

Space Microdatacenters

Nathaniel Bleier
nbleier3@illinois.edu

University of Illinois at Urbana-Champaign
Urbana, Illinois, USA

Gary R. Swenson

University of Illinois at Urbana-Champaign
Urbana, Illinois, USA

Muhammad Husnain Mubarak

University of Illinois at Urbana-Champaign
Urbana, Illinois, USA

Rakesh Kumar

University of Illinois at Urbana-Champaign
Urbana, Illinois, USA

ABSTRACT

Earth observation (EO) has been a key task for satellites since the first time a satellite was put into space. The temporal and spatial resolution at which EO satellites take pictures has been increasing to support space-based applications, but this increases the amount of data each satellite generates. We observe that future EO satellites will generate so much data that this data cannot be transmitted to Earth due to the limited capacity of communication that exists between space and Earth. We show that conventional data reduction techniques such as compression [126] and early discard [41] do not solve this problem, nor does a direct enhancement of today's RF-based infrastructure [133, 153] for space-Earth communication. We explore an unorthodox solution instead - moving to space the computation that would have happened on the ground. This alleviates the need for data transfer to Earth. We analyze ten non-longitudinal RGB and hyperspectral image processing Earth observation applications for their computation and power requirements and discover that these requirements cannot be met by the small satellites that dominate today's EO missions. We make a case for space microdatacenters - large computational satellites whose primary task is to support in-space computation of EO data. We show that one 4KW space microdatacenter can support the computation need of a majority of applications, especially when used in conjunction with early discard. We do find, however, that communication between EO satellites and space microdatacenters becomes a bottleneck. We propose three space microdatacenter-communication co-design strategies - k -list-based network topology, microdatacenter splitting, and moving space microdatacenters to geostationary orbit - that alleviate the bottlenecks and enable effective usage of space microdatacenters.

KEYWORDS

Computational satellite, Micro datacenter, Compute in space

ACM Reference Format:

Nathaniel Bleier, Muhammad Husnain Mubarak, Gary R. Swenson, and Rakesh Kumar. 2023. Space Microdatacenters. In *56th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO '23)*, October 28-November 1, 2023, Toronto, ON, Canada. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3613424.3614271>

1 INTRODUCTION

The ability to launch satellites into space and then control them to accomplish a wide variety of tasks such as navigation [43], communication [51], forecasting [115], early warning [107], reconnaissance [102], broadcasting [99], scientific research [33], signals intelligence [104, 154], weapons delivery [62], and Earth observation [123] has been one of the most wondrous achievements of humankind. These satellites have different volumes (0.01 m^3 to 916 m^3) and weights (1.26 kg to 420 000 kg) and are placed into outer space at different altitudes above the Earth (274 km to 35 786 km) in different orbits (low Earth orbit [34], geostationary orbit [139], sun-synchronous orbit (SSO) [17], etc.) using launch vehicles [23, 37]. These satellites have different sources of power generation (none - for passive satellites [122], solar panels [119], radioisotopic thermoelectric generators [118], etc.) to support their functionality, use transponders [42] for communication to Earth-based ground stations [85], and work either alone or together as a group (often called a *constellation* [145]).

Earth observation (EO) has been a key task for satellites since inception. EO satellites image the Earth using camera [123], radar [55], lidar [112], photometer [137], or atmospheric instruments [21] in order to support a variety of scientific [19], military [104, 154], and commercial [42] applications. As imaging satellites, they are often placed in low Earth orbit for high data resolution (though some EO satellites are placed in a geostationary orbit [139] for uninterrupted coverage or in a SSO for consistent lighting during imaging [18]), and transmit their images to Earth-based ground stations for further processing. Following Sputnik-1 [113], the first satellite ever launched, thousands of EO satellites have been placed in space to support different applications [93]. A vast number of future satellite launches are also devoted to Earth observation [84] to support a fast growing Earth observation industry [84].

A key parameter for an EO satellite is the resolution at which it takes its pictures. Increasingly Earth observation space missions are being planned with aggressive goals of spatial and temporal resolution (Section 3) to support emerging EO applications such as forest fire detection [147], realtime video [131], conflict zone monitoring [10], tasking [27], warning systems for early responders [157],


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 the author(s) 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.

MICRO '23, October 28-November 1, 2023, Toronto, ON, Canada

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0329-4/23/10...\$15.00

<https://doi.org/10.1145/3613424.3614271>



(a) A conventional EO Constellation. At (b) Compression and Early Discard reduce (c) An EO Constellation with a S μ DC. high resolutions, the amount of EO data to the amount of data to be downloaded, but S μ DCs process in-space the data generated by EO Constellations be downloaded becomes prohibitive. it is still prohibitive at high resolutions.

Figure 1: S μ DCs reduce or eliminate the need to satellite download data to Earth-based ground stations.

and tracking of events such as Earthquakes [160], hurricanes [38], and tornadoes [30], as well as objects such as aircraft [66] and missiles [14]. Even traditional EO applications such as flood monitoring [155], traffic monitoring [75], mapping [31], etc., seek higher resolutions requirements now. Mapping a narrow path in a dense urban area easily requires sub-meter resolution [159], for example. Fig. 2 shows how spatial resolution of EO satellites has improved over the decades.

In this paper, we observe that the amount of data that future high resolution Earth observation satellites will generate will be so massive that data cannot simply be transmitted to the Earth considering present or projected ground station capacity (Section 3). The limited number of ground stations on the Earth limit the total amount of data that can be transmitted. At current costs, the monetary cost of transmission will also be prohibitive (Section 3).

We first evaluate two techniques (Section 4) that have been previously proposed to reduce the amount of data transmitted to the Earth - compression [126] and early discard [41] - to address the problem (Fig. 1b). We show that compression or early discard may not provide sufficient data reduction for many high resolution space missions either alone or in conjunction. We also consider (Section 4) if today’s RF-based communication infrastructure can be enhanced to support high resolution space missions. We show that practical RF-based satellite antennas may not support the needs for many such missions. The number of channels needed to be supported on the ground may also be unrealistic.

We explore an unorthodox solution instead (Section 5) - whenever possible, move the computation that would have happened on the ground to space itself. If we are able to perform the computation in space itself, only insights, not raw sensor data, may need to be transmitted to the ground alleviating the need for massive data transfer to the ground for high resolution applications.

We analyze ten emerging non-longitudinal RGB and hyperspectral image processing Earth observation applications that process high resolution satellite data. We estimate for these applications their computation and power requirements at different resolutions. We find that small satellites which dominate Earth observation today, cannot support many of these applications, especially at high resolutions, as these satellites cannot generate enough power to support the power requirements of these applications. While early discard helps reduce the power requirements, the reduction is not enough to support many of these applications.

With the above in mind, we make a case for space microdatacenters (S μ DCs) for high resolution Earth observation space missions (Section 6). A S μ DC (Fig. 1c) is a relatively large computational satellite whose primary task is to support in-space computation on data generated by the observation satellites. The power generation capability for the S μ DC is commensurate with the amount of computation supported by the S μ DC. Inter-satellite links (ISLs) are used to offload the data generated by the observation satellites to the S μ DC.

We consider the S μ DC requirements for a 64-satellite constellation of Earth observation satellites for 4kW S μ DCs based on NVIDIA RTX 3090-class processors. We show (Section 6) that one 4kW S μ DC can support the computation needs for a majority of our applications for most resolutions, especially when used in conjunction with early discard.

We do find, however, that communication between the observation satellites and the S μ DCs becomes a bottleneck (Section 7). We propose three S μ DC-communication co-design strategies - *k-list*-based network topology, S μ DC splitting, and moving S μ DCs to geostationary orbit - to alleviate this bottleneck and effectively use these S μ DCs (Section 8). Finally, we analyze the impact of placement and chip architecture on S μ DC design and performance.

This paper makes the following contributions:

- We show that future high resolution Earth observation missions will generate so much data that the generated data cannot be transmitted to the Earth considering present or projected ground station capacity or considering the transmission costs.
- We show that compression, early discard, or antenna scaling have limited effectiveness at addressing the problem.
- We explore moving the Earth-based computation that computes on EO data into space and show that this computation cannot be performed on the typically small EO satellites since these satellites cannot meet the corresponding power requirements.
- We make a quantitative case for S μ DCs that are designed to run the Earth-based computation in space. We show that a 4kW S μ DC can support a majority of the applications if communication bottlenecks can be alleviated.
- We present multiple S μ DC-communication co-design strategies (new connection topologies, S μ DC splitting, moving S μ DCs to geostationary orbit) that alleviate the communication bottlenecks of S μ DCs.

2 BACKGROUND AND RELATED WORK

In this work, we will focus on Low-Earth orbit (LEO) Earth observation satellites — satellites with orbital periods of < 128 min and with low eccentricity (i.e., near-circular orbits), resulting in altitude < 2000 km. We focus on LEO EO satellites because a) EO satellites are often placed in LEO orbit in order to improve the spatial resolution of the generated imagery, and b) The number and size of LEO EO satellite constellations has been increasing [35], in large part due to significant decreases in LEO satellite launch costs [46, 72] as well as due to emergence of new EO applications (Section 5).

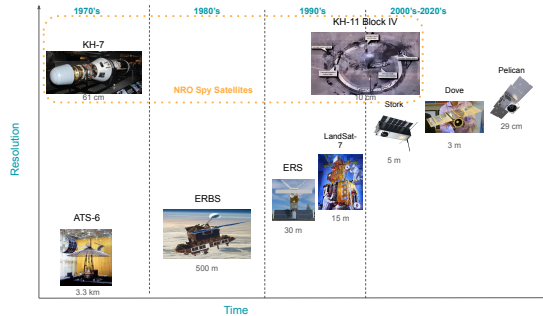


Figure 2: Improvements in imaging technology have led to increasing EO satellite spatial resolutions. The massive and expensive Key Hole line of National Reconnaissance Office (NRO) spy satellites greatly outperform commercial and scientific EO satellites, but also see significant improvements in spatial resolution over time.

Unlike the EO data generation rate, which is increasing rapidly, there is limited opportunity to increase RF downlink capacity [121]. As such, downlink data rates have increased less than data generation rates (Fig. 3). Several approaches have been proposed to deal with this downlink deficit. Lossless and high quality lossy compression can be used to decrease the number of bits needed to represent each pixel downlinked. More aggressively, data can be discarded — either not downlinked or not even generated. This is done commonly in practice (e.g., Dove does not image the ocean); prior work [41] also propose to do it via image processing (e.g., detect and discard images occluded by clouds). Our work does not focus on reducing the amount of EO data to be sent to the applications running on Earth; we move the applications themselves to space.

The closest related work is the deployment of HPE’s SpaceBorne and SpaceBorne-2 computers to the International Space Station (ISS). These computers have been used to compute on data generated in space which had historically been slow to downlink. For example, astronauts have used these computers to monitor their DNA for mutation due to radiation exposure. This decreased the amount of time needed to analyze astronaut DNA from 12 h (mostly in downlink time) to 6 min [142]. Unlike our work, the HPE ISS computers do not process EO data from EO satellites.

Another closely related work is by Orbits Edge [106], a start-up that is trying to build frames to send servers to outer space. Limited information is available about their design.

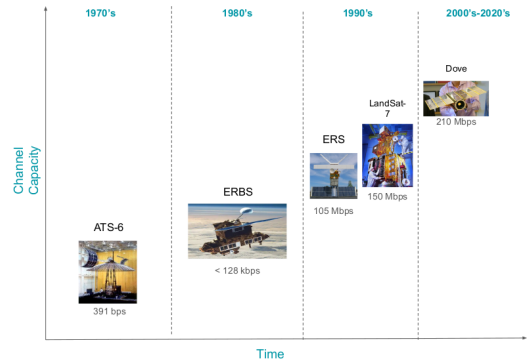


Figure 3: Satellite downlink capacity has grown over time, due to improvements in communication system design and changes in frequency bands used, however, RF downlink capacity is limited by bandwidth constraints.

To the best of our knowledge, no prior work makes a quantitative case for S μ DCs. Ours is also the first work to analyze the computation requirements for a S μ DC, the associated communication bottlenecks, and the S μ DC-communication co-design approaches to address the bottlenecks.

3 DATA REQUIREMENTS OF HIGH RESOLUTION EO SPACE MISSIONS

Earth observation (EO) space missions are increasingly being planned with aggressive goals of spatial and temporal resolution. Table 1 lists some of the current and planned LEO EO constellations - spatial resolution targets are now routinely sub-meter. Satellites have similarly started emerging with continuous imaging goals (Earthnow). These aggressive goals have been the result of impressive advances in addressing challenges in imaging at fine resolutions, including the diffraction limit of telescopes, dispersion due to diffraction by the atmosphere, orbiter motion compensation of ~ 8 km s^{-1} , etc. On larger satellites, such as the NRO’s KH-11, a 2.4 m mirror has a diffraction-limited resolution of 0.05 arcseconds, or, at a 250 km altitude, a spatial resolution of 0.6 cm [52]. Smaller satellites can produce 10 cm resolution imagery by processing multiple coarse-resolution (e.g., 40 cm) images [60].

Considering the aggressive resolution targets, the amount of data these missions will generate will be massive. Fig. 4a shows the data generation rate at different resolutions assuming a global coverage target (i.e., $\frac{\text{surface area of Earth}}{\text{spatial res.}} \cdot \frac{1}{\text{temporal res}}$): at fine spatial resolutions, tens of Tbit s^{-1} , and at fine spatial and temporal resolutions, tens of Pbit s^{-1} of data needs to be generated.

Today’s LEO EO constellations use RF downlinks to transmit data from orbit to Earth ground stations. Using Planet’s Dove constellation’s 96 MHz X-band channels [42] as a baseline, Fig. 4b shows the number of concurrent, continuous Dove-like channels needed to transmit all of the data from space to Earth. At fine resolutions, this is many orders of magnitude more channels than can currently be supported by Earth’s ground stations. Table 2 shows the number and continental locations of ground stations operated by commercial Ground Station as a Service providers. While many of these ground stations can support multiple simultaneous channels,


Table 1: LEO EO Constellations

Company Name	Constellation Name	# Satellites	Form Factor	Imaging	Spatial Resolution	Temporal Resolution	“Goals” or Mission
SatRev	Stork	14	3U	RGB+Near Infrared	5 m	6 Hour	Hosted Payload Missions
SatRev	REC	1024	6U	RGB	50 cm	30 Min	Insurance, land survey, precision farming, smart cities, imagery intelligence, early warning systems support, assistance of missile homing systems
Planet	Dove	159	3U	RGB-Hyperspectral	3 m	24 Hour	Daily imaging of Earth’s land
Planet	SkySat	21	100 kg	RGB-Hyperspectral	50 cm	<24 Hour	Sub-daily high resolution imaging of any point on Earth, stereo video for up-to 90 seconds
Spacestry	Spacestry SAR	56	185 kg	C-Band SAR	1 m	“High Frequency Revisits”	Real-time SAR imagery of every point on Earth, day and night, rain or shine.
Chang Guang	Jilin-1	300	225 kg	Color Video, PAN, MSI	1-1.3 m (video), 75 cm (PAN), 3-4 m (MSI)	2-3.3 days	
Spacestry	ADASPACE	192	185 kg	RGB, hyperspectral	1 m (RGB), 4 m (hyperspectral)	< 24 h	A global, minute-level updated Earth image data network
Space JLTZ	Gemini	378	6U	Multispectral	4 m	10 Min	
Planet	Pelican	32	150 kg to 200 kg	RGB	0.29 m	30 Min	Provide responsive, rapid, very-high resolution imagery
Airbus	EarthNow	300	230 kg	Color Video	1 m	Continuous	Hurricane monitoring, fisheries management, forest fire detection, crop-health monitoring, conflict zone observation
LeoStella	BlackSky	18	50 kg	RGB Imagery	1 m	1h	Hourly revisit time for most major cities
Earth-i	Vivid-i	15	100 kg	RGB Color Video	60 cm, 1 m	< 12h	First constellation to provide full-color video

Table 2: Number and location of Ground Station as a Service (GSaaS) providers’ ground stations.

Service	Ground Stations					Total
	N. America	S. America	Africa	Europe/MENA	Asia/Pacific	
AWS Ground Station [71]	2	1	1	3	4	11
Azure Ground Stations [11]	4	1	3	6	5	19
KSat Ground Network Services [83]	4	2	4	9	6	26
Viasat Real-Time Earth [149]	4	1	2	4	3	14
US Electrodynamics Inc [146]	2	0	0	0	0	2
Swedish Space Corporation [140]	3	2	0	2	3	10
Atlas Space Operations [8]	4	0	1	3	5	13
Leaf Space [86]	1	0	1	8	4	14
RBC Signals [135]	12	2	3	18	16	51

they are ultimately limited by both number of antennas (typically < 100, e.g., KSat’s Antarctica ground station has 20 antennas [59], and its entire network of 26 ground stations has only 270 antennas [83]) and limited S-band and X-band bandwidth. Thus, even with a planned doubling of the number of ground stations over the next 3 years [153], the number of downlink channels is orders of magnitude too low to support high resolution LEO EO missions.




(a) Constellation Data Generation Rates. (b) Constellation Downlink Requirements


Figure 4: At fine resolutions, data generation of LEO EO constellations becomes prohibitive. The number of concurrent Dove-like 220 Mbit s⁻¹ [42] channels needed to support constellation data generation rates grows enormously.

The scarcity of ground stations also leads to high prices for ground stations. At the price-points of three leading services (AWS, Azure, and KSat), which charge \$3 per minute per channel, the cost of downlinks to support a fine resolution LEO EO constellation would be in the millions of dollars *per minute!* Thus, at current costs, the monetary cost of transmission of high resolution EO data will also be prohibitive.

Another view of this phenomenon is presented in Fig. 5. In 5a, we present the ‘downlink deficit’ (DD), or portion of generated data which must be discarded due to downlink capacity limitations, as a function of the number of downlink channels available to a satellite per orbital revolution. As number of channels per revolution increases, downlink deficit decreases. Different curves represent different spatial resolutions (for a given satellite, these curves



(a) Downlink Deficit (DD)



(b) Time spent downlinking (per satellite per revolution)

Figure 5: Downlink Deficit and time spent downlinking for an EO satellite at different spatial resolutions, assuming an 220 Mbit/s] downlink channel [42]. These figures assume a 95% early discard rate (i.e., only 1 in 20 images is downlinked to Earth), as in [41].

are invariant with respect to temporal resolution). 5b depicts the amount of time each satellite spends downlinking each revolution. As this time increases, so too, does the monetary cost of transmitting data to Earth. The results show that the amount of data generated by high resolution EO missions lead to prohibitively high downlink deficit or high cost or both.

4 EFFECTIVENESS OF DATA REDUCTION AND TRANSMISSION TECHNIQUES

We first calculate the effective compression ratio (ECR) required to support various imaging resolutions of an EO mission for a given downlink capacity. ECR, in this case, is the data reduction ratio achieved by combination of early discard and image compression. Let’s optimistically assume that sufficient downlink capacity exists for 3 m-1 d resolution RGB imagery of all of Earth – Planet’s current Dove constellation (Table 1) provides 3 m-1 d resolution RGB imagery of Earth’s *land*. Fig. 6 shows the ECR needed to support various resolutions using this downlink capacity. The results show that fine resolutions require ECRs in the the thousands to hundreds of thousands. Such ECRs are likely unachievable in most settings.

Table 3 shows achievable rates of early discard and their associated effective compression ratios (ECR) for several types of early discard. These rates have been calculated using gross Earth characteristics (50% images correspond to night, 70% images correspond

to ocean, 10% of Earth is uninhabited, built-up areas - those areas which contain vertical construction - account for 2% of Earth’s area, $\frac{2}{3}$ of Earth is covered by cloud) for optimistic assumption orbital dynamics (e.g., a non-dawn/dusk circular orbit for night data, equatorial orbit instead of polar orbit for habitation data, etc.). As the results show, the achievable early ECRs are far lower than the required ECRs reported in Figure 6. Note that some forms of early discard may be combined (e.g., imaging only built-up areas during the day) to achieve higher ECRs. However, this is limited by conditional dependencies (e.g., cloud distribution is dependent on land vs ocean, uninhabited implies non-built up, etc.).

We similarly estimate realistic ECR values for EO data when data compression algorithms are used. We ran different data compression algorithms on the Crowd AI Mapping Challenge [98] dataset of satellite imagery of built-up areas of Earth. One thousand images were randomly selected from the dataset after removing images which did not display a full scene¹. Analysis of compression of SAR imagery used the xView3 validation dataset [108]. Table 4 shows the results. The results show that achievable compression ratios using lossless image compression are limited to $< 4\times$ for RGB imagery. This is in line with previous studies on lossless compression of satellite imagery [152]. High-quality ‘quasi-lossless’ lossy compression, also results in compression ratios of only 10–20 \times [12]. These numbers are off from the required ECRs by several orders of magnitude.

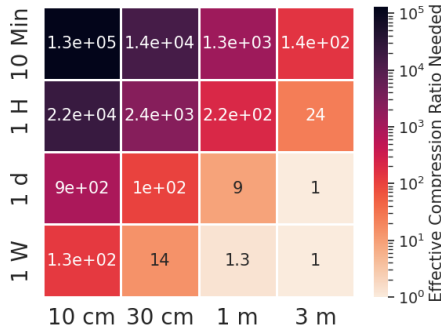


Figure 6: Effective compression ratio needed to downlink data from different target resolutions given downlink sufficient for 3 m-1 day spatial and temporal resolution.

Assuming independence of early discard and compression (e.g., by discarding images of non-built-up areas and images at night), the ECR of combined compression and early discard is $\leq 4 \times 100 = 400$. This best case ECR is still up to 3.5 orders of magnitude short of the ECR needed for some of the fine resolution targets. These results show that compression and early discard have limited effectiveness at addressing the problem of too much data that will be generated by future high resolution EO missions.

Another way to decrease the downlink deficit is to increase the amount of information which can be moved from space to Earth. While number of ground stations is anticipated to double from 2021

¹the dataset contains images which are padded to full resolution by a black background. These images were removed as they can achieve unrealistically high compression ratios.

Table 3: Achievable early discard Rates and their associated ECRs

Metric	None	Night	Ocean	Uninhabited[128]	Non-Built-Up[45]	Cloudy[80]
Early Discard Rate	0	0.5	0.7	0.9	0.98	0.67
ECR	1	2	3.4	10	50	3

to 2026 [153], doubling the number of ground stations leads to no more than proportional increase in downlink capacity.

Satellite designers can also increase downlink capacity by modifying the design of their satellites RF communications. RF downlinks are modeled as additive white Gaussian noise communication channels [133], and are thus subject to the Shannon-Hartley theorem [130], which relates channel capacity, C , (in bit s^{-1}) to the channel bandwidth, B , (in Hz), and the signal-to-noise ratio (SNR) at the ground-station:

$$C = B \log_2 (1 + \text{SNR}) .$$

Note that $\frac{\partial C}{\partial B} = \log_2 (1 + \text{SNR})$ and $\frac{\partial C}{\partial \text{SNR}} \propto \frac{1}{\text{SNR} \log_2(2) + \log_2(2)}$. Thus, when $\text{SNR} \gg 0$, C scales linearly with B , but with the reciprocal of the SNR. This is called a ‘bandwidth limited’ regime, and satellite downlinks are squarely in this regime (e.g., Dove’s ground stations experience $\text{SNR} \sim 19[42]$).

As the electromagnetic spectrum is a limited natural resource, satellites cannot simply scale their bandwidth, which is allocated to them by national and international governing bodies, such as the Federal Communications Commission (USA), and the International Telecommunication Union. Thus, satellite designers can only increase RF channel capacity by increasing signal strength. This can be achieved in one of two ways: 1) increase the power output by the antenna, and 2) increase the directionality, or gain, of the antenna. Increasing antenna output power requires increasing the input power, while increasing the antenna gain requires increasing the antenna’s aperture size, and thus increasing the physical size of the antenna for common satellite antenna types (i.e., patch antennas, helical antennas, and parabolic antennas).

Fig. 7 shows the infeasibility of meeting fine spatial resolution targets through scaling of RF downlinks in a bandwidth limited communications regime. Both a 2 kW antenna input power and a 30 m antenna fall far short of meeting the downlink capacity requirements of a 1 m resolution target, let alone < 1 m resolutions.




Figure 7: Increasing channel capacity by modifying satellite antennas requires exponential increases in power consumption and antenna size.

The number of channels needed to be supported on the ground may also be unrealistic. Fig. 4b shows that the number of channels needed is many orders of magnitude greater than the number of

Table 4: Compression ratios achieved using lossless compression techniques for RGB satellite imagery from the Crowd AI Mapping Challenge [98], and SAR imagery from the xView3 [108] dataset.

Imagery	JPEG2000	LZW	Zip	RLE	PNG	CCSDS [151]
RGB	3.93	2.14	2.38	1.	2.49	1.88
SAR	966	469	2436	64	924	9.89

channels which can be supported by current or near future ground stations.

5 MOVING GROUND-BASED COMPUTATION INTO SPACE

We explore an unorthodox solution instead - whenever possible, move the data processing computation that would have happened on the ground to space. If we are able to perform the computation in space, only insights, not raw sensor data, may need to be transmitted to the ground alleviating the need for massive data transfer to the ground for high resolution applications.

To determine the feasibility of moving ground-based computation into space, we analyzed a collection of representative Earth-based applications that process high resolution satellite image data. We consider only ‘memoryless’ applications which process a single frame at a time using only the single frame’s data. Some EO imagery applications are *longitudinal* – they assess changes to a location over days, months, or even years, and thus require significant data storage.

Table 5 lists our applications. Air Pollution Prediction (APP) is used to monitor urban areas and other areas where air pollution is a concern. Missions supported by NASA and the California Air Resources Board use satellite imagery to predict air pollution. Satellite imagery-based crop monitoring (CM) is used to identify how much of a crop is grown in a region, which is important information for commodities markets, and to monitor crop growth and performance on a macro scale. Satellite imagery is used to perform Flood Detection (FD) and flood severity estimation. Satellites can provide timely identification of fast moving flash floods [144]. In Forage Quality Estimation (FQE), satellite imagery is used to estimate the quality (quantity) of animal forage for use by ranchers, shepherds, etc. Urban Emergency Detection (UED) is a multifaceted application which attempts to identify emergent life threatening phenomenon in built-up and urban areas, enabling timely emergency response and public awareness. Processing in space enables low latency detection, an important metric for this application. Aircraft Detection (AD) enables detecting and classifying aircraft from satellite imagery. While 3 m resolution is sufficient for commercial airliners and large, manned combat aircraft, < 1 m resolutions are likely required to detect and classify small drones and loitering munitions which have played impactful roles in recent battlefields [150]. LEO satellites provide benefit over aircraft for this role in that they do not violate restricted or contested airspace. Panoptic Segmentation (PS) [81] is an emerging machine vision application which attempts to perform both semantic segmentation of an image, as well as identification of individual objects within the segments. It can be used to support numerous other applications. In Oil Spill Monitoring (OSM), waterways are monitored for signs of spills of oil and refined petroleum products. As oil is often

shipped via intercontinental shipping lanes, satellites offer timely and inexpensive (relative to aircraft) monitoring of sea-lanes. Traffic Monitoring (TM) detects moving vehicles due to the offset of different wavelengths that moving objects produce, causing a specific reflectance relationship in RGB channels. This enables effective vehicle detection with very low compute overhead. Land Surface Clustering (LSC) is an unsupervised machine learning technique which attempts to segment imagery to detect changes in a landscape over time. Satellites, which periodically revisit locations with little to no additional cost per revisit (unlike aircraft), are thus a good fit for this application. The majority of the applications are machine learning based, with most using deep learning. The variety of kernels and architectures lead to a wide spread in computational complexity, with over $10^5\times$ difference in floating point operations per pixel between aircraft detection and traffic monitoring.

To estimate the performance and power requirements of these application, we run them on Jetson AGX Xavier (32GB AGX) that features an NVIDIA Volta GPU with eight streaming multiprocessors. The Jetson AGX Xavier has been proposed for use in cubesat-class EO satellites [6] due to its good radiation tolerance [120]. We installed JetPack 5.0.1 with L4T 34.1.1, which supports CUDA version 11.4.315. To maintain compatibility with the hardware and software environment, we installed the appropriate cuDNN version 8.6.0.166 and TensorFlow version 2.11. This configuration allowed us to successfully run all but one of our applications on the Jetson AGX Xavier. We ran the inference 100 times, for different batch sizes, and employed the TegraStats tool to measure the average GPU utilization. To approximate the GPU power consumption, we used the utilization data along with the reported maximum power of Jetson AGX Xavier, an accepted technique for estimating embedded GPU power consumption [20]. Table 6 shows the performance and power of our applications on Jetson AGX Xavier, including pixels processed $s^{-1} W^{-1}$.

We use the performance and power numbers of applications on Jetson AGX Xavier to determine how much compute and power generation a satellite must support to run a given application in space. Fig. 8 shows these requirements for a single satellite at 0.10 m to 3 m resolutions and 0-99% early discard rates. As in [41], each ground frame at 3 m is represented by a single 4K RGB image; scaling resolution holds the ground frame size constant by increasing the number of pixels per frame. Thus, as resolution becomes finer, the number of pixels needed to be processed each second increases. The horizontal lines in the graph represent the number of pixels per second needed to be processed to run the applications at a given resolution and early discard rate. The curves (lines with non-zero slope) represent the number of pixels per second that can be supported for a given power budget (x -axis) with power efficiency equal to a Jetson AGX Xavier. Where a curve intersects a horizontal bar gives the amount of power needed to support the application in a satellite. We assume computational complexity scales linearly with number of pixels, as is decidedly the case in TM, and is often the case in deep learning based image processing [57].

The results in Fig. 8 show that only one application can be supported at 3 m resolution with a power budget typical of a small satellite (Table 7) without a high early discard rate. No application

Table 5: Applications which consume satellite imagery

Application	Description	Imagery	Kernel	FLOPs/Pixel	Application Users / Providers
Air Pollution Prediction (APP)	Predict air pollution levels using CNN	RGB	Inception-ResNet	3317	NASA [141], CARB [26]
Crop Monitoring (CM)	Identify type and quality of crops	Hyperspectral	Inception v3	67113	Ministry of Agriculture of China [156], ESA [48],
Flood Detection (FD)	Identify floods and assess flood severity	RGB	DenseNet	178969	GDACS [54], NASA [100]
Aircraft Detection (AD)	Identify stationary and moving aircraft from satellite images using CNN	RGB	Custom 4-layer CNN	7387714	Orbital Insights [105], militaries
Forage Quality Estimation (FQE)	Estimate forage quality for use in agriculture and animal husbandry	RGB	EfficientNet based	8491	USDA [50], UN [89]
Urban Emergency Detection (UED)	Fire, Traffic Accident, Building Collapse Detection	RGB	MobileNet v3	4484	NASA [25], UN [114]
Panoptic Segmentation (PS)	Simultaneous detection of countable objects and backgrounds	RGB	Mask RCNN	6874279	Crop Monitoring [110], urban classification [40], environmental monitoring [39]
Oil Spill Monitoring (OSM)	Deep water environmental monitoring	Hyperspectral	VGg19	390625	KSAT, NOAA, ESA
Traffic Monitoring (TM)	Detect moving vehicles via blue reflectance	RGB	Custom DSP algo using ratios between channels	51	DST [94], ESA [108]
Land Surface Clustering (LSC)	Unsupervised segmentation of land/land cover change detection	Hyperspectral	K-Means ($K = 4$)	15984	NASA [7], ESA [16]

Table 6: Application results for RTX 3090 and Jetson AGX Xavier. Results are for optimal batch sizes. PS could not be mapped to Jetson AGX Xavier.

App Name	RTX 3090				Jetson AGX Xavier			
	Pow (W)	Util (%)	Infer time (s)	kPixel s ⁻¹ W ⁻¹	Pow (W)	Util (%)	Infer time (s)	kPixel s ⁻¹ W ⁻¹
Air Pollution	119	25	0.59	1168	4.04	27	3.07	825
Crop Monitoring	222	42	1.57	395	12.5	84	16.0	86
Flood Detection	325	88	5.53	307	13.8	92	78.4	64
Aircraft Detection	124	6	0.26	74	2.62	18	17.5	39
Forage Quality Estimation	129	27	0.56	843	5.13	34	3.29	449
Urban Emergency Detection	266	72	2.04	569	12.6	17	17.4	177
Oil Spill Monitoring	347	98	3.84	231	14.6	97	80.2	33
Traffic Monitoring	19	< 1	2.72	2597	1.00	< 1	0.05	9630
Land Surface Clustering	108	2	0.35	2175	2.21	1	0.6	5792
Panoptic Segmentation	160	80	7.81	20	X	X	X	X

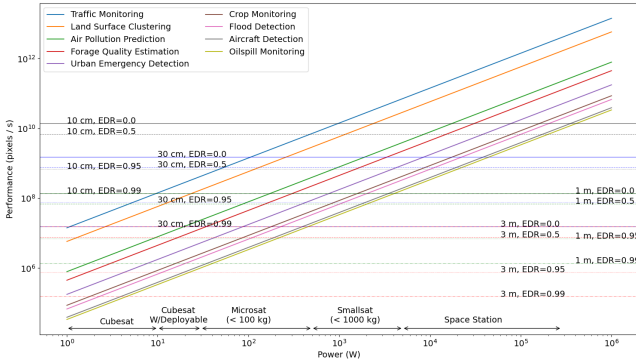


Figure 8: Power needed to meet compute requirements for EO applications on a given satellite assuming zero early discard and Jetson AGX Xavier as the on-satellite computer architecture. Horizontal bars represent the performance needed to meet various spatial resolution targets: 3 m, 1 m, 30 cm, and 10 cm, at 0%, 50%, 95%, and 99% early discard rate (early discard rate).

can be supported by a small satellite at fine resolutions where they require hundreds to hundreds of thousands of watts.

Even with aggressive early discard (99%), many applications still require hundreds of Ws at fine resolutions. Aircraft detection requires > 400 W of compute per satellite at 30 cm. At 99% early discard rate, several applications cannot be supported at 1 m on a cubesat or cubesat with deployable solar panel. At 10 cm, several applications cannot even be supported on a typical 100 kg micro-satellite. Further, for many applications, 99% early discard rate is unrealistic, as applications such as OSM, CM, AD, FD, FQE, etc, may be interested in large portions of Earth’s surface.

To summarize, the above results show that while moving Earth-based computation (that computes on EO data) into space may be promising, this computation cannot be performed on the typically small EO satellites themselves (Table 7) since the corresponding power requirements cannot be met by most of these satellites. While some large satellites may be able to natively support some of

Table 7: Satellite capabilities by weight class. Applications supported at 10 cm spatial resolution for 0% and 95% early discard rates (in parentheses).

Satellite Class	Examples	Power Generation	Apps Supported at all res. (at 0.95 ED)
Cubesat	Swarm Technologies	1 W to 10 W [4, 127, 148]	TM (APP, UED)
Cubesat (Deployable Panels)	Dove, REC, Stork, Gemini	10 W to 30 W [44, 138]	(FQE, LSC)
< 100 kg	SkySat, BlackSky	55 W to 210 W [77, 124]	APP, UED, FQE, LSC, CM, FD
< 1000 kg	Vivid-i, EarthNow, ADASPACE, Jilin-1, Spacety	200 W to 6600 W [36, 68, 70]	CM, FD (OM)
Space Station	ISS	240 kW [53]	OM, AD, FS

the applications, many of the emerging LEO EO constellations are based on microsat and cubesat class satellites, including ones with < 1 m spatial resolutions (Table 1), as well as the *largest* current and planned EO constellations. As such, an alternate approach to computing in space must be developed.

6 A CASE FOR SPACE MICRODATACENTERS

With the above in mind, we make a case for space microdatacenters (S_μDCs) for high resolution Earth observation space missions. A S_μDC is a relatively large computational satellite whose primary task is to support in-space computation on data generated by the observation satellites. The power generation capability for the S_μDC is commensurate with the amount of computation supported by the S_μDC. Inter-satellite links (ISLs) are used to offload the data generated by the observation satellites to the S_μDC.

To support in-space computation of Earth-based applications, one could also simply make each EO satellite much bigger (i.e., increase its power generation and computation capability). However, a LEO EO constellation supported by S_μDCs offers several advantages over a homogeneous constellation of EO satellites large enough to support the applications natively. First, by concentrating compute onto S_μDCs, EO satellite design – satellite bus design, heat dissipation, power generation and power management, etc, – is simplified allowing continuing low mission costs [41] which is critical for growth of the EO industry. Second, changes in computational requirements (e.g., an improved neural network model, increased accuracy requirements, change in application, etc) would be hard to support in a homogeneous constellation, while they can be supported by adding additional S_μDCs in our model. Third,

S μ DCs act as data integrators, minimizing the impact of variation in data generation (not all EO satellites within their constellation would generate the same amount of data - e.g., land vs ocean, day vs night, cloudy vs not). Thus, average case design for S μ DCs would be more effective than average case design for a homogeneous constellation. Finally, S μ DCs may also be used to provide space-based cloud computing, supporting excess compute requirements of multiple constellations, including from multiple organizations.

Fig. 9 shows the number of 4 kW S μ DCs needed to support a constellation of 64 EO satellites for various resolutions and early discard rates. We assume a 4 kW S μ DC for this study since Orbits Edge SATFRAME 445 [106] uses a 19-inch server rack, which can easily support up to 4 kW of compute² We assume no images are downlinked to Earth – all are processed in space. These results have been generated using measured power and delay numbers on a RTX 3090 (Table 6). RTX 3090 is a state-of-the-art GPU that provides high energy efficiency for image processing workloads [74], and support high productivity programming paradigms [61]. We used CUDA version 11.7 along with cuDNN version 8.9.0 and the supported TensorFlow version 2.12. For the Panoptic Segmentation application, which uses Mask R-CNN, we employed the Mask R-CNN [63] implementation available in the application’s repository. To run this specific application on the RTX 3090, we used TensorFlow-GPU version 1.14, ensuring compatibility with the provided Mask R-CNN implementation. For all of the DNNs, we performed inference 100 times, for different batch sizes, and utilized the Python NVML (pynvml) library to measure the average GPU utilization and average GPU power. In addition, we used the timing library to measure the inference time. We ported the TM workload from a CPU implementation to a CUDA implementation. We implemented LSC using k-means in CUDA. Batch-sizes which maximize energy efficiency (maximize pixels $W^{-1} s^{-1}$) are used. We assume a ground track frame period of 1.5 s, meaning each satellite in the constellation generates one image every 1.5 s. We assume that early discard is applied uniformly over all generated images.

The results show that one 4 kW S μ DC can support the computation needs for a majority of our applications for most resolutions, especially when used in conjunction with early discard. For example, only a single 4 kW S μ DC is needed to support all but one application at 1 m with 95% early discard rate. At finer resolutions and low early discard rates, multiple 4KW S μ DCs may be needed. In some cases, S μ DCs may need to be significantly larger (e.g., 256 kW “Space Station” class S μ DCs). While the number of S μ DCs needed to support some applications at fine resolutions is high, the costs of downlinking data to Earth are *prohibitive*. Even with 99% early discard, downlink at current commercial rates would cost the constellation operator over \$1000 per minute at 10 cm resolution. But at that early discard rate, eight out of ten applications can be supported with only a small number of S μ DCs computing in space. Launching these S μ DCs, especially at projected future launch costs, will invariably be cheaper than paying significant recurring costs for data downlink.

²The power budget of 4KW is for computing, not the entire satellite. Other components that draw significant power include ISLs, ground communication, flywheel-based attitude control, radiation-hardened flight controller, battery heating, propulsion, and active thermal management. We estimate that these components may add up to 1KW more for an overall power dissipation to be <5KW.




Figure 9: The number of RTX 3090-based 4 kW S μ DCs needed to support applications for various resolutions and early discard rates.




Figure 10: Two S μ DCs in ‘ring’ network topologies

7 COMMUNICATION AS A SPACE MICRODATACENTER BOTTLENECK

The above analysis did not consider communication of EO data from the satellites to an S μ DC. Fig. 10 shows a small constellation supported by two S μ DCs in ‘ring’ topology [76]. Data from distant EO satellites is relayed to the S μ DC by more proximal EO satellites. Thus, the number of connected EO satellites is potentially limited by the capacity of the ISL between the S μ DC and the closest EO satellites.

A ring topology in which the S μ DCs are part of the same orbit as the EO satellites has clear benefits. By flying the S μ DCs in formation with the EO satellites and using a fixed ring topology, ISLs are also fixed. This is important when ISLs are optical, since optical ISLs can take seconds or even minutes to orient [15, 64]. Small satellites which contain optical ISLs often orient the ISL by rotating the entire satellite [64]. This means that the satellite cannot perform simultaneous communication and imaging. However, by using a ring topology with fixed distance and angle between satellites, satellite designers can design the ISL and camera such that they are usable simultaneously.

If one S μ DC can support the computation of n satellites, but the capacity limitation of the ISL between the S μ DC and the closest EO satellites dictates that the S μ DC can only receive data from $m < n$ satellites, then the number of clusters (and thus S μ DCs) needed is $\frac{64}{m} > \frac{64}{n}$. In this case, the constellation is ISL-bottlenecked. If $m \geq n$, the constellation is ISL-unconstrained.

For lightweight applications, the minimal number of S μ DCs that are needed in a ring topology may not be set by the total amount of computation required, but rather by the number required to mitigate the ISL bottleneck. Table 8 shows how many EO satellites an S μ DC can support before becoming ISL-bottlenecked at various data rates and for several ISL capacities based on RF [92, 134] and optical [28, 78] LEO to LEO ISLs. This table assumes that a base (at

3 m) 4K RGB image is generated every 1.5 s on each EO, and transmitted via ISL to the μ DC. As resolution improves, so does the number of pixels in the image (i.e., the imaged area remains constant). In a ring topology, the limiting links are the ones between the μ DC and its adjacent EO satellites. Thus, for example, at 3 m resolution and 1 Gbit s⁻¹ ISL capacity, each ISL can support transmitting over four images every 1.5 s. Since the μ DC has two ISL receivers, it can support up to nine EO satellites.

The results show that < 100 Gbit s⁻¹ ISLs are often insufficient to support even a single EO satellite for high data rates. Even 100 Gbit s⁻¹ ISLs fail at 10 cm resolutions. On the other hand, a single μ DC can support a large number of EO satellites at low data generation rates (i.e., coarse resolution and high early discard rates) – more than what would realistically be placed into a single orbital plane. This table data, combined with Fig. 9 indicates when ISL-bottlenecks or computational requirements dictate the number of μ DCs needed. Fig. 11 shows that the number of clusters, and thus μ DCs, is set by the ISL-bottlenecks for many applications – especially for high-power μ DCs. As ISL capacity increases, the bottleneck goes away, and the number of clusters required matches the number of μ DCs needed to support the computation, as in Fig. 9.

In general, it is preferable for a constellation to be ISL unconstrained, as an ISL bottlenecked constellation means more μ DCs are launched than are strictly required based on computational power requirements. This increases constellation equipment, launch, and management costs.

Our results also show that ISLs considerations can have important influence on μ DC design for lightweight applications – high power μ DCs are more likely to be ISL-bottlenecked than a low power μ DCs. They also suggest that ISL network topology may play an important role in enabling high μ DC utilization. Thus ISL considerations will impact EO/ μ DC satellite constellation design.

Table 8: The number of EO satellites supportable by a single μ DC using ring topology for various data generation rates for ISLs with capacity 1 Gbit s⁻¹, 10 Gbit s⁻¹ and 100 Gbit s⁻¹.

Early Discard Rate	Resolution	1 Gbit s ⁻¹	10 Gbit s ⁻¹	100 Gbit s ⁻¹	Resolution	1 Gbit s ⁻¹	10 Gbit s ⁻¹	100 Gbit s ⁻¹
0	3 m	9	98	992	1 m	1	10	110
0.5		18	198	1986		2	22	220
0.95		198	1986	19868		22	220	2206
0.99		992	9934	99340		110	1102	11036
0		0	0	8		0	0	0
0.5	30 cm	0	0	18	10 cm	0	0	2
0.95		0	18	198		0	2	22
0.99		8	98	992		0	10	110
0		0	0	0		0	0	0

8 SPACE MICRODATACENTER AND COMMUNICATION CO-DESIGN

One way to mitigate an ISL-bottleneck in context of μ DCs is to modify the network topology to increase the amount of data onboarded onto the μ DCs. Figure 12a shows how, by adding more receivers to a μ DC, the cluster topology can be changed from a ring, or ‘2-list’, to a ‘4-list’, or, more arbitrarily, a ‘ k -list’ for even k . While this may not help RF communication-based constellations due to limited available bandwidth, tremendous amounts of bandwidth is available in the optical frequencies, allowing linear growth in incoming data rate with the number of optical receivers [136]. Thus, for optical ISLs, k -lists for $k > 2$ can be used to increase

the μ DC’s incoming data rate at the cost of additional optical receivers on the μ DC and additional transmit power.

As k increases, the link distance between relay satellites grows. Optical ISL transmit power grows quadratically with distance [88], meaning a 4-list’s ISLs consume 4× the power of a 2-list (while also transmitting 2× the data). Also, this distance can eventually grow to such an extent that the ISL must aim through significant amounts of atmosphere. This results in atmospheric turbulence induced fading of the optical signal [162], degrading the channel capacity. If the distance is large enough, then the Earth’s landmass will directly block the signal. The specific value of k for which distance becomes a concern is dependent on the constellation’s formation: for evenly distributed – ‘orbit spaced’ – formations, maximum k is smaller than for tightly packed formations in which satellites are relatively close to one another.

Alternatively, μ DCs can be *split* Figure 12b – increasing the number of clusters in a ring-topology without increasing the compute power of the μ DCs in aggregate. By using smaller split μ DCs, costs associated with higher cluster counts (e.g., launch cost, booster fuel requirements, etc.) are mitigated. μ DC splitting is effectively a form of disaggregation, and thus can lead to increased *total* launch costs and constellation management costs. However, μ DC splitting is effective for all constellation formations, including orbit-spaced constellations which may see limited benefit from k -list topologies.

μ DC splitting and k -list topologies can be used in conjunction. Their benefits are orthogonal, and the increase in aggregate data rate into μ DCs scales multi-linearly with number of clusters (from splitting), and number of links into each μ DC (from k -lists). That is, the number of EO satellites supported by a k -list topology cluster is $\frac{k}{2}$ times those shown in Table 8, while μ DC splitting multiplies the number of clusters. Figure 13 shows that k -lists combined with μ DC splitting leads to significantly increased ISL communication capacity (the rate at which data can be transmitted from an EO satellite to an μ DC) in a frame-spaced constellation.

9 IMPACT OF CHIP ARCHITECTURE, RADIATION, PLACEMENT, ETC.

Section 5 looks at commodity GPUs as the computer architecture in a μ DC since these are easily programmable, abundant, and perform well on image processing workloads. However, it is unclear that GPUs are the best architecture for μ DCs. GPUs are designed for applications with latency considerations (e.g., customer satisfaction [1], real-time video processing [5], etc.), while many satellite applications do not have stringent latency requirements (e.g., TM, APP, AD, CM, LSC, FQE). For such applications, energy efficiency is the most important metric. As such, alternate architectures (i.e., accelerators) or computing regimes (e.g., near-threshold computing) may be attractive. In the recent MLPerf [117] v3.0 competition [96, 97], the Qualcomm Cloud AI 100 was the most energy efficient architecture for offline batch image processing inference tasks – > 2.5× better than the NVIDIA A100 and nearly 2× better than the NVIDIA H100. We also compare Qualcomm Cloud AI 100 to RTX 3090 – the accelerator is 18.25× more energy efficient on MLPerf workloads. With this energy efficiency, Figure 14 estimates the the number of 4KW μ DCs needed to support EO

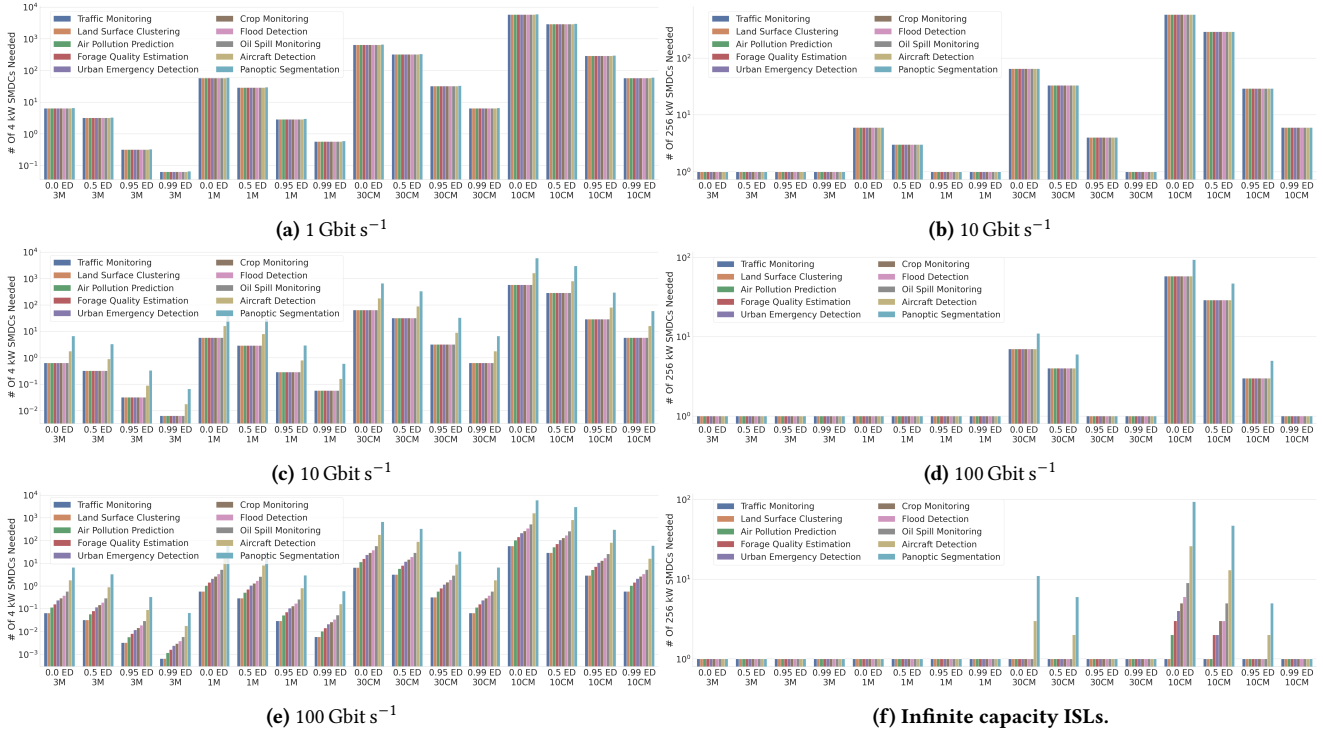



Figure 11: For 4 kW SμDCs, ISL bottlenecks occur for many applications for ISL channels < 100 Gbit s⁻¹ (Left). For 256 kW SμDCs, ISL bottlenecks persist, even with high capacity ISLs in a ring topology (Right).



(a) A '4-list' topology. (b) SμDC Splitting.

Figure 12: An ISL-bottleneck can be resolved by (a) increasing the number of incoming edges to each SμDC, or (b) increasing the number of SμDCs (while proportionally decreasing the computational power of each SμDC).

satellite applications at various resolutions and early discard rates using a Qualcomm Cloud AI 100 rather than a RTX 3090 (Fig. 9). The results show that AI 100 enables more applications to be supported at finer resolutions and with lower early discard rates with only a small number of SμDCs. The significant benefits relative to even the most state-of-the-art GPUs (A100, H100) suggests that SμDCs would benefit from incorporating an architecture focused on energy efficiency. In general, unlike compression-based or RF communication-based approaches, SμDCs scale with technology and architectural developments, and thus the number and complexity of applications they can support will grow in the future.




Figure 13: The total ISL communication capacity, defined as the rate at which data is transmitted from EO satellites to SμDCs, and total ISL transmission power consumption for various k -list topologies and SμDC splitting factors. Values are normalized against a 2-list (ring) with out splitting SμDCs.

We assume that SμDCs will use radiation-hardened SBCs such as the BAE Systems' RAD750 [13] for flight control. For data processing, however, we note that SμDCs are based in LEO. LEO satellites experience much less radiation (e.g., 1 krad/year of gamma radiation [101]) than MEO and GEO, except in the South Atlantic Anomaly (SAA) [65], so use of COTS hardware, instead of radiation-hardened hardware, is common [58, 129], especially for CubeSats [24] - an ITAR regulated hardened component that tolerates 300 krad [9] is significant overdesign for LEO.

We believe that SμDCs can also use non-hardened COTS GPU and accelerator hardware, especially since many COTS hardware, including NVIDIA Xavier studied in the paper, already show good

radiation tolerance [120]³. NASA’s Ingenuity helicopter on Mars already uses non-hardened COTS Snapdragon 801 GPU SoC to perform in-flight image processing and guidance tasks [143], NASA is already using non-hardened EdgeTPUs for CubeSats [58]. The effects of the SAA can be mitigated either through pausing computation while the S μ DC passes through the SAA, or by applying additional software-based hardening while in the SAA (ISS’s SBC-2 [47] and MISSE-7 mission [158] show feasibility of software-based radiation hardening on COTS hardware). There is already a large body of research on software based hardening for GPUs [22, 56, 111, 116].

We note also that our workloads are largely ML workloads. Due to inherent algorithmic resilience to radiation-induced bit-flip errors for such workloads [125], the overhead of software-based soft-error hardening will be low: 18% [3], <5% for convolution layers [132].

Finally, radiation exposure can be reduced by bus design – this is OrbitsEdge’s approach with their SatFrame [106].




Figure 14: Non-GPU architectures which focus on energy efficiency are highly effective on the massively parallel, latency unconstrained workloads in EO satellite applications.

The orbital placement of the S μ DC can impact its cost and effectiveness. Since LEO EO constellations typically place many satellites in the same orbital plane at the same altitude, perhaps the most obvious orbit to give the S μ DC is the same orbit as the satellites it supports. This also enables a fixed ‘ring’ or k -list network topology, as discussed above. However, the satellites need significant boosting [90] at lower altitude to prevent atmospheric drag from causing them to crash into Earth. Thus, a second possible placement for an S μ DC is in the same orbital plane, but at a higher altitude, to limit the amount of boosting needed for the larger S μ DC. Unfortunately, angular velocity of a satellite’s orbit decreases with altitude, meaning a static network topology is no longer possible. This may not be a significant issue for RF-based ISLs, which can be quickly aimed via beamforming [109], but high-capacity optical ISLs require pointing which can take seconds or even minutes to perform [15, 64]. Another possible location for S μ DCs is in geostationary (GEO) orbit. A GEO orbit is a 36 000 km equatorial orbit in which the satellite is located above a fixed location on Earth (i.e., the angular velocity of the satellite is equal to the angular velocity of Earth’s rotation about its axis). This means a collection of three S μ DCs in GEO could provide continuous support for LEO EO satellites. Drawbacks of GEO positioning include greater launch cost due to higher altitude, and need for increased radiation hardening,

³GPU4S team [82] specifically states that “LEO [missions] will likely adopt COTS GPUs first, due to the more limited exposure to radiation.”

since GEO orbits are in Earth’s outer Van Allen belt, which contains higher energy radiation than the inner Van Allen belt [29].

Placement also impacts the design of S μ DCs. Lower altitude (circular) orbits spend more time in an eclipsed region, where access to sunlight is blocked by Earth. LEO satellites spend $\sim \frac{1}{3}$ of their time eclipsed, while GEO satellites spend most of the year without eclipse – GEO satellites experience small amounts ($< 45 \text{ min d}^{-1}$) of eclipsed time in an eclipse for several weeks before and after the two equinoxes. Thus, S μ DCs in LEO must support greater power generation than S μ DCs in GEO in order to support the same computational workload. Likewise, S μ DCs in LEO will need boosting capabilities in order to extend lifetime. GEO requires less boosting than LEO as discussed above meaning lifetime of S μ DCs in GEO may be considerably longer than in LEO. This means that S μ DCs in GEO will need additional radiation hardening and hardware redundancy. Additionally, the ‘sun-setting’ or retirement of satellites differs between the two. GEO satellites are retired by increasing their altitude by $\sim 300 \text{ km}$ into a ‘graveyard’ orbit, while LEO satellites are retired by lowering their altitude into a ‘disposal’ orbit in which atmospheric friction destroys the satellites. In GEO, since satellites can stay in orbit indefinitely, back-up hardware is also used to extend the lifetime of a satellite [91], especially as GEO satellite mission lifetimes are often longer than commodity hardware lifetimes ($\sim 15 \text{ years}$ for GEO satellites [67] vs as low as four years for commodity hardware [95]).

Also, S μ DCs in GEO have the benefit of lowered ISL-bottleneck. A GEO-based ISL-bottleneck mitigation strategy uses the increased optical ISL count of k -lists, but changes the network topology to three dynamic star-clusters in which the S μ DCs are in GEO, as seen in Fig.15. Numerous works demonstrate high capacity, low power LEO-GEO ISLs [2, 64]. By using three S μ DCs spaced 120° apart, each EO satellite is guaranteed to have line of sight with at least one S μ DC at all times. Benefits to this approach are that very large S μ DCs can be used with mitigated ISL capacity concerns.




Figure 15: Three S μ DCs are placed in GEO, ensuring all EO satellites have LOS with at least one S μ DC at all times.

A S μ DC will produce large amounts of heat waste. As such, dissipation of heat is an important S μ DC design consideration. Some heat may be dissipated using thermocouples, as in RTG powered systems, and as has been argued for use in terrestrial datacenters [161]. Excess heat is also an issue in space for infra-red cameras and other heat sensitive payloads, and thus numerous thermal control subsystem design approaches are in use, including radiating surfaces, heat transport and heat conductive lines [73], thermoelectric coolers [69], and design via evolutionary algorithm [49].

The discussion in this paper assumes a monolithic implementation of an S μ DC. A *disaggregated* S μ DC may also be interesting. In

Table 9: Comparison of downlink deficit mitigation strategies.

	S μ DCs	Homogeneous Compute	Compression	RF Comms
Scales to Future Resolution Targets	Yes	Yes	No	No
High Power	Yes	Yes	No	Yes
Requires ISLs	Yes	No	No	No
Adaptive to Mission Changes	Yes	No	No	No

a disaggregated spacecraft design [32], a large satellite is divided into sub-components. These sub-components are placed on their own buses and launched in close proximity to one another, leading to a single logical satellite composed out of several physical satellites. The physical satellites communicate with one another over high capacity, short range ISLs, and can even perform wireless power transfer with high efficiency [87]. With a disaggregated design, existing S μ DCs can be augmented with additional compute hardware and power generation as needed to support changes to the EO constellation mission. Disaggregated design can also help to increase satellite resilience, and also can lower cost when subsystems malfunction, since only a replacement for the subsystem must be launched, rather than a full satellite. Disaggregated design, though, has higher costs, than aggregated design, since design complexity (and thus testing requirements) and total design mass are increased. In the case of an S μ DC, this may be acceptable, due to its high power requirements. While compute hardware lasts for years (and may be outdated before it malfunctions), solar panels last for decades [79]. For S μ DCs, disaggregating power generation and compute may make sense, especially for large S μ DCs in GEO, where compute hardware is susceptible to damage from high energy outer van Allen belt radiation.

Table 9 compares the different techniques for dealing with downlink deficit. Unlike S μ DCs, techniques based on early discard and compression, increasing RF channel capacities, or homogeneous constellations with more powerful EO satellites do not scale to future resolution targets, or are not adaptive to changes in EO satellite mission or algorithmic developments.

Figure 16 shows the impact of radiation hardening overhead. In 16a, software-based radiation hardening with overhead of 20% [3] is used. In 16b and 16c, 2 \times and 3 \times redundancy is used, leading to high overheads. At coarse resolution, radiation hardening has negligible impact – for most applications, the number of S μ DCs needed is unchanged. The same is true even at 1 m with high early discard rate. However, at fine resolutions, and especially with low early discard rates, the impact is significant, especially for redundancy based hardening. For example, at 30 cm and 50% early discard, 3 \times S μ DCs are needed to support a constellation. For software based hardening, this number is unchanged, but for double and triple redundancy, it increases to 5 and 8, respectively. Thus, we anticipate software-based radiation hardening will be the most common solution to radiation hardening in S μ DCs.

10 SUMMARY AND CONCLUSION

In this work, we observe that future EO satellites will generate so much data that this data cannot be transmitted to Earth due to limited capacity of communication that exists between space and Earth. We showed that conventional data reduction techniques such as compression [126] and early discard [41] do not solve this

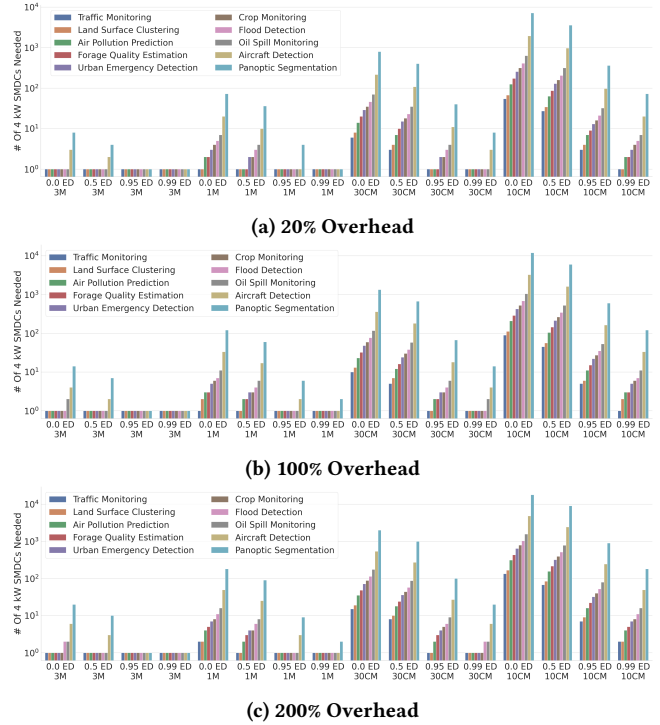


Figure 16: The impact of radiation tolerance overhead on S μ DC equipped constellations.

problem, nor does a direct enhancement of today’s RF-based infrastructure [133, 153] for space-Earth communication. We explored an unorthodox solution instead - moving to space the computation that would have happened on the ground. This alleviates the need for data transfer to Earth. We analyzed ten non-longitudinal RGB and hyperspectral image processing Earth observation applications for their computation and power requirements and discovered that these requirements could not be met by the small satellites that dominate today’s EO missions. We made a case for space microdatacenters (S μ DCs) - computational satellites tasked to support in-space computation of EO data. We showed that one 4KW space microdatacenter can support the computation need of a majority of applications. To address the communication bottleneck between EO satellites and S μ DCs, we proposed three space microdatacenter-communication co-design strategies – k -list-based network topology, microdatacenter splitting, and moving space microdatacenters to geostationary orbit. These techniques enable effective usage of S μ DCs.

ACKNOWLEDGMENTS

We thank Michael Lembeck, Rick Eason, Stanley Wu, and Ramakrishna Kanungo, as well as the anonymous reviewers, for their feedback and suggestions.

REFERENCES

- [1] Mohammad Aazam and Eui-Nam Huh. 2015. Fog Computing Micro Datacenter Based Dynamic Resource Estimation and Pricing Model for IoT. In *2015 IEEE 29th International Conference on Advanced Information Networking and Applications*. 687–694. <https://doi.org/10.1109/AINA.2015.254>
- [2] Maan M. Abdulwahid and Sefer Kurnaz. 2023. The channel WDM system incorporates of Optical Wireless Communication (OWC) hybrid MDM-PDM for higher capacity (LEO-GEO) inter satellite link. *Optik* 273 (2023), 170449. <https://doi.org/10.1016/j.jiloe.2022.170449>
- [3] Geancarlo Abich, Jonas Gava, Rafael Garibotti, Ricardo Reis, and Luciano Ost. 2021. Applying Lightweight Soft Error Mitigation Techniques to Embedded Mixed Precision Deep Neural Networks. *IEEE Transactions on Circuits and Systems I: Regular Papers* 68, 11 (2021), 4772–4782. <https://doi.org/10.1109/TCSI.2021.3097981>
- [4] Abdelkader Aoudeche, Xiangyang Zhao, and Kamel D. Kerrouche. 2018. Design of a High Performance Electrical Power System for an Earth Observation Nano-Satellite. In *Proceedings of the 2018 International Conference on Electronics and Electrical Engineering Technology (Tianjin, China) (EET '18)*. Association for Computing Machinery, New York, NY, USA, 140–146. <https://doi.org/10.1145/3277453.3286093>
- [5] Atakan Aral, Ivona Brandic, Rafael Brundo Uriarte, Rocco De Nicola, and Vincenzo Socca. 2019. Addressing application latency requirements through edge scheduling. *Journal of Grid Computing* 17 (2019), 677–698.
- [6] Antonio Arcidiacono, Daniele Finocchiaro, Riccardo De Gaudenzi, Oscar del Rio-Herrero, Stefano Cioni, Marco Andrenacci, and Riccardo Andreotti. 2021. Is Satellite Ahead of Terrestrial in Deploying NOMA for Massive Machine-Type Communications? *Sensors* 21, 13 (2021), 4290.
- [7] Arset - change detection for land cover mapping 2018. Arset - change detection for land cover mapping. <https://appliedsciences.nasa.gov/join-mission/training/english/arset-change-detection-land-cover-mapping>
- [8] ATLAS Space 2023. ATLAS Space. <https://atlasspace.com/>
- [9] Atmel. 2008. Rad Hard 16 MegaBit 3.3V SRAM Multi-Chip Module. https://ww1.microchip.com/downloads/en/DeviceDoc/at68166h_ds.pdf
- [10] Thiri Shwesin Aug. 2021. Satellite analysis of the environmental impacts of armed-conflict in Rakhine, Myanmar. *Science of the Total Environment* 781 (2021), 146758.
- [11] Azure GS 2023. Azure Orbital Ground Station - Satellite Comms: Microsoft Azure. <https://azure.microsoft.com/en-au/products/orbital/>
- [12] Pascal Bacchus, Renaud Fraisse, Aline Roumy, and Christine Guillemot. 2022. Quasi Lossless Satellite Image Compression. In *IGARSS 2022 - 2022 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, Kuala Lumpur, 1532–1535. <https://doi.org/10.1109/IGARSS46834.2022.9883135>
- [13] BAE Systems 2020. *RAD750 3U CompactPCI Single Board Computer*. BAE Systems. <https://www.baesystems.com/en-media/uploadFile/20210404070637/1434555679066.pdf>
- [14] Djedjiga Belfadel and Isaac Macwan. 2022. Ballistic missile tracking in the presence of decoys using space base IR sensors. In *Signal Processing, Sensor/Information Fusion, and Target Recognition XXXI*, Ivan Kadar, Erik P. Blasch, and Lynne L. Grewe (Eds.), Vol. 12122. International Society for Optics and Photonics, SPIE, 1212209. <https://doi.org/10.1117/12.2628140>
- [15] Edoardo Benzi, Daniel C Troendle, Ian Shurmer, Mark James, Michael Lutzer, and Sven Kuhlmann. 2016. Optical inter-satellite communication: the Alphasat and Sentinel-1A in-orbit experience. In *14th International Conference on Space Operations*. 2389.
- [16] Pablo Blanco and Malcolm Davidson. 2018. https://seom.esa.int/page_project017.php
- [17] Ronald J Boain. 2004. AB-Cs of sun-synchronous orbit mission design. (2004).
- [18] L. Boldt-Christmas. 2023. Polar and sun-synchronous orbit. https://www.esa.int/ESA_Multimedia/Images/2020/03/Polar_and_Sun-synchronous_orbit
- [19] Dana Bolles. 2023. Trace. <https://science.nasa.gov/missions/trace>
- [20] Robert A. Bridges, Neena Imam, and Tiffany M. Mintz. 2016. Understanding GPU Power. A Survey of Profiling, Modeling, and Simulation Methods. *Comput. Surveys* 49, 3 (9 2016). <https://doi.org/10.1145/2962131>
- [21] RA Brown. 1983. On a satellite scatterometer as an anemometer. *Journal of Geophysical Research: Oceans* 88, C3 (1983), 1663–1673.
- [22] Fredrik C Bruhn, Nandinbaatar Tsog, Fabian Kunkel, Oskar Flordal, and Ian Troxel. 2020. Enabling radiation tolerant heterogeneous GPU-based onboard data processing in space. *CEAS Space Journal* 12, 4 (2020), 551–564.
- [23] Alice Brunello, Andrea Valmorbidia, Enrico C Lorenzini, Stefania Cantoni, Mario De Stefano Fumo, Alberto Fedele, Roberto Gardi, and Raffaele Votta. 2021. Deorbiting small satellites from the ISS using a tether system. *CEAS Space Journal* 13 (2021), 217–230.
- [24] Jan Budroweit and Hagen Patscheider. 2021. Risk assessment for the use of COTS devices in space systems under consideration of radiation effects. *Electronics* 10, 9 (2021), 1008.
- [25] Emre Çolak and Filiz Sunar. 2020. The importance of ground-truth and crowd-sourcing data for the statistical and spatial analyses of the NASA FIRMS active fires in the Mediterranean Turkish forests. *Remote Sensing Applications: Society and Environment* 19 (2020), 100327. <https://doi.org/10.1016/j.rsase.2020.100327>
- [26] California Air Resources Board 2023. California Air Resources Board. <https://ww2.arb.ca.gov/resources/documents/air-quality-research-using-satellite-remote-sensing>
- [27] Davide Castelletti, Gordon Farquharson, Craig Stringham, Michael Duersch, and Duncan Eddy. 2021. Capella space first operational SAR satellite. In *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*. IEEE, 1483–1486.
- [28] Aizaz U. Chaudhry and Halim Yanikomeroglu. 2021. Free Space Optics for Next-Generation Satellite Networks. *IEEE Consumer Electronics Magazine* 10, 6 (2021), 21–31. <https://doi.org/10.1109/MCE.2020.3029772>
- [29] Yue Chen, Geoffrey D Reeves, and Reiner HW Friedel. 2007. The energization of relativistic electrons in the outer Van Allen radiation belt. *Nature Physics* 3, 9 (2007), 614–617.
- [30] Alexander Chernokulsky, Andrey Shikhov, Alexey Bykov, and Igor Azhigov. 2020. Satellite-based study and numerical forecasting of two tornado outbreaks in the Ural Region in June 2017. *Atmosphere* 11, 11 (2020), 1146.
- [31] Emilio Chuvieco, Florent Mouillot, Guido R Van der Werf, Jesús San Miguel, Mihai Tanase, Nikos Koutsias, Mariano García, Marta Yebra, Marc Padilla, Ioannis Gitas, et al. 2019. Historical background and current developments for mapping burned area from satellite Earth observation. *Remote Sensing of Environment* 225 (2019), 45–64.
- [32] Air Force Space Command. 2013. Resiliency and disaggregated space architectures. *white paper* (2013).
- [33] MA Cordiner, H Linnartz, NLJ Cox, J Cami, F Najjarro, CR Proffitt, R Lallement, P Ehrenfreund, BH Foing, TR Gull, et al. 2019. Confirming interstellar C60+ using the hubble space telescope. *The Astrophysical Journal Letters* 875, 2 (2019), L28.
- [34] Nicholas H Crisp, Peter CE Roberts, Sabrina Livadiotti, Vitor Toshiyuki Abrao Oiko, Steve Edmondson, SJ Haigh, Claire Huyton, LA Sinpetru, KL Smith, SD Worrall, et al. 2020. The benefits of very low earth orbit for earth observation missions. *Progress in Aerospace Sciences* 117 (2020), 100619.
- [35] Giacomo Curzi, Dario Modenini, and Paolo Tortora. 2020. Large constellations of small satellites: A survey of near future challenges and missions. *Aerospace* 7, 9 (2020), 133.
- [36] Thibery Cussac, Marie-Anne Clair, Pascale Ultré-Guerard, François Buisson, Gérard Lassalle-Balier, Michel Ledu, Christian Elisabelar, Xavier Passot, and Nelly Rey. 2006. The DEMETER microsatellite and ground segment. *Planetary and Space Science* 54, 5 (2006), 413–427.
- [37] Lance A Davis. 2018. Falcon heavy. *Engineering* 4, 3 (2018), 300–300.
- [38] Muhammad Dawood, Amina Asif, and Fayyaz ul Amir Afsar Minhas. 2020. Deep-PHURIE: deep learning based hurricane intensity estimation from infrared satellite imagery. *Neural Computing and Applications* 32 (2020), 9009–9017.
- [39] Osmar Luiz Ferreira de Carvalho, Osmar Abilio de Carvalho Júnior, Anesmar Oolino de Albuquerque, Nickolas Castro Santana, Dibio Leandro Borges, Argelia Saiaika Luiz, Roberto Arnaldo Trancoso Gomes, and Renato Fontes Guimar aes. 2022. Multispectral panoptic segmentation: Exploring the beach setting with worldview-3 imagery. *International Journal of Applied Earth Observation and Geoinformation* 112 (2022), 102910. <https://doi.org/10.1016/j.jag.2022.102910>
- [40] Osmar Luiz Ferreira de Carvalho, Osmar Abilio de Carvalho Júnior, Cristiano Rosa e Silva, Anesmar Oolino de Albuquerque, Nickolas Castro Santana, Dibio Leandro Borges, Roberto Arnaldo Trancoso Gomes, and Renato Fontes Guimarães. 2022. Panoptic Segmentation Meets Remote Sensing. *Remote Sensing* 14, 4 (2022). <https://doi.org/10.3390/rs14040965>
- [41] Bradley Denby and Brandon Lucia. 2020. Orbital edge computing: Nanosatellite constellations as a new class of computer system. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*. 939–954.
- [42] Kiruthika Devaraj, Ryan Kingsbury, Matt Ligon, Joseph Breu, Vivek Vittaldev, Bryan Klofas, Patrick Yeon, and Kyle Colton. 2017. Dove high speed downlink system. (2017).
- [43] Timothy H Dixon. 1991. An introduction to the Global Positioning System and some geological applications. *Reviews of Geophysics* 29, 2 (1991), 249–276.
- [44] Yi Duann, Loren C Chang, Chi-Kuang Chao, Yi-Chung Chiu, Rong Tsai-Lin, Tzu-Ya Tai, Wei-Hao Luo, Chi-Ting Liao, Hsin-Tzu Liu, Chieh-Ju Chung, et al. 2020. IDEASat: A 3U CubeSat mission for ionospheric science. *Advances in Space Research* 66, 1 (2020), 116–134.
- [45] Brian Dunbar. 2005. Satellites and the city: NASA scientists encourage satellite observations to improve our understanding of urban effects on climate and weather. <https://www.eurekalert.org/news-releases/587468>
- [46] Brian Dunbar. 2017. Advanced Space Transportation Program Fact sheet. <https://www.nasa.gov/centers/marshall/news/background/facts/astp.html>
- [47] Emily Dunkel, Jason Swope, Zaid Towfic, Steve Chien, Damon Russell, Joseph Sauvageau, Douglas Sheldon, Juan Romero-Cañas, Jose Luis Espinosa-Aranda, Léonie Buckley, Elena Hervas-Martin, Mark Fernandez, and Carrie Knox. 2022.

- Benchmarking Deep Learning Inference of Remote Sensing Imagery on the Qualcomm Snapdragon And Intel Movidius Myriad X Processors Onboard the International Space Station. In *IGARSS 2022 - 2022 IEEE International Geoscience and Remote Sensing Symposium*. 5301–5304. <https://doi.org/10.1109/IGARSS46834.2022.9884906>
- [48] NASA Earth Science Data Systems. 2023. Normalized difference vegetation index (NDVI). <https://www.earthdata.nasa.gov/topics/biosphere/vegetation/vegetation-index/normalized-difference-vegetation-index-ndvi>
- [49] Emanuel Escobar, Marcos Diaz, and Juan Cristóbal Zagal. 2016. Evolutionary design of a satellite thermal control system: Real experiments for a CubeSat mission. *Applied Thermal Engineering* 105 (2016), 490–500.
- [50] Frank Forcella, Roberto L. Benech Arnold, Rudolfo Sanchez, and Claudio M. Ghersa. 2000. Modeling seedling emergence. *Field Crops Research* 67, 2 (2000), 123–139. [https://doi.org/10.1016/S0378-4290\(00\)00088-5](https://doi.org/10.1016/S0378-4290(00)00088-5)
- [51] Carl E Fossa, Richard A Raines, Gregg H Gunsch, and Michael A Temple. 1998. An overview of the IRIDIUM (R) low Earth orbit (LEO) satellite system. In *Proceedings of the IEEE 1998 National Aerospace and Electronics Conference. NAECON 1998. Celebrating 50 Years (Cat. No. 98CH36185)*. IEEE, 152–159.
- [52] D. L. Fried. 1966. Optical Resolution Through a Randomly Inhomogeneous Medium for Very Long and Very Short Exposures. *J. Opt. Soc. Am.* 56, 10 (Oct 1966), 1372–1379. <https://doi.org/10.1364/JOSA.56.001372>
- [53] Mark Garcia. 2017. International Space Station Solar Arrays. https://www.nasa.gov/mission_pages/station/structure/elements/solar_arrays.html#.VvIBj2Ou90I
- [54] Global Disaster Alert and Coordination System 2010. Global Flood Detection System - Version 2. <https://www.gdacs.org/flooddetection>
- [55] Richard M Goldstein, Howard A Zebker, and Charles L Werner. 1988. Satellite radar interferometry: Two-dimensional phase unwrapping. *Radio science* 23, 4 (1988), 713–720.
- [56] Marcio M Goncalves, Jose Rodrigo Azambuja, Josie ER Condia, Matteo Souza Rorda, and Luca Sterpone. 2020. Evaluating software-based hardening techniques for general-purpose registers on a GPGPU. In *2020 IEEE Latin-American Test Symposium (LATS)*. IEEE, 1–6.
- [57] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep learning*. MIT press.
- [58] Justin Goodwill, Gary Crum, James MacKinnon, Cody Brewer, Michael Monaghan, Travis Wise, and Christopher Wilson. 2021. NASA spacecube edge TPU smallsat card for autonomous operations and onboard science-data analysis. In *Proceedings of the Small Satellite Conference*. AIAA.
- [59] Kongsberg Group. 2021. KSAT to build six antennas at the troll station. <https://www.ksat.no/news/news-archive/2021/ksat-to-build-six-antennas-at-the-troll-station/>
- [60] Toppheer Haddad. 2021. Approved for 10cm. <https://albedo.com/post/approved-for-10cm>
- [61] Tianyi David Han and Tarek S. Abdelrahman. 2009. HiCUDA: A High-Level Directive-Based Language for GPU Programming. In *Proceedings of 2nd Workshop on General Purpose Processing on Graphics Processing Units (Washington, D.C., USA) (GPGPU-2)*. Association for Computing Machinery, New York, NY, USA, 52–61. <https://doi.org/10.1145/1513895.1513902>
- [62] David C Hardesty. 2005. Space-based weapons: long-term strategic implications and alternatives. *Naval War College Review* 58, 2 (2005), 45–69.
- [63] Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross Girshick. 2017. Mask R-CNN. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- [64] F. Heine, H. Kämpfner, R. Czichy, R. Meyer, and M. Lutzer. 2010. Optical inter-satellite communication operational. In *2010 - MILCOM 2010 MILITARY COMMUNICATIONS CONFERENCE*. 1583–1587. <https://doi.org/10.1109/MILCOM.2010.5680175>
- [65] James R Heirtzler. 2002. The future of the South Atlantic anomaly and implications for radiation damage in space. *Journal of Atmospheric and Solar-Terrestrial Physics* 64, 16 (2002), 1701–1708.
- [66] Henning Heiselberg. 2019. Aircraft and ship velocity determination in Sentinel-2 multispectral images. *Sensors* 19, 13 (2019), 2873.
- [67] Caleb Henry. 2023. Companies are flying old satellites longer, study finds. <https://spacenews.com/companies-are-flying-old-satellites-longer-study-finds/>
- [68] Sung-Soo Jang and Jaeho Choi. 2008. Energy balance analysis of small satellite in Low Earth Orbit (LEO). In *2008 IEEE 2nd International Power and Energy Conference*. IEEE, 967–971.
- [69] Zhang Jianzhong, Wang Tiemin, Liu Xueming, and Tong Guanghui. 1998. Application of the thermoelectric cooler in the seventeenth Chinese retrievable satellite. In *1998 5th International Conference on Solid-State and Integrated Circuit Technology. Proceedings (Cat. No.98EX105)*. 865–868. <https://doi.org/10.1109/ICSICT.1998.786476>
- [70] Chance Johnson, Lauchie Scott, and Stefan Thorsteinson. 2021. Comparing Photometric Behavior of LEO Constellations to SpaceX Starlink using a space-based optical sensor. (2021).
- [71] Daniel Jolliffe, Liane Davidson, and Jocelyn Robert. 2003. Ground station. <https://aws.amazon.com/ground-station/>
- [72] Harry Jones. 2018. The recent large reduction in space launch cost. 48th International Conference on Environmental Systems.
- [73] Albert J. Juhasz and George P. Peterson. 1994. *Innovative Radiator Technologies*. National Aeronautics and Space Administration, Office of Management, Scientific and Technical Information Program.
- [74] Sunguk Jung, Hyeonbeom Heo, Sangheon Park, Sung-Uk Jung, and Kyungjae Lee. 2022. Benchmarking Deep Learning Models for Instance Segmentation. *Applied Sciences* 12, 17 (2022). <https://doi.org/10.3390/app12178856>
- [75] Lynn H Kaack, George H Chen, and M Granger Morgan. 2019. Truck traffic monitoring with satellite images. In *Proceedings of the 2nd ACM SIGCAS Conference on Computing and Sustainable Societies*. 155–164.
- [76] Nikos Karafolas and Stefano Baroni. 2000. Optical Satellite Networks. *J. Lightwave Technol.* 18, 12 (Dec 2000), 1792. <https://opg.optica.org/jlt/abstract.cfm?URI=jlt-18-12-1792>
- [77] Matsuaki Kato, Shinichiro Takayama, Yosuke Nakamura, Keisuke Yoshihara, and Hidekazu Hashimoto. 2005. Road map of small satellite in JAXA. In *56th International Astronautical Congress of the International Astronautical Federation, the International Academy of Astronautics, and the International Institute of Space Law*. B5–6.
- [78] S. Kaur. 2019. Analysis of inter-satellite free-space optical link performance considering different system parameters. *Opto-Electronics Review* 27, 1 (2019), 10–13. <https://doi.org/10.1016/j.opelre.2018.11.002>
- [79] Jaewon Kim, Matheus Rabelo, Siva Parvathi Padi, Hasnain Yousuf, Eun-Chel Cho, and Junsin Yi. 2021. A Review of the Degradation of Photovoltaic Modules for Life Expectancy. *Energies* 14, 14 (2021). <https://doi.org/10.3390/en14144278>
- [80] Michael D. King, Steven Platnick, W. Paul Menzel, Steven A. Ackerman, and Paul A. Hubanks. 2013. Spatial and Temporal Distribution of Clouds Observed by MODIS Onboard the Terra and Aqua Satellites. *IEEE Transactions on Geoscience and Remote Sensing* 51, 7 (2013), 3826–3852. <https://doi.org/10.1109/TGRS.2012.2227333>
- [81] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. 2019. Panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9404–9413.
- [82] Leonidas Kosmidis, Iván Rodriguez, Alvaro Jover-Alvarez, Sergi Alcaide, Jérôme Lachaize, Olivier Notebaert, Antoine Certain, and David Steenari. 2021. GPU4S: Major Project Outcomes, Lessons Learnt and Way Forward. In *2021 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. 1314–1319. <https://doi.org/10.23919/DAT51398.2021.9474123>
- [83] KSAT 2023. The KSAT Global Ground Station Network. <https://www.ksat.no/ground-network-services/the-ksat-global-ground-station-network/>
- [84] Erik Kulu. 2021. Satellite constellations-2021 industry survey and trends. (2021).
- [85] JA Lawton. 1987. Numerical method for rapidly determining satellite-satellite and satellite-ground station in-view periods. *Journal of Guidance, Control, and Dynamics* 10, 1 (1987), 32–36.
- [86] Leaf Space 2023. Leaf Space - ground segment as Service. <https://leaf.space/>
- [87] Ying Li and Vikram Jandhyala. 2012. Design of Retrodirective Antenna Arrays for Short-Range Wireless Power Transmission. *IEEE Transactions on Antennas and Propagation* 60, 1 (2012), 206–211.
- [88] Jintao Liang, Aizaz U Chaudhry, Eylem Erdogan, and Halim Yanikomeroglu. 2022. Link Budget Analysis for Free-Space Optical Satellite Networks. In *2022 IEEE 23rd International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*. IEEE, 471–476.
- [89] Han Liu, Randy A. Dahlgren, Royce E. Larsen, Scott M. Devine, Leslie M. Roche, Anthony T. O’Geen, Andy J.Y. Wong, Sarah Covello, and Yufang Jin. 1970. Estimating rangeland forage production using remote sensing data from a small unmanned aerial system (SUAS) and PlanetScope satellite. <https://agris.fao.org/agris-search/search.do?recordID=CH2023188188>
- [90] Jinyu Liu, Xingqun Zhan, Gangqiang Li, Qingtao Wang, and Shizhuang Wang. 2020. Dynamics of orbital boost maneuver of low Earth orbit satellites by electrodynamic tethers. *Aerospace Systems* 3 (2020), 189–196.
- [91] W.Q. Lohmeyer and K. Cahoy. 2013. Space weather radiation effects on geostationary satellite solid-state power amplifiers. *Space Weather* 11, 8 (2013), 476–488. <https://doi.org/10.1002/swe.20071> arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1002/swe.20071
- [92] Jinling Ma and Ruimin Lu. 2011. The Study on RF Inter-Satellite Links. In *2011 7th International Conference on Wireless Communications, Networking and Mobile Computing*. 1–4. <https://doi.org/10.1109/wicom.2011.6040269>
- [93] Salvo Maruccio, Silvia Ullò, Marco Carminati, and Olfa Kanoun. 2019. Smaller satellites, larger constellations: Trends and design issues for earth observation systems. *IEEE Aerospace and Electronic Systems Magazine* 34, 10 (2019), 50–59.
- [94] Mark McCord, Carolyn Merry, and Prem Goel. 1998. Incorporating Satellite Imagery in Traffic Monitoring Programs. In *North American Travel Monitoring Exhibition and Conference*. 1–18.
- [95] Justin Meza, Qiang Wu, Sanjeev Kumar, and Onur Mutlu. 2015. Revisiting Memory Errors in Large-Scale Production Data Centers: Analysis and Modeling of New Trends from the Field. In *2015 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*. 415–426. <https://doi.org/10.1109/DSN.2015.57>

- [96] MLCommons Datacenter 2023. MLCommons V3.0 Datacenter Results. <https://mlcommons.org/en/inference-datacenter-30/>
- [97] MLCommons Edge 2023. MLCommons V3.0 Edge Results. <https://mlcommons.org/en/inference-edge-30/>
- [98] Sharada Prasanna Mohanty, Jakub Czakon, Kamil A Kaczmarek, Andrzej Pyskir, Piotr Tarasiewicz, Saket Kunwar, Janick Rohrbach, Dave Luo, Manjunath Prasad, Sascha Fleer, et al. 2020. Deep Learning for Understanding Satellite Imagery: An Experimental Survey. *Frontiers in Artificial Intelligence* 3 (2020).
- [99] Shaun Moores et al. 1996. *Satellite television and everyday life: articulating technology*. John Libbey Media.
- [100] NASA - Flood Detection and Mapping 2010. Near real-time flood detection and mapping. <https://appliedsciences.nasa.gov/what-we-do/projects/near-real-time-flood-detection-and-mapping>
- [101] NASA Space Radiation Effects 1999. Space Radiation Effects on Electronic Components in Low-Earth Orbit. <https://llis.nasa.gov/lesson/824>
- [102] Directorate of Special Projects. 1967. *Summary Analysis of Project 206 (GAMBIT)*. AF Unit Post Office. 1–52 pages.
- [103] Directorate of Technical, Operational Support, Erling Kristiansen, Claude Loisy, and Willem v.d. Bosch. 2003. . Vol. 115. European Space Agency. 52–57 pages.
- [104] National Reconnaissance Office. 2008. *Review and Redaction Guide* (2008 ed.). NRO. 1–190 pages.
- [105] Orbital Insight - Aircraft Detection 2023. Aircraft detection. <https://up42.com/marketplace/blocks/processing/orbital-pleiades-aircraft>
- [106] OrbitsEdge. 2022. SATFRAME 445 Le - an overview. <https://orbitsedge.com/satframe>
- [107] Alexandre Paleologue. 2005. Early warning satellites in Russia: What past, what state today, what future?. In *Modeling, Simulation, and Verification of Space-based Systems II*, Vol. 5799. SPIE, 146–157.
- [108] Fernando Paolo, Tsu-ting Tim Lin, Ritwik Gupta, Bryce Goodman, Nirav Patel, Daniel Kuster, David Kroodma, and Jared Dunnmon. 2022. xView3-SAR: Detecting Dark Fishing Activity Using Synthetic Aperture Radar Imagery. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- [109] Jinseok Park, Seong-Mo Moon, Junhan Lim, Soonyoung Hong, Dongpil Chang, and Byung-Sun Lee. 2022. A Ka-Band CMOS Beamforming Front-End IC for RF Inter Satellite Link System. In *2022 27th Asia Pacific Conference on Communications (APCC)*. IEEE, 184–186.
- [110] Marian Petruk. 2022. *Multi-temporal Satellite Imagery Panoptic Segmentation of Agricultural Land in Ukraine*. Ph. D. Dissertation. Ukrainian Catholic University.
- [111] L. L. Pilla, P. Rech, F. Silvestri, C. Frost, P. O. A. Navaux, M. Sonza Reorda, and L. Carro. 2014. Software-Based Hardening Strategies for Neutron Sensitive FFT Algorithms on GPUs. *IEEE Transactions on Nuclear Science* 61, 4 (2014), 1874–1880. <https://doi.org/10.1109/TNS.2014.2301768>
- [112] Sorin C Popescu, Kaiguang Zhao, Amy Neuenschwander, and Chinsu Lin. 2011. Satellite lidar vs. small footprint airborne lidar: Comparing the accuracy of above-ground biomass estimates and forest structure metrics at footprint level. *Remote Sensing of Environment* 115, 11 (2011), 2786–2797.
- [113] Adam Clayton Powell III. 1998. Satellite Imagery: The ethics of a new technology. *Journal of Mass Media Ethics* 13, 2 (1998), 93–98.
- [114] B. Quayle, R. Sohlberg, and J. Descloitres. 2004. Operational remote sensing technologies for wildfire assessment. In *IGARSS 2004. 2004 IEEE International Geoscience and Remote Sensing Symposium*, Vol. 3. 2245–2247 vol.3. <https://doi.org/10.1109/IGARSS.2004.1370809>
- [115] P Krishna Rao, Susan J Holmes, Ralph K Anderson, Jay S Winston, and Paul E Lehr. 1990. Weather satellites: Systems, data, and environmental applications. (1990).
- [116] P. Rech, C. Aguiar, C. Frost, and L. Carro. 2013. An Efficient and Experimentally Tuned Software-Based Hardening Strategy for Matrix Multiplication on GPUs. *IEEE Transactions on Nuclear Science* 60, 4 (2013), 2797–2804. <https://doi.org/10.1109/TNS.2013.2252625>
- [117] Vijay Janapa Reddi, Christine Cheng, David Kanter, Peter Mattson, Guenther Schmuelling, Carole-Jean Wu, Brian Anderson, Maximilien Breughe, Mark Charlebois, William Chou, Ramesh Chukka, Cody Coleman, Sam Davis, Pan Deng, Greg Diamos, Jared Duke, Dave Fick, J. Scott Gardner, Itay Hubara, Sachin Idgunji, Thomas B. Jablin, Jeff Jiao, Tom St. John, Pankaj Kanwar, David Lee, Jeffrey Liao, Anton Lokhmotov, Francisco Massa, Peng Meng, Paulius Micikevicius, Colin Osborne, Gennady Pekhimenko, Arun Tejusve Raghunath Rajan, Dilip Sequeira, Ashish Sirasao, Fei Sun, Hanlin Tang, Michael Thomson, Frank Wei, Ephrem Wu, Lingjie Xu, Koichi Yamada, Bing Yu, George Yuan, Aaron Zhong, Peizhao Zhang, and Yuchen Zhou. 2020. MLPerf Inference Benchmark. In *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*. 446–459. <https://doi.org/10.1109/ISCA45697.2020.00045>
- [118] Fred Ritz and Craig E Peterson. 2004. Multi-mission radioisotope thermoelectric generator (MMRTG) program overview. In *2004 IEEE aerospace conference proceedings (IEEE Cat. No. 04TH8720)*, Vol. 5. IEEE, 2950–2957.
- [119] Julie A Rodiek and Henry W Brandhorst. 2008. Solar array reliability in satellite operations. In *2008 33rd IEEE Photovoltaic Specialists Conference*. IEEE, 1–4.
- [120] Ivan Rodriguez-Ferrandez, Maris Tali, Leonidas Kosmidis, Marta Rovituro, and David Steenari. 2022. Sources of Single Event Effects in the NVIDIA Xavier SoC Family under Proton Irradiation. In *2022 IEEE 28th International Symposium on On-Line Testing and Robust System Design (IOLTS)*. 1–7. <https://doi.org/10.1109/IOLTS56730.2022.9897236>
- [121] J Rosello, A Martellucci, R Acosta, J Nessel, LE Bråten, and C Riva. 2012. 26-GHz data downlink for LEO satellites. In *2012 6th European Conference on Antennas and Propagation (EUCAP)*. IEEE, 111–115.
- [122] JL Ryerson. 1960. Passive satellite communication. *Proceedings of the IRE* 48, 4 (1960), 613–619.
- [123] Mike Safyan. 2020. Planet’s dove satellite constellation. *Handbook of Small Satellites: Technology, Design, Manufacture, Applications, Economics and Regulation* (2020), 1–17.
- [124] Hirobumi Saito, Atsushi Tomiki, Prilando Rizki Akbar, Takashi Ohtani, Kunitoshi Nishijo, Jiro Hirokawa, and Makoto Ando. 2013. Synthetic aperture radar compatible with 100kg class piggy-back satellite. In *Conference Proceedings of 2013 Asia-Pacific Conference on Synthetic Aperture Radar (APSAAR)*. 88–91.
- [125] Fernando Fernandes dos Santos, Pedro Foletto Pimenta, Caio Lunardi, Lucas Draghetti, Luigi Carro, David Kaeli, and Paolo Rech. 2019. Analyzing and Increasing the Reliability of Convolutional Neural Networks on GPUs. *IEEE Transactions on Reliability* 68, 2 (2019), 663–677. <https://doi.org/10.1109/TR.2018.2878387>
- [126] Lucana Santos, Ana Gomez, and Roberto Sarmiento. 2019. Implementation of CCSDS standards for lossless multispectral and hyperspectral satellite image compression. *IEEE Trans. Aerospace Electron. Systems* 56, 2 (2019), 1120–1138.
- [127] Jake Schaffner. 2002. The electronic system design, analysis, integration, and construction of the Cal Poly State University CP1 CubeSat. (2002).
- [128] Science Daily 2008. Urbanization: 95% of the world’s population lives on 10% of the land. <https://www.sciencedaily.com/releases/2008/12/081217192745.htm>
- [129] Pratih Shah and Anthony Lai. 2021. Cots in space: From novelty to necessity. In *35th Annual Small Satellite Conference*.
- [130] C.E. Shannon. 1949. Communication in the Presence of Noise. *Proceedings of the IRE* 37, 1 (1949), 10–21. <https://doi.org/10.1109/JRPROC.1949.232969>
- [131] Jia Shao, Bo Du, Chen Wu, and Lefei Zhang. 2019. Tracking objects from satellite videos: A velocity feature based correlation filter. *IEEE Transactions on Geoscience and Remote Sensing* 57, 10 (2019), 7860–7871.
- [132] Uzair Sharif, Daniel Mueller-Gritschneider, Rafael Stahl, and Ulf Schlichtmann. 2023. Efficient Software-Implemented HW Fault Tolerance for TinyML Inference in Safety-critical Applications. In *2023 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. 1–6. <https://doi.org/10.23919/DATES6975.2023.10137207>
- [133] Bo Shi and Guohua Zhai. 2018. Performance Evaluation of Gbps X-Band Satellite Downlink with Transmitter Imperfections. In *TENCON 2018 - 2018 IEEE Region 10 Conference*. 1110–1113. <https://doi.org/10.1109/TENCON.2018.8650054>
- [134] K. Sidibeh and T. Vladimirova. 2006. IEEE 802.11 Optimisation Techniques for Inter-Satellite Links in LEO Networks. In *2006 8th International Conference Advanced Communication Technology*, Vol. 2. 1177–1182. <https://doi.org/10.1109/ICACT.2006.206181>
- [135] RBC Signals. 2020. Global Ground Station Network for Satellite Operations. <https://rbcsignals.com/partners/global-ground-station-network/>
- [136] Mehtab Singh and Jyoteesh Malhotra. 2020. A high-speed long-haul wavelength division multiplexing-based inter-satellite optical wireless communication link using spectral-efficient 2-D orthogonal modulation scheme. *International Journal of Communication Systems* 33, 6 (2020), e4293.
- [137] Stanley C Solomon, PB Hays, and Vincent J Abreu. 1984. Tomographic inversion of satellite photometry. *Applied optics* 23, 19 (1984), 3409–3414.
- [138] Sua Song, Hongrae Kim, and Young-Keun Chang. 2018. Design and implementation of 3U CubeSat platform architecture. *International Journal of Aerospace Engineering* 2018 (2018), 1–17.
- [139] Erik Mattias Soop. 1994. *Handbook of geostationary orbits*. Vol. 3. Springer Science & Business Media.
- [140] SSC Space 2023. Our stations. <https://sscspace.com/services/satellite-ground-stations/our-stations/>
- [141] Kurtis Thome. 2007. Satellite observations. <https://terra.nasa.gov/citizen-science/air-quality/part-ii-track-pollution-from-space>
- [142] Daisy Thornton. 2023. Edge computing enables DNA sequencing, lower risk of contaminants on International Space Station. <https://federalnewsnetwork.com/technology-main/2023/01/edge-computing-enables-dna-sequencing-lower-risk-of-contaminants-on-international-space-station/>
- [143] Zaid Towfic, Dennis Ogbe, Joe Sauvageau, Douglas Sheldon, Andre Jongeling, Steve Chien, Faiz Mirza, Emily Dunkel, Jason Swope, Mehmet Ogut, Vlad Cretu, and Chris Pagnotta. 2022. Benchmarking and Testing of Qualcomm Snapdragon System-on-Chip for JPL Space Applications and Missions. In *2022 IEEE Aerospace Conference (AERO)*. 1–12. <https://doi.org/10.1109/AERO53065.2022.9843518>
- [144] Do Ngoc Tuyen, Tran Manh Tuan, Le Hoang Son, Tran Thi Ngan, Nguyen Long Giang, Pham Huy Thong, Vu Van Hieu, Vassilis C Gerogiannis, Dimitrios Tzimos, and Andreas Kanavos. 2021. A novel approach combining particle swarm optimization and deep learning for flash flood detection from satellite images. *Mathematics* 9, 22 (2021), 2846.

- [145] Yuri Ulybyshev. 2008. Satellite constellation design for complex coverage. *Journal of Spacecraft and Rockets* 45, 4 (2008), 843–849.
- [146] USEI 2022. <https://www.usei-teleport.com/>
- [147] K Vani et al. 2019. Deep learning based forest fire classification and detection in satellite images. In *2019 11th International Conference on Advanced Computing (ICoAC)*. IEEE, 61–65.
- [148] Ivo Vertat and Ales Vobornik. 2014. Efficient and reliable solar panels for small CubeSat picosatellites. *International Journal of Photoenergy* 2014 (2014).
- [149] ViaSat. 2023. Ground stations as a Service. <https://www.viasat.com/space-innovation/space-and-networking-technology/ground-network/>
- [150] Mark Voskuijl. 2022. Performance analysis and design of loitering munitions: A comprehensive technical survey of recent developments. *Defence Technology* 18, 3 (2022), 325–343.
- [151] Hongqiang Wang and Hongqiang Wang. 2008. An implementation of CCSDS 122.0-B-1 recommended standard. *University of Nebraska-Lincoln* (2008).
- [152] Ze Wang, Tianxu Zhang, Luxin Yan, and Cheng Gong. 2010. A High Performance Fully Pipelined Architecture for Lossless Compression of Satellite Image. In *2010 International Conference on Multimedia Technology*. 1–4. <https://doi.org/10.1109/ICMULT.2010.5631429>
- [153] Shaun Waterman. 2021. Renting the ground: The growing future of ground segment as a Service. https://interactive.satellitetoday.com/via/november-2021/renting-the-ground-the-growing-future-of-ground-segment-as-a-service/_fragment.html
- [154] Tim Weiner. 1993. Titan lost payload: Spy-satellite system worth 800 million. <https://www.nytimes.com/1993/08/04/us/titan-lost-payload-spy-satellite-system-worth-800-million.html?src=pm>
- [155] Marc Wieland and Sandro Martinis. 2019. A modular processing chain for automated flood monitoring from multi-spectral satellite data. *Remote Sensing* 11, 19 (2019), 2330.
- [156] Shuo Yang. 2007. <http://www.stats.gov.cn/english/icas/papers/p020071017422083591595.pdf>
- [157] Sami Yenice, Mehmet Aktaş, Fazıl Yaraşlı, Tunç Durmuş Medeni, Ahmet Savaş, and Halil Yeşilçimen. 2019. A New Emergency Notification System Developed by TÜRKSAT as Part Of EMYNOS Project Turkish Field Trials. (2019).
- [158] Grace T Yi, Kim K de Groh, Bruce A Banks, Athena Haloua, Emily C Imka, and Gianna G Mitchell. 2013. Overview of the MISSE 7 polymers and zenith polymers experiments after 1.5 years of space exposure. In *International Symposium on Materials in the Space Environment (ISMSE-12)*.
- [159] Ju Zhang, Qingwu Hu, Jiayuan Li, and Mingyao Ai. 2020. Learning from GPS trajectories of floating car for CNN-based urban road extraction with high-resolution satellite imagery. *IEEE Transactions on Geoscience and Remote Sensing* 59, 3 (2020), 1836–1847.
- [160] Xiangwei Zhao, Shun Pan, Zhongchang Sun, Huadong Guo, Lei Zhang, and Kang Feng. 2021. Advances of satellite remote sensing technology in earthquake prediction. *Natural Hazards Review* 22, 1 (2021), 03120001.
- [161] Haojie Zhou, Tong Tian, Xinyue Wang, and Ji Li. 2023. Combining looped heat pipe and thermoelectric generator module to pursue data center servers with possible power usage effectiveness less than 1. *Applied Energy* 332 (2023), 120539.
- [162] Xiaoming Zhu and J.M. Kahn. 2002. Free-space optical communication through atmospheric turbulence channels. *IEEE Transactions on Communications* 50, 8 (2002), 1293–1300. <https://doi.org/10.1109/TCOMM.2002.800829>