

Deciphering the Enigma of Satellite Computing with COTS Devices: Measurement and Analysis

Ruolin Xing[♦], Mengwei Xu[♦], Ao Zhou[♦], Qing Li[◇]
Yiran Zhang[♦], Feng Qian[▲], Shangguang Wang[♦]

[♦]Beijing University of Posts and Telecommunications, [◇]Peking University, [▲]University of Southern California

ABSTRACT

In the wake of the rapid deployment of large-scale low-Earth orbit satellite constellations, exploiting the full computing potential of Commercial Off-The-Shelf (COTS) devices in these environments has become a pressing issue. However, understanding this problem is far from straightforward due to the inherent differences between the terrestrial infrastructure and the satellite platform in space. In this paper, we take an important step towards closing this knowledge gap by presenting the first measurement study on the thermal control, power management, and performance of COTS computing devices on satellites. Our measurements reveal that the satellite platform and COTS computing devices significantly interplay in terms of the temperature and energy, forming the main constraints on satellite computing. Further, we analyze the critical factors that shape the characteristics of onboard COTS computing devices. We provide guidelines for future research on optimizing the use of such devices for computing purposes. Finally, we have released the datasets to facilitate further study in satellite computing.

CCS CONCEPTS

• **Networks** → *Network measurement*; • **Hardware** → *Testing with distributed and parallel systems*.

KEYWORDS

Satellite Networking, Earth Observation, Satellite Computing, Network Measurement, Dataset

ACM Reference Format:

Ruolin Xing[♦], Mengwei Xu[♦], Ao Zhou[♦], Qing Li[◇], Yiran Zhang[♦], Feng Qian[▲], Shangguang Wang[♦]. 2024. Deciphering the Enigma

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACM MobiCom '24, September 30–October 4, 2024, Washington D.C., DC, USA

© 2024 Association for Computing Machinery.

ACM ISBN 979-8-4007-0489-5/24/09...\$15.00

<https://doi.org/10.1145/3636534.3649371>

of Satellite Computing with COTS Devices: Measurement and Analysis. In *International Conference On Mobile Computing And Networking (ACM MobiCom '24), September 30–October 4, 2024, Washington D.C., DC, USA*. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3636534.3649371>

1 INTRODUCTION

The resurrection of aerospace technologies is making low-Earth orbit (LEO) satellites a promising frontier in mobile edge computing [4, 10, 12, 30]. The Commercial Off-The-Shelf (COTS) computing devices are becoming core units of in-orbit computing. They offer the unparalleled advantage of reusing advanced embedded systems deployed on Earth, thereby facilitating rapid development and deployment of computational capabilities. Several initiatives have successfully employed COTS computing devices to boost LEO satellite computational power. For instance, ESA's PhiSat-1 [20] using Intel's specialized image processing chip to achieving high-speed, low-power in-orbit image data processing, CubeSat Missions [40, 66], and even more commercial projects [22, 56, 64, 65, 67].

Utilizing COTS devices in space for computing purposes is challenging due to the harsh, distinct space environment compared to those on Earth. This is primarily manifested in two aspects: (1) Thermal Control: The vacuum of space prevents heat dissipation through air convection [42]. The frequent alternation between daylight and eclipse zones leads to significant variations in the thermal conditions external to the satellite systems [63]. (2) Energy Management: Satellites can only capture and store a very limited amount of solar energy [29]. The energy is subject to internal contention within the satellite systems [78], necessitating prioritization for the normal operation of the satellite platform [3]. Additionally, radiation effects can induce damage or anomalies in COTS semiconductor devices. Such effects are generally not pronounced at LEO around 500 km. The impact of radiation effects is briefly discussed in §6.

Within the satellite industry, pre-launch methods [1, 2, 6, 77] are used to address these issues by incorporating *sufficient design margins*: (1) For thermal control, the design [42] of thermal structures is informed by results from thermal analysis [2], with appropriate thermal materials [63] selected to facilitate insulation, conduction, and dissipation

processes, culminating in tests such as thermal vacuum experiments. Across these critical steps, margins are incorporated to ensure that, the thermal load on various satellite components does not exceed threshold values under normal conditions. For instance, if the design temperature range for the satellite interior is -10 to 30°C , the active and passive thermal management systems typically limit temperatures to a range of 0 to 20°C . (2) Energy management primarily involves the acquisition, storage, and distribution of energy, corresponding to the solar cells, batteries, and power control/distribution boards, respectively. These key components undergo simulation, selection, design, and testing to ensure their functionality [1, 35, 54]. To accommodate the satellite's overall power consumption budget, energy management systems also incorporate margins. For example, the remaining battery capacity in satellites is usually maintained above 70% to ensure sufficient energy availability for unexpected demands, while also mitigating battery degradation.

Despite the rigorous processes to ensure operational integrity, a pivotal question remains unanswered: to what extent do thermal control and energy management impact the *computing capability* and *reliability* of COTS computing devices in actual satellite systems? Furthermore, how can computing tasks be scheduled to both comply with the constraints of the satellite platform and maximize computing capabilities?

To this end, the satellite community has developed mathematical models and simulation tools for analyzing temperature [2, 8, 15, 42, 63] and energy [1, 14, 35, 54] variations mainly in the pre-launch stage, aiding in the theoretical assessment of temperature and energy constraints. The computer science community has devised simulation platforms [32, 33, 51] tailored for satellite networks and in-orbit computing, capable of conducting pure or hardware-in-the-loop simulations of satellite computing workflows. Building on this foundation, *real-world data* would enable the synergistic integration of both strands of work, facilitating simulations of satellite computing that mirror authentic conditions and allowing for the effective scheduling of actual computing tasks on satellites.

However, a genuine in-orbit testing of COTS computing devices presents a plethora of challenges that extend beyond mere technical considerations. First, the construction of a real satellite system is an expensive and intricate process. Even for a small LEO satellite, the costs can be ranging from tens to hundreds of millions of dollars [46]. The "one-shot" nature of satellite make it impossible to maintain the system [48]. Second, COTS computing devices are exposed to a highly dynamic environment, making experimental design challenging when trying to eliminate the irrelevant factors [6]. For instance, the eclipse and daylight periods, the thermal structure of the payloads, differences in payloads, and

variations in work loads can all influence temperature and energy consumption measurements. These multiple factors converge, creating a complex system landscape with various variables and extensive experimentation. This complexity makes direct comparisons and clear conclusions challenging.

To overcome these challenges, we invest two years and over one million dollars in preparation and testing, creating a realistic satellite system. The satellite has been operating normally in orbit for over 1 year, completing more than 6000 circuits around the Earth. We have successfully conducted various experiments on typical COTS computing devices, amounting to over 1000 hours of investigative work. Simultaneously, we develop a terrestrial testbed, replicating the satellite structure for COTS computing devices in a 1:1 ratio to minimize the impact of unrelated factors. We design a series of experiments for key factors that affect temperature and energy, classifying and contrasting each major influencing factor. After 6-month experimental data collection, we amass 10,000,000+ lines of telemetry information and over 22GB of trimmed text data. The findings from these experiments can be summarized as follows:

- **Temperature—the Critical Bottleneck for In-orbit Computing.** The heating of COTS device chips on satellites may lead to frequency throttling, resulting in an up to 10% reduction in computational performance. A computation task lasting 10 hours, operating at approximately 9 Watts, causes the surface temperature to rise beyond the operational limit (e.g. 30°C in our case) [77], leading to instability. This phenomenon can be attributed to the passive heat dissipation mechanism [6]. Limited by the satellite's volume and weight budget, active overheating control is generally unacceptable. In comparison, the terrestrial systems can leverage air convection and cooling devices to quickly stabilize temperature. The implication is a need for computational task planning in orbit to avert prolonged computation and prevent overheating.

- **The Influence of Eclipse and Daylight on Computing.** The eclipse or daylight periods can slightly affect payload chip and surface temperatures, with variations mostly within $\pm 5^{\circ}\text{C}$. However, regardless of whether the cycle begins during the eclipse or daylight periods, the temperatures at the beginning and ending times remain roughly similar, with a difference of less than 2°C . This phenomenon results from the Sun being the primary external heat source, with the change in solar exposure affecting COTS device cooling [6]. For long-duration satellite computing tasks, eclipse or daylight launch consideration is unnecessary, while short-duration computing tasks may benefit from launching in eclipse period. This contrasts with ground edge computing devices, where active temperature control renders environmental temperature changes insignificant.

• **Computing Dominance on Short-term Power Consumption.** Even over short periods (e.g. within a single day), the power consumption of COTS computing devices can have a significant impact on the battery’s Depth of Discharge (DoD), potentially causing it to exceed 30%, even 50%. The satellite’s steady-state power consumption is relatively stable, but an increase in the average DoD from 25% to 30% can reduce the battery’s lifespan by approximately 25% [16], consequently affecting the satellite’s operational longevity. Unlike terrestrial edge computing devices, where battery maintenance or replacement is feasible, the satellite computing tasks must softly control the maximum DoD within 30% for short-term intensive computing.

• **Potential Energy Efficiency Bonus for Computing.** We observe a cyclical pattern in satellite energy collection over two distinct periods: one orbit around Earth and one natural day. But approximately 6% of the converted solar energy remains unutilized. On the other hand, we note that there is a 30% to 40% energy consumption occupied by the computing tasks while the communication power consumption and others power consumption remain stable. This discrepancy arises from an imperfect alignment between the scheduling of satellite computations and the periodic energy collection [77]. When the battery reaches full capacity, the residual solar energy is not consumed by the computational tasks. Unlike typical terrestrial edge devices with on-demand energy supply, in-orbit computing scheduling is better to consider these intricate cycles, optimizing the utilization of solar energy that cannot be stored.

The above results unequivocally establish that *both temperature and energy are significant limiting factors that directly impede the capability and even reliability of onboard COTS computing devices*. In extreme conditions such as overheating, computing tasks, as well as COTS computing devices, may be subjected to an involuntary shutdown. Furthermore, we identify surface temperature and DoD as the key metrics for quantitatively evaluating the constraints imposed by the satellite platform. Given these insights, we advocate for the community to increase focus on the challenges of computing task scheduling under platform constraints to finally meet the goal of long-term stable execution of in-orbit computing tasks. In summary, our main contributions are as follows:

- We conduct extensive measurements on a real satellite equipped with COTS computing devices, filling the gap in empirical evaluation of satellite computing.
- We provide insightful guidance for future research work related to hardware-in-the-loop simulation or task scheduling in the context of satellite computing.
- We have released the detailed datasets at <https://github.com/TiansuanConstellation/MobiCom24-SatelliteCOTS>.




Figure 1: An overview of the measurement setup.

2 MEASUREMENT METHODOLOGY

2.1 Real-world Satellite System

The BUPT-1 is a 12U small satellite platform, with performance parameters slightly exceeding mainstream 6U and smaller CubeSats [68]. The platform can provide robust support for COTS computing devices, i.e. the payloads. As indicated in Table 1, it operates in a Sun-synchronous orbit, a common choice for experimental satellites [62]. Therefore, our experiments on BUPT-1 regarding temperature control and power management can reflect the mutual interactions and impacts between satellite COTS computing devices and the satellite platform, without losing generality.

We use two ground stations to control the satellite (see Fig. 1). The telemetry, tracking, and control (TT&C) station enables loading command sequences onto the satellite for execution and controls. The data transmission station is responsible for receiving the telemetries [18] and payload data generated by the satellite platform and COTS computing devices respectively.

2.2 Onboard Devices and Experiments

Hardware. Two types of COTS computing devices, Raspberry Pi 4B and Atlas 200 DK are selected for the experiments. The Raspberry Pi 4B, already employed in several successful in-orbit missions [13, 66], has proven its feasibility and stability. Featuring an ARM Cortex-A72 chip with four cores, it is well-suited for general-purpose in-orbit processing such as photo control, image processing, and AI inference. The Atlas 200 DK integrates the Ascend 310 AI processor [24] specifically for AI applications. This processor has built-in circuits capable of handling tasks like image and video encoding/decoding, AI classification, and AI inference. Furthermore, Atlas 200 DK offers four adjustable computational levels, i.e. Low, Mid, High, Full, corresponding to varying power consumption and computing performance.

As shown in Fig. 1, we deploy two Raspberry Pi 4B and two Atlas 200 DK. These COTS computing devices were labeled for identification purposes, with the Raspberry Pi denoted as

Table 1: Basic Information of BUPT-1

Parameter	Value
ORCID	55261
Orbit	Sun-synchronous Orbit
Orbital Altitude	487.607 km to 494.651 km
Orbital Inclination	97.3710°
Mass	17.44 kg
Volume	434.5 mm × 346.6 mm × 340.8 mm
Internal Thermal Range	-10°C to 30°C
Solar Panel Size	1199.3 mm × 868.4 mm × 423.4 mm
Battery Capacity	115 Wh × 2
TT&C	4.8 kbps uplink, 9.6 kbps downlink
Data Transmission	1 Mbps uplink, 100 Mbps downlink

Pi-A and Pi-B, and the Atlas 200 DK denoted as Atlas-A and Atlas-B. Pi-A and Atlas-B were primarily utilized for most of the experiments, while the other two devices served as control groups for comparing.

Experiments. On Raspberry Pi 4B, we deploy three kinds of experiments, i.e. stress tests, image segmentation, and image inference. The stress tests are divided into four groups, with each group respectively invoking 1 to 4 CPU cores (denoted as level 1 to 4). The image inference utilized 4 different models, i.e. SSD-MV1, YOLO-Fastest, YOLOv3, and YOLOv3-Lite. On the Atlas 200 DK, we also deploy experiments in three main categories. The first category is image encoding (denoted as 'jpege'). The second is image classification (denoted as 'imgcl'), utilizing the ResNet50 model. The third is object detection (denoted as 'od'), utilizing the YOLOv3 model for image or video stream inference. On Atlas, the inference tasks can be offloaded to the AI processor by 1 or 4 threads (denoted as 1T or 4T).

Characteristics. The satellite computing tasks generally have the following traits:

- (1) The data generation rate far exceeds the processing capability [45, 52]. This implies the strong motivations of satellite computing.
- (2) The tasks are stateless, meaning previous computations do not affect current ones. This is applicable to satellite-generated raw data, like images. Common image processing tasks like segmentation [17], classification [40], or object detection [31, 53] fit the stateless definition.
- (3) The computing tasks could tolerate a certain degree of latency. Low-latency processing is a key metric for satellite tasks [71], often demanding more computing resources in a short time. However, within the constraints of the satellite platform, the focus is not on real-time requirements but more on the computational volume [12]. In this paper, we focus on the latter case, aiming at maximizing the processed data volume.

Table 2: Collected Data in Telemetries

Date Name	Unit	Sample Rate
Platform Total Voltage/Current	mV/mA	1s
Battery Voltage/Current	mV/mA	4s
MPPT Input Voltage/Current	mV/mA	3s
MPPT Output Current	mA	3s
Atlas-A/B Current	mA	1s
Pi-A/B Current	mA	1s
X-Band Transceiver Current	mA	1s
TT&C Module Current	mA	1s
Surface Temperature of Atlas-A/B	°C	4s
Surface Temperature of Pi-A/B	°C	4s

2.3 Data Collection and Processing

Payload and Telemetry Data. The data collected from BUPT-1 can be divided into payload and telemetry data. Payload data is generated by the onboard COTS computing devices. It includes measurements like chip temperature, CPU frequency, CPU usage per process, memory occupancy, image inference latency, etc. Telemetry data is produced by the satellite platform. It mainly encompasses categories of temperature, energy consumption, satellite attitude, orbit, and more. We list the essential data types in Table 2. From the table, the term *MPPT* stands for maximum power point tracking, a technique employed in photovoltaic solar systems to optimize the power output.

Data Collection. The BUPT-1 is successfully launched into orbit on January 15th. After approximately two months of stability testing, we conduct official experiments from March 22nd to July 30th, resulting in over 22GB of telemetry data and over 30GB of payload experiment data. The datasets are available at <https://github.com/TiansuanConstellation/MobiCom24-SatelliteCOTS>.

Representativeness. Influences introduced by irrelevant factors may destroy the representativeness of our study. We make the following 3 efforts to eliminate irrelevant factors in terms of the thermal conditions, the time variations and the program correctness:

- We build a testbed on the ground for the two COTS computing devices with the completely same thermal alteration (see Fig. 1).
- We keep the time of the onboard COTS computing devices synchronized with the standard UTC time to avoid errors introduced by time variations.
- We develop a program to generate identical command sequence both on satellites and the ground, minimizing the error introduced by the software to the greatest extent.

3 TEMPERATURE RESULTS

For satellites equipped with COTS computing devices, the typical temperature control process involves [6, 63]: (1) heat




Figure 2: Temperature Overview.

generation from the computing tasks; (2) the heat transfer to a spreader through thermal pads and an aluminum alloy enclosure; (3) the subsequent conduction of heat from the spreader to the satellite’s external structure for dissipation through radiation. Fig. 1 illustrates the thermal structure of the satellite, where the deep red area represents the heat spreader and the light red area indicates the satellite’s external structure utilized for heat dissipation. Table 1 presents the internal temperature design range for the satellite, set between -10 to 30°C . Therefore, the temperatures of the aluminum alloy enclosures and the heat spreader must not exceed this range, as it could potentially impact the satellite platform equipment and lead to overall system instability.

Uncontrolled temperature variations can have catastrophic consequences for a satellite. If the temperature surpasses its operational limit [77], components may degrade or fail, leading to diminished performance or total system failure [63]. In extreme instances, this could culminate in irreparable damage to the spacecraft, underscoring the paramount importance of maintaining thermal regulation [42]. Although COTS computing devices are subjected to thermal analysis, design, and testing in the pre-launch stage, there remains a potential risk for excessive internal satellite temperatures during prolonged durations under heavy workloads. This risk arises primarily from two factors:

- (1) COTS computing devices and their accompanying cooling structures occupy a significant volume and weight on small satellites (see §2.1). Under the passive cooling architecture, once the heat generation rate of the COTS computing devices surpass the heat dissipation rate of the satellite’s external structure, the surface temperature will rapidly increase. Given the relatively low efficiency of radiative cooling, prolonged high-load computing can easily lead to overheating within the satellite’s interior.
- (2) The thermal vacuum testing ensures that COTS computing devices can function at sufficiently high power in a vacuum environment, but only for a relatively short duration. The design of thermal isolation can largely isolate

the heat exchange between COTS computing devices and the external environment as well as critical internal components of the satellite [42]. However, these testings are unable to guarantee the safety of the real satellite system once the internal temperature exceeds the designed safe range.

To study the actual impact of COTS computing devices on temperature, we define the *surface temperature* to indicate the temperature of the aluminum alloy enclosure and the heat spreader. We define the unified temperature metric for the two components because the heat spreader and enclosure rapidly achieve thermal equilibrium. The surface temperature should be within the range of -10 to 30°C . Furthermore, the rise in enclosure temperature originates from the increase in the chip temperature of the COTS computing devices.

Consequently, we also emphasize the *chip temperature* as the indicator of heat generation. The limitations on chip temperature are specified in their respective documentation [19, 25]. Finally, we use *external temperature* to indicate the temperature of the satellite’s external structure.

Additionally, the alternation between daylight and eclipse zones have a direct impact on the external temperature. Recall that the aluminum alloy enclosure of COTS computing devices ultimately dissipates heat through radiation via the satellite external structures (see Fig. 1). And variations in external temperature can affect the cooling efficiency. Furthermore, due to the contact thermal conduction between the aluminum alloy enclosure and the COTS computing devices themselves, the impact on surface temperature and chip temperature may also be studied.

Preliminary Results. In Fig. 2, we record the surface temperature variations for 9 hours and 2 days respectively. Fig. 2a illustrates the variation of surface temperature over 2 days. During 90% of this time, the surface temperature remains at or below 4°C . This suggests that typical COTS computing tasks exert normal influence on the surface temperature. Fig. 2b depicts the surface temperature over a 9-hour span, a period without computing and communication tasks. In this state, the surface temperature is observed to lie within a range of 0 to 1°C . Only slight variations occur in tandem with the transition between the eclipse zone and the daylight zone.

3.1 Chip/Surface Temperature Variations

Experiment Methodology. There are 4 principal factors influencing both the surface temperature and the chip temperature: device power consumption, computing load, task duration, and the alternating daylight and eclipse periods. Among these, the computing load can directly affect device

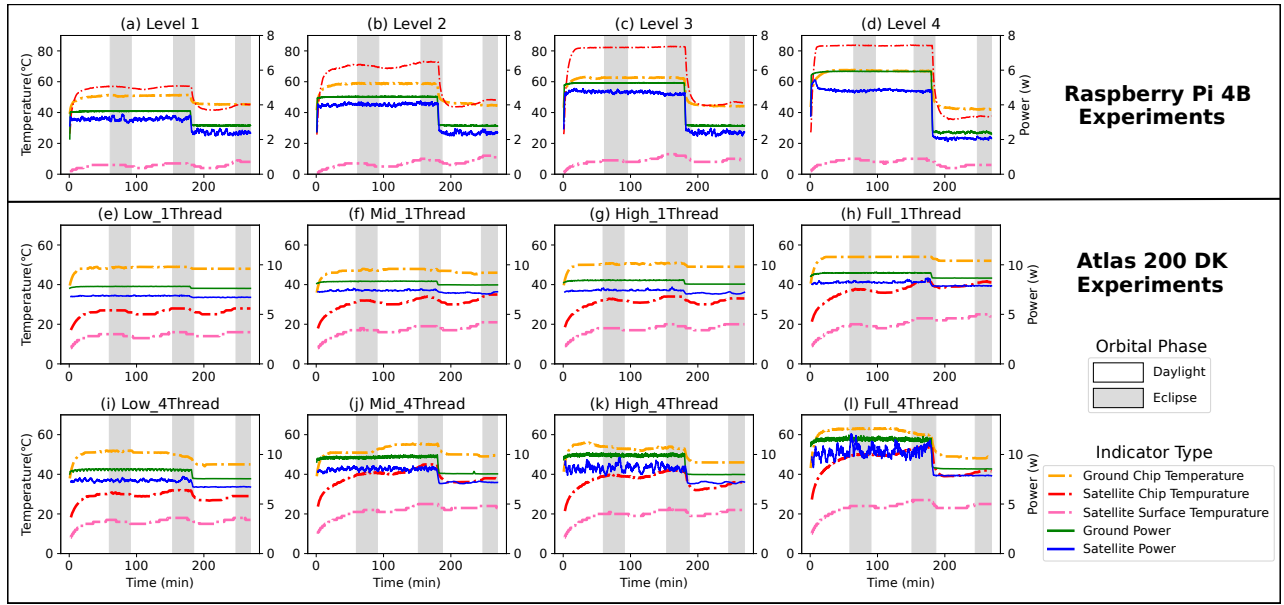


Figure 3: Temperature and Power Variations for Atlas and Pi in Satellite-Terrestrial Environments.

power consumption. To investigate these variables, we design 12 experiments. Each computing experiment is controlled to last 5 hours, with the first 3 hours (approximately equivalent to two orbital periods) devoted to computing, followed by 2 hours of halted computing, reducing the load to an idle state. For the Raspberry Pi, we measure the variations in both temperatures when utilizing 1 to 4 cores for computation (see Fig. 3a to Fig. 3d). For the Atlas, we execute image inference tasks with both single-thread and 4-thread configurations, recording the temperature variations at four power levels for each scenario (see Fig. 3e to Fig. 3l).

Characteristic Difference. The variation in chip temperature and surface temperature when COTS computing devices execute computing tasks exhibits significant disparities for different devices.

For the Raspberry Pi, the upper-plot on Fig. 3 shows that they are more influenced in chip temperature by computing, with minimal effects on surface temperature. Under level 3 and level 4 workloads (see §2.2), the Raspberry Pi reaches a saturation temperature of approximately 82°C. The saturation temperature refers to the point at which the COTS computing devices and satellite’s external structure reach a thermal equilibrium. This equilibrium leads to a stable thermal state. However, its surface temperature rarely exceeds 15°C, far below the operational limit, i.e. 30°C.

For the Atlas, conversely, the computing impact the surface temperature more and the chip temperature less. When Atlas operates at Full capacity with 4 threads (see Fig. 3i), after approximately 3 hours, the surface temperature reaches a maximum of about 24°C. This result approaches the platform’s daily operational temperature design limit of 30°C.

Comparative ground experiments reveal that as Atlas computational level increases from Low to Full, the power level and temperature on the satellite gradually approach the values obtained under the same computing tasks on the ground.

Analysis. The observed differences in temperature behavior between the two devices can be attributed to several factors. Firstly, the dissimilar thermal designs employ for the two devices on the BUPT-1 contribute to this phenomenon (see Fig. 1). The Atlas, with its higher power consumption, is equipped with an enhanced thermal design during the pre-launch stage. it also includes an addition of thermal conductive plates to its aluminum alloy enclosure for heat dissipation. The Raspberry Pi, in contrast, is enveloped in a simpler aluminum alloy enclosure, resulting in inferior heat dissipation.

Secondly, the nature of the COTS computing devices themselves plays a role. Satellite heat dissipation relies on contact cooling. The efficiency of temperature conduction depends on factors such as the thermal conductivity, thickness, and surface area of the contact material. Both the Atlas and Raspberry Pi utilize aluminum alloy and thermal pads for cooling, with similar thermal conductivity and thickness. However, they differ in chip surface area, with the Raspberry Pi’s being smaller (14mm × 14mm) and the Atlas’s larger (52.6mm × 38.5mm). Consequently, the reduced heat dissipation area on Raspberry Pi restricts the degree of chip cooling. Besides these two factors, the variance in power consumption levels between the two devices may also be considered.

Implications. The experiments conducted reveal critical insights into the varying emphasis required for different COTS computing devices, particularly concerning the design of thermal structures, the scheduling of computing




Figure 4: Frequency Throttling of Pi in Space.




Figure 6: Atlas Surface Overheating.

task duration, and the magnitude of computing loads. (1) For low-power computing devices such as the Raspberry Pi, *the utilization of materials with higher thermal conductivity* is advisable. Such an approach can slow down the time taken to reach saturation temperature and accelerate the heat dissipation process; (2) Conversely, for high-power, high-performance devices like the Atlas, attention must be paid to *controlling both the load and duration of computing tasks*. Prolonged operation must be avoided to prevent excessive surface temperature, which could potentially compromise the stability of the entire satellite platform.

3.2 Overheating

In the experiments in §3.1, overheating phenomena occur in both the aluminum alloy enclosure and the chips: the surface temperature of Atlas approaches 30°C after 3 hours of computation; the chip temperature of Raspberry Pi escalates to above 80°C. This observation raises two pertinent questions:

- **Q1:** Does Atlas induce enclosure overheating with extended, high-load computing tasks?
- **Q2:** What impact does elevated chip temperature have on the computing capabilities of the Raspberry Pi?

To investigate the **Q1**, we design two experiments. The Experiment 1 (denoted as "Single On" in Fig. 6) extends the task duration to 9 hours with a configuration of 4-thread and level Full on Atlas-B to monitor the surface temperature changes. The second experiment (denoted as "Double On" in Fig. 6), simultaneously activates Atlas-A and Atlas-B. Atlas-B executes image inference with the same configuration in Experiment 1, and Atlas-A performs a 4-thread stress test task using only CPU. We deploy a lower load on Atlas-A, because




Figure 5: Atlas Latency Variations.

the satellite does not support both Atlases running high-load tasks concurrently in the ground temperature testing. A target computing time of 3 hours is set for both devices.

Enclosure Overheating. Experiment 1 (see the lower-plot on Fig. 6) reveals that Atlas-B reaches approximately 17°C after 50 minutes of computation and 30°C, the operational limit, after 540 minutes. Atlas-B ceases its computing task normally at this point, and 15 minutes later, its surface temperature returns to around 17°C. Experiment 2 (see upper-plot on Fig. 6) indicates that Atlas-B's surface temperature reaches 30°C around 40 minutes in, persisting for about 100 minutes. Anomalously, Atlas-A is forcibly halted around 50 minutes (process killed by OS), and Atlas-B eventually gets shut down around 110 minutes. This phenomenon underscores that the operational limit of surface temperature is a key constraint on the computing duration and load of onboard COTS computing devices.

To explore the **Q2**, we deploy image inference applications using various models (see §2.2) on the Raspberry Pi. All the tasks employ the same task duration as in §3.1 to evaluate the impact of chip overheat on computing performance of the Raspberry Pi.

Chip Overheating. The chip overheat leads to a mild frequency throttling of approximately 5% in the onboard Raspberry Pi. While an identical task running on the ground Raspberry Pi exhibits no such frequency throttling. For Raspberry Pi, the impact of this frequency throttling varies across different models. In the case of YOLOv3 and YOLOv5, the average inference latency increases by roughly 10% after the temperature reaches saturation. Simultaneously, for inference tasks other than image segmentation, the absolute inference time rises by around 5% to 8%. For Atlas, the increase in chip temperature has almost no effect on computing performance. As the Fig. 5 shows, the difference between satellite and ground computing performance is less than 1%. Furthermore, under 4 different power levels, Atlas exhibits some variation in performance. The Mid and High power levels improve approximately 45% over Low. The Full power level has around a 55% performance boost over the Mid and High power levels.

Analysis on Enclosure Overheating. For the Experiment 1, *the task duration* plays the key role in causing the

Table 3: Temperature: Differences of Space and Ground

	Average Power (W)		Ascending Time (min)	
	Space	Ground	Space	Ground
Pi (Idle)	2.36	2.80	None	None
Pi (1 Core)	3.16	3.63	161	67
Pi (2 Core)	3.99	4.47	155	41
Pi (3 Core)	4.66	5.26	125	38
Pi (4 Core)	4.82	5.92	86	39
Atlas (Idle, Low)	6.67	7.61	None	None
Atlas (1T, Low)	6.81	7.81	60	27
Atlas (4T, Low)	7.31	8.46	147	35
Atlas (Idle, Mid)	7.08	8.03	None	None
Atlas (1T, Mid)	7.35	8.32	155	53
Atlas (4T, Mid)	8.45	9.71	169	154
Atlas (Idle, High)	7.07	8.03	None	None
Atlas (1T, High)	7.37	8.42	148	40
Atlas (4T, High)	8.54	9.91	170	36
Atlas (Idle, Full)	7.75	8.63	None	None
Atlas (1T, Full)	8.17	9.17	164	145
Atlas (4T, Full)	10.19	11.48	147	88

aluminum alloy enclosure of Atlas slightly exceeds the operational limit. In the Experiment 2, both of computing tasks on Atlas-A and Atlas-B are passively terminated due to temperature elevation. It further illustrates that multiple devices running together intensify the rate of surface temperature increase. This phenomenon primarily results from the thermal structure design. On BUPT-1, the two computing devices share an aluminum surface to conserve the internal space. Thus the computing load leads to rapid heat release when the both of the two Atlas are on. To prevent such problem, one solution involves thermally isolating COTS computing devices. But the cost of such isolation is typically high, making it generally applicable only to critical satellite components like batteries [6].

Analysis on Chip Overheating. Raspberry Pi employs Dynamic Frequency Scaling (DVFS) as its CPU frequency control method. This mechanism causes Raspberry Pi to actively frequency throttling when the chip temperature exceeds 80°C. Given the cooling capabilities on BUPT-1, Raspberry Pi reaches an active frequency throttling zone of over 80°C within 50 minutes under load levels above 75%. Atlas, on the other hand, does not have a frequency control mechanism similar to DVFS; instead, its power consumption level is adjusted at fixed levels across 4 power levels.

Implications. *Overheating phenomena limit the duration and load level of satellite computing and might degrade the computing performance of the SoC.* (1) In Experiment 1, Atlas possibly reaches the 30°C operational limit in surface temperature after 9 hours of computing. This achievement occurs at an 8.5W power consumption level under the thermal conditions specific to Atlas after modification. It highlights the need for considering satellite platform temperature constraints for specific COTS computing devices under high-load computing tasks. *Excessive computing time becomes the primary limiting factor in this case.* It necessitates the reduction of the maximum task duration for a single experiment

and the setting of reasonable intervals for cooling; (2) In Experiment 2, parallel computing with two Atlas devices leads to the structure temperature reaching the 30°C operational limit in a short time. The computing tasks on both devices are eventually forcibly terminated due to overheating. This further illustrates that onboard COTS computing devices cannot operate for extended periods under excessive load. *Excessive computing load becomes the main limiting factor in this case.* Furthermore, even without excessive load, overheating may still cause a decline in computing performance for specific devices, such as the Raspberry Pi.

The experiments in this section provide insights for *modeling or conducting hardware-in-the-loop simulations* of the capabilities of onboard COTS computing devices. Computing with COTS on a satellite requires full consideration of thermal constraints. Constraints must be made on task duration, computing load, and computing performance according to the specific computing device. Additionally, it is essential to design more efficient, space-saving, and cost-effective cooling systems for the deployment of multiple COTS computing devices on the satellite.

3.3 Heating Rate

In the experiments of §3.1, differences are observed in both the power and the temperature rise rate when COTS computing devices execute identical tasks under satellite and ground conditions. However, two problems still confront us: how to conduct a quantitative assessment of the differences in these two aspects; and how to clearly delineate the variations in the impact of space and terrestrial environments on both facets. To clarify on these two points, we calculate the average power and times to reach maximum temperature for different computing tasks from the experiments in §3.1. Table 3 illustrates the average power consumption and temperature rise time for COTS computing devices operating under different computing tasks in both satellite and ground environments. The temperature rise time quantifies the time taken by the COTS computing devices from task initiation to the attainment of maximum temperature.

Rate of Temperature Increase. Generally, the chip temperature rises more smoothly in space than on the ground. However, for both space and ground conditions, the chip temperatures on the Raspberry Pi reach a stable state more quickly than those on the Atlas (see Fig. 3). In Table 3, the time required for the satellite-based COTS computing devices to heat to saturation temperature is longer, amounting to 2-4 times the duration on the ground. Moreover, within a computing span of 3 hours, the satellite experiments in most cases reach their maximum values in the final 20% of the time.




Figure 7: Surface Temperature and Power Variations of COTS Computing Devices Starting at Eclipse or Daylight.




Figure 8: Pi Latency.

Figure 9: Atlas Latency.

Analysis. Two factors may contribute to the disparity in heating time between satellite and ground. First, the vast difference between the vacuum of space and ground environment matters. Although we make sure that the both cases of space and ground utilize the same passive contact cooling mechanism (see §2.3). Convection through the air allows the COTS computing devices on the ground to achieve saturation more quickly, while the vacuum in space doesn't have the same effect [6]. Second, eclipse and daylight zones exert certain influences. for Raspberry Pi operating under level 3 and level 4 (see Fig. 3c and Fig. 3d), saturation temperature can be reached within 30 minutes. But the peak temperature is only attained after 120 minutes, largely owing to a slight temperature increase caused by passing through the daylight zone.

Implications. This characteristic of satellite temperature can be mitigated by intermittent computing [39] to reduce the average temperature over a period. This strategy leverages the slower heating rate in space, allowing for a more controlled thermal environment and potentially enhancing the overall system's reliability and performance.

3.4 Daylight and Eclipse Zones

From previous experiments, it is understood that the periodic variations of the daylight and eclipse zones exert a certain influence on both the surface temperature and chip temperature (see §3.1 and §3.3).

Due to the alternating cycles of eclipse and daylight in Sun-synchronous orbits, satellite computing tasks can typically be scheduled to start either in the eclipse zone or in the daylight zone. Thus we aim to further explore the extent to which initiating in the eclipse zone or in the daylight zone impacts satellite computing. To elucidate this issue, we compare two scenarios for COTS computing devices: *starting*

in the eclipse zone (denoted as Test2 in Fig. 7) and starting in the daylight zone (denoted as Test1 in Fig. 7). Fig. 7 shows the two The computing tasks, lasting approximately 60 minutes, run the same computing program on both of the Raspberry Pi and Atlas.

Impact of Eclipse and Daylight. Figure 7 illustrates that for the Raspberry Pi, initiating the process in either the eclipse zone or daylight zone has almost no impact on both the chip temperature and surface temperature. However, for the Atlas, the groups starting in the daylight zone ultimately reach a temperature difference of around one hour. In Figure 7, the chip temperature displays a final difference of approximately 8°C, and the surface temperature shows a difference of around 5°C. In Fig. 8 and Fig. 9, there is minimal difference in the performance of different tasks during the aforementioned process.

Analysis. Solar irradiation serves as the main external heat input for the satellite. The passive cooling structure on satellites mainly dissipates heat through its external structure [63]. Thus the impact caused by the eclipse and daylight zones primarily results from uneven irradiation from the Sun toward the satellite external structure during its movement. The difference in thermal conductivity remains the main cause of the distinct temperature performances between these two devices.

Implications. (1) The differences in starting within the daylight or eclipse zone have minimal impact on computing performance. This finding implies that for hardware-in-the-loop simulations targeting satellite computing, this effect can be simplified; (2) On the other hand, in the scheduling of satellite computing tasks, the effect of computing in the eclipse zone or first in the daylight zone on temperature tends to become consistent as the total time increases.

4 ENERGY RESULTS

The Electrical Power Subsystem (EPS) is responsible for handling power generation, storage, and distribution within LEO satellites. The satellite energy recycling relies primarily on solar panels [61], rechargeable lithium batteries [16], and power distribution/control board. As of 2021, approximately 85% of nanosatellite power systems apply these setups [68].

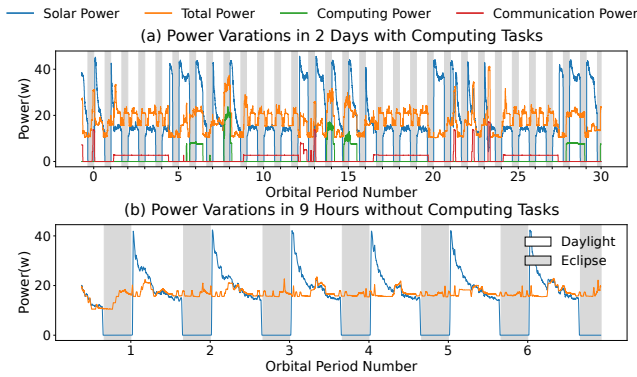


Figure 10: Power Overview.

The EPS efficiently oversees energy conversion, distribution, and regulation, while preserving battery life and catering to the rigors of space and the fluctuating power needs of various mission phases [1, 6].

In terms of energy collection, the deciding factor is the periodic change in sunlight conditions caused by satellite motion. Regarding energy usage, we concentrate on the power consumption of COTS computing devices. As seen in the preliminary measurement results in Fig. 10, computing tasks occupy a significant portion of the non-routine energy consumption. For energy storage, the DoD is a key index determining satellite life. We aim to control the BUPT-1's DoD level below 30%. The general design target is approximately 25% DoD[77], so 30% is not considered excessive.

Preliminary Results. The variation of harvested solar power and individual power changes in computing and communication devices over two days are shown in Fig. 10a. Power consumption pattern reveals four cycles in a day where the average power consumption reaches 40W, with other times approximately 20W. We make further study about this in §4.1. Fig. 10b presents power changes over 10 hours, a period devoid of computing and communication tasks, with the corresponding power consumption being zero. The harvested solar power and the total power consumption of the satellite platform maintain dynamic balance.

Metrics. In Fig. 11, the *Solar Power* refers to the power generated by MPPT (see §2.3) to be used to charge the battery and supply COTS computing devices or other components. The *Total Power* refers to the power consumed by the satellite, including all kinds of consumed power. The *Communication Power*, *Computing Power*, and *Others Power* refers to the consumed power by the communication components, COTS computing devices, and other components on the satellite platform. In Fig. 12 and Fig. 13, we use the corresponding terms, i.e. *Solar*, *Communication*, *Computing*, and *Others*, to denote the energy generated or consumed by the corresponding components.

4.1 Energy Harvesting Patterns

Short-term Power Variation. Figure 11 analyzes a typical cycle featuring a data transmission task. At around 60 minutes, communication power rises to about 18W due to the activation of communication components. At the same time, computing power rises for the activation of Atlas-B. At around 78 minutes, the peak is caused by the beginning of the data transmission. The solar power reaches a peak of approximately 40W within about 20 minutes. It subsequently declines to around 10W, before finally entering the shadowed region and reducing to zero. During the initial 30 minutes depicted in Figure 11, the batteries are nearly fully charged. The solar power exceeds the total power of the satellite platform, resulting in about 3Wh of the solar energy not being utilized by the satellite.

Periodic Energy Variation. In Fig. 12, we adjusted the granularity of time, observing energy variation on the orbital periods. It can be observed that BUPT-1 exhibits periodicity in energy collection at the granularity of an orbital period. Approximately every 15 orbital periods, corresponding to roughly 24 hours, there are two peaks and valleys depicted in Fig. 12. This pattern reveals that the power generation capacity of BUPT-1's solar panels initiates a new cycle approximately every 24 hours.

In Fig. 13, we observe over a specific 60-min time interval in daylight per day. Despite the cyclical nature of energy harvesting, the energy yield remains relatively stable. Within each round, the energy collected by the solar panels predominantly clusters around the average value of 16Wh, without manifesting pronounced peaks and valleys.

Analysis. The satellite's energy collection shows regular patterns over two different time scales: the orbital period and 24 hours. When the cycle is defined by the orbital period, the periodic variation in energy collection primarily stems from the relative motion between the satellite and the Sun, leading to periodic changes in sunlight exposure. In the eclipse zone, the solar irradiance is zero, while in the daylight zone, the satellite's power undergoes an initial increase followed by a decrease. This pattern results mainly from the alteration in the angle of the satellite's solar panels during motion, thereby inducing corresponding fluctuations in the energy received by the solar panels.

When the cycle spans 24 hours, the formation of daily energy collection cycles is attributable to the working mode of the EPS. The EPS can be configured with a power profile to augment the energy reception of the solar panels at specific intervals[77]. However, this enhancement is not without constraints. An increase in the power of the solar panels would elevate their temperature, which must not surpass the tolerance threshold. The continuous cycles of collected energy depicted in Fig. 12 are illustrative of BUPT-1's EPS




Figure 11: Power Variation.

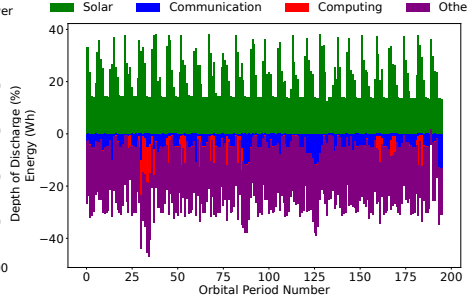


Figure 12: Energy Variation per Orbit.




Figure 13: Energy Variation per Day.

working mode. Out of approximately 15 cycles a day, 10 exceed 14Wh, while the remaining five cycles hover around 14Wh.

Implications. (1) *Stability and Predictability of Energy Collection:* The periodic nature of the satellite’s energy collection ensures that the acquired energy is both stable and predictable. From the perspective of factors determining satellite energy collection across two time dimensions, they demonstrate relative stability over longer time cycles. On one hand, the satellite’s orbital changes and its ability to harness solar energy can be predicted within a small margin of error; on the other hand, the working mode of the EPS is fundamentally determined at the time of the satellite’s fabrication, so its impact on energy harvesting is also well predictable; (2) *Maximizing Utilization of Collected Energy with Satellite Computing:* As depicted in Fig. 11, the energy collected by the satellite is not fully utilized. Given the characteristics of satellite computing (see §2.2), employing the unutilized energy to perform computing tasks helps processing the excessive data. The predictability of solar power can also help the scheduling of computing tasks to avoid energy under-utilization.

4.2 Energy Expenditure Breakdown

Short-term Power Variation. In Fig. 11, at around the 62nd minute, the satellite activates the computing and communication devices, preparing for data transmission. At around the 78th minute, the satellite commences the data transmission, lasting for around 10 minutes, with the communication power staying 38W. Meanwhile, the power consumption of the COTS computing device remains around 14W throughout. Overall, short-term communication and computation power consumption have a significant impact on the battery’s DoD, consuming 20% within 30 minutes.

Periodic Energy Variation. In Fig. 12, the energy consumption of communication remains relatively steady, basically maintaining around 5Wh per round most of the time. But the short-term share of computing energy can even reach 30% to 40%. From Fig. 13, energy collection and consumption maintain balance in 24-hour intervals, with absolute differences generally not exceeding 20%.

Depth of Discharge and Available Energy. Fig. 14 illustrates the variations in DoD and the available energy for each cycle, corresponding to the same time frame shown in Fig. 12. The available energy is calculated by subtracting the energy consumed from the energy collected. The figure also marks the 30% DoD limit and the equilibrium situation where available energy is 0. During the two-week testing period, the satellite maintains an average 16% DoD level. However, during the time period from the second to the third day, the average DoD rises to around 35%, which can severely affect battery degradation. Within a short timeframe, the DoD can even exceeded 50%. From the perspective of available energy, the expected value is 0, representing energy balance within a cycle. But on one hand, there is still 6% excess energy converted from solar energy that is not utilized when the battery is full. On the other hand, there is insufficient energy during compute-intensive tasks.

Analysis. Communication power can rise rapidly in a short time. On a broader time scale, the communication power is more gradual. This is mainly because communication time is relatively limited. Meanwhile, computing energy consumption can greatly exceed the energy collected by the solar panels as shown in Fig. 12. Fig. 13 shows the share of computing task energy consumption remains high. According to the information in Fig. 12 for the corresponding time frame, the COTS device is performing continuous computing tasks from the second to the third day. It results in the DoD level exceeding the 30% limit during that time frame as shown in Fig. 14. Additionally, the consumed energy other than communication and computing is relatively stable.

Implications. (1) *Dominance of Computing Energy Consumption:* The computing tasks are the primary consumers of the variable energy, considering both of power magnitude and duration. Within an orbital period, communication tasks dominate most of the power consumption if present. But over extended time frames exceeding a day, the proportion of energy consumed by communication is not substantial; (2) *Necessity of Reasonable Scheduling for Satellite Computing:* Onboard COTS computing devices must carefully schedule computing tasks to fully utilize energy and to avoid exceeding the DoD limit.




Figure 14: DoD and Available Energy Variation.

4.3 Computational Energy Efficiency

The power supply for onboard COTS computing devices is exclusively provided by the satellite platform. To investigate whether computing efficiency in space differs from that on Earth, we have deployed various types of computing tasks (see §2.2) both on the satellite and on the ground. The satellite is powered by solar panels and batteries, while the ground uses a programmable power monitor. This comparative study aims to uncover potential disparities and understand the underlying factors that might contribute to any observed differences in computing performance.

Comparisons. Tables 4 and 5 present the energy consumption for processing 100 images on Atlas and Raspberry Pi, respectively, as well as the number of images processed per 60-minute. From the tables, we can conclude that: (1) Computing tasks on the satellite are more energy-efficient than those on the ground. For the same task on Raspberry Pi, the difference in energy consumption between the ground and space is trivial (around 2%). For the same task on Atlas, the difference is about 10%; (2) Computing tasks on the ground are more performant than those on satellites. When comparing the number of inference tasks, the ground is greater than or equal to the space. But this is at the cost of more power consumed (see Table 3) in the same time frame; (3) As the computing load increases, the efficiency of energy utilization becomes higher. Tables 4 shows that when Atlas is at the Full level, the energy consumption for processing 100 images is the least.

Analysis. Conclusions (1) and (2) share a common root cause. This root cause pertains to the instability in the power supply of COTS computing devices onboard satellites, leading to a slight decline and consequently a lower overall power profile. However, under these reduced power levels, the performance degradation of COTS computing devices remains minimal. Two factors contribute to this outcome. First, in the daylight zone, the satellite’s attitude affects the stability of the solar power. Second, in the eclipse zone, the voltage of batteries may decline due to either an increase in the DoD or

Table 4: Energy Consumption on Atlas.

Task	Level	Sat Energy (J)	Gnd Energy (J)	Sat Quantity	Gnd Quantity
od	1T, Low	3983.7	4407.3	1846	1912
	1T, Mid	2448.2	2613.8	3242	3439
	1T, High	2428.2	2644.2	3279	3440
	1T, Full	2151.7	2281.1	4098	4342
	4T, Low	180.4	203.6	33	34
	4T, Mid	103.3	117.4	67	68
	4T, High	104.2	120	68	68
imgcl	Mid	1138	1173	2310	2641
	Full	1024	1023	2832	3319
jpeg	Mid	645	665	4030	4601
	Full	550	566	5207	5896

Table 5: Energy Consumption on Pi.

Models	Sat Energy (J)	Gnd Energy (J)	Sat Quantity	Gnd Quantity
yolo-fastest	21	21.38	1640	1742
yolo-v5lite	207.6	212.9	168	176
yolov3	335	345.2	105	113

battery aging. Conclusion (3) aligns with intuitive reasoning. As the computational load escalates, the additional power is exclusively channeled towards computing, hence, under higher computational load, computing efficiency improves.

Implications. (1) *On hardware-in-the-loop simulations:* Overall, the computing efficiency under satellite power conditions proves superior. This insight implies that terrestrial hardware-in-the-loop simulations may overestimate the energy measurements, while the actual satellite computing efficiency likely experiences a moderate increase; (2) *On satellite computing scheduling:* Given that higher loads yield superior efficiency ratios, COTS computing devices should ideally operate under elevated loads. However, this high-load scenario elevates the risk of thermal constraints, necessitating a comprehensive consideration of computing limitations during task scheduling.

5 RELATED WORK

Satellite System and Measurement. Owing to advancements in rocket technology and the deployment of COTS computing devices [13, 44], the costs of LEO satellite systems have significantly reduced [26]. In recent years, an emerging trend centers around system and measurement works grounded in real satellites and authentic data. These efforts can be broadly classified into the following categories: (1) *Real Satellite Systems* [47, 58, 59, 74]; (2) *Satellite System Measurements* [27, 43, 50]; (3) *Satellite Systems Based on Authentic Data* [55, 70, 71]. The majority of existing work focuses on integrated terrestrial-satellite Internet systems

[27, 43, 50, 58, 59, 70, 71]. While a minor fraction is devoted to specific computing tasks [74] or applications [47, 55]. The trend toward the convergence of satellite system networking and computing is becoming increasingly pronounced, facilitating the emergence of new practical application scenarios [4, 34].

Orbital Edge Computing. Due to the rapid data generation on satellites, which significantly outpaces the satellites' data transmission capabilities [45, 52], Onboard Edge Computing (OEC) [4, 10–12, 38] uses an architecture for executing in-orbit computing tasks using COTS devices. The prevailing demand for in-orbit processing emanates primarily from the remote sensing domain [9]. Through in-orbit processing, valuable insights are swiftly and real-time extracted from voluminous remote sensing imagery, or conversely, irrelevant data are discarded. In this context, a multitude of novel algorithms and models have emerged [31, 53, 60, 69, 75], underpinned by COTS computing platforms and massive image datasets. Concurrently, limitations such as energy consumption of COTS computing platforms receive increasing attention within the OEC architecture [10, 12, 37, 75]. Additionally, as the scenarios and requirements for in-orbit processing proliferate [4, 34], task scheduling in in-orbit computing gains in significance [36, 37, 73, 76].

Testbed for Satellite Research. Historically constrained by the high costs and complexity of real systems, research in the domain of satellite networking and computing primarily relies on theoretical modeling and simulation platforms [4, 5, 12, 23, 28, 33, 34, 51, 73]. Significant advancements have occurred in theoretical analysis [7, 36]. But a lack of in-depth system work and empirical validation in real-world environments often results in conclusions that diverge considerably from actual scenarios. This growing divergence underscores the increasing emphasis on the development of hardware-in-the-loop simulation platforms [32] and even fully physical experimental setups [49, 59, 74]. Platforms like StarryNet [32] utilize a hybrid simulation approach to offer the advantage of simulation flexibility while partially maintaining system realism and credibility. Fully physical experimental platforms require substantial initial investment but excel in reproducing the challenges and issues inherent to real-world scenarios. As the costs of satellite manufacturing and launching continue to decline, we envision an increasing proportion of physical components in more experimental platforms [32, 49, 59, 74].

6 DISCUSSION ON RADIATION EFFECTS

In the context of space environmental effects on spacecraft, radiation effects refer to the energy transfer by charged particles [77]. The conduction or ionization caused by a charged

particle can produce various effects on semiconductor structures, with Single Event Effects (SEE) being a common category [41], including Single Event Upsets (SEU) and Single Event Latchups (SEL) [57]. SEUs are errors caused by bit flips in memory or registers during computation. Specifically, the affected area may contain data or instructions, and such errors can be detected and corrected at the software level. SEL represents a more severe error, such as a sudden short circuit in a charged transistor, where prolonged high current can lead to device burnout, causing irreversible damage to the spacecraft.

Protection against SEEs involves multiple dimensions, including reliability designs such as Triple Modular Redundancy (TMR) [21] and Byzantine algorithms [72] at both software and hardware levels. Through integrated software and hardware system design [79], terrestrial high-performance COTS computing devices are also expected to meet aerospace-grade reliability standards.

Due to the protective effect of the Earth's magnetic field, SEEs are less pronounced in low Earth orbits. In our 6-month computational experiment, no computational errors caused by SEEs were observed, despite our computing devices not undergoing special radiation hardening. This demonstrates that COTS computing devices selected with simple electromechanical and thermal modifications are broadly reliable in addressing radiation effects at LEO around 500KM.

7 CONCLUSION

To conclude, in this paper, we build a real satellite system to study the impact of the thermal control and power management on the COTS computing devices in terms of computing constrains and scheduling for computing tasks. We will release the comprehensive datasets and hope to help the community make further progress on the research of the satellite computing system.

ACKNOWLEDGMENTS

We sincerely thank our shepherd and all anonymous reviewers for their valuable feedback. This work was supported by National Natural Science Foundation of China under Grant (NSFC No. 62032003, U21B2016, 62302055, and 62302015) and Postdoctoral Innovation Talents Support Program (No.BX20230011). Ao Zhou is the corresponding author of this work.

REFERENCES

- [1] Japan Aerospace Exploration Agency. 2009. *DESIGN STANDARD POWER SUB SYSTEM*. Industrial Standard JAXA-JERG-2-214. Japan Aerospace Exploration Agency. https://sma.jaxa.jp/TechDoc/Docs/E_JAXA-JERG-2-214.pdf
- [2] Japan Aerospace Exploration Agency. 2009. *SPACECRAFT THERMAL CONTROL SYSTEM*. Industrial Standard JAXA-JERG-2-310. Japan

- Aerospace Exploration Agency. https://sma.jaxa.jp/TechDoc/Docs/E_JAXA-JERG-2-310.pdf
- [3] Vladimir S Aslanov and Dmitry A Sizov. 2021. Chaos in flexible CubeSat attitude motion due to aerodynamic instability. *Acta Astronautica* 189 (2021), 310–320.
 - [4] Debopam Bhattacharjee, Simon Kassing, Melissa Licciardello, and Ankit Singla. 2020. In-orbit Computing: An Outlandish thought Experiment?. In *HotNets*. ACM, 197–204.
 - [5] Vaibhav Bhosale, Ketan Bhardwaj, and Ada Gavrilovska. 2020. Toward Loosely Coupled Orchestration for the LEO Satellite Edge. In *HotEdge*. USENIX Association.
 - [6] Sonja Caldwell. 2022. *State-of-the-Art of Small Spacecraft Technology*. NASA. <https://www.nasa.gov/smallsat-institute/sst-soa>
 - [7] Long Chen, Feilong Tang, and Xu Li. 2021. Mobility-and load-adaptive controller placement and assignment in leo satellite networks. In *IEEE INFOCOM 2021-IEEE Conference on Computer Communications*. IEEE, 1–10.
 - [8] Zhezhen Chen, Haibo Yang, Yongjun Yu, and Songlin Yang. 2023. Thermal control system design and analysis for a micro-nano satellite stays on target satellite. In *Journal of Physics: Conference Series*, Vol. 2472. IOP Publishing, 012048.
 - [9] Mingmin Chi, Antonio Plaza, Jon Atli Benediktsson, Zhongyi Sun, Jinsheng Shen, and Yangyong Zhu. 2016. Big data for remote sensing: Challenges and opportunities. *Proc. IEEE* 104, 11 (2016), 2207–2219.
 - [10] Bradley Denby, Krishna Chintalapudi, Ranveer Chandra, Brandon Lucia, and Shadi A. Noghahi. 2023. Kodan: Addressing the Computational Bottleneck in Space. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3, ASPLOS 2023, Vancouver, BC, Canada, March 25-29, 2023*, Tor M. Aamodt, Natalie D. Enright Jerger, and Michael M. Swift (Eds.). ACM, 392–403. <https://doi.org/10.1145/3582016.3582043>
 - [11] Bradley Denby and Brandon Lucia. 2019. Orbital edge computing: Machine inference in space. *IEEE Computer Architecture Letters* 18, 1 (2019), 59–62.
 - [12] Bradley Denby and Brandon Lucia. 2020. Orbital Edge Computing: Nanosatellite Constellations as a New Class of Computer System. In *ASPLOS*. ACM, Lausanne, Switzerland, 939–954.
 - [13] Steve Guertin Edward J. Wyrwas. 2021. *NEPP Processor Enclave: Post COVID Update*. NASA Electronic Parts and Packaging (NEPP) Program. https://nepp.nasa.gov/workshops/etw2021/talks/16-JUN-21_Wed1030_a_Wyrwas_Guertin-NEPP-Processor-Enclave-20210017075.pdf
 - [14] Justin T Ehren. 2019. Energy Analysis and Orbit Simulation of Actuating CubeSat Solar Arrays. (2019).
 - [15] Assal Farrahi and Isabel Pérez-Grande. 2017. Simplified analysis of the thermal behavior of a spinning satellite flying over Sun-synchronous orbits. *Applied Thermal Engineering* 125 (2017), 1146–1156.
 - [16] JP Fellner, GJ Loeber, SP Vukson, and CA Riepenhoff. 2003. Lithium-ion testing for spacecraft applications. *Journal of Power Sources* 119 (2003), 911–913.
 - [17] Tharindu Fernando, Clinton Fookes, Harshala Gammulle, Simon Denman, and Sridha Sridharan. 2023. Towards On-Board Panoptic Segmentation of Multispectral Satellite Images. *IEEE Transactions on Geoscience and Remote Sensing* (2023).
 - [18] Consultative Committee for Space Data Systems (CCSDS). 2021. *CCSDS 132.0-B-3: Lossless Data Compression*. Technical Report. Consultative Committee for Space Data Systems. <https://public.ccsds.org/Pubs/132x0b3.pdf>
 - [19] Raspberry Pi Foundation. 2023. Raspberry Pi 4 Model B Specifications. <https://www.raspberrypi.com/products/raspberry-pi-4-model-b/specifications/>. Accessed: 2023-04-17.
 - [20] Gianluca Giuffrida, Luca Fanucci, Gabriele Meoni, Matej Batič, Léonie Buckley, Aubrey Dunne, Chris van Dijk, Marco Esposito, John Hefebe, Nathan Vercauysen, et al. 2021. 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 (2021), 1–14.
 - [21] Robért Glein, Florian Rittner, Andreas Becher, Daniel Ziener, Jürgen Fricke, Jürgen Teich, and Albert Heuberger. 2015. Reliability of space-grade vs. COTS SRAM-based FPGA in N-modular redundancy. In *2015 NASA/ESA Conference on Adaptive Hardware and Systems (AHS)*. IEEE, 1–8.
 - [22] Spire Global. 2023. *Lemur-2 Satellites*. Spire Global. <https://www.eoportal.org/satellite-missions/spire-global>
 - [23] Yannick Hauri, Debopam Bhattacharjee, Manuel Grossmann, and Ankit Singla. 2020. "Internet from Space" without Inter-satellite Links. In *HotNets*. ACM, Virtual Event, USA, 205–211.
 - [24] Huawei. 2023. Ascend 310 AI Processor. <https://e.huawei.com/mx/products/cloud-computing-dc/atlas/ascend-310> Accessed: 2023-05-12.
 - [25] Huawei. 2023. Atlas 200 AI Accelerator Module. <https://e.huawei.com/en/products/computing/ascend/atlas-200>. Accessed: 2023-03-21.
 - [26] Harry Jones. 2018. The recent large reduction in space launch cost. 48th International Conference on Environmental Systems.
 - [27] Mohamed M Kassem, Aravindh Raman, Diego Perino, and Nishanth Sastry. 2022. A browser-side view of starlink connectivity. In *Proceedings of the 22nd ACM Internet Measurement Conference*. 151–158.
 - [28] Simon Kassing, Debopam Bhattacharjee, André Baptista Águas, Jens Eirik Saethre, and Ankit Singla. 2020. Exploring the "Internet from space" with Hypatia. In *Proceedings of the ACM Internet Measurement Conference*. 214–229.
 - [29] Vaclav Knap, Lars Kjeldgaard Vestergaard, and Daniel-Ioan Stroe. 2020. A review of battery technology in CubeSats and small satellite solutions. *Energies* 13, 16 (2020), 4097.
 - [30] Vivek Kothari, Edgar Liberis, and Nicholas D. Lane. 2020. The Final Frontier: Deep Learning in Space. In *HotMobile*. ACM, Austin, TX, USA, 45–49.
 - [31] Andrzej S Kucik and Gabriele Meoni. 2021. Investigating spiking neural networks for energy-efficient on-board ai applications. a case study in land cover and land use classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020–2030.
 - [32] Zeqi Lai, Hewu Li, Yangtao Deng, Qian Wu, Jun Liu, Yuanjie Li, Jihao Li, Lixin Liu, Weisen Liu, and Jianping Wu. 2023. {StarryNet}: Empowering Researchers to Evaluate Futuristic Integrated Space and Terrestrial Networks. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. 1309–1324.
 - [33] Zeqi Lai, Hewu Li, and Jihao Li. 2020. Starperf: Characterizing network performance for emerging mega-constellations. In *2020 IEEE 28th International Conference on Network Protocols (ICNP)*. IEEE, 1–11.
 - [34] Zeqi Lai, Weisen Liu, Qian Wu, Hewu Li, Jingxi Xu, and Jianping Wu. 2022. SpaceRTC: Unleashing the Low-latency Potential of Mega-constellations for Real-Time Communications. In *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications, London, United Kingdom, May 2-5, 2022*. IEEE, 1339–1348. <https://doi.org/10.1109/INFOCOM48880.2022.9796887>
 - [35] Qiang Li, Yu Jiang, Yan Zhang, Yu Gao, Xiaohong Guo, Ruifei Cui, Sichen Wang, Wenming Guo, Tiexin Lv, Lifeng Cai, et al. 2021. Optimal attitude control for solar array orientation. *Open Astronomy* 30, 1 (2021), 73–82.
 - [36] Qing Li, Shangguang Wang, Xiao Ma, Qibo Sun, Houpeng Wang, Suzhi Cao, and Fangchun Yang. 2021. Service coverage for satellite edge computing. *IEEE Internet of Things Journal* 9, 1 (2021), 695–705.
 - [37] Qing Li, Shangguang Wang, Xiao Ma, Ao Zhou, and Fangchun Yang. 2021. Towards Sustainable Satellite Edge Computing. In *2021 IEEE International Conference on Edge Computing (EDGE)*. IEEE, 1–8.

- [38] Brandon Lucia, Brad Denby, Zachary Manchester, Harsh Desai, Emily Ruppel, and Alexei Colin. 2021. Computational nanosatellite constellations: Opportunities and challenges. *GetMobile: Mobile Computing and Communications* 25, 1 (2021), 16–23.
- [39] Brandon Lucia and Benjamin Ransford. 2015. A simpler, safer programming and execution model for intermittent systems. *ACM SIGPLAN Notices* 50, 6 (2015), 575–585.
- [40] Abhas Maskey and Mengu Cho. 2020. CubeSatNet: Ultralight Convolutional Neural Network designed for on-orbit binary image classification on a 1U CubeSat. *Engineering Applications of Artificial Intelligence* 96 (2020), 103952.
- [41] Joseph E Mazur, Joseph F Fennell, James L Roeder, Paul T O'Brien, Timothy B Guild, and Justin J Likar. 2011. The timescale of surface-charging events. *IEEE Transactions on Plasma Science* 40, 2 (2011), 237–245.
- [42] Jianyin Miao, Qi Zhong, Qiwei Zhao, Xin Zhao, Jianyin Miao, Qi Zhong, Qiwei Zhao, and Xin Zhao. 2021. Design of Spacecraft Thermal Control Subsystem. *Spacecraft Thermal Control Technologies* (2021), 45–64.
- [43] François Michel, Martino Trevisan, Danilo Giordano, and Olivier Bonaventure. 2022. A first look at starlink performance. In *Proceedings of the 22nd ACM Internet Measurement Conference*. 130–136.
- [44] Ferdinando Tonicello Mikko Nikulainen. 2021. *Utilization of COTS in ESA Missions*. NASA Electronic Parts and Packaging (NEPP) Program. https://nepp.nasa.gov/workshops/etw2021/talks/17-JUN-21_Thur/1045_Nikulainen_Tonicello-Utilisation-of-COTS-in-ESA-Missions.pdf
- [45] Doug Mohnhey. 2020. *Terabytes From Space: Satellite Imaging is Filling Data Centers*. Data Centers Frontier. <https://www.datacenterfrontier.com/internet-of-things/article/11429032/terabytes-from-space-satellite-imaging-is-filling-data-centers>
- [46] NanoAvionics. 2022. *How Much Do CubeSats and SmallSats Cost?* <https://nanoavionics.com/blog/how-much-do-cubesats-and-smallsats-cost/> Accessed: 2023-01-02.
- [47] Sujay Narayana, R Venkatesha Prasad, Vijay Rao, Luca Mottola, and T Venkata Prabhakar. 2020. Hummingbird: Energy efficient gps receiver for small satellites. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. 1–13.
- [48] NASA. n.d. *CubeSat 101: Basic Concepts and Processes for First-Time CubeSat Developers*. Technical Report. NASA. https://www.nasa.gov/sites/default/files/atoms/files/nasa_csli_cubesat_101_508.pdf
- [49] Kathryn O'Donnell, Meghan Weber, Joy Fasnacht, Jeff Maynard, Margaret Cote, and Shayn Hawthorne. 2023. Extension of cloud computing to small satellites. (2023).
- [50] Daniel Perdices, Gianluca Perna, Martino Trevisan, Danilo Giordano, and Marco Mellia. 2022. When satellite is all you have: watching the internet from 550 ms. In *Proceedings of the 22nd ACM Internet Measurement Conference*. 137–150.
- [51] Tobias Pfandzelter and David Bermbach. 2022. Celestial: Virtual Software System Testbeds for the LEO Edge. In *Middleware*. ACM, Quebec, QC, Canada, 69–81.
- [52] H. K. Ramapriyan. 2015. *The Role and Evolution of NASA's Earth Science Data Systems*. NASA. <https://ntrs.nasa.gov/citations/20150018076>
- [53] Markus Reichstein, Gustau Camps-Valls, Bjorn Stevens, Martin Jung, Joachim Denzler, Nuno Carvalhais, and fnm Prabhat. 2019. Deep learning and process understanding for data-driven Earth system science. *Nature* 566, 7743 (2019), 195–204.
- [54] Ahmad Zammir Ribah and Sri Ramayanti. 2019. Power produced analysis of solar arrays in nadir pointing mode for low-earth equatorial micro-satellite conceptual design. In *IOP Conference Series: Earth and Environmental Science*, Vol. 284. IOP Publishing, 012048.
- [55] João GP Rodrigues and Ana Aguiar. 2019. Extracting 3D maps from crowdsourced GNSS skyview data. In *The 25th Annual International Conference on Mobile Computing and Networking*. 1–15.
- [56] Mike Safyan. 2020. Planet's dove satellite constellation. *Handbook of Small Satellites: Technology, Design, Manufacture, Applications, Economics and Regulation* (2020), 1–17.
- [57] Felix Siegle, Tanya Vladimirova, Jørgen Ilstad, and Omar Emam. 2015. Mitigation of radiation effects in SRAM-based FPGAs for space applications. *ACM Computing Surveys (CSUR)* 47, 2 (2015), 1–34.
- [58] Vaibhav Singh, Akarsh Prabhakara, Diana Zhang, Osman Yagan, and Swarun Kumar. 2021. A community-driven approach to democratize access to satellite ground stations. In *MobiCom*. ACM, New Orleans, Louisiana, 1–14.
- [59] Ankit Singla. 2021. Satnetlab: a call to arms for the next global internet testbed. , 28–30 pages.
- [60] SmartSat. 2021. *Machine Learning Onboard Satellites*. Technical Report. SmartSat CRC. <https://smartsatcrc.lbcdn.io/uploads/Technical-report-machine-learning-onboard-satellites-PUB-1.pdf>
- [61] H Socolovsky, S Muoz, D Raggio, and C Bolzi. 2017. Development and Testing of Solar Panels for Small Satellite Applications at CNEA. In *Proceedings of the 31st Annual AIAA/USU Conference on Small Satellites, SSC17-P1-18*.
- [62] Michael Swartwout. 2013. The first one hundred cubesats: A statistical look. *Journal of small Satellites* 2, 2 (2013), 213–233.
- [63] Sumitaka Tachikawa, Hosei Nagano, Akira Ohnishi, and Yuji Nagasaka. 2022. Advanced Passive Thermal Control Materials and Devices for Spacecraft: A Review. *International Journal of Thermophysics* 43, 6 (2022), 91.
- [64] Arkedged Space Team. 2023. *Arkedged Space*. Arkedged Space. <https://arkedgedspace.com/en>
- [65] Exo-Space Team. 2023. *Exo-Space*. Exo-Space. <https://www.exo-space.com/>
- [66] Get Away Special Team. 2022. *Get Away Special Passive Attitude Control Satellite*. Utah State University. <https://www.usu.edu/physics/gas/projects/gaspacs>
- [67] Ubotica Team. 2023. *Ubotica*. <https://ubotica.com/>
- [68] Union of Concerned Scientists 2022. *UCS Satellite Database*. Union of Concerned Scientists. <https://www.ucsusa.org/resources/satellite-database>
- [69] Burak Uz Kent, Christopher Yeh, and Stefano Ermon. 2020. Efficient object detection in large images using deep reinforcement learning. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*. 1824–1833.
- [70] Deepak Vasisht and Ranveer Chandra. 2020. A distributed and hybrid ground station network for low earth orbit satellites. In *Proceedings of the 19th ACM Workshop on Hot Topics in Networks*. 190–196.
- [71] Deepak Vasisht, Jayanth Shenoy, and Ranveer Chandra. 2021. L2D2: low latency distributed downlink for LEO satellites. In *SIGCOMM*. ACM, Virtual Event, 151–164.
- [72] Giuliana Santos Veronese, Miguel Correia, Alysso Neves Bessani, Lau Cheuk Lung, and Paulo Verissimo. 2011. Efficient byzantine fault-tolerance. *IEEE Trans. Comput.* 62, 1 (2011), 16–30.
- [73] Peng Wang, Hongyan Li, Binbin Chen, and Shun Zhang. 2022. Enhancing Earth Observation Throughput Using Inter-satellite Communication. *IEEE Transactions on Wireless Communications* (2022).
- [74] Shangguang Wang, Qing Li, Mengwei Xu, Xiao Ma, Ao Zhou, and Qibo Sun. 2021. Tiansuan Constellation: An Open Research Platform. In *EDGE*. IEEE, Chicago, IL, USA, 94–101.
- [75] Shangguang Wang, Qiyang Zhang, Ruolin Xing, Fei Qi, and Mengwei Xu. 2023. The first verification test of space-ground collaborative intelligence via cloud-native satellites. *arXiv preprint arXiv:2311.06078* (2023).
- [76] Yu Wang, Min Sheng, Weihua Zhuang, Shan Zhang, Ning Zhang, Runzi Liu, and Jiandong Li. 2018. Multi-resource coordinate scheduling for

- earth observation in space information networks. *IEEE Journal on Selected Areas in Communications* 36, 2 (2018), 268–279.
- [77] James Richard Wertz, Wiley J Larson, Douglas Kirkpatrick, and Donna Klungle. 1999. *Space mission analysis and design*. Vol. 8. Springer.
- [78] Mark White. 2017. *Commercial Off-The-Shelf (COTS) Parts Risk and Reliability User and Application Guide*. Technical Report.
- [79] Xin Yuan, Ruolin Xing, Shangguang Wang, and Mengwei Xu. 2022. Towards Robust Intelligence in Space. In *2022 IEEE Smartworld, Ubiquitous Intelligence & Computing, Scalable Computing & Communications, Digital Twin, Privacy Computing, Metaverse, Autonomous & Trusted Vehicles (SmartWorld/UIC/ScalCom/DigitalTwin/PriComp/Meta)*. IEEE, 706–713.