler
NASA Ames Research Center
M/S 244-19, Moffett Field, CA, USA 94035 – 650/604-3155
marcus.s.murbach@nasa.gov

Alejandro J. Salas, Kwabena Boateng, Gregoire Marty
MEIS LLC/Axient Corporation
M/S 244-19, Moffett Field, CA, USA 94035
alejandro.j.salas@nasa.gov

Avery Brock
KBR Wyle Services LLC
M/S 244-19, Moffett Field, CA, USA 94035
avery.d.brock@nasa.gov

Malachi Mooney-Rivkin, Stanley M. Krześniak
Metis Technology Services LLC
M/S 244-19, Moffett Field, CA, USA 94035
malachi.mooney-rivkin@nasa.gov

ABSTRACT

Space missions have become more ambitious with exploration targets growing ever distant while simultaneously requiring larger guidance and communication budgets. These conflicting desires of distance and control drive the need for in-situ intelligent decision making to reduce communication and control limitations. While ground based research on Artificial Intelligence and Machine Learning (AI/ML) software modules has grown exponentially, the capacity to experimentally validate such software modules in space in a rapid and inexpensive format has not. To this end, the Nano Orbital Workshop (NOW) group at NASA Ames Research Center is performing flight evaluation tests of 'commercially' available bleeding-edge computational platforms via what is programmatically referred to as the BrainStack on the TechEdSat (TES-n) flight series. Processors selected as part of the BrainStack are of ideal size, packaging, and power consumption for easy integration into a cube satellite structure. These experiments have included the evaluation of small, high-performance GPUs and, more recently, neuromorphic processors in LEO operations. Additionally, it is planned to measure the radiation environment these processors experience to understand any degradation or computational artifacts caused by long term space radiation exposure on these novel architectures. This evolving flexible and collaborative environment involving various research teams across NASA and other organizations is intended to be a convenient orbital test platform from which many anticipated future space automation applications may be initially tested.

THE NEED FOR ADVANCED ON-ORBIT COMPUTING AND THE BRAINSTACK LABORATORY PAYLOAD

Recent advances in spacecraft technology have greatly improved system miniaturization, cost reduction, and have enabled distributed science[1,2] via spacecraft swarms. Enhanced by these technologies, low-cost missions can now easily create large, valuable data products, but they are still faced by constrained communication budgets. A lowered cost barrier has also enabled novel mission concepts visiting the Moon, Mars, or beyond, with projected data products in the terabytes per day range potentially generated across not a single spacecraft, but a large swarm of distributed vehicles requiring coordination. Direct and complete Earth-based control of swarms or deep-space vehicles, and Earth-based processing of raw data, is simply becoming infeasible given the volume of data and communication distances. The remedy for these problems is to increase onboard computational capabilities and leverage modern advancements in artificial intelligence (AI) and machine learning (ML) models to autonomously manage missions and perform statistical analysis on data to reduce communication requirements.

Currently, for single vehicles or small swarms, artificial intelligence (AI) and machine learning (ML) models executed on general-purpose processors operate acceptably[4,5], but general-purpose is insufficient for more advanced models or mass data processing. The software models required to operate a complete mission and its host spacecraft fully autonomously beyond the reach of real-time communications, or to coordinate hundreds of distributed sensors across space are simply too complex for general computing platforms. The inclusion of AI/ML tailored co-processors, optimized FPGA fabrics, or the inclusion of micro-GPUs on standard microprocessors and even microcontrollers have improved the ability to perform most basic mission operations autonomously[3,4], but data processing and automation desires are still outpacing space-rated computational capabilities.

A challenge in space-based AI/ML development is the impact of radiation and the general space environment on electronic hardware. Commercial and government off-the-shelf (COTS/GOTS) hardware products designed for AI/ML applications tend to be poorly characterized for space operation due to the cost and time required to perform such characterization. The pace of hardware development has outpaced the time required for space environmental testing and characterization, making most testing financially inadvisable to any manufacturer despite the recent growth in the space industry. This leaves the space performance of a new processor up to the customer to evaluate. Simply put, the way new processor architectures, silicon structures, and AI/ML models perform in the space environment is largely unknown except for specific configurations used by specific missions. However, it is understood that mass market-driven technologies, while less radiation tolerant, will generally outpace space-targeted, radiation tolerant, AI/ML devices.

The Nano Orbital Workshop's BrainStack payload seeks to expand knowledge in this area of advanced orbital computing through extensive spaceflight testing and characterization of advanced processing hardware and software models, as well as through collaboration with government, academia, and commercial industry to find a path forward to enable prolific usage of substantial AI/ML models for LEO and beyond-Earth applications. The BrainStack payload series creates an orbital laboratory where both novel hardware and software models can be tested and evaluated in real-time, with failure made an acceptable risk with no impact to other mission requirements. To date, the BrainStack payload has flown two AI/ML accelerator-capable platforms: the NVIDIA Tegra X2 general-purpose graphics processing unit (GPGPU), and the Intel Loihi-1 neuromorphic processor (NP). This paper discusses the current and notional BrainStack platform architecture, concept of operations, recent results from the ongoing TES-13 mission, and future architectures and system designs of BrainStack.

Figure 1: ISS Deployment of TES-10

INITIAL PROTO-BRAINSTACK FLIGHTS: TES-8 AND TES-10

An early BrainStack concept was first tested on the TES-8 and TES-10 spacecraft, both deployed from the International Space Station (ISS) in 2019 and 2020, respectively. The deployment of TES-10 is shown in Figure 1. These first tests of what would become the BrainStack payload consisted solely of an NVIDIA® Tegra X2 (Jetson™ TX2) GPGPU. For these first two experiments, the TX2 was fed video from a set of stereoscopic cameras mounted to the aft of the spacecraft. The experiment consisted of recording a sixty second video from both cameras, then compressing the recording, then transfering the compressed file to the TES 'Lunar Radio' S-Band SDR via an internal Wi-Fi network. The video would then be downlinked when able with the resulting video footage providing a view of the spacecraft moving away from the ISS after deployment.

TES-8 and TES-10 BrainStack Development

To facilitate this experiment, several key technologies had to be built up to support the use of the TX2 in the heritage TechEdSat core avionics stack. Prior to TES-8, TES spacecraft had not flown a high-bandwidth radio, something required to retrieve meaningful amounts of data from a test of the TX2. The TES team developed an S-band SDR platform from COTS parts called the Lunar Radio after its potential capability to support a lunar CubeSat mission. On the TX2 GPGPU side of development, schedule and budget pushed the use of COTS parts, with an Orbitty Carrier being used to interface with the TX2 as it would easily fit within the TES-bus avionics stack form factor.

Initial development of the TX2 experiment hardware and software was performed by the San Jose State University (SJSU) Computer Vision Club, a partnership arising out of the TechEdSat group's longstanding relationship with the university. SJSU students and NASA interns selected the initial hardware components and performed initial development and ground testing of the flight software to be executed on the TX2 on orbit.

Supporting vehicle integration, the TES team began developing the flight avionics hardware and software needed to facilitate the novel use of on-vehicle Wi-Fi to transfer the large video file between the TX2 experiment and the TES Lunar Radio SDR. This originally point-to-point Wi-Fi network was expanded to mesh all major avionics processors and radios together to allow for a flexible internal communications network capable of easily transferring large files between experiments, data storage, and radios, a design topology that would become integral to all subsequent TES missions and a notable feature of TES avionics. Both TES-8 and TES-10 flew nearly identical avionics stacks consisting of the TX2 experiment module and the TES Lunar Radio SDR, as shown above in Figure 2.

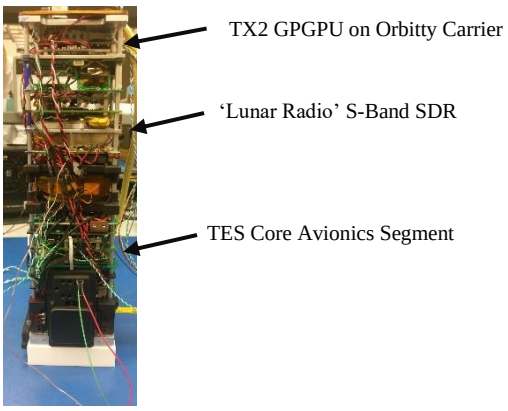

Figure 2: TES-8 Stack with Integrated TX2 GPGPU and S-Band SDR Radio

TES-8 and TES-10 Results

Both versions of the payload produced partially successful results; it was confirmed the TX2 payload activated, recorded a video, compressed the file, and wirelessly transferred the data to the S-band radio via Wi-Fi, determined through event logs and verified onboard file transfers. Unfortunately, in both cases the Lunar Radio did not perform as expected, preventing recovery of the full video file as an S-band downlink was never successfully maintained. However, due to the flexible nature of the onboard Wi-Fi network, TES-10 was able to transmit several frames of video by re-routing data from the Lunar Radio to the several Iridium short-burst data (SBD) modems in the core avionics stack. One of these frames is shown in Figure 3 on the following page. However, the payload did demonstrate mastery of incorporating the TX2 in a nanosatellite system, paving the way for more complex experiments. The experience of design and integration of the TX2 on TES-8 and 10 proved invaluable for future BrainStack payloads, as a complete electrical, mechanical, software, firmware, and programmatic approach methods were developed to support the operation of complex novel processors and software packages as experimental payloads, which paved the way for increasingly complex BrainStack payloads.

Figure 3: TES-10 Still Frame Showing Deployment of Exo-Brake Drag Device Payload on Orbit, Demonstrating Functional Data Pathway Architecture

TES-13 NEUROMORPHIC BRAINSTACK PAYLOAD

The first 'true' BrainStack payload was flown on TES-13, which was launched on Virgin Orbit's *Above the Clouds* mission on January 13, 2022. Considerably more complex than the prior TX2 experimental payload, this experiment featured the first reported flight of a neuromorphic processor in space. The neuromorphic processor used was the Intel® Lab's Loihi-1 in the Kapoho Bay USB form factor. The Loihi is a 14nm, 128-core Spiking Neural Network (SNN) able to support on-chip training. The goal of this payload was to demonstrate Loihi-1 functionality in orbital conditions by running simple test applications over TES-13's two-year operational period[2]. The Loihi-1 processor is attractive for aerospace applications because it's computational efficiency is much greater than that of a traditional CPU or GPU when running learning models, which is of great benefit to low-power, thermally sensitive aerospace applications. Flexible learning models are being studied for use as intelligent systems monitoring and optimization, and real-time system debugging and fault recovery. For example, a model integrated into an ADCS controller could hypothetically learn to recognize degradation of ADCS mechanical components and optimize a compensation algorithm through system monitoring and feedback. Such a model would require more power and hardware resources running on a GPU or CPU than on a neuromorphic processor topology optimized to run such learning models.

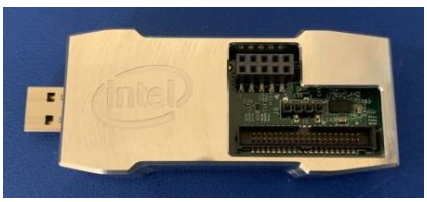

Figure 4: Intel Loihi Packaged in a Kapoho Bay USB Module

Supporting A Neuromorphic Processor as a Payload

Several hardware, electrical, and software design challenges were faced in accommodating this unique payload. As the Kapoho Bay is designed to be a USB-interfaced neuromorphic accelerator/co-processor, it requires an interface computer acting as a host to program its neural network and to provide and receive data that it processes. Being an Intel® product in early development, the Kapoho Bay development environment is only compatible, at present, with x86 processors able to execute the low-level instructions supported by the Loihi-1 processor. This prohibited the use of the TX2 or a more common Raspberry Pi as a host computer. As such, an UP Squared single board computer (SBC) with an Intel® Pentium™ CPU was selected to host the Kapoho Bay.

Packaged together, the UP Squared host computer and Kapoho Bay were then treated much like the TX2 flown on TES-8 and 10, where the UP Squared was linked to the new version of the TES Lunar Radio via Wi-Fi to facilitate large data downlink. Execution of experimental trials and pass/fail results were routed via Iridium constant-coverage SBD modems. The UP Squared communicated with an intermediate TES flight computer called 'Crayfish', which parsed and formed the SBD packets and stored data prior to transmission. This Crayfish avionics computer had a hardline serial connection to the UP Squared host computer to transmit run and Wi-Fi initialization commands. The physical arrangements of these core components are shown below in Figure 5, as flown on TES-13.

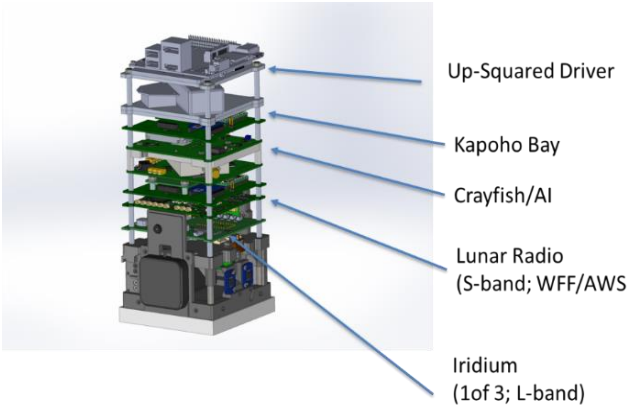

Figure 5: Intel Loihi BrainStack and Support Modules as Assembled for TES-13

Neuromorphic Payload Operations

The unknown operational behavior of the Loihi on orbit created a challenge when developing the operation methodology and experimental design of this BrainStack payload. The power and thermal load produced by the Kapoho Point and UP Squared module in the space environment was unknown, as was the reliability of the newly redesigned S-band Lunar Radio. To reduce the risk of latch-up or program hangs, it was decided that the payload would be executed for commanded time durations with all execution halting and the payload being powered off when the core avionics determined the payload had run its allowed time. As such, the UP Squared host SBC was responsible for scheduling experiments after being given a set allowable run time. Once the computational experiment concluded, either after a successful run or timeout, two additional time commands allowed for maximum periods during which the UP Squared board could transfer small 100-byte summary result packets to the Crayfish avionics module for immediate Iridium SBD downlink, and another period for full, several megabyte compressed log files to transfer to the Lunar Radio SDR for later downlink during a ground pass. These three sequential states, experimentation, summary reporting, and log transfer, could be executed individually at any time on orbit should the prior time allotment have been insufficient, or should a fault have occurred causing system reset or brownout. In this way, data would not be lost in the event the experiment forced a halt during any point in the data gathering or transfer process though either excessive heating, power consumption, or other error. Previously un-transferred data would simply be transferred during the next time allocation commanded.

This concept of operations required software to be broken into three main components, as seen on the following page in Figure 6, with additional detail in Figure 7. The TES Flight Software executing on the Lunar Radio and Crayfish Processor managed command and control of the UP Squared SBC and downlinking of experimental results. The UP Squared SBC software was broken into two parts, the scheduler and experimental data manager, and the actual software experiment payloads.

To summarize, the concept of operations was to run the payload for a commanded duration, with the UP Squared SBC scheduling test software execution and logging all data generated, then transferring brief status updates to Crayfish for immediate SBD downlink, then transferring full log files to the Lunar Radio for S-band passes.

Software Modules as a Payload

The UP Squared SBC flight software consisted of two software modules, one to control the Loihi processor in the Kapoho Bay, and another to interface with the TechEdSat avionics software running on the Crayfish module. Control of the Kapoho Bay was implemented using a software scheduler on the UP Squared SBC, which would choose between seven experimental applications based on the maximum allowed execution time. The maximum execution duration of each experiment was known, allowing the scheduler to only run experiment applications that would execute within the allowable experiment duration. Not all experiment software applications utilized the Loihi, some were designed to directly compare performance between the CPU and Neuromorphic processing architectures. A summary table of tasks is shown below in Table 1. The software payloads were developed by Michael Mercury of the Exploration Institute under NASA SBIR, and Tarek M. Taha of the University of Dayton, and as such will not be detailed in this publication. The SBC management software and the scheduler managing the software payloads was developed by the NASA Ames Intelligent Systems Division.

Table 1: BrainStack Software Payloads

ID	Processor	Description
A	Kapoho Bay	Base Model executed on Kapoho Bay
B	Kapoho Bay	Cognitive radio online learning
C	Kapoho Bay	Cognitive radio on simulated data
O	Kapoho Bay	Online learning on spacecraft anomaly data
T	Kapoho Bay	Online learning with spacecraft health
W	UP CPU	Validates any data generated from task T
X	UP CPU	Base Model run without simulated SNN
Y	UP CPU	Base Model ran on simulated SNN on a CPU
Z	Kapoho Bay	SNN anomaly detection on test spacecraft data

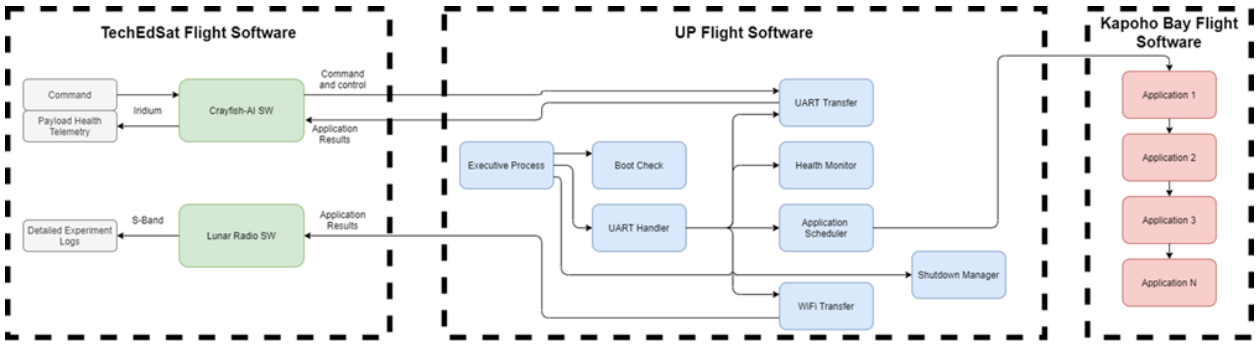

Figure 6: TES-13 BrainStack Software Modules

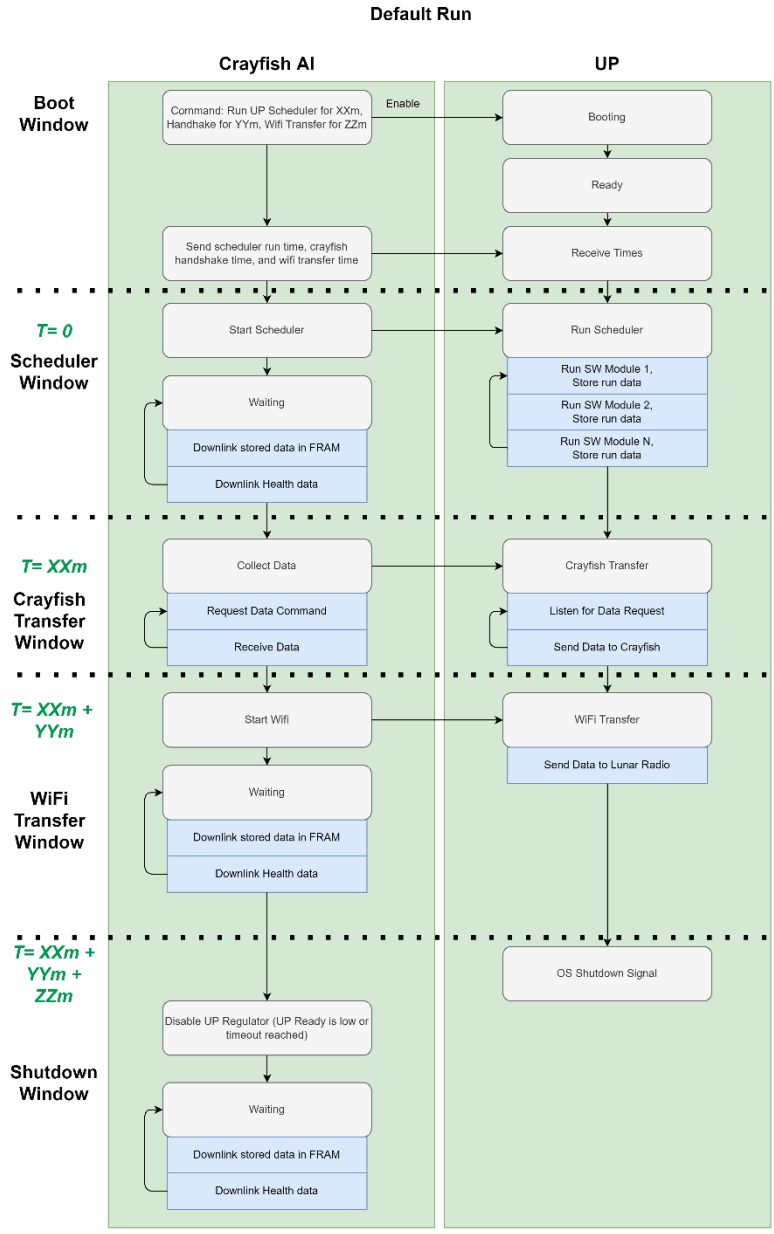

Figure 7: BrainStack UP Squared SBC Software Interaction with TES Crayfish Avionics Computer with Scheduler Time Parameters

TES-13 BrainStack Platform Results

Overall, the Loihi-1 neuromorphic processor has performed well on board TechEdSat-13 and has achieved comprehensive success per the initial mission criteria. It should be noted this experiment is ongoing, with TES-13 still active and on orbit as of August 2023, with testing now focused on extended experimental runs during the solar events occurring over the summer of 2023. As of this writing, no significant SEE events have been noted, even with this complex-architecture COTS 14nm processor only protected by two millimeters of aluminum.

Summary packets stating pass/fail results and top-level system health logs have been received via Iridium SBD, and full BrainStack execution log files have been received through now successful Lunar Radio S-Band downlink opportunities. As confidence in this platform has grown, the run time has been incrementally increased, with experiments now allowed to operate continuously for over two full orbits.

A summary of past tests, run times, and application execution success results below are shown below in Table 2. The table shows a breakdown of the applications executed and which applications executed successfully. One of the key observations from this testing is the Loihi-1 processor does not perform well when limited to a shorter run time. The reason for this result was determined to be that due to the experimental nature of the Kapoho Bay and Loihi-1 silicon design, the processor does not operate properly at colder temperatures and needs several minutes to self-heat before it begins to execute successfully. TES-13's nominal interior temperature hovers around 0°C in its 500km, 45° circular orbit, and the Kapoho Bay was determined to not properly operate until reaching 15°C. Additionally, experiment runs exceeding 220 minutes have resulted in overheating conditions in the processor, resulting in automatic self-shutdown. However, these behaviors were a known operational artifact of the Loihi-1, revision 'A' fabrication.

Table 2: Summary of TES-13 BrainStack Software Experiments

Test ID	Date	Run Time	Total Packets	H Packets	A Packets	B Packets	C Packets	O Packets	T Packets	W Packets	X Packets	Y Packets	Z Packets	Success Rate
1	1/14/2022	4	2	-	-	-	-	-	-	-	1	-	-	50%
2	1/14/2022	6	5	-	-	-	-	-	-	-	1	-	-	20%
3	1/16/2022	4	7	-	-	-	-	-	-	-	1	-	-	14%
4	1/17/2022	30	12	-	-	-	-	-	-	-	1	1	-	17%
5	1/18/2022	15	6	1	-	-	-	-	-	-	-	-	-	0%
6	1/19/2022	19	9	2	-	-	-	-	-	-	-	-	-	0%
7	1/21/2022	41	10	3	-	-	-	-	-	-	-	-	-	0%
8	1/24/2022	45	21	3	-	1	1	-	-	-	-	-	-	11%
9	1/25/2022	65	19	5	-		1	-	-	-	-	-		7%
10	1/30/2022	65	40	4	1	5	2	-	2	2	-	-	2	39%
11	2/7/2022	95	66	7	2	6	4	-	13	13	-	-	6	75%
12	2/10/2022	96	35	-	3	6	5	-	-	-	6	4	6	86%
13	2/12/2022	180	67	-	5	12	9	5	-	-	10	8	11	90%
14	2/22/2022	225	159	15	4	18	17	-	34	34	-	-	14	84%
15	3/4/2022	226	61	-	5	5	5	2	-	-	9	8	6	66%

FUTURE FLIGHTS OF THE NEXT BRAINSTACK

The next two BrainStack payloads will fly the next generations of the NVIDIA® TX2 and Kapoho Bay. TechEdSat-11, targeted to launch in late 2023, will fly the NVIDIA® TX2-I GPU, a more rugged version of the TX2. In an updated experiment, the GPU will execute a learning model on a sample image set while in orbit and compare processor and model performance and results with that of a ground control. This new software experiment is being developed by Caleb Adams of the NASA Ames Intelligent Systems Division, with hardware development and integration performed by the TES team.

TechEdSat-16, projected to launch in early 2024, will fly the Kapaho Bay Revision 'B' module, which is designed to alleviate several of the issues found with the original experimental model. This hardware-improved module uses the same first generation Loihi as the TES-13 experiment, but the Kapaho Bay USB host is designed to be more tolerant to temperature fluctuations compared to the previous revision.

Operating under constrained budget and schedule given the highly experimental nature of the BrainStack experiments, only devices readily available to the TES team have been flown, or those under active study by affiliated NASA partner groups. From an experimental standpoint, and to continue the purpose of the BrainStack concept, there is a larger list of novel processors from which to draw, with potential near-term options shown on the following page in Table 3. Of great interest is the development of NASA's High Power Space Computer (HPSC), intended for deep space missions that require radiation tolerance and long-term reliability. However, delivery is not expected for another two years and the HPSC is not focused purely on AI/ML execution.

The intent of the BrainStack architecture is to reduce the effort required to integrate different AI/ML devices and related software experiments into small spacecraft through the development of a standardized hardware and software solution. The TES-n Common Software Interface that will run various potential AI/ML platforms and experiment menus is shown below in Figure 8. It is an evolution and expansion of the interface currently operational on TES-13. The interface system would be commanded to schedule a particular set of experiments on a platform in the BrainStack through an Iridium command line packet. A particular processor payload (shown as n) would be selected, and the initialization sequence initiated, starting the run of the selected experiment on the desired processor. Once completed or after the allowed duration, the data is then fetched and transferred to the S-band radio module's memory and awaits a downlink command schedule. The experiment is then shutdown, and the BrainStack can commence with a new experiment on the same or a different hardware element. This shift from simple time-allocation to a configuration and command period will allow on-orbit configuration of customized experiments, enabling different hardware and software development teams to execute different unique experiments as desired.

Additionally, the flexible physical and electronic standard of the TES-n avionics bus enables a wide range of options for hosting new technologies, as shown in the following Figure 9. The unique manufacturing and design approach of TES allows comparatively rapid customization for a particular mission. These range from standard 2U (e.g., TES-7) to 12U (e.g., TES-14,16) spacecraft configurations, with 150 to 300Wh of power storage capability, and sustained power delivery up to 200W, depending on configuration. The 2U or greater width of larger formats allows for non-standard size components and instruments to be easily integrated, from which the physical BrainStack can then be coupled to science instruments, permitting testing of early AI/ML techniques to rapidly evaluate data and minimize the sizeable downlink data burden.

Figure 8: Next-Generation TES Common Software Interface Flow Diagram

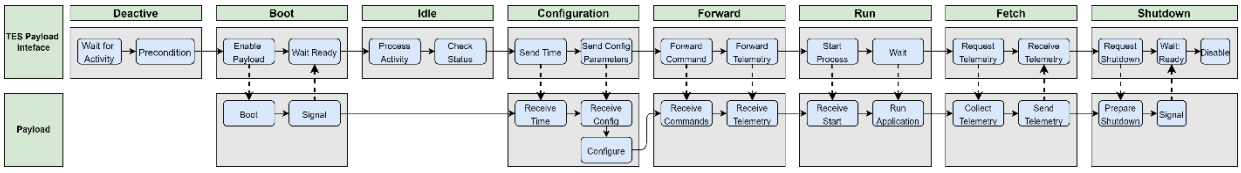

Table 3: Potential BrainStack Payload Processors

AI/ML Architectures					
Processor	Developer	Dimensions	Power	Supported APIs	Comments
Tegra X2	NVIDIA	256-core NVIDIA Pascal™ GPU architecture with 256 NVIDIA CUDA cores	15W	TensorFlow (TF), CUDA	Flown on TES-8/10 Ready for TES-11 flight
Loihi 1 Kapoho Bay	Intel	128K neurons per chip in 2D mesh of 128 neural cores	<1.5W	NxSDK	TES-13 in orbit (first test flight) TES-12 in dev.
Loihi 2	Intel	1M+ neurons per chip in 3D mesh of 128 neural cores	<1W	LAVA	TES-17 in dev. Core capacity significantly higher than Loihi 1
Akida	Brainchip	1.2 million neurons, 256 nodes	<4W	TF, Keras, BrainChip MetaTF	Minimal CPU intervention needed, mini PCIe board with Brainchip
Movidius™ Myriad™ X Vision	Intel	16 SHAVE cores (916 billion operations per second)	1.5W TDP	Flic Hub SDK	14mm x 14mm x 0.84mm 105°C max & -40 °C min
Coral TPU	Google	85x56mm	2 TOPS/W	TFLite	Low power usage
Apple A16 Bionic	Apple	16 Apple neuron Engine cores	17 TOPS 8W TDP	Swift	iPhone GPU, not tested for flight, very small size
HPSC	NASA	8x X280 at 4.6TOPS/c 4x TBD RISC-V cores	7W	TFLite, ROS	Scalability: less than 1W or up to 10 cores 2025 delivery

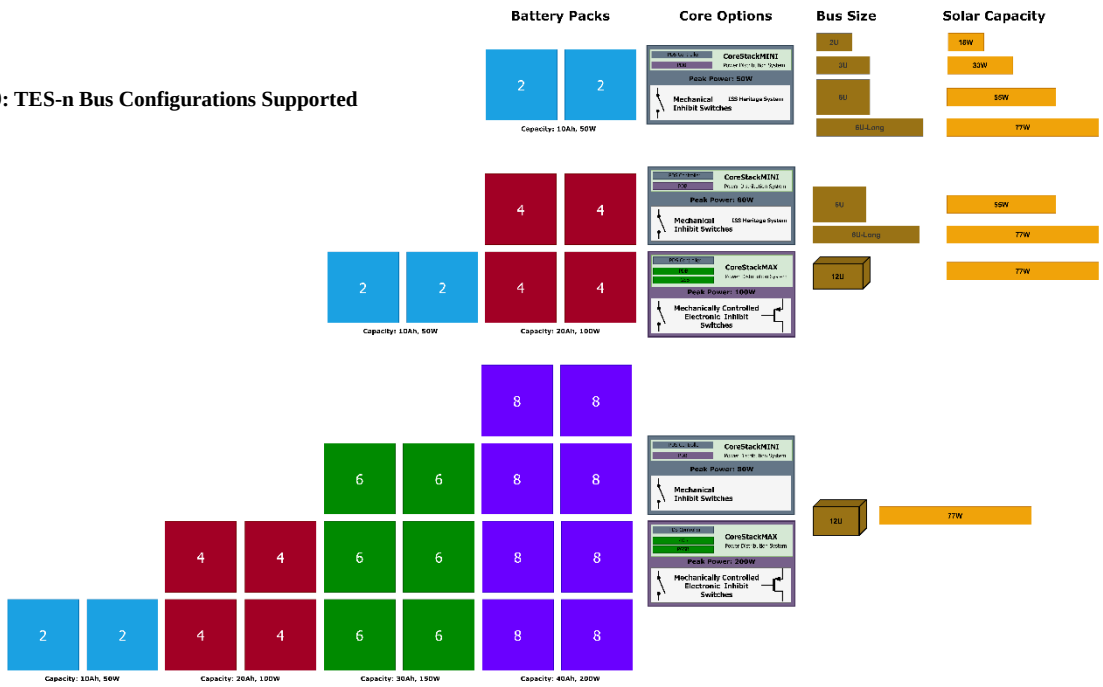

Figure 9: TES-n Bus Configurations Supported

CONCLUSION

The major challenge with any deep-space mission, whether distributed across a swarm or monolithic, is not only to survive and gather data, but to collect scientifically useful information with minimal human interaction. The ability to execute large AI/ML models may present a solution to the communications and data processing limitations faced by the space sciences field. By analyzing large amounts of data as it is generated, scientifically meaningful results can be picked out or statistically generated and compressed for transmission, greatly reducing the required communications budget. This solution is not as simple as improving software, as such models require substantial processing capabilities, and a change in how they are treated compared to research models in labs. The size, weight, power, processor efficiency, algorithm type, figures of merit for determining effectiveness, and many other constraints limit AI/ML operations in an environment where it can be difficult to monitor and adjust models, let alone increase hardware resources as a model grows.

While a considerable amount of experimentation is occurring in the cloud or on large, specialized AI/ML platforms, the space community has been behind in accessing and implementing these new technologies. In part, this has been due to the perception of demanding power requirements, implications of non-radiation tolerant devices, the incompatibility of physical form factors, or perhaps general accessibility in a silicon shortage. Building on earlier successes with GPUs, and more recently, a neuromorphic processor flight test, the notion of a collaborative BrainStack orbital AI/ML laboratory module is presented. The intention is to be able to perform experiments on multiple hardware and software AI/ML elements on the same flight with different collaborative teams. The TES-n Common AI/ML Software Interface will permit a menu driven set of experiments across individual elements. The next set of TES BrainStack experiments will host combinations of GPUs and neuromorphic processors, with flexibility to support upcoming novel systems and their unique interfaces. Such a collaborative BrainStack system will greatly expand the use of these remarkable new tools and methods in the space sector.

REFERENCES

1. Hanson, J. and Chartres, J. and Sanchez, H. and Oyadomari, K., "The EDSN Intersatellite Communications Architecture", 28th Annual Small Satellite Conference, paper SSC14-WK-2, 2014.
2. Cannon, H., "Starling – Objectives and Preparation for Launch", 36th Annual Small Satellite Conference, NASA Short Talk, 2022.
3. Jonsson, A. and Morris, R. A. and Pedersen, L., "Autonomy in Space: Current Capabilities and Future Challenge", AI Magazine, 28(4), 27. doi:10.1609/aimag.v28i4.2066, 2007.
4. Tipaldi, M. and Glielmo, L., "A Survey on Model-Based Mission Planning and Execution for Autonomous Spacecraft", IEEE Systems Journal, 12(4), pp 3893-3905, 2018.
5. Labrèche, G., and Alvarez, C. G., "SaaSyML: Software as a Service for Machine Learning On-Board the OPS-SAT Spacecraft", 2023 IEEE Aerospace Conference, Big Sky, MT, USA, 2023.
6. Murbach. M., Salas, A., Lowry, M., Barszcz, E., Briones, J., Schemmel, P., Mercury, M., Taha, T., Priscal, C., Ntone, R., Zuniga, S., Krześniak, S., Boateng, K., Schisler, S., Stone, T., Alena, R., Gannon, A., Dudukovich, R., Downey, J., Doxley, C., Chelmins, D. "TechEdSat-13: The First Flight of a Neuromorphic Processor". CubeSat Developer's Workshop, 2022.
7. Murbach. M., Salas, A., Lowry, M., Barszcz, E., Briones, J., Schemmel, P., Mercury, M., Taha, T., Priscal, C., Ntone, R., Zuniga, S., Krześniak, S., Boateng, K., Schisler, S., Stone, T., Alena, R., Gannon, A., Dudukovich, R., Downey, J., Doxley, C., Chelmins, D. "TechEdSat-13: On the importance of nano-satellites and the rapid advancement of artificial intelligence and machine learning for space applications". Proceedings of the 1st Artificial Intelligence in Action Conference, San José State University, San José, CA, USA 95112 and Athens, Greece, 2022.
8. Murbach, M., Stone, T., Guarneros-Luna, A., Priscal, C., Salas, A., Kanniainen, H., Ntone, R., Williams, N., Brock, A., Tanner, S., Dono Perez, A. "Topics in Advanced Communication and Design in the TES-n Nanosatellite Flight Series – Use of Iridium® as a Primary Encrypted Command/Control Gateway". NASA Small Spacecraft Virtual Institute Webinar Series, NASA, 2020.