

**Spatial and temporal analysis of the Polar
Radiant Energy in the Far-InfraRed Experiment
(PREFIRE) intersections**

Natasha Vos

A thesis submitted in partial fulfillment of the requirements

for the degree of

Master of Science

Atmospheric and Oceanic Sciences

at the

UNIVERSITY OF WISCONSIN-MADISON

Acknowledgements

This document describes the culmination of my graduate research that began in Fall 2021. From the outset, I have had the good fortune of encountering numerous benefactors in the form of mentors, collaborators and colleagues as well as supportive family and friends, whose contributions to my thesis have been paramount to its success. I extend considerable gratitude to those whom this applies and would like to personally acknowledge the following individuals:

My advisor, Tristan L'Ecuyer, who has advised me with unfailing patience and enthusiasm. Thank you for affording me the opportunity to take part in the momentous PREFIRE mission as a student researcher. Above all, I value your passion for pioneering novel approaches to Earth science and your penchant for sharing this passion with others, especially early career scientists.

My thesis readers, Andrea Lopez Lang and David Henderson, whose thoughtful feedback has empowered me to communicate my research with greater efficacy. Thank you for the time you invested in reviewing this paper and helping me achieve more clarity and purpose in my work.

Tim Michaels, Kyle Mattingly, Nathan Miller and Aronne Merrelli, essential members of the PREFIRE science team whose insights have guided me through numerous challenges. Thank you for helping me navigate and interpret the files underpinning my research. An extra special thank you to Tim, who provided the simulated orbit files on which this work relied and who has “put out” countless figurative fires to make many lives easier, mine included.

Rudradutt Thaker, Hamish Prince and Stephanie Orland, three exceptional early career scientists I have had the privilege of meeting in this program. You have each reliably helped me when needed, and you have done so without hesitation. Thank you for enduring me and inspiring me to be a better scientist.

Most importantly, my dad (Steven), mom (Sharon), brother (Shawn), and partner (Mark). Thank you for being the best part of my life. Language is insufficient to express how much I love and value you.

Abstract

The Polar Radiant Energy in the Far-InfraRed Experiment (PREFIRE) will use two 6U CubeSats to continuously measure spectral far-infrared (FIR) emissions for the first time in the modern satellite era. By strategically operating two CubeSats in separate sun-synchronous orbits, PREFIRE will achieve frequent orbit resampling, or intersections, that afford insights into the underlying polar processes that modulate FIR emissions. These orbit intersections are integral to PREFIRE science and will likely feature prominently in future CubeSat missions, motivating methods to characterize resampling distributions. Accordingly, this thesis develops new methods to locate orbit intersections and extract co-located pixels between crossovers. Such methods are then applied to simulated PREFIRE orbits to characterize the spatial and temporal distribution of hypothetical PREFIRE intersections and identify a subset with minimal revisit times that can be used for inter-calibration during the mission. This analysis confirms that hundreds of PREFIRE intersections are anticipated each day, with the majority (>75%) occurring poleward of 66.5° latitude. Notably, the addition of a second CubeSat nearly quadruples the number of sub-daily polar intersections. Spatial and temporal resampling coverage is clearly enhanced when more than one CubeSat is enlisted, securing greater latitudinal representation and more diverse time differences between crossovers compared to a scenario in which only a single CubeSat is deployed. Additionally, when two CubeSats are placed at different altitudes, we find that their intersections exhibit time-varying, cyclic coverage. This orbit configuration further yields simulated intersections between latitudes 72° and 78° with time differences sufficiently short for inter-calibration. Finally, we consider three conceptual cases that indicate PREFIRE intersections will be valuable for capturing processes like sea ice melt and high-latitude cloud cover change as well as for inter-calibrating the PREFIRE sensors.

Table of Contents

| | |
|---|----|
| 1. Introduction | 1 |
| 1.1 Polar Radiant Energy in the Far-InfraRed Experiment (PREFIRE) | 1 |
| 1.2 PREFIRE satellites and measurements | 5 |
| 1.3 Classifying PREFIRE intersections | 10 |
| 2. Methodology | 12 |
| 2.1 Identifying intersections..... | 12 |
| 2.1.1 Simulated PREFIRE orbits..... | 12 |
| 2.1.2 Fourth-scene method | 13 |
| 2.1.3 Quality testing..... | 16 |
| 2.1.4 Invalid resampling | 17 |
| 2.2 Extracting co-located pixels | 18 |
| 2.2.1 Intersection vertices..... | 19 |
| 2.2.2 Interior pixels..... | 21 |
| 2.2.3 “Common pixels” | 23 |
| 2.3 Methods for conceptual cases | 24 |
| 2.3.1 Arctic heat wave case methods..... | 25 |
| 2.3.2 Cloud dissipation case methods..... | 32 |
| 2.3.3 Calibration case methods..... | 35 |
| 3. Spatial and temporal analysis of PREFIRE intersections | 39 |
| 3.1 Numerical distribution of simulated intersections | 39 |
| 3.2 Spatial distribution of simulated intersections | 43 |
| 3.3 Temporal distribution of simulated intersections..... | 49 |
| 3.4 Simulated short-term revisits | 59 |
| 4. Conceptual cases | 61 |
| 4.1 Arctic heat wave case results | 61 |
| 4.2 Cloud dissipation case results | 73 |
| 4.3 Calibration case results..... | 82 |
| 5. Discussion | 88 |
| 6. References | 96 |

1. Introduction

1.1 Polar Radiant Energy in the Far-InfraRed Experiment (PREFIRE)

Thermal emissions at wavelengths between 15 and 100 μm , generally referred to as the far-infrared (FIR), dominate outgoing longwave radiation (OLR) in the poles, peaking at over 60% in Antarctica during austral winter (L'Ecuyer et al. 2021). This suggests that outgoing FIR characterizes much of the polar energy budget. FIR emissions also account for a considerable fraction of OLR in middle and low latitudes, particularly in elevated regions like the Tibetan Plateau as well as regions associated with persistent deep convection (L'Ecuyer et al 2021). Harries et al. (2008) showed that first-principle calculations yield emissions that peak in the FIR for terrestrial temperatures, suggesting that emitted FIR is important to the energy budget everywhere on Earth. Contemporary research further reveals that outgoing FIR exceeds 40% of outgoing thermal energy at virtually every point on Earth (L'Ecuyer et al. 2021). Averaged globally, approximately half (49.1%) of terrestrial OLR in 2019 comprised wavelengths greater than 15 μm , as shown by L'Ecuyer et al. (2021). It therefore follows that upwelling FIR is a major constituent of thermal emissions, especially in the frigid poles.

Despite their large contribution to the global energy budget, FIR emissions have not been systematically constrained, owing to past observational challenges (e.g., Harries et al. 2008; Feldman et al. 2014; Palchetti et al. 2020; L'Ecuyer et al. 2021). This stands in marked contrast to outgoing mid-infrared (MIR), taken to be between 4 and 15 μm , which is tightly constrained following decades of direct MIR observations (e.g., Huang et al. 2016). Since 1999 and 2002, the Moderate Resolution Imaging Spectroradiometer (MODIS; Justice et al. 1998) aboard Terra and Aqua, respectively, has collected global multispectral radiances between 0.4 and 14.4 μm . Additionally, hyperspectral MIR retrievals have been achieved using high spectral resolution

sounders like the Atmospheric Infrared Sounder (AIRS; Pagano et al. 2010) and the Cross-track Infrared Sounder (CrIS; Glumb et al. 2003), which provide fine resolution data for wavelengths up to 15.4 μm . Conversely, spectrally resolved FIR estimates have been stymied by a lack of direct observations, thus resulting in knowledge gaps with respect to both global and polar energy budgets (e.g., Harries et al. 2008; L'Ecuyer et al. 2021).

Terrestrial FIR measurements from space have generally been limited to broadband observations. Two notable broadband missions include the Earth Radiation Budget Experiment (ERBE) and the Clouds and the Earth's Radiant Energy System (CERES). Beginning in 1984 with the launch of the Earth Radiation Budget Satellite, ERBE facilitated broadband measurements using three detectors that respectively collected shortwave, longwave and total outgoing energy, with the longwave contingent measuring up to about 25 μm (Barkstrom et al. 1989). Its well-known successor, CERES, has collected broadband data since 1997 using a three-channel scanning broadband radiometer (Wielicki et al. 1996). In addition to a shortwave channel and a channel dedicated to the atmospheric window, CERES provides broadband radiances between 0.2 and 100 μm (Smith et al. 2011). Although missions like ERBE and CERES have granted us valuable insights into Earth's OLR, their data do not compensate for the near-complete lack of narrow-band FIR observations.

Remarkably, almost no terrestrial FIR missions have been attempted since NASA's 1970–71 Nimbus-4, which achieved the first global narrow-band FIR measurements from space (Huang et al. 2014). It is worth noting that the InfraRed Interferometer Spectrometer (IRIS) aboard Nimbus-4 measured only up to 25 μm (Hanel et al. 1972). Subsequent Soviet missions Meteor-25 and Meteor-28 that were launched in 1976 and 1977, respectively, also measured wavelengths up to 25 μm (Kempe et al. 1980), but there have been no continuous narrow-band observations of

terrestrial FIR in the modern satellite era. Critically, emissions at longer FIR wavelengths ($> 25 \mu\text{m}$) are altogether uncharted.

This observational shortcoming can be attributed to several challenges unique to FIR spectra. First, the high sensitivity needed to narrowly resolve outgoing FIR previously required expensive supercooled detectors (L'Ecuyer et al. 2021). Additionally, the low photon energy associated with FIR frequencies traditionally mandated long integration times (Harries et al. 2008). These peculiar features of FIR emissions have behaved as natural barriers to their systematic measurement via satellite. Another primary challenge to direct FIR observations is water vapor continuum absorption (e.g., Clough et al. 1989). That is, strong H_2O rotational absorption dominates much of the FIR band, making low-altitude observations from space impractical under most conditions (Davis 1993), with evidence suggesting that total column water vapor must be less than 1 mm for strong FIR emissions to reach the upper atmosphere from the surface (Palchetti et al. 2020). Historically, such challenges have been central to the persistent lack of narrow-band measurements in this important portion of thermal energy. Incomplete spectral measurements ultimately result in lingering uncertainties with respect to FIR fluxes, and, by extension, the atmospheric greenhouse effect (AGHE) to which FIR fluxes are major constituents. Uncertainties are moreover confounded by persistent knowledge gaps concerning FIR emissivities.

Recall that, for any surface, thermal emissions directly depend not only on surface temperature but also on the spectral emissivity intrinsic to the surface (e.g., Schmugge et al. 2002). Unfortunately, our inability to constrain terrestrial emissivities in the FIR has necessitated making assumptions about emissivity across these spectra, and such assumptions remain a source of error in radiative transfer models (Chen et al. 2014). For example, numerical models have generally assumed that terrestrial surfaces exhibit blackbody emissivity in the FIR; however, radiative

transfer studies suggest that the blackbody assumption results in non-negligible discrepancies in calculated top-of-atmosphere (TOA) OLR, with polar oceans and low-latitude deserts exhibiting the greatest sensitivities (Huang et al. 2016). In the latter regions, peak differences in outgoing longwave were on the order of 10 Wm^{-2} , while OLR differed by roughly 2 Wm^{-2} over polar oceans when realistic emissivity values were used rather than assuming an emissivity of unity (Huang et al. 2016). Thus, the extent to which such discrepancies impact energy budget analyses may be substantial, especially in sensitive areas such as the poles where surface flux estimates are plagued with uncertainties. For instance, comparing estimates of decadal mean annual surface flux from 2000-2009 reveals uncertainties ranging between $35\text{-}70 \text{ Wm}^{-2}$ and $70\text{-}80 \text{ Wm}^{-2}$ in the Arctic and Antarctic, respectively (L'Ecuyer et al. 2021). These findings strongly imply that energy fluxes at either pole are far from settled and that reducing overall uncertainty will require, among other things, systematic measurements in the FIR to a degree not yet seen in past Earth System observations.

In response to this longstanding observational gap, a collaborative effort led by the University of Wisconsin-Madison and NASA's Jet Propulsion Laboratory has culminated in the Polar Radiant Energy in the Far-InfraRed Experiment (PREFIRE; L'Ecuyer et al. 2021). PREFIRE will pioneer passive, moderate-resolution FIR measurements from space by leveraging nascent sensor technology deployed on two nanosatellites called CubeSats, which are detailed in the next section. Together, the PREFIRE CubeSats will retrieve near-global measurements of TOA radiances up to $54 \mu\text{m}$. With its aim decidedly polar, the fundamental goal of PREFIRE is to improve our understanding of the AGHE in the poles where outgoing FIR dominates. Moreover, the mission intends to constrain FIR spectral emissivities for common polar surfaces. To pursue either end, it will be critical to achieve frequent observations of localized surface and atmospheric

changes. PREFIRE was therefore designed so that its duo CubeSats frequently resample high latitudes, yielding two-dimensional regions in which the fields of view (FOVs) of two unique orbits are approximately co-located on various sub-daily timescales. Such regions are hereafter referred to as the PREFIRE intersections, and they lend themselves as the focus of this work.

1.2 PREFIRE satellites and measurements

PREFIRE is among the first Earth science missions to rely on CubeSats to meet its primary objectives (Tristan et al. 2021). CubeSats, which constitute modular miniaturized satellites, are noted for their light and inexpensive designs vis-à-vis traditional research satellites (e.g., Shiroma et al. 2011). For reference, CubeSats generally do not exceed 10 kg whereas traditional satellite platforms commonly weigh more than 2000 kg (Shiroma et al. 2011). Conceived in 1999 by Jordi Puig-Suari at California Polytechnic State University, San Luis Obispo and Bob Twiggs at Stanford's Space Systems Development Laboratory, CubeSats have become standardized space vehicles with a cubic base unit denoted 1U, which are frequently combined in 2U, 3U or 6U configurations (Lee 2017). Each 1U cube has dimensions of 10 cm^3 and mass on the order of 1 kg (Puig-Suari et al. 2001). Their simple, standardized design means they can be built on relatively short timescales, with production times ranging from four years to less than two (Gill et al. 2010). This represents a rapid turnaround time compared to traditional platforms and hence promotes the use of CubeSats as testbeds for new payloads, among other novel uses.

Since their debut in academia 25 years ago, CubeSats have become ideal candidates for experimental missions undertaken by national space agencies, colleges and universities, and industry the world over (e.g., Buckhard & Weston 2021). In fact, modern CubeSat missions have been featured in myriad fields including Earth science, biology, global navigation, communications, deep space exploration, and a host of other emerging technological and other

industries (Saeed et al. 2020). As detailed in Buckhard & Weston (2021), however, CubeSats were seldom leveraged in science missions before 2011 and have only recently been incorporated into broader Earth System endeavors.

As the applicability of CubeSats expanded, the intrinsic research potential of dispatching multiple CubeSats has been swiftly recognized, gaining momentum as a research topic over the last decade (Buckhard & Weston 2021). Several modern missions have experimented with CubeSat networks. In 2017, for example, an international research collective pioneered a constellation of 36 2U and 3U CubeSats for QB50, a mission that sought in-situ measurements of the mid-lower Thermosphere (Malan et al. 2017). During preliminary planning, a network of traditional satellites had been dismissed outright since their use in a mission of this nature would be financially impractical (Gill et al. 2010). QB50 thus served not only as a science mission of a kind but also a concept mission aimed at demonstrating the revolutionary utility of the CubeSat platform. More recently, NASA launched the Time-Resolved Observations of Precipitation structure and storm Intensity with a Constellation of Smallsats (TROPICS; Blackwell et al. 2018). This tropical mission deployed four 3U CubeSats in 2023, with each TROPICS satellite making passive microwave measurements in the low latitudes (Leslie et al. 2023). Such measurements will help reveal the temperature, moisture and precipitation profiles during the life cycle of tropical cyclones. By strategically combining CubeSats, TROPICS anticipates unprecedented temporal coverage of tropical cyclones (Blackwell et al. 2018; Leslie et al. 2023). Notably, PREFIRE is among the newest space missions to leverage the enhanced temporal coverage permitted by CubeSats to advance our understanding of the complex Earth System. In this case, PREFIRE is a first-of-its-kind mission to use the CubeSat platform to secure frequent observations in polar regions.

By June 2024, PREFIRE successfully deployed two 6U CubeSats into separate sun-synchronous orbits via two launches, where “separate” implies a different mean local time of the ascending node (MLTAN) between satellites. The MLTAN offset follows from the orbital trade study undertaken by Drouin et al. (2022), which found that an equatorial separation between PREFIRE CubeSats would yield favorable high-latitude coverage. For completeness, the author notes that a sun-synchronous orbit implies an inclination angle of approximately 98° , though the exact inclination of a sun-synchronous satellite is impacted by drift and other perturbations (Jordan et al. 1993). Estimated PREFIRE orbit parameters following the respective launch of each satellite, presently referred to as simply CubeSat 1 and CubeSat 2, are listed in Table 1.

| | |
|---------------------------|-----------------|
| Orbit type | Sun-synchronous |
| CubeSat 1 altitude | 530.9 km |
| CubeSat 2 altitude | 530.6 km |
| CubeSat 1 MLTAN | 03:38 UTC |
| CubeSat 2 MLTAN | 08:03 UTC |

Table 1: Estimated PREFIRE orbit parameters following launch. Note that altitude and MLTAN are subject to change over time.

Over the course of the mission, the satellites will continuously measure TOA radiances via ambient-temperature Thermal Infrared Spectrometers (TIRS), which occupy only about half the space aboard their respective 30 cm by 20 cm by 10 cm CubeSat. Together, CubeSat 1 and CubeSat 2 will yield observations on timescales ranging from sub-daily to seasonal, with a spectral resolution of about $0.86 \mu\text{m}$ and coverage extending from roughly 82°S to 82°N . It should be noted that the central wavelengths of the TIRS channels are not identical across detectors, and their approximate pre-launch values are listed separately in Table 2 for reference. Note that either detector has spectral gaps near $7 \mu\text{m}$, $15 \mu\text{m}$ and $30 \mu\text{m}$ (Miller et al. 2023).

| TIRS1 channel centers (μm) | TIRS2 channel centers (μm) |
|---|---|
| 2.95 | 4.20 |
| 3.80 | 5.05 |
| 4.64 | 5.89 |
| 5.48 | 6.73 |
| 8.02 | 9.27 |
| 8.86 | 10.11 |
| 9.70 | 10.95 |
| 10.55 | 11.80 |
| 11.39 | 12.64 |
| 12.24 | 13.48 |
| 13.08 | 14.33 |
| 15.61 | 16.86 |
| 16.45 | 17.70 |
| 17.30 | 18.55 |
| 18.14 | 19.39 |
| 18.99 | 20.24 |
| 19.83 | 21.08 |
| 20.67 | 21.92 |
| 21.52 | 22.77 |
| 22.36 | 23.61 |
| 23.20 | 24.45 |
| 24.05 | 25.30 |
| 24.89 | 26.14 |
| 25.74 | 26.99 |
| 26.58 | 27.83 |
| 27.42 | 28.67 |
| 28.27 | 29.52 |
| 30.80 | 32.05 |
| 31.64 | 32.89 |
| 32.49 | 33.74 |
| 33.33 | 34.58 |
| 34.17 | 35.42 |
| 35.02 | 36.27 |
| 35.86 | 37.11 |
| 36.71 | 37.96 |
| 37.55 | 38.80 |
| 38.39 | 39.64 |
| 39.24 | 40.49 |
| 40.08 | 41.33 |
| 40.92 | 42.17 |
| 41.77 | 43.02 |
| 42.61 | 43.86 |
| 43.46 | 44.71 |

| | |
|-------|-------|
| 44.30 | 45.55 |
| 45.14 | 46.39 |
| 45.99 | 47.24 |
| 46.83 | 48.08 |
| 47.68 | 48.92 |
| 48.52 | 49.77 |
| 49.36 | 50.61 |
| 50.21 | 51.46 |
| 51.05 | 52.30 |
| 51.89 | 53.14 |
| 52.74 | 53.99 |

Table 2: Approximate central wavelengths for each TIRS channel. Values shown are derived from pre-launch estimates.

Each uncooled detector, described in greater detail by L’Ecuyer et al. (2021), collects observations via near-nadir pushbroom scanning. As the CubeSats move forward in time, they complete discrete orbits in roughly 95 minutes, resulting in approximately 15 orbits per day. An example of simulated PREFIRE orbits is shown in Figure 1. Swaths nominally exceed 200 km across and are composed of eight scenes that extend roughly perpendicular to the motion of the subtrack. When projected onto Earth’s surface, scenes form footprints that are roughly 12 km wide (Miller et al. 2023). In the across-track direction, scenes are separated from one another by a space approximately double their width, while scenes overlap in the along-track direction so that roughly every six successive scenes are partially co-located.




Figure 1: Simulated PREFIRE orbits (or “granules”), with pink and blue denoting CubeSat 1 and CubeSat 2, respectively, and corresponding arrows denoting the approximate scan direction. For reference, an arbitrary cross-track segment has been extracted from each granule and is shown in isolation. Polygons represent the eight scenes that compose the swath of either satellite and are numbered accordingly. Note that the actual PREFIRE orbits may differ.

The mission’s compact hardware combined with the absence of spacecraft fuel contributes to the low-cost, lightweight design characteristic of CubeSat technology. Although such features naturally limit the mission duration overall, they also enable more experimental Earth observations, including the strategic use of orbit intersections.

1.3 Classifying PREFIRE intersections

As described in section 1.1, the orbit parameters selected for PREFIRE were chosen in part to facilitate frequent orbit resampling, particularly in high latitudes (Drouin et al. 2022). Orbit

resampling has previously featured in other satellite applications, including sensor calibration where satellite revisits are commonly recognized as near-nadir overpasses (e.g., Cao et al. 2004). Resampling is therefore not unique to PREFIRE; however, the mission is among the first of its kind to leverage resampling to investigate high-latitude OLR. By targeting orbits that yield recurring ground tracks in either pole, as shown by Drouin et al. (2022), the resulting satellite revisits will provide continuous time-lapsed observations of polar processes—for example, ice sheet melt or Arctic cloud cover changes—on a variety of sub-daily and longer timescales, thereby revealing the spectral signatures of many key features that modulate polar emissions. Resampling is therefore essential to the mission and constitutes the so-called *PREFIRE intersections* on which this paper will focus.

Consider the scenario in which a single CubeSat revisits a region it previously sampled, a type of resampling we refer to as *self-intersections* in the context of PREFIRE. This intersection type avails itself to either CubeSat, thereby yielding *SAT1 self-intersections* and *SAT2 self-intersections*, where *SAT1* and *SAT2* are shorthand references to CubeSat 1 and CubeSat 2, respectively. From the outset of this analysis, we assumed there could be meaningful differences between SAT1 self-intersections and SAT2 self-intersections and as such investigated them separately. Note that self-intersections are not possible on timescales shorter than the duration of one orbit since sampling will never repeat during a single orbit.

Alternatively, a CubeSat may revisit a region previously sampled by the *other* CubeSat, resulting in time-lapsed, co-located FOVs across satellites. When resampling occurs between CubeSats, we describe this as a cross-satellite intersection, and going forward they are collectively referred to as *SAT1-SAT2 intersections*. Unlike self-intersections, intersections between CubeSats have no theoretical lower bound with respect to time, such that they can occur at any timescale,

including sub-orbit timescales. This is a meaningful distinction between the intersection types and will be explored in greater detail in subsequent sections.

Guided by the premise that the intersections will serve an integral role in the mission as a matter of providing insights into dynamic local conditions, we sought to characterize them ahead of the launch. As such, this research investigated their potential spatial and temporal pattern in the context of the targeted orbit configurations. Additionally, a conceptual analysis was undertaken to demonstrate the utility of the intersections as a tool for both Earth System observations and inter-calibration.

2. Methodology

2.1 Identifying intersections

2.1.1 Simulated PREFIRE orbits

Absent real orbit data, simulated orbits were generated using Simplified General Perturbations No. 4 to facilitate an analysis of plausible intersections. Roughly two continuous months of orbits were simulated for both satellites, a duration taken to be sufficiently long to capture longer-term patterns, if any. Each orbit was assigned a hypothetical timestamp over the duration extending from June 28 to August 28, 2021. For every orbit simulated, the date, UTC time, and latitude and longitude for each scene was stored in a single NetCDF file (Rew & Davis 1990). Each discrete simulated orbit is hereafter referred to as a “granule.”

It is worth emphasizing that this thesis was an undertaking that preceded the launch of either PREFIRE satellite, such that the orbits the mission eventually achieved were not known during the planning or implementation of this work. Consequently, the prescribed orbit parameters that feature in this thesis are not identical to those characterized in Table 1, and they are in fact modeled after preliminary targets that were later amended as the launch date neared. As

summarized in Table 3, this analysis represented the PREFIRE CubeSats by simulating separate sun-synchronous orbits and prescribing an MLTAN of 08:00 UTC and 16:00 UTC to CubeSat 1 and CubeSat 2, respectively. Additionally, we assumed the launches may not secure identical altitudes for both satellites, an assumption nurtured by the lack of fuel with which to raise or lower either spacecraft. The synthetic granules were therefore created by selecting different altitudes for each simulated satellite, with CubeSat 1 and CubeSat 2 granted an altitude of 538 km and 525 km, respectively. This arbitrary altitude difference was chosen to model a launch with relatively large errors, guiding us to a presumably conservative estimate of the potential resampling distribution. Drag was also continuously applied to both CubeSats.

| | |
|---------------------------|-----------------|
| Orbit type | Sun-synchronous |
| CubeSat 1 altitude | 538 km |
| CubeSat 2 altitude | 525 km |
| CubeSat 1 MLTAN | 08:00 UTC |
| CubeSat 2 MLTAN | 16:00 UTC |

Table 3: The orbit configuration prescribed in the simulated orbits analyzed in this thesis. Note that the values shown are different from the actual PREFIRE orbits summarized in Table 1.

2.1.2 Fourth-scene method

Identifying intersections is a task that, sans simplifications, concerns hundreds of thousands of pixels for every pair of orbits. To optimize this process, the method used for this analysis considers only the fourth scene of each granule and neglects the remaining seven. This follows from noting that the fourth and fifth scenes constitute the interior-most scenes of a swath. Consequently, we can approximate the center of each intersection by determining where either the fourth or fifth scenes of two orbits are roughly co-located. Since there is no obvious advantage to using the fourth scene over the fifth or vice versa, the fourth was chosen arbitrarily. By limiting

the scope of our analysis to only the fourth scene, the process of determining the relative center of simulated intersections was streamlined.

It is sufficient to compile general sampling statistics by treating each scene as an idealized point represented by its center. Ergo, unless otherwise indicated, the coordinates of each scene correspond to the latitude and longitude of its central point. Additionally, owing to the fact scenes partially overlap in the along-track direction, computation time was reduced by a factor of six by skipping every six along-track segments for any pair of orbits.

With the simplifications described above, our objective was to approximate the location at which the fourth scenes of two granules intersect. Prior to testing the proximity of the fourth scenes, some preliminary steps were taken to eliminate regions that were obviously lacking intersections. Every granule was first divided into global quadrants, resulting in northwestern, northeastern, southwestern, and southeastern segments. This manner of partitioning granules facilitated more targeted, and therefore more efficient, proximity testing. For each regional segment, pixels were then binned every 10° longitude, following from the observation that the longitudinal coverage within the four regions varies by orbit. Thus, regions for which intersections are not possible can be quickly eliminated by subselecting segments and comparing pixels in mutual longitude bins. For every pair-wise comparison between granules, scenes falling outside of mutual longitudes for either granule were immediately discarded, and only the remaining pixels were tested for proximity.

Note that identifying intersections does not require the coordinates of two unique fourth scenes to completely converge. Since the swath extends some 200 km across, an intersection is properly detected when the fourth scenes of two granules are *sufficiently close*. The threshold at which we allow two interior scenes to be separated during proximity testing is therefore somewhat

arbitrary. At the upper bound, the distance threshold is constrained by the dimensions of the swath. In practice, however, a much smaller threshold is desired to approximate the center of an intersection rather than simply any point contained therein. On the opposite end, the smaller the selected threshold, the greater the likelihood of failing to identify an intersection altogether. This is particularly concerning since intersection orientation and size is not fixed. To navigate between either extreme, the approach taken in this thesis was to immediately discard every pair of fourth scenes greater than 25 km apart. Then, the distance between each retained pair was calculated using a function that converts coordinates in decimal degrees to physical distance, and the pair yielding the smallest distance was selected as a probable intersection. To avoid duplicates, each pair selected in this manner was then compared to all previous intersections identified for the same pair of granules, if any. If the coordinates of the newly flagged pair were found to occur within 500 km of any intersection previously identified for that granule pair, then it was discarded as a duplicate intersection. Otherwise, the selected pair of fourth scenes was taken to be the approximate center of an intersection.

Since a pair of fourth scenes is retained as an intersection center when they are approximately co-located, retaining the scene coordinates from both granules is therefore redundant. Consequently, for every intersection, the coordinates for only the second granule are retained and taken to be the relative intersection center. In addition to the latitude and longitude corresponding to the intersection, the simulated timestamp for either granule is retrieved and the time difference calculated and retained. For every intersection, a so-called “intersection ID” was then generated that contained CubeSat type and granule number as well as the along-track position for the flagged scene in either granule. Figure 2 illustrates how our simple method identifies the approximate center of an intersection over interior Antarctica.




Figure 2: The fourth-scene method for an arbitrary Antarctic intersection. The intersection is defined as the region bounded by four vertices.

Using the fourth-scene method, we identified all intersections within a rolling 48-hour (32-granule) window over the two-month period, starting with the granule timestamped at +48 hours of the first available granule and ending with the granule timestamped at -48 hours of the last available granule. Hereafter, the elapsed time between sampling is called the *time difference* of an intersection and is denoted Δt . Note that a time difference is relative to the granule being analyzed, so that a positive (negative) time difference implies that the granule being considered completed the earlier (later) crossover. This sign convention is arbitrary and may change during the mission.

2.1.3 Quality testing

Systematic quality assurance was required to ensure the accuracy of our results. This process relied largely on the important observation that intersections come in pairs. To leverage this observation, the number of intersections between every combination of two orbits was enumerated. If an odd number of intersections was identified for a particular orbit pair, then that pair of orbits was flagged as having had either a duplicate or missed intersection. Additional testing

was undertaken to determine which scenario—either a duplicate or missed intersection—impacted the flagged pairs.

Missed intersections were rare and primarily occurred in the vicinity of the prime meridian. To identify those missed, an iterative process was undertaken to repeat the fourth-scene method with more lenient requirements. For the first iteration, the distance threshold was increased from 25 km to 40 km, a threshold determined to be sufficiently large to identify most missed intersections. Any remaining missed intersections were then identified by repeating the process again but testing every along-track segment rather than skipping every six as before. Extending the distance threshold and considering every along-track segment, though computationally much slower, was sufficient to identify all remaining missed intersections.

Conversely, duplicate intersections were more common and occurred in the vicinity of the equator. This follows from the fact that equatorial intersections occur near the terminus of granules, such that it is possible to identify a single equatorial intersection once near the end of one granule and then again near the beginning of the next successive granule. Recall that real-time duplicate testing was undertaken by comparing the intersections per granule pair and therefore did not account for the scenario described above. Thus, duplicates that survived initial testing were identified by retrieving the time difference and coordinates for all intersections. For intersections with common time differences and coordinates, a simple test was done to determine whether the timestamps of the corresponding granules were less than 90 minutes apart, the approximate duration of a single orbit. Doing so revealed intersections that had been detected for successive granules, which were then flagged as duplicates and removed from the final analysis.

2.1.4 Invalid resampling

There were numerous orbit combinations that yielded partially or fully repeating ground tracks. This phenomenon, described in previous literature (e.g., Rees 1992), may yield valuable information content, but repeating ground tracks were excluded from this thesis. Instead, we required that resampling regions exhibit approximate diamond shapes (i.e., a clearly defined polygon bounded by four vertices) in order to be classified as intersections. Moreover, we excluded resampling regions that were relatively elongated in favor of restricting intersections to more compact and homogenous regions. All resampling regions that failed to satisfy these criteria were classified as invalid intersections and were discarded accordingly.

To identify invalid intersections, all combinations of orbits were tested for proximity by measuring the distance between the granules near the equator. This approach follows from the relationship between a) the equatorial proximity of two orbits and b) the spatial extent of their resulting intersections, where a) and b) are positively correlated. Thus, if the equatorial fourth scenes of two orbits was less than 700 km apart, a threshold chosen following qualitative testing, then the separation between them was considered too small and the granule pair was removed from this thesis. Valid intersections are therefore those that are compactly bounded by the tracks of the first and eighth scenes of two unique orbits.

2.2 Extracting co-located pixels

The following subsections present methods for extracting additional co-located pixels between orbit pairs. Operationally, analyses of spectral changes captured by PREFIRE intersections will critically rely on the careful discrimination of scenes within a resampled region. For reference, Figure 3 summarizes four categories of pixels that characterize PREFIRE intersections. Intersection vertices, interior pixels and common pixels will be addressed in section 2.2.1, 2.2.2 and 2.2.3, respectively.

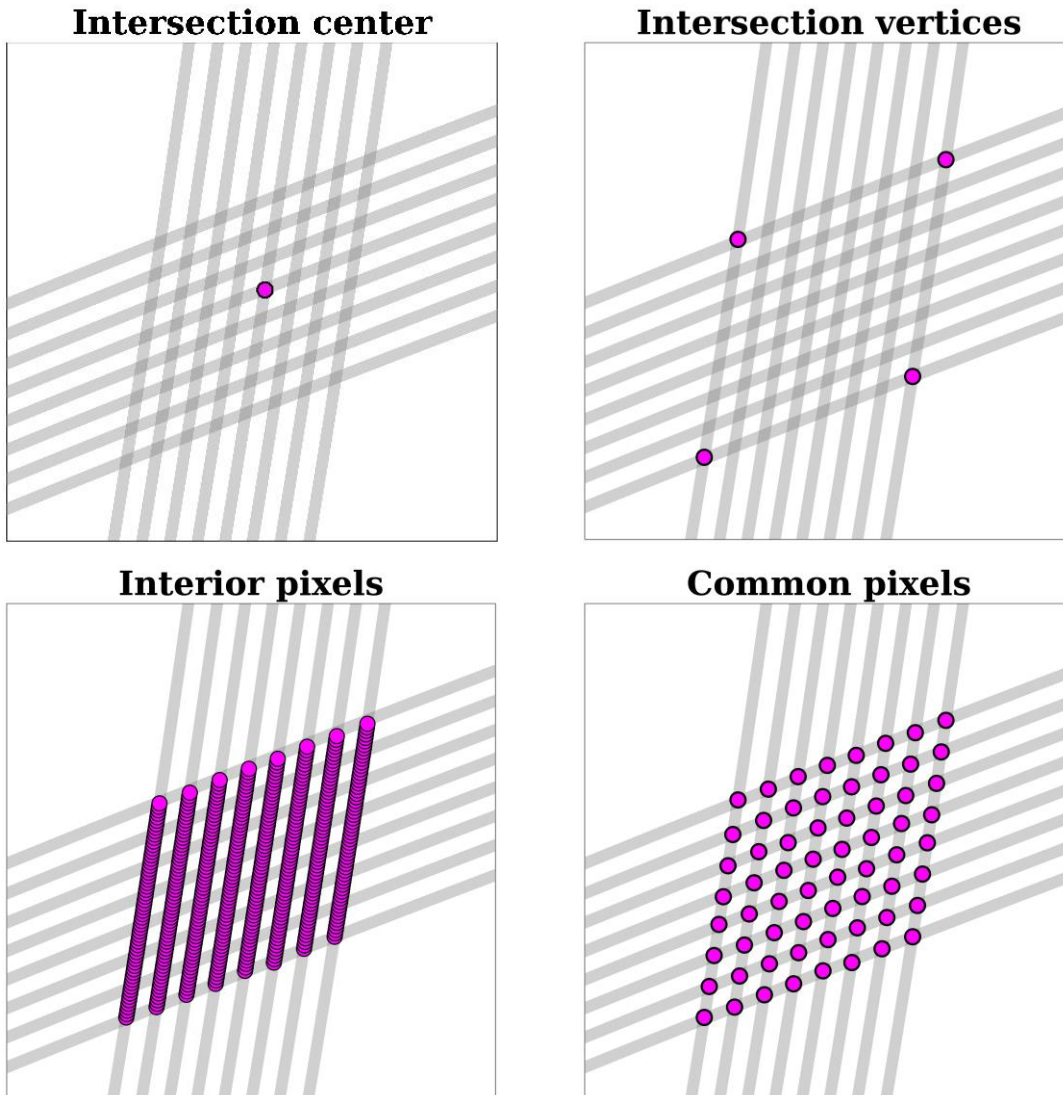


Figure 3: The anatomy of an intersection. Each panel represents one of the four categories of pixels that characterize the PREFIRE intersections.

2.2.1 Intersection vertices

Having established a process for approximating intersection centers, we continue with a method to extract all remaining co-located pixels from a resampling region. Our method employs the same simple distance calculations underpinning the fourth-scene method, beginning with pixels that form an intersection's four vertices. Note that each vertex represents a pair of exterior pixels that are approximately co-located, such that one vertex corresponds to an overlap between the first scenes of two granules, another corresponds to the overlap between their eighth scenes, a

third constitutes the overlap between the first scene of one granule and the eighth scene of the other, and vice versa for the remaining vertex. To identify vertices for any arbitrary intersection, we first retrieve the swath-relative position that corresponds to its center, taken with respect to both constituent granules. This is achieved by extracting the along-track information embedded in the intersection ID. Recall that the intersection center is, according to our definition, a pair of fourth scenes of two orbits that are approximately co-located. Here we select the along-track positions of each fourth scene in either orbit. That is, we identify the position of the across-track segment containing the intersection center by identifying where it occurs in the along-track direction, and we do so for each granule forming an intersection. In other words, if a granule consists of N across-track segments that follow sequentially in the along-track direction, the intersection center is located in an across-track segment ranging from the first to the N^{th} segment. This position is explicitly given by the intersection ID generated earlier and is taken to be the “target track” of an intersection, such that it will be targeted in the steps that follow.

After retrieving the target tracks for both intersecting orbits, we subselect the first and eighth scenes of each granule and subsequently reduce them into smaller segments centered about the target tracks. The lengths of the smaller segments are determined by orbit proximity, as it can be shown that orbits positioned close together near the equator tend to form elongated intersections and therefore require extended segments. Depending on proximity, the first and eighth scene segments range from ± 150 across-tracks to ± 600 across-tracks, again taken relative to the target tracks. Such a process results in four lists, one subselected latitude list and one subselected longitude list per granule. In similar fashion to the fourth-scene method, each list is subselected again to include only mutual latitudes and longitudes and then binned with respect to the latter. As before, distances between pixels in common bins are then calculated between granules. A distance

threshold of 15 km was used to eliminate pixels that obviously did not converge, such that all pairs located within 15 km of each other were flagged as possible vertices and remaining pairs discarded.

In the remaining steps, we determine four pairs of pixels that are separated by an arbitrarily small distance while ensuring they represent a unique vertex of the intersection region. To do so, flagged pairs are sorted in ascending order of distance, and we thereafter proceed in the following manner: For any intersection, the pair of pixels with the absolute minimum separation was always taken to be one of the vertices. The remaining three vertices were identified by iterating over each flagged pair in ascending order of distance and comparing their location to other confirmed vertices. This is a reasonable approach since vertices are located at separate regions of an intersection and should therefore never be near one another. It is worth clarifying that there are two distances being considered here: The distance between the flagged pixels of the orbit pair and the distance between a flagged pixel and a confirmed vertex. Possible vertices are identified using the former distance while true vertices are confirmed using the latter. If a flagged pixel is located near a confirmed vertex, then it is merely a pixel in the periphery of that vertex rather than an additional vertex. Testing confirmed that true vertices are located at least 25 km apart from any other vertex, and this threshold was used accordingly. In sum, vertices are identified by iterating over each flagged pixel and retaining the pixel closest to that of the other granule, provided that it does not land within 25 km of a confirmed vertex. Once four pairs were identified, one vertex for each orbit, the process concludes, and each pair is taken to represent the four vertices of a particular intersection region.

2.2.2 Interior pixels

With the vertices identified, the next objective is to extract the remaining co-located pixels. In theory, we can employ existing geometry packages to achieve this, provided that any

combination of vertices yields a “valid” polygon. In reality, however, available packages perform unreliably owing to the highly variable orientation of the intersections. It was found during this analysis that some intersections are not perceived as valid polygons and interior pixels were inaccurately extracted. A unique method to find interior pixels was therefore sought.

Our chosen method begins with a few observations. First and most obvious, each vertex corresponds to a specific orbit track. This observation lends itself to another: The perimeter of an intersection comprises segments of the first and eighth scenes that connect two vertices, such that each orbit contributes two bounding segments. It then naturally follows that all interior pixels are contained within the sequence of tracks bounded by the lower and upper tracks of the perimeter. Although this holds irrespective of the criteria used to define the pixels that are “interior,” the across-track discontinuities imply, rather critically, that a number of pixels that are technically within the perimeter are not truly co-located between orbits. Consequently, the pixels we take to be “interior” to and thus part of an intersection are subjective. In this thesis, we begin with an inclusive definition in which any pixel landing within the perimeter is accepted as an interior pixel. This concept guides the method discussed in this section, whereas a method to extract pixels that are strictly co-located is considered in the next section.

Obliging with the more inclusive definition, our principal strategy calls for extracting a single pixel per scene that approximately intersects a perimeter segment of the opposite granule. Since the method described in section 2.2.1 yields the segments of the first and eighth scenes that form the intersection boundary, our remaining task need only consider the remaining six scenes, namely scenes 2-7. As before, we can simply subselect each of the remaining six scenes, this time between the lower and upper track of the perimeter, and then repeat the process of calculating distances with respect to the first and eighth scenes of the opposite granule. In other words, we

iterate over the interior six scenes of one granule and find where each interior scene intersects both exterior scenes (i.e., the first and eighth scenes) of the other, again using the same simple distance calculations as before. The full intersection region is then represented by the segments of each scene, starting from the nearest pixel corresponding to scene 1 of the other granule and extending to the nearest pixel corresponding to scene 8, including the segments between vertices that were found previously.

2.2.3 “Common pixels”

Precision may require us to narrow the scope of our interior pixels to include only pixels that are strictly co-located. Hereafter, such pixels are referred to as the “common pixels” of an intersection. Common pixels help facilitate a more direct comparison between sampling since they eliminate complications arising from the discontinuities across each swath.

Our method to identify common pixels relies on the same general concept as before: Identify the pixels that roughly converge across satellites for each scene. Unlike our initial method, however, we no longer restrict this computation to only the perimeter scenes of one granule. Instead, we extract common pixels by extending the distance computations to *every* scene. This requires us to calculate, via the subselection processes and the simple distance calculations cited earlier, where each of the eight scenes of the first granule approximately intersect the eight scenes of the other, resulting in 64 pixels per orbit that share approximately the same coordinates. Note that some pairs were identified during preceding steps and this reduces the computational burden as such. For reference, Figure 4 summarizes our approach and specifies the scene-wise overlaps that define each intersection feature.

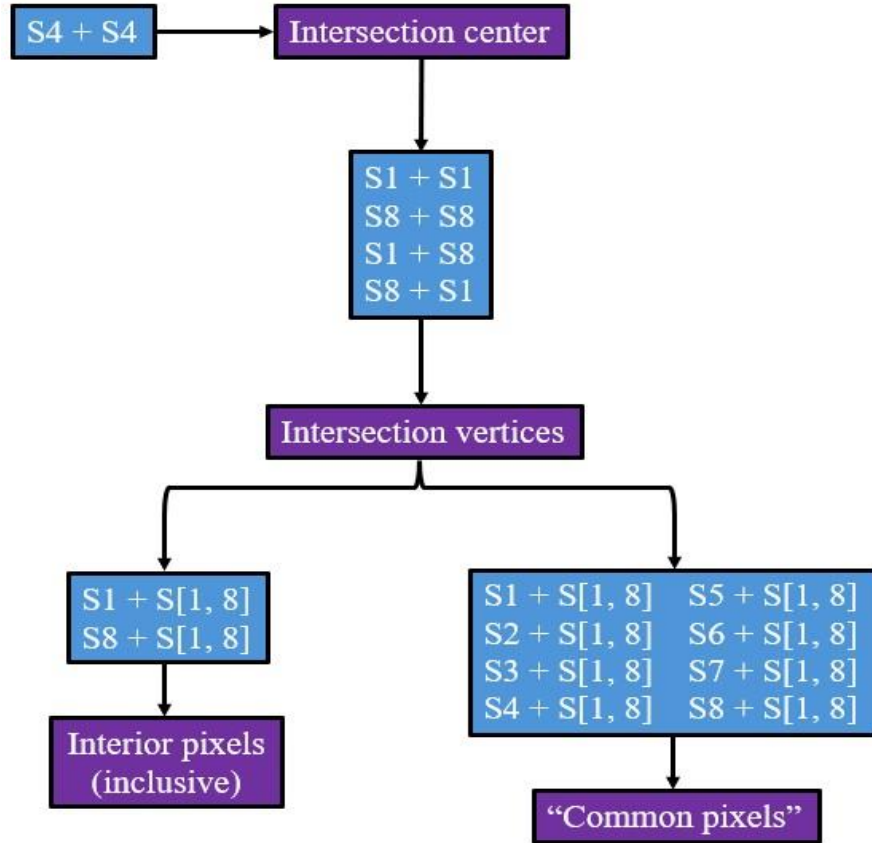


Figure 4: The sequential method to extract intersection pixels used in this paper. Purple boxes list intersection features, whereas blue boxes specify the corresponding scene (“S”) combination(s) used to derive them. In the context of scene combinations, it is implied that we are referring to their point of convergence. The format $SX + SY$ is used to denote overlaps between scene X of one granule and scene Y of the other. Alternatively, $SX + S[Y, Z]$ designates overlaps between scene X of one granule and all scenes ranging from Y to Z of the other. For either the interior or common pixels, some combinations are accounted for in preceding steps.

2.3 Methods for conceptual cases

This thesis sought to answer not only specific questions concerning where and when intersections will plausibly occur but also broader questions concerning their utility. To the latter end, we present three illustrative examples from the simulated orbits introduced earlier. Each case, detailed below, behaves as a conceptual exercise and is therefore intended to demonstrate broad, open-ended applications rather than rigid interpretations or techniques.

For each conceptual case, radiances were simulated by running state variables in the Principal Component-based Radiative Transfer Model (PCRTM), described in greater detail in Liu

et al. (2006). In summary, PCRTM is a high spectral-resolution forward model with coverage from about 3.6 to 200 μm and 101 fixed vertical pressure levels (Miller et al. 2023). The model requires vertical temperature and moisture variables as input, as well as surface properties like surface pressure and spectral emissivity. Atmospheric composition can also be explicitly adjusted in the model. Additionally, PCRTM allows for prescribed clouds and accepts a number of cloud inputs, including cloud phase (a binary choice between liquid or ice), cloud pressure level and cloud optical depth.

In the hypothetical cases analyzed in the sections that follow, temperature, moisture and pressure conditions were based on hourly data from the Goddard Earth Observing Systems for Instrument Teams (GEOS-IT), provided courtesy of the Global Modeling and Assimilation Office at NASA Goddard Space Flight Center. Scene conditions were approximated by interpolating GEOS-IT data to TIRS scenes and TIRS crossover times. Surface emissivity values were based on Huang et al. (2016), and constituent gases like CO_2 and CH_4 were assigned standard contemporary values. Clear sky radiances were simulated in the first and third case, as described in section 2.3.1 and 2.3.3, respectively. Conversely, the cloud profile varied in the second case and is detailed in section 2.3.2.

2.3.1 Arctic heat wave case methods

Our first example features a coastal SAT2 self-intersection centered near 71.01°N , 164.42°E in the East Siberian Sea. Figure 5 depicts the constituent granules. Note that this resampled region permits us to consider multiple polar surface types concurrently, a scenario that may present frequently during operation. Temporally, the elapsed time between crossovers is approximately 33 hours, a duration sufficiently long to capture notable scene changes on daily timescales.




Figure 5: The constituent granules forming the intersection featured in the first case. Purple and orange denote the swaths of the first and second crossover, respectively. The intersection region is demarcated in gray.

Archived melt analyses from the National Snow and Ice Data Center guided us to a period of relatively rapid Arctic warming and sea ice melt from June 29 to July 4, 2021. Hourly, single-level European Centre for Medium-Range Weather Forecasts Re-Analysis 5 (ERA5; Hersbach et al. 2020) data confirmed regional surface warming and sea ice attrition across the Arctic Basin from late June to early July 2021. In particular, sharp changes in ERA5 skin temperature and total column water vapor (TCWV) were noted in Eastern Siberia. We selected the intersection described above after confirming that it adequately captured the temperature and moisture changes impacting the broader region. This confirmation was achieved by retrieving hourly GEOS-IT data and then interpolating the data to each scene of either granule. In addition to spatial interpolation, GEOS-

IT data were also interpolated with respect to time in order to represent sub-hourly conditions at the approximate crossover times. Figure 6 shows the interpolated GEOS-IT skin temperature and TCWV for the first and second orbit of the selected intersection.

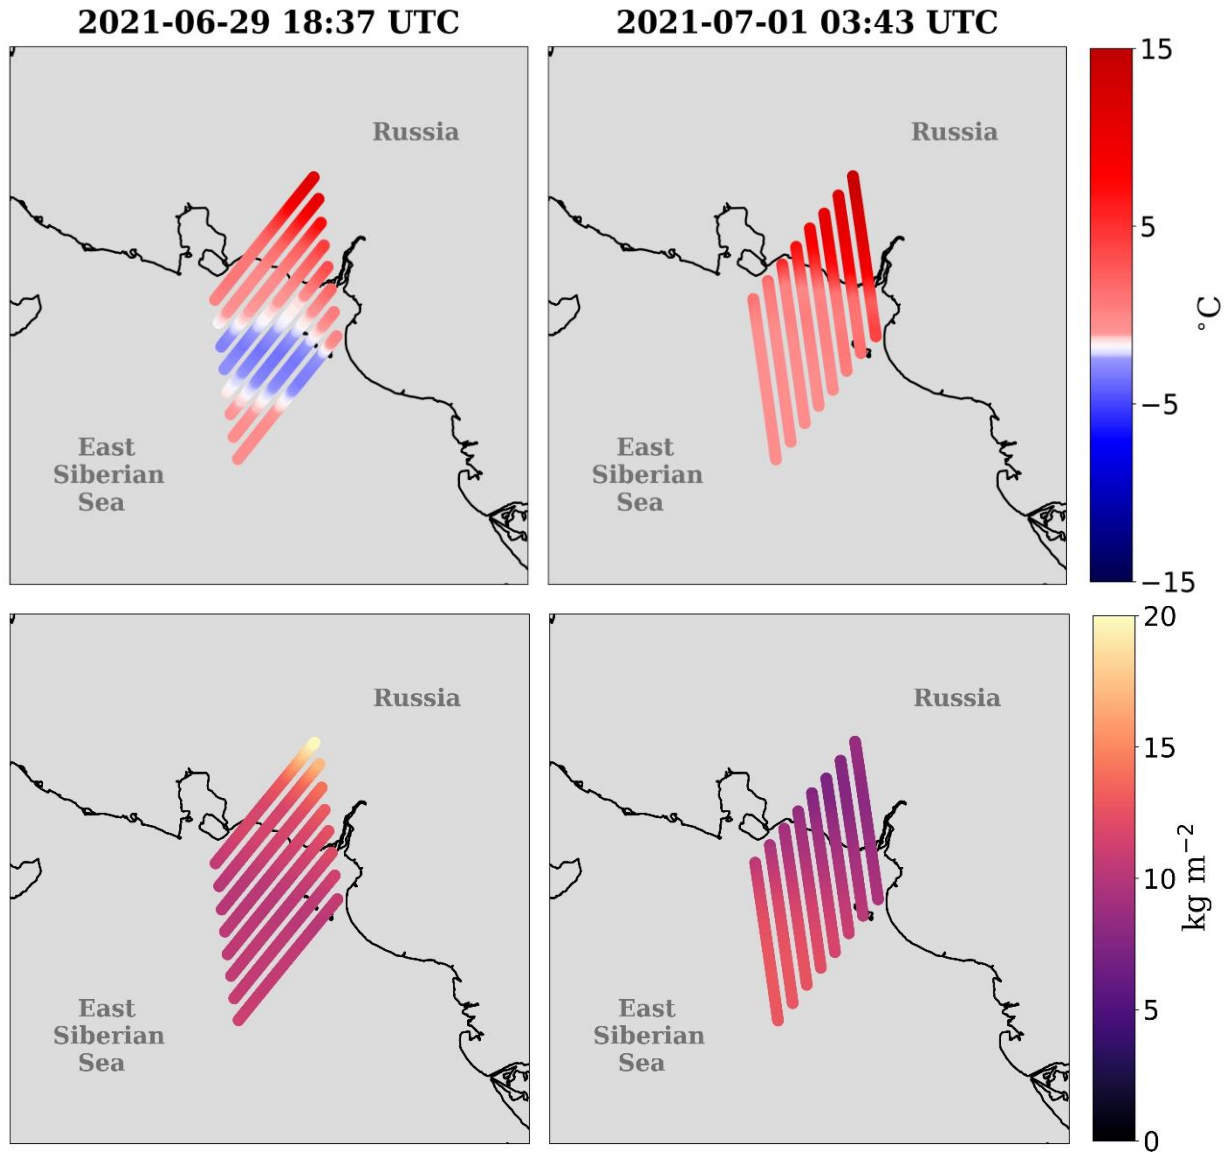


Figure 6: GEOS-IT skin temperatures in degrees Celsius (top) and TCWV in kg m⁻² (bottom) interpolated to each scene of either granule during the first (left) and second (right) crossover for the first case.

For our analysis, we sought a straightforward comparison between crossovers and as such neglected pixels that were not strictly co-located. All remaining pixels, the so-called “common

pixels” discussed in section 2.2.3, are shown with interpolated skin temperature and TCWV in Figure 7. These common pixels are the subject of the analysis unless otherwise noted.

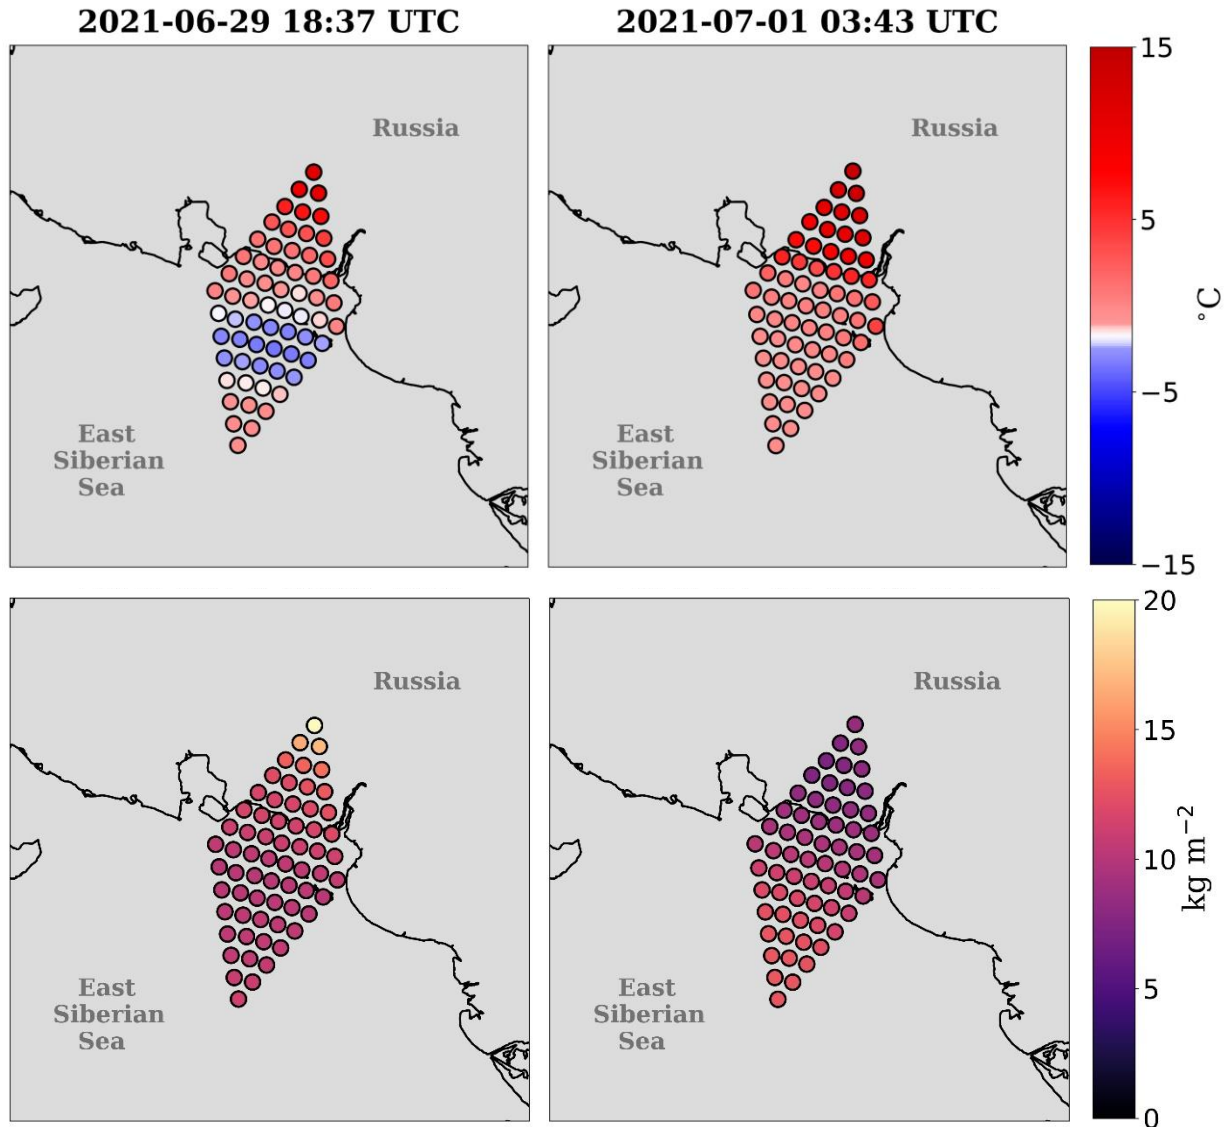


Figure 7: The same as Figure 6 but in the context of the common pixels for each granule.

A quick review of the intersection area confirms that much of the region extends offshore, with a smaller fraction overlaying coastal land. Among the offshore sector, twenty pixels transitioned from below to above freezing—taken with respect to sea water, -1.8°C —between the first and second crossover, respectively. The rest of the offshore pixels remained above freezing during either crossover time. Although reanalysis confirmed local sea ice melt, the degree of

attrition was variable across the resampled region. To facilitate a more robust example, we imposed a binary sea ice mask using GEOS-IT skin temperature as a proxy for surface type. That is, offshore pixels were taken to be 100% sea ice if they exhibited a skin temperature less than the freezing point of sea water (-1.8°C). Conversely, offshore pixels above freezing were treated as 100% ocean, while the remaining onshore pixels were regarded as uniform land. The three surface types are shown in Figure 8. Note that our treatment of offshore pixels introduces an element of idealization to the case; however, our approach lends insight into OLR sensitivities with respect to sea ice melt, an end to which our idealization is a worthwhile means. Additionally, our fidelity to the temperature and moisture profiles helps ensure a reasonably realistic case.

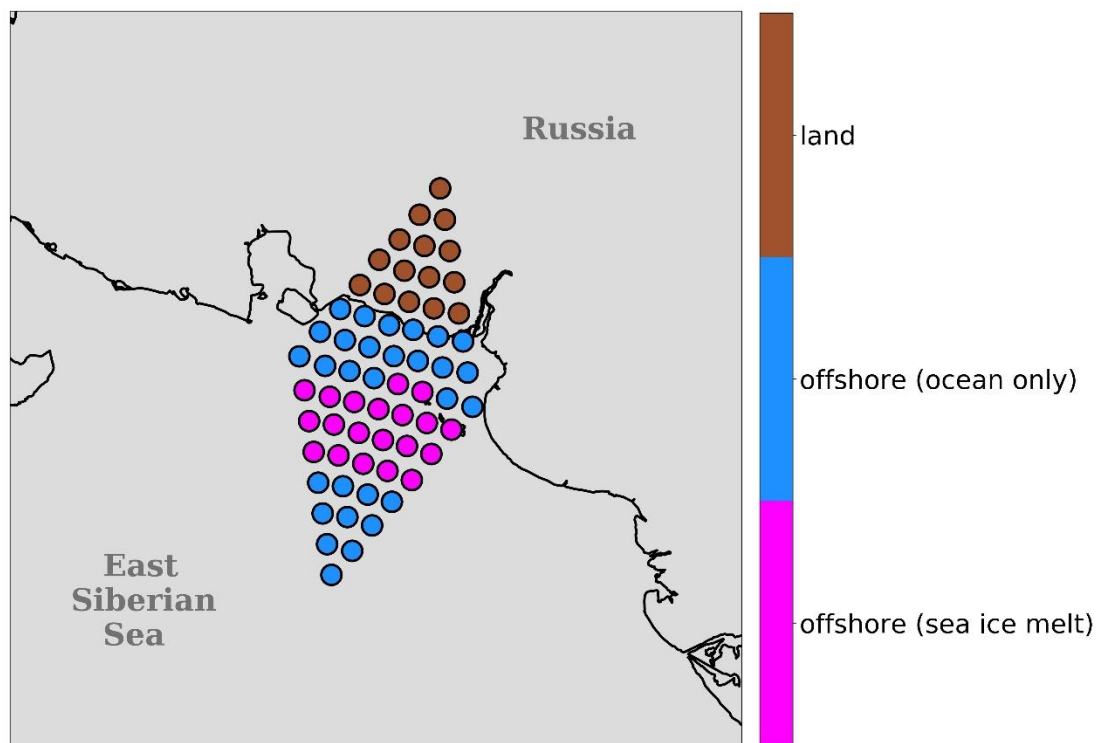


Figure 8: The three surface types considered in the first case. For simplicity, land and sea ice are taken to be snow-free.

Surface emissivity values prescribed for each surface type are shown in Figure 9 and were based on Huang et al. (2016). For simplicity, sea ice was assumed to be snow-free, such that its

emissivities were based on values for ice only. Land was also taken to be snow-free and granted emissivity values consistent for grass (Huang et al 2016).




Figure 9: Spectral emissivities for each of the three surface types considered. Values were based on Huang et al. 2016.

Figure 10 shows the difference in interpolated GEOS-IT skin temperature and TCWV between crossovers in the context of the common pixels. We note variable warming across the diamond-shaped area, with some onshore pixels warming by nearly 8°C. It also strikes us that shifts in moisture are highly variable across the resampled region, with TCWV decreasing by more than 11 kg m⁻² for some onshore pixels while many offshore pixels exhibited subtle increases in TCWV.

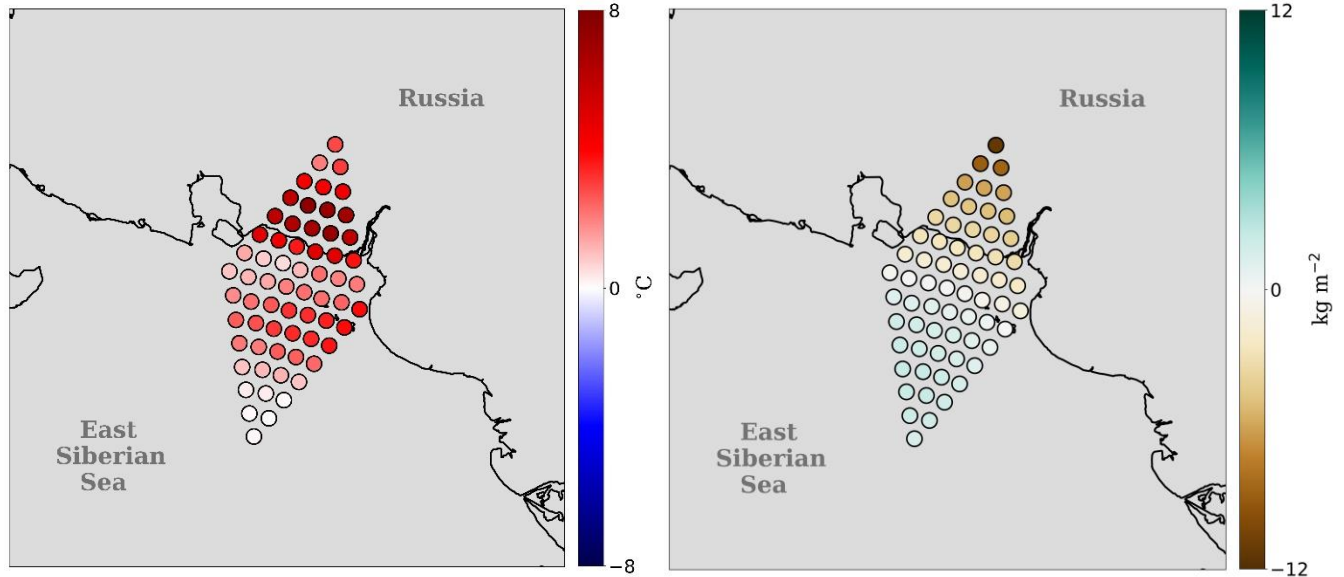


Figure 10: Changes in skin temperature in degrees Celsius (left) and TCWV in kg m⁻² (right) captured by the hypothetical intersection in the first case.

The noted temperature and moisture changes, combined with surface type changes where applicable, were interpreted in the context of OLR by simulating radiances for each TIRS2 channel. Although each TIRS has some 60 nominal channels, this analysis focuses on six channels centered on 6.73, 10.95, 11.8, 19.39, 26.99 and 39.64 μm , where 6.73 μm was chosen to represent a MIR channel with strong water vapor sensitivity; 10.95 μm and 11.88 μm were selected to sample two important MIR channels in the atmospheric window; 19.39 μm was intended to represent a “dirty” window channel in the FIR, such that it exhibits some sensitivity to water vapor; 26.99 μm was chosen to represent a FIR channel with greater sensitivity to water vapor than the preceding FIR channel; and 39.64 μm was included to represent a longer-wave FIR channel deep in the water vapor continuum.

By sampling this variety of wavelengths, we expect to gain insight into the vertical structure of the resulting spectral changes captured by the intersection, if any. This follows from the differences in weighting functions, where weighting functions generally peak higher in the vertical with increasing sensitivity to water vapor. Put another way, channels that are less sensitive

to water vapor (i.e., window channels) are commonly said to be transparent, a reference to the fact that their weighting functions peak near the surface. As a result, transparent channels can reveal near-surface conditions. In contrast, channels sensitive to water vapor are commonly described as opaque, such that their weighting functions tend to peak in the vicinity of water vapor above the surface. The degree of opacity is somewhat relative, where a channel that may be fully opaque in the moist tropics may be more transparent in dry regions, including the poles.

Readers should bear in mind that the spectral response functions are wavelength- and scene-dependent. As a reminder, clouds were removed before simulating radiances in this case to ensure the inclusion of surface processes in the final results. The resulting clear sky radiances are discussed in section 4.1.

2.3.2 Cloud dissipation case methods

In our second example, we simulate a transition from cloudy to clear conditions in a frigid polar environment. By doing so, we can assess the impact of high-latitude cloud dissipation on spectral OLR. The intersection featured in this case is a SAT1 self-intersection centered near 80.7°S, 73.15°E in interior East Antarctica. Both granules this intersection comprises can be seen via Figure 11. Although Antarctic clouds can persist over extended timescales (e.g., Silber et al. 2019), we opted for an intersection with a moderate time difference between crossovers. This follows from our efforts to assess cloud forcing, an endeavor met with greater difficulty if conditions change substantially between orbits. We therefore selected an intersection with a duration of about 4.5 hours, a timescale on which cloud cover changes could realistically be observed sans dramatic changes in the underlying conditions.




Figure 11: The constituent granules forming the intersection featured in the second case. Purple and orange again denote the swaths of the first and second crossover, respectively, with the intersection region demarcated in gray.

As in Case 1, we opt to preserve realistic scene conditions by interpolating GEOS-IT data between hourly timesteps and then interpolating to the scenes of either granule. Figure 12 shows the resulting skin temperature and TCWV profiles observed during each crossover. We note the extremely low skin temperature consistent with our expectations for this location during austral winter, with values ranging from roughly -70°C to -56°C during either crossover. Our expectations of a dry wintertime environment are also met, evident by the low TCWV values ranging from 0.16 to 0.27 kg m^{-2} .

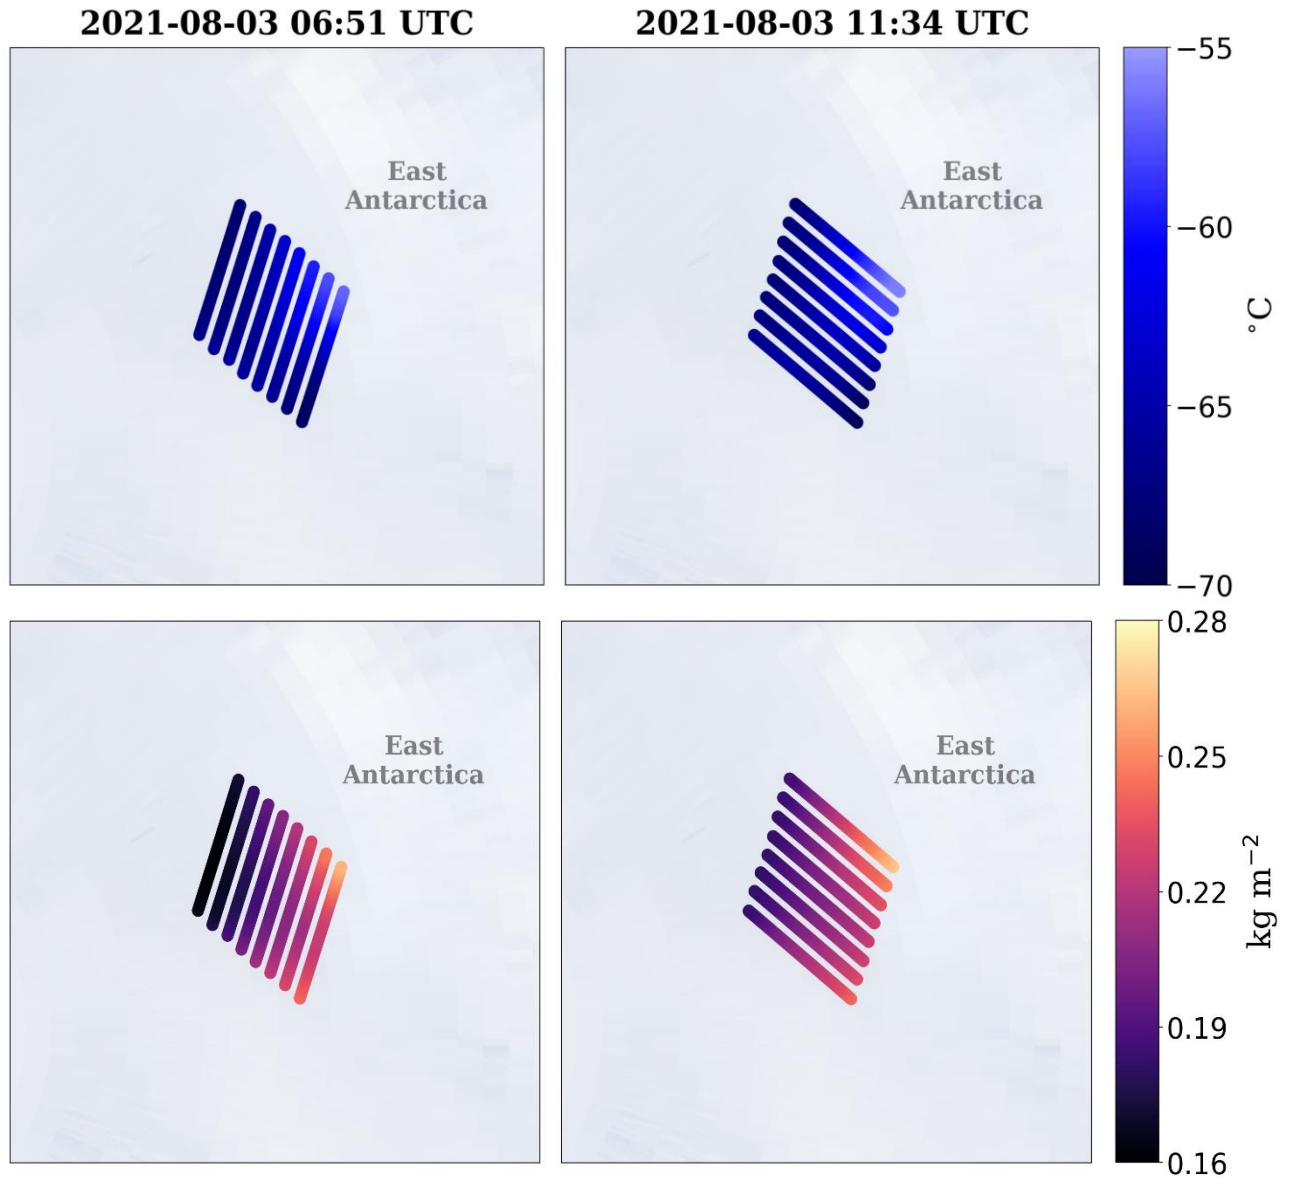


Figure 12: GEOS-IT skin temperatures in degrees Celsius (top) and TCWV in kg m⁻² (bottom) interpolated to each scene of either granule during the first (left) and second (right) crossover of the second case.

To simplify our analysis, we removed all native clouds from the first granule and added a uniform layer of ice clouds near 400 hPa with a prescribed optical depth of 1. We then removed all clouds from the second crossover to simulate a transition from cloudy to clear conditions. To target the impact of cloud optical thickness on OLR, we repeated this process for optical depths 2

and 4. Finally, for each iteration, we generated simulated radiances for each TIRS1 channel. Our results are detailed in section 4.2.

2.3.3 Calibration case methods

Our third case addresses inter-calibration. In the event of instrument drift, we require an in-orbit technique to identify impacted measurements. The chosen technique should reliably flag spectral differences between CubeSats that result from drift rather than changes in underlying conditions. Two unique factors must guide our approach: Foremost, recall that the channels for TIRS1 are centered on different wavelengths than TIRS2. Additionally, the spectral response functions are both instrument- and scene-dependent. It therefore follows that, provided identical conditions, each CubeSat will register different spectral profiles even in the absence of instrument drift. As such, our technique must anticipate a baseline offset between CubeSats.

An intuitive approach to inter-calibration is one that employs static conditions to achieve two ends: First, we want to establish the baseline offset between the eight scenes of TIRS1 and the eight scenes of TIRS2. In other words, we want to benchmark the differences in channel radiances per satellite per scene while holding conditions constant. Second, once this is established, we can verify whether we identify *differences in the differences* over time while again controlling for scene conditions. Differences in the differences imply drift and can be flagged accordingly.

If we assume conditions are more likely to remain static as our time difference between sampling decreases, then it follows that our shortest intersections are ideal candidates for inter-calibration. Thus, our third case selects a SAT1-SAT2 intersection with a relatively rapid time difference of less than three minutes. As shown in Figure 13, this intersection is centered near 74.88°S, 107.35°E in East Antarctica. We caution that the lower bound on resampling time may be greater during operation than that which we observed in this thesis. Nonetheless, this case

demonstrates a general procedure that will guide calibration, and it holds regardless of the degree of fidelity between the simulated orbits considered here and those we achieve in operation.




Figure 13: The constituent granules forming the intersection featured in the third case. Purple and orange again denote the swaths of the first and second crossover, respectively, with the intersection region demarcated in gray.

In addition to a scene that is sufficiently static with time, calibration also favors a scene that is approximately homogeneous with respect to underlying conditions. Figure 14 shows GEOS-IT skin temperature interpolated to each scene (top, in the context of the common pixels), and it indicates a relatively small range of skin temperatures across the resampling region, varying from -74°C to -68°C over the intersection pixels. TCWV, shown in the same figure (bottom), also exhibits very little variation across the resampling region.

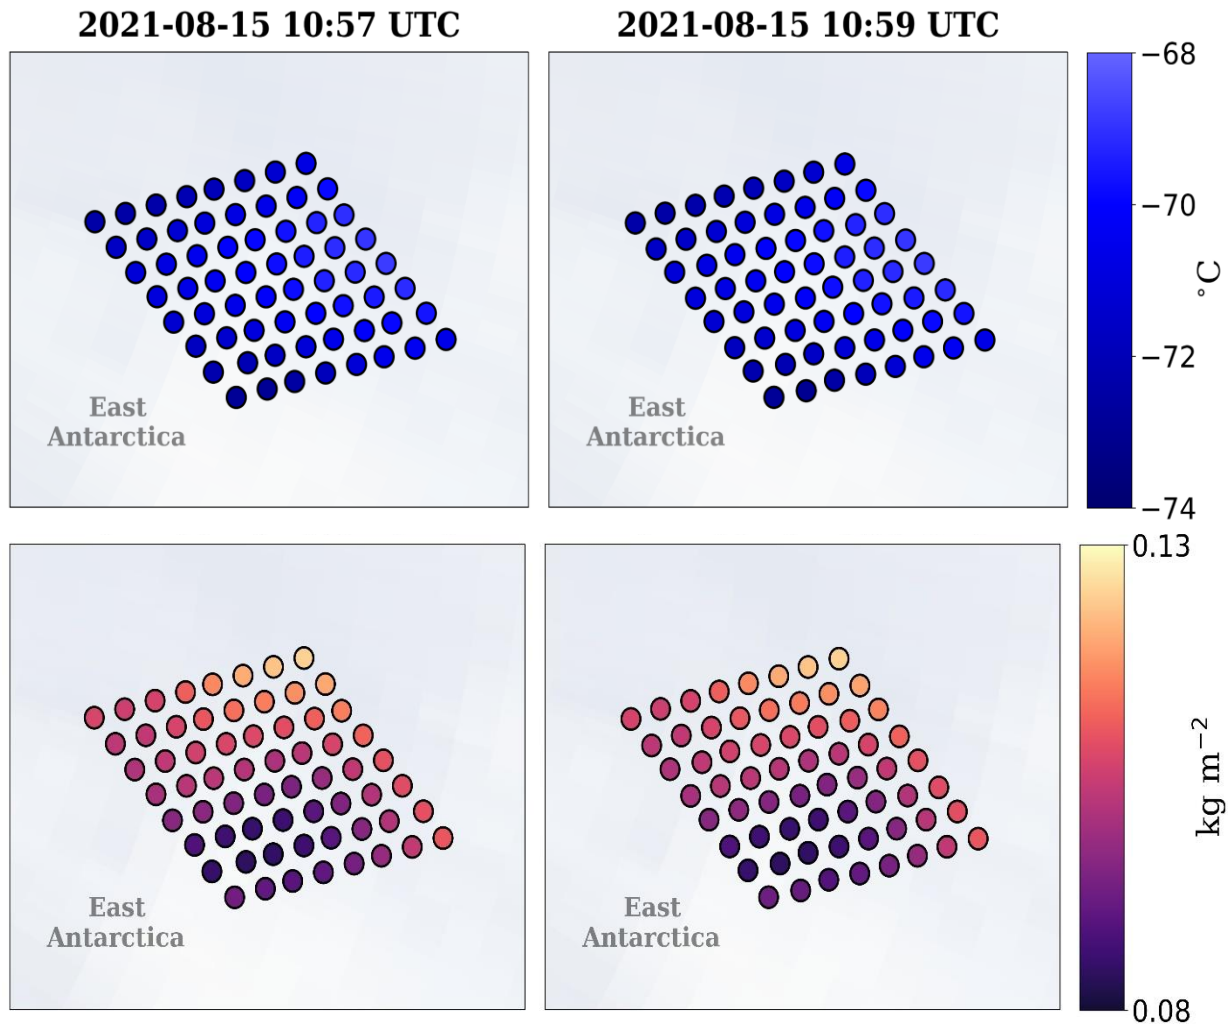


Figure 14: GEOS-IT skin temperatures in degrees Celsius (top) and TCWV in kg m^{-2} (bottom) interpolated to each scene of either granule during the first (left) and second (right) crossover of the third case. Here we show only the common pixels.

By design, the selected intersection should exhibit negligible scene differences between the first and second crossover. We confirm this by plotting the skin temperature and TCWV for every common pixel per granule. This is respectively shown in the top and bottom of Figure 15. Strong consistency between the pixels of the first and second crossover for either variable confirms conditions are approximately static between observations.

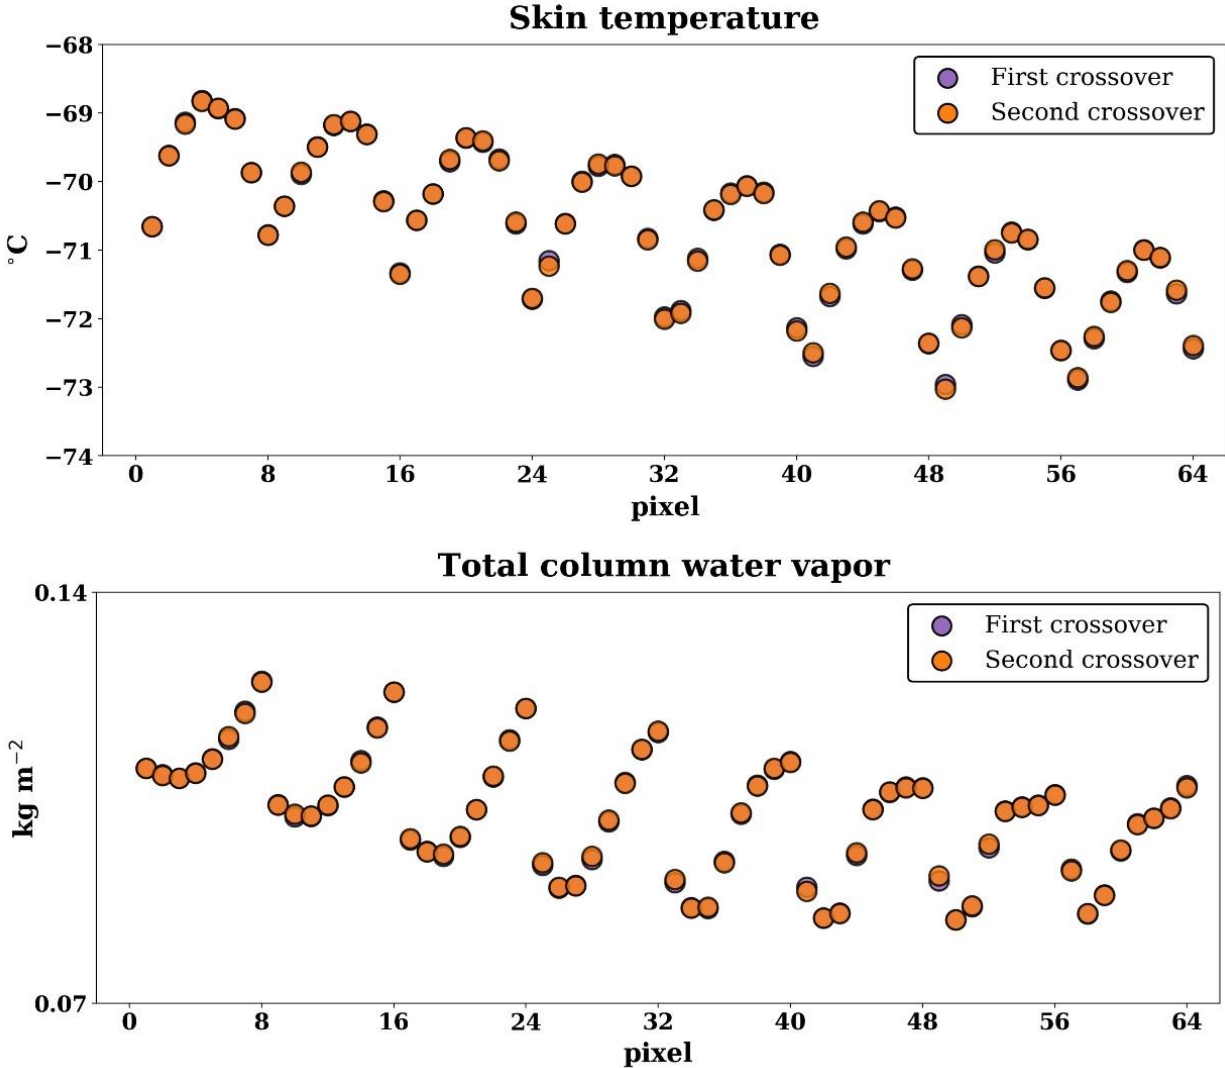


Figure 15: A comparison of skin temperature (top) and TCWV (bottom) for the common pixels during the first and second crossover for the granules constituting the third case. Note that values for the first crossover are hardly visible owing to the strong agreement between sampling.

When we recall that each scene has a unique spectral response function, it becomes clear that our calibration technique must consider them individually. This condition is met in our proposed method, which is demonstrated using the simulated clear sky radiances for both granules of the intersection noted above. Fundamentally, our goal is to achieve a comparison between the scenes of TIRS1 and the scenes of TIRS2, a comparison by which we can characterize their baseline spectral differences. We demonstrate this conceptually by first computing spectral radiances (I_λ) for every 1 μm between 5 and 55 μm . Note that channels are not centered at discrete

integers and we thus select the channels closest to every integer wavelength. For each scene of TIRS1, we then calculate ΔI_λ with respect to each of the eight scenes of TIRS2. It should be obvious that we are comparing radiances for channels that are not identically equal; however, by comparing channels that are closest to an objective threshold—in this case, every 1 μm —we are facilitating a non-arbitrary comparison between instruments. Finally, after generating baseline spectral differences per scene, we provide a simple conceptual example of drift by arbitrarily increasing I_λ by 10% for a single TIRS2 scene and recalculating the spectral differences. Our case concludes by computing the differences in the differences after drift is applied. Our results are discussed in section 4.3.

3. Spatial and temporal analysis of PREFIRE intersections

After identifying all intersections within a rolling 48-hour window, as described in section 2.1.2, we then categorized them by latitude and Δt . From the outset, we considered the three intersection types separately and classified their respective intersections by time using thresholds of ± 48 hours, ± 24 hours and ± 12 hours. For completeness, we will discuss numerical results for each of the three time classes over the simulated period; however, we anticipate shorter-term resampling will distinguish PREFIRE from other missions and we therefore give priority to intersections in the shortest time class ($|\Delta t| < 12$ hours) for much of this paper. The discussion that follows begins by enumerating our hypothetical intersections per type in section 3.1. We then describe the spatial distribution of our simulated intersections in section 3.2, followed by a discussion of their temporal characteristics in section 3.3. Finally, in section 3.4, we shift our focus to a brief summary of the short-term revisits (i.e., Δt is arbitrarily small) observed in the simulated orbits, an important feature for inter-calibration.

3.1 Numerical distribution of simulated intersections

We start by considering the number of intersections over the two-month simulation for each of the three time classes. For intersections ± 48 hours, our analysis identified more than 200,000 combined intersections over this period, with SAT1 self-intersections, SAT2 self-intersections and SAT1-SAT2 intersections contributing 53,402, 55,398 and 109,790, respectively. Not surprisingly, we see that for intersections ± 24 hours, the number of SAT1 self-intersections, SAT2 self-intersections and SAT1-SAT2 intersections is roughly halved at 26,476, 26,566 and 54,960 intersections, respectively. This pattern holds for intersections ± 12 hours, where we again see an approximate halving of the number of intersections per type compared to the preceding time bin. That is, for resampling less than 12 hours, we identified 14,072 SAT1 self-intersections, 14,120 SAT2 self-intersections and nearly double the number of SAT1-SAT2 intersections at 27,538 over the two-month period. Our results are synthesized in Figure 16. As a reminder, our discussion will focus primarily on intersections ± 12 hours going forward.

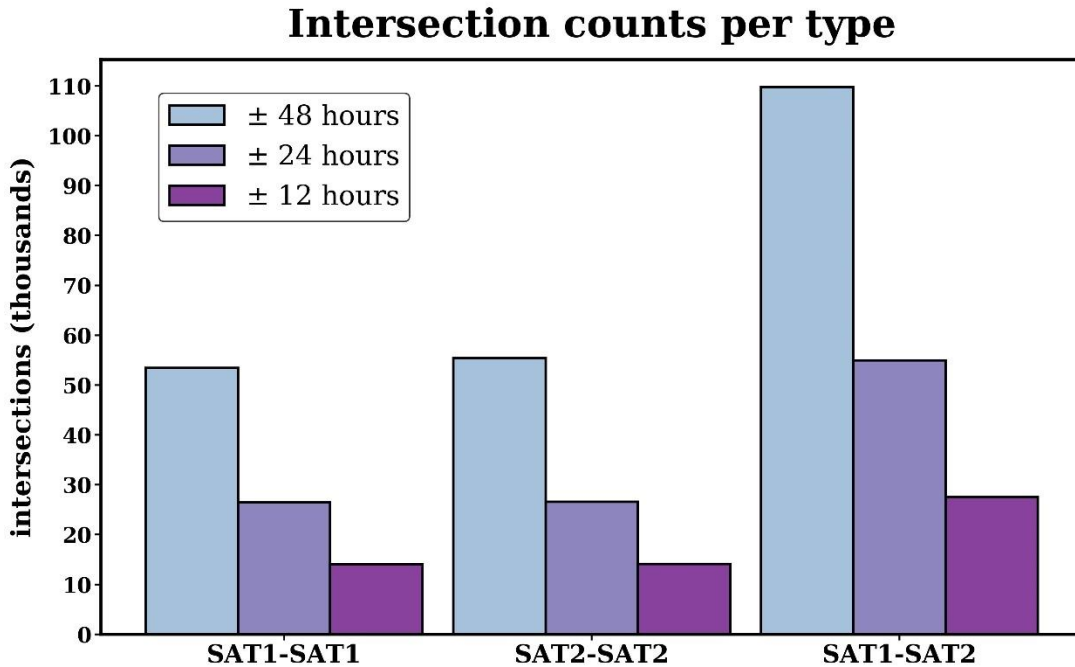


Figure 16: Histogram showing the number of intersections over the simulated two-month period for each intersection type and for three temporal categories. Counts may differ during the mission.

Taken together, we can make several notes about the overall distribution. First, we see there is very little difference in the number of self-intersections for one CubeSats over the other, particularly when we only consider shorter-term intersections. As a percentage of all intersections ± 12 hours, self-resampling for either CubeSat accounts for approximately one-quarter (25.3%) of the two-month total. Second, the number of SAT1-SAT2 intersections is virtually double the number of self-intersections for a single CubeSat, constituting roughly half (49.4%) of all intersections ± 12 hours. Third, this distribution implies that we would lose almost three-quarters (74.7%) of intersections if we had one CubeSat rather than two. Figure 17 visualizes the percentages by type as well as the intersections we would forfeit if we lost one CubeSat.

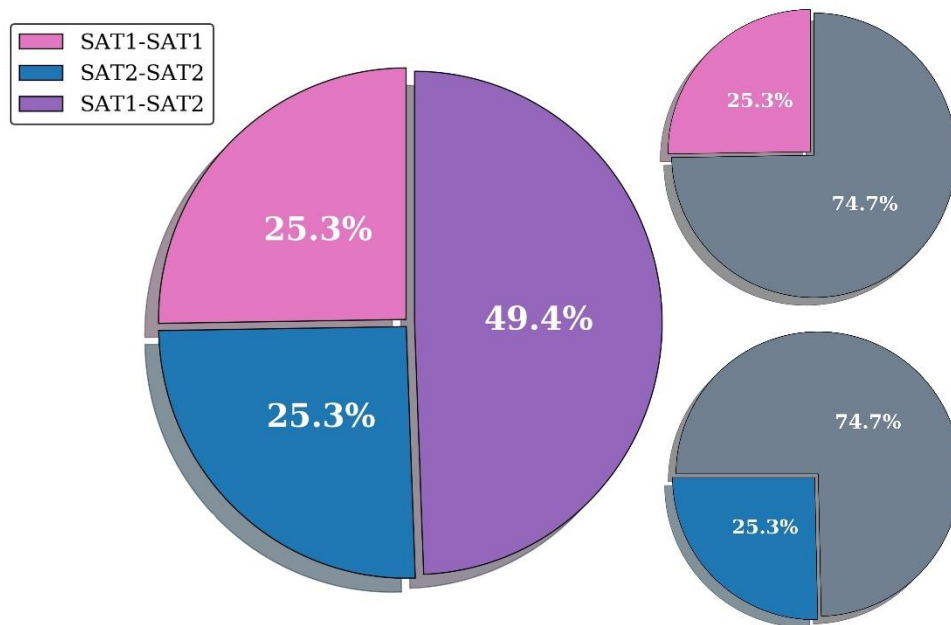


Figure 17: Relative percentages of ± 12 -hour intersections per intersection type (left). The right-most charts represent a hypothetical loss of CubeSat 1 (top right) or CubeSat 2 (bottom right), where shaded regions signify the percent of lost intersections compared to the original scenario. Percentages may differ during the mission.

Since the two-month period is arbitrary, we sought to determine a plausible number of intersections *per day* given the orbit parameters prescribed. Shifting to this shorter timescale, we calculated the mean, minimum and maximum number of daily intersections as well as the standard

deviation per day. Our results are synthesized via Table 4. If we consider self-intersections together, a reasonable approach given their similar numerical distributions, we see that we would expect just over 240 global intersections daily on average for either self-intersection type, with a daily minimum and maximum of 240 and 247, respectively. SAT1 self-intersections had a standard deviation of 2.1 global intersections per day, slightly less than SAT2 self-intersections, which had a standard deviation of 2.5 global intersections per day.

| | Mean daily | Min daily | Max daily | Std. dev. |
|------------------|-------------------|------------------|------------------|------------------|
| SAT1-SAT1 | 242 (181) | 240 (180) | 247 (186) | 2.1 (2.1) |
| SAT2-SAT2 | 243 (182) | 240 (180) | 247 (186) | 2.5 (2.6) |
| SAT1-SAT2 | 474 (362) | 451 (330) | 523 (403) | 21.4 (15.9) |

Table 4: Mean, minimum and maximum daily intersections as well as the standard deviation for each intersection type. Values indicate global (polar) intersections, where 66.5° is used as the polar cutoff for either hemisphere. The number of daily intersections may differ during the mission.

Consistent with our earlier findings, we see that the daily number of SAT1-SAT2 intersections again diverges from the former two types. For resampling between satellites, we average over 470 global intersections per day, nearly double the number of global self-intersections for a single CubeSat. Moreover, the minimum and maximum number of global intersections per day stands at 451 and 523, respectively, with a standard deviation of 21.4 globally. This implies that the number of daily intersections between CubeSats exhibits much greater variability than self-intersections, with the latter showing very little variability overall insofar as the numerical distribution.

In addition to the number of global intersections per day, we also sought to isolate daily statistics for high-latitude resampling. For this exercise, we consider any intersection centered at or poleward of $\pm 66.5^\circ$ to be a “polar intersection.” Our results indicate that, on average, we would

expect just over 180 polar intersections per day for either self-intersection type, with a minimum and maximum of 180 and 186, respectively. SAT1 self-intersections had a standard deviation of 2.1 polar intersections per day, unchanged from its global standard deviation, whereas SAT2 self-intersections also had nearly the same standard deviation compared to their global counterparts with 2.6 polar intersections per day. Conversely, we found an average of just over 360 polar intersections per day for SAT1-SAT2 intersections, ranging from 330 polar intersections to 403 polar intersections per day. Just as we encountered for the global resampling, SAT1-SAT2 polar intersections exhibit greater variability compared to self-intersections, with a standard deviation of 15.9 polar intersections per day.

3.2 Spatial distribution of simulated intersections

The preceding section concluded by isolating high-latitude intersections from the global collective, thereby giving us cursory insight into the hypothetical spatial profile of our intersections. Although such an exercise is broadly useful, this thesis sought to approximate the locations of our resampling regions with greater precision, an objective for which the following two steps were taken: First, intersections were classified as either *Northern Hemisphere* or *Southern Hemisphere*, depending on the location of their respective centers. Second, for each hemisphere, intersection centers were binned every 5° latitude, starting from the equator and extending poleward to $\pm 85^\circ$. The last theoretical bin containing latitudes higher than $\pm 85^\circ$ is not applicable since this range exceeds the maximum sampling latitude. To facilitate comparison, we grouped bins into low-, mid-, and high-latitude sectors, where the low latitudes are taken to range from 0° to 30° (-0° to -30°) in the Northern (Southern) Hemisphere; mid-latitudes extend from 30° to 60° (-30° to -60°) in the Northern (Southern) Hemisphere; and alas, the remaining latitudes represent the high latitudes relevant to PREFIRE.

Our results inform us that intersection coverage varies sharply depending on whether we consider self-intersections or intersections between CubeSats. Starting with SAT1 self-intersections, the focus of Figure 18, we find no meaningful difference in the hypothetical distribution between hemispheres, with either yielding practically identical percentages of intersections per sector. In both hemispheres, about one-eighth (12.5%) of intersections were documented in the low latitudes, while another one-eighth (12.5%) were centered in the mid-latitudes. Thus, we encounter the majority of intersections in the remaining sector, totaling approximately three-quarters (roughly 75%) overall. In other words, most SAT1 self-intersections were identified in the high latitudes, up to our maximum sampling latitude near $\pm 82^\circ$, with only about one-quarter of intersections located in low- and mid-latitude sectors combined.

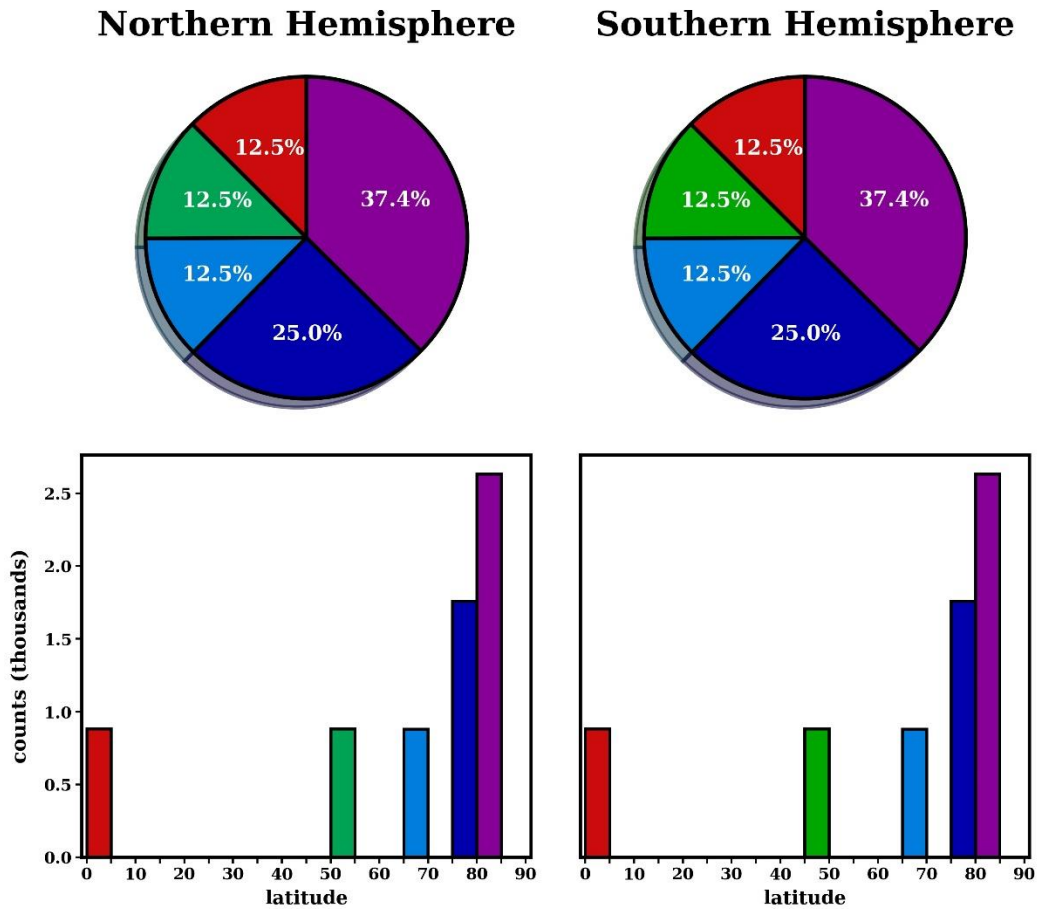


Figure 18: Latitudinal distribution of SAT1 self-intersections based on the simulated orbits analyzed in this thesis. For completeness, the Northern and Southern Hemisphere are considered

separately. Colors signify regional sectors, where warm colors denote low-latitude bins; greens denote mid-latitude bins; blues denote high-latitude bins, excluding the bin containing the maximum sampling latitude; and purple denotes the polar-most bin containing the maximum sampling latitude for either hemisphere.

To the critical observer, it might be obvious that the distribution of SAT1 self-intersections is not *strictly* identical between hemispheres. Mid-latitude intersections in particular appear to be slightly offset when we compare one hemisphere to the other, with Southern Hemisphere intersections in the mid-latitudes located equatorward relative to mid-latitude intersections in the Northern Hemisphere. Further inspection reveals, however, that this distinction lacks a meaningful difference; in reality, intersection centers are nominally located within the same latitude bands in either hemisphere, with any offset only marginal and resulting from the fact that some intersections are centered near the boundaries of latitude bins. In this case, some SAT1 self-intersections tend to occur near $\pm 45^\circ$, a latitude dividing two adjacent bins. Going further, the fourth-scene method yields approximate, rather than true, centers of each intersection. This means that the difference between documented and true centers will inevitably produce some degree of offset when intersections are viewed in the Northern versus Southern Hemisphere. We can therefore suggest with confidence that the spatial distributions between hemispheres is approximately identical with respect to latitude.

Recall that we investigated each self-intersection type independently, having assumed that different prescribed altitudes could give rise to unique self-resampling for each CubeSat. Yet the results for SAT2 self-intersections, shown in Figure 19, reveal that any difference in self-intersection coverage is unremarkable between satellites. Consistent with our earlier results, simulated SAT2 self-intersections numbered approximately one-eighth (12.5%) in the low latitudes; another one-eighth (12.5%) were centered in the mid-latitudes; and as before, the remaining three-quarters (about 75%) are located in high latitudes, up to the maximum sampling

latitude near $\pm 82^\circ$. These percentages are approximately static across hemispheres, just as they were for SAT1 self-intersections, with the same caveat cited earlier—any discrepancy between hemispheres is a matter of exactitude rather than an indicator of a true statistical difference.

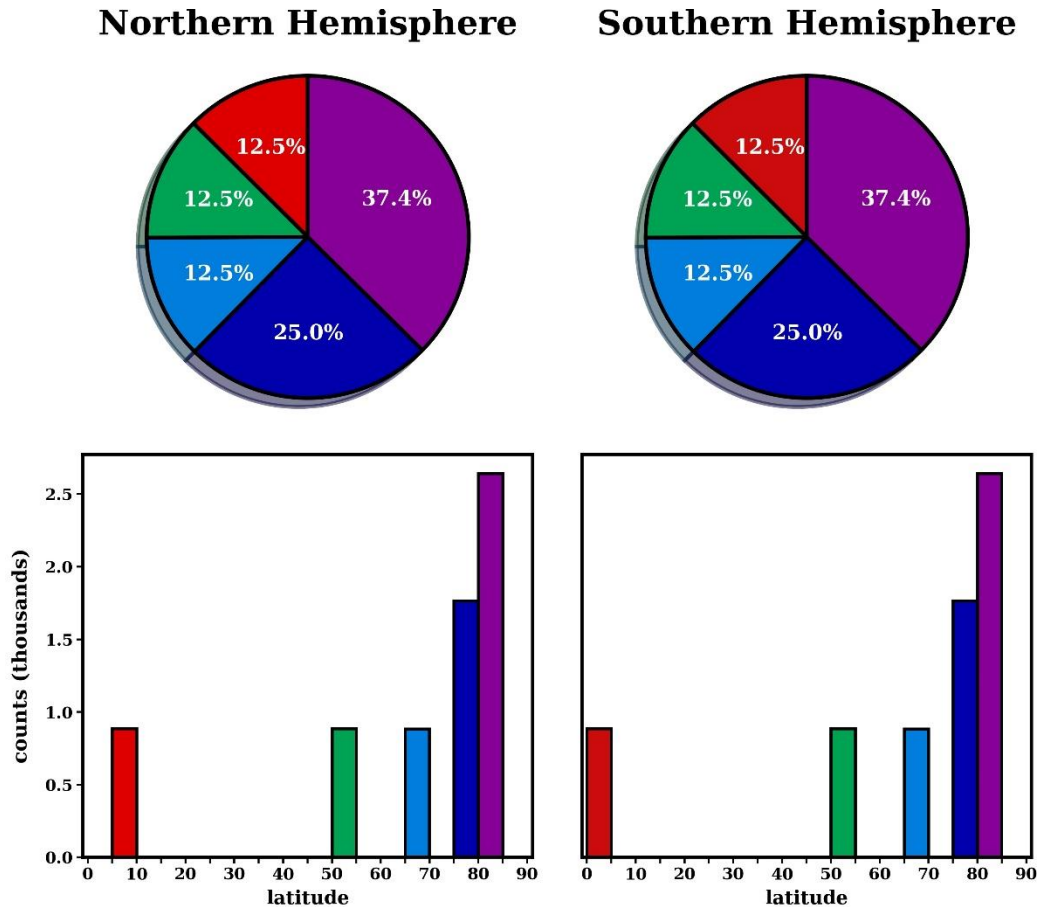


Figure 19: The same as Figure 18 but for SAT2 self-intersections. As before, warm colors denote low-latitude bins; greens denote mid-latitude bins; blues denote high-latitude bins, excluding the bin containing the maximum sampling latitude; and purple denotes the polar-most bin containing the maximum sampling latitude for either hemisphere.

To the extent the spatial distribution for self-intersections is not identical across satellites, they diverge only when we consider individual bins within each latitudinal sector. We note that, in some low- and mid-latitude cases, SAT2 self-intersections are centered marginally poleward with respect to SAT1 self-intersections. For example, mid-latitude intersections are contained between $\pm 50^\circ$ and $\pm 55^\circ$ for SAT2 self-intersections, whereas SAT1 self-intersections were

tracked between -45° and -50° in the Southern Hemisphere. A similar pattern is borne out for self-intersections near the equator: For CubeSat 2, self-intersections exhibit a slight polar offset in the Northern Hemisphere low latitudes compared to CubeSat 1 self-intersections near the vicinity. Taken together, our comparison suggests that, although the broad distribution of self-intersections is roughly the same for either satellite, they are not strictly co-located. This implies that altitude exerts a modest but nonzero impact on self-resampling.

The preceding figures confirm what the previous observations tacitly reveal: Self-intersections are anticipated within discrete latitude bands, particularly in the low- and mid-latitudes where we note an extensive gap in coverage stretching across numerous low- and mid-latitude bins. Such a pattern is indicative of the polar-centric distribution targeted by the mission (as described in Drouin et al. 2022). In other words, our simulation yields more robust self-intersection coverage in high latitudes, and this is true both in terms of the relative percent of intersections identified as well as the increasingly narrow gaps in coverage as we move toward either pole, up to the maximum sampling latitude. This finding, which applies to self-intersections regardless of CubeSat, is consistent with our expectations.

Up to this point, we have spatially characterized self-intersections at the exclusion of intersections *between* CubeSats. We now endeavor to estimate where the latter are located in the context of our simulated orbits. Figure 20 illustrates the hypothetical spatial distribution of SAT1-SAT2 intersections per hemisphere, which immediately underscores the unique nature of this type of resampling. Whereas self-intersections occupy discrete latitude bands, SAT1-SAT2 intersections exhibit continuous coverage from the equator to our maximum sampling latitude. Despite this continuity, it remains that a majority of SAT1-SAT2 intersections are centered in the high latitudes, just as we found earlier for self-intersections. In fact, over 75% of SAT1-SAT2

intersections are found at or poleward of $\pm 60^\circ$, with the number of intersections increasing approximately monotonically in the direction of either pole. Intersections are particularly dense near the maximum sampling latitude, where we encounter roughly twice as many SAT1-SAT2 intersections compared to self-intersections for a single CubeSat.

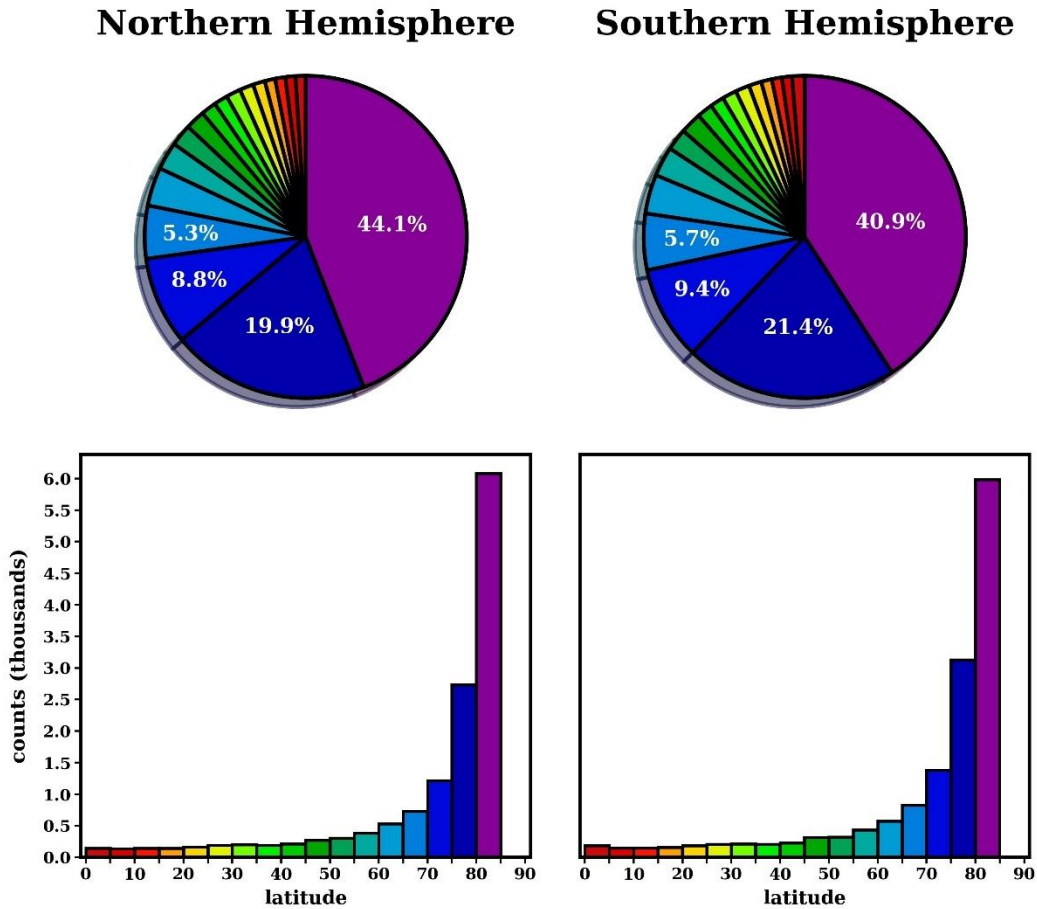


Figure 20: The same as Figures 18 and 19 but for SAT1-SAT2 intersections. Note the y-axis is nearly double compared to the previous figures. The color convention remains the same as before, such that warm colors denote low-latitude bins; greens denote mid-latitude bins; blues denote high-latitude bins, excluding the bin containing the maximum sampling latitude; and purple denotes the polar-most bin containing the maximum sampling latitude for either hemisphere.

Percentages are not strictly identical across hemispheres for intersections between satellites, with bin-wise differences increasing in magnitude as we move to higher latitudes. Although by now we have the impression that SAT1-SAT2 intersections behave differently than self-resampling, the same caveat that applies to hemispheric differences for self-intersections also

applies here. That is, we expect that some latitudinal offset between hemispheres may be the result of the imperfect approximation of true centers. In any case, we note that the general spatial pattern is statistically similar whether considering SAT1-SAT2 intersections in the Northern or Southern Hemisphere, with dense high-latitude coverage and greater spatial resolution overall compared to self-intersections.

3.3 Temporal distribution of simulated intersections

The observational value we derive from PREFIRE intersections will depend only in part on where they occur. Integral to this thesis of the mission's resampling is not only the spatial profile we can resolve given the prescribed orbits but also the temporal profile. Thus, following our spatial results described in the preceding section, we then compiled statistics to characterize our intersections with respect to Δt . Recall that Δt denotes the elapsed time between any two crossovers forming an intersection.

As before, we begin our discussion with the results for SAT1 self-intersections, but now we start by looking at a random CubeSat 1 granule. This is shown schematically in Figure 21, where the centers of every self-intersection for the selected granule have been plotted, up to ± 12 hours. Our results confirm that the number of SAT1 self-intersections per orbit is constant over the simulated period, totaling 32 self-intersections per CubeSat 1 granule. When we review the temporal behavior of the 32 intersections, illustrated chromatically in the same figure, we encounter an obvious pattern: For SAT1 self-intersections, Δt decreases as we move toward either pole. In other words, Δt gets progressively shorter with increasing latitude for this intersection type. Near the equator, Δt hovers near our imposed upper limit of ± 12 hours. If we move poleward we see a gradual decrease in Δt , reaching the theoretical minimum of about 1.5 hours near the maximum sampling latitude. Additionally, SAT1 self-intersections are found in ascending-

descending pairs. This implies that when we identify a discrete Δt on an ascending segment of an orbit, we can find a matching Δt on the descending segment of the same orbit, located in roughly the equal latitude in the opposite hemisphere. In an effort to avoid redundancy, we will note here that this ascending-descending pairing is observed for the other intersection types as well.

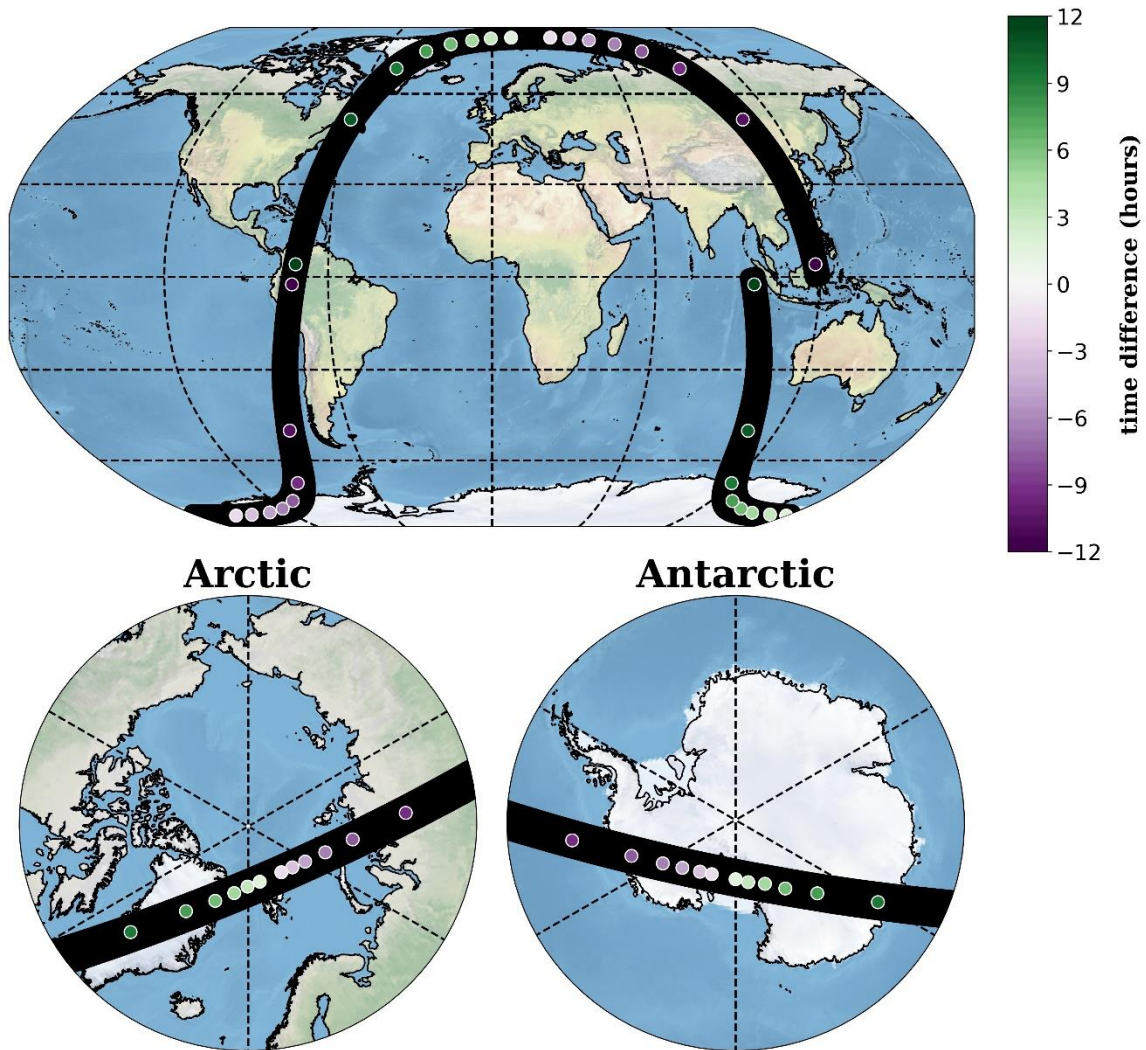


Figure 21: An arbitrary CubeSat 1 granule and the centers of all ± 12 -hour self-intersections specific to that orbit. Each CubeSat 1 orbit yielded a constant number of self-intersections over the simulated period, totaling 32 as shown. Center colors represent the elapsed time between crossovers.

Enriched with these insights, we now consider all SAT1 self-intersections ± 12 hours over the two-month period, represented by their centers in Figure 22. The pattern that emerges is one of time-invariant spatial coverage. In either hemisphere, we encounter eight distinct latitude bands

containing intersections, with each band differentiated by a unique Δt . Our impression that intersections get shorter temporally as we climb to the highest latitudes is confirmed, with the longest and shortest intersections wrapped tightly in latitude bands near the equator and maximum sampling latitude, respectively. Combining our observations, we affirm that SAT1 self-intersections exhibit constant coverage with respect to latitude and time. Moreover, Δt is a discrete function of latitude, with the two inversely proportionate to each other. Our confidence in the results must be tempered by the caveat that realistic changes in the orbit on timescales exceeding the simulated period may result in some departure from the results shown here.

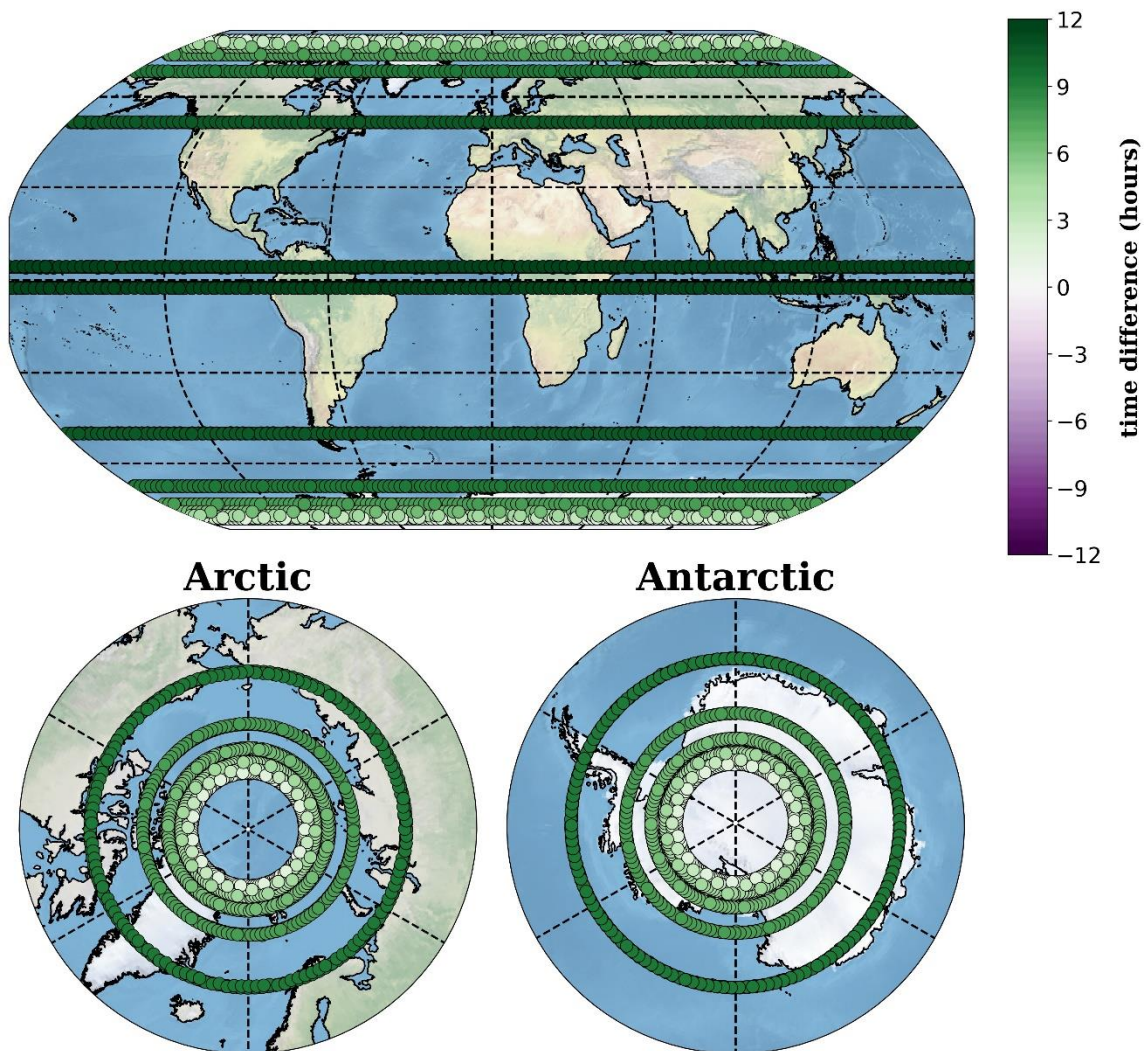


Figure 22: The centers of all SAT1 self-intersections observed over the simulated two-month period. Center colors represent the elapsed time between crossovers.

For completeness, our discussion continues by briefly reviewing results for SAT2 self-intersections. Previously, we found that SAT1 self-intersections and SAT2 self-intersections are practically identical when we consider their regional distributions. We might in turn expect their temporal patterns to tend toward agreement as well, an expectation that is met accordingly. Focusing on the random CubeSat 2 granule shown in Figure 23, we encounter for a second time a fixed number of self-intersections per granule, again totaling 32 per CubeSat 2 orbit. More importantly, we find mutual agreement in the inverse relationship between latitude and Δt , such that SAT2 self-intersections also become temporally shorter as we shift poleward. Here again we find intersections on the shortest theoretical timescales near the highest latitudes sampled in either hemisphere, increasing to ± 12 hours just north or south of the equator.




Figure 23: An arbitrary CubeSat 2 granule and the centers of all ± 12 -hour self-intersections specific to that orbit. Like CubeSat 1, CubeSat 2 orbits yielded a constant number of self-intersections over the simulated period, totaling 32 as shown. Center colors represent the elapsed time between crossovers.

Figure 24 represents the centers of all SAT2 self-intersections ± 12 hours over the simulated period. A complementary analysis confirms the presence of similar latitude bands as SAT1 self-intersections. Owing to the modest influence of altitude described earlier, the latitude bands for SAT2 self-intersections do not strictly overlap the latitude bands of SAT1 self-intersections, but their locations are regionally consistent. Ultimately, we anticipate nearly identical, time-invariant spatial and temporal coverage for both self-intersection types, and this holds for timescales at least as long as two months, the duration of this thesis.

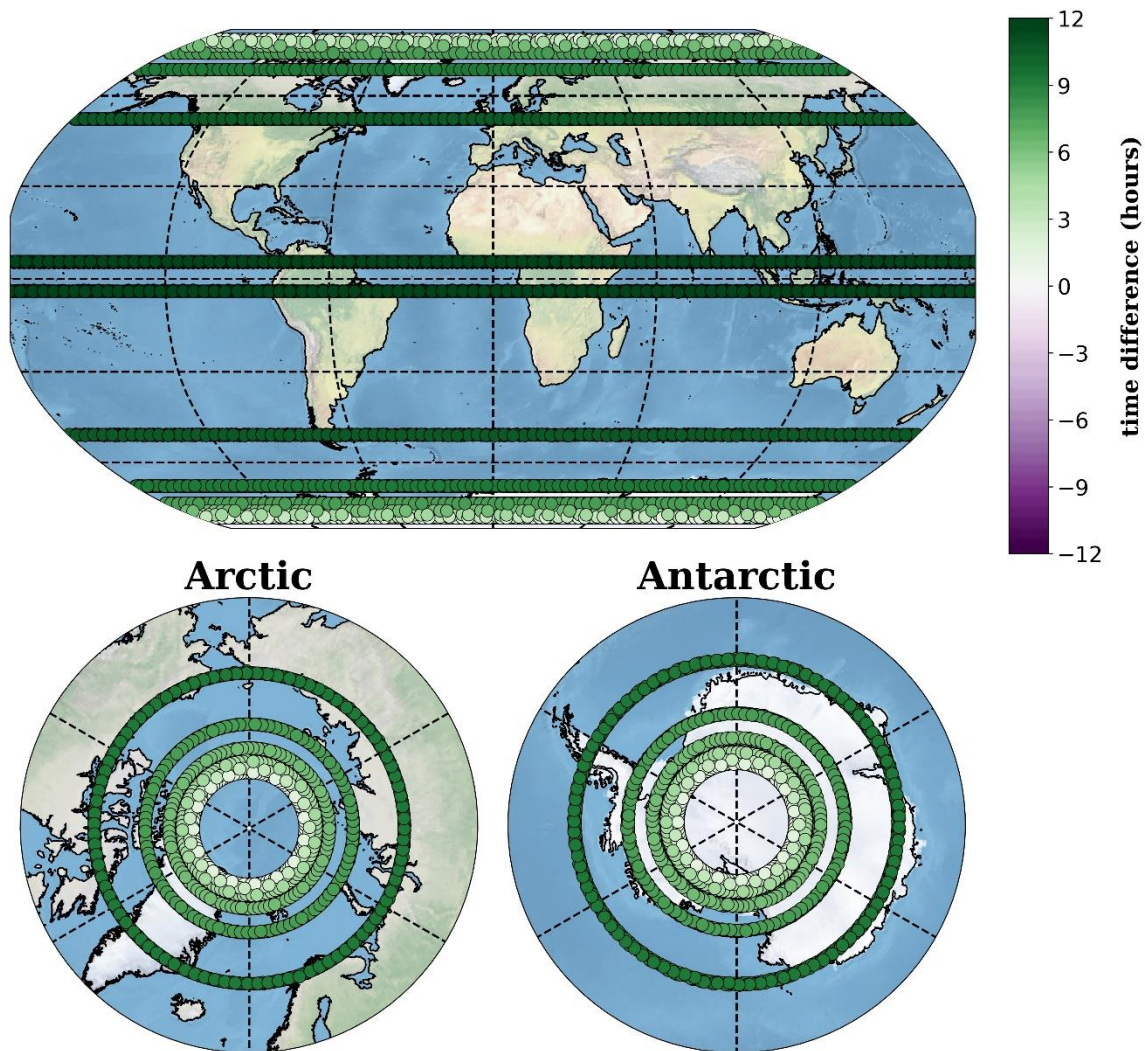


Figure 24: The centers of all SAT2 self-intersections observed over the simulated two-month period. Colors represent the elapsed time between crossovers.

Among other distinguishing features, we previously uncovered that SAT1-SAT2 intersections have relatively high variability on daily timescales, ranging by as many as 72 intersections per day globally. To understand this large spread, we narrow our scope and consider the number of SAT1-SAT2 intersections per granule, just as we did previously. Unlike self-intersections, the number of SAT1-SAT2 intersections per orbit is not fixed and varies from 30 to 34 intersections per granule. We plot a random CubeSat 1 granule that has 34 intersections with CubeSat 2, up to ± 12 hours, thereby illustrating the upper bound per orbit. This is shown in Figure

25, with intersections again represented by their centers. Little more than a cursory glance is required to see notable temporal differences compared to self-intersections. Perhaps most obvious, Δt no longer decreases monotonically as we traverse from low to high latitudes. Conversely, we find a staggering of time differences as we move poleward from the equator. Starting with intersections at lower latitudes, we witness Δt near the median of our imposed 12-hour range. In the highest sampling latitudes, our attention is drawn to manifold Δt values, which signify fairly diverse temporal coverage in this vicinity. Notably, it is within these polar regions that we encounter Δt ranging from rapid timescales up to our maximum of ± 12 hours. The former will be discussed in section 3.4, but for now we simply emphasize that resampling between CubeSats will yield intersections on sub-orbit timescales, a departure from self-intersections. In general, we see little agreement between the temporal coverage for self-intersections and the coverage we observe in the simulated SAT1-SAT2 intersections.




Figure 25: An arbitrary CubeSat 1 granule and the centers of all ± 12 -hour intersections it encountered with CubeSat 2. The number of SAT1-SAT2 intersections per granule varied between 30 and 34 over the simulated period. This example illustrates the upper bound of SAT1-SAT2 intersections, totaling 34. Center colors represent the elapsed time between crossovers.

We can better qualify the differences in coverage by reviewing all SAT1-SAT2 intersections over the two-month period. The resulting mosaic of intersection centers, captured in Figure 26, reveals temporal patterns dissimilar to any encountered previously. In place of the discrete latitude bands we witnessed for self-intersections, we now find a cyclic patchwork of resampling. Coverage is no longer time-invariant and this is evident from the resonating network of Δt . To the degree that our simulations produce realistic insights, we form the impression that SAT1-SAT2 intersections will yield periodic spatial and temporal coverage. Critically, dense and

varied intersections in the high latitudes imply that we expect polar regions to be generously resampled between CubeSats.




Figure 26: The centers of all SAT1-SAT2 intersections observed over the simulated two-month period. Colors represent the elapsed time between crossovers.

Finally, we draw our temporal analysis to a close by binning intersection centers with respect to time. It becomes obvious momentarily that our selected bin widths are not constant, such that bins were designed to emphasize shorter-term intersections. For this exercise, each intersection was assigned to one of the following bins depending on its corresponding Δt , with units in hours: $[0, 0.5)$, $[0.5, 1)$, $[1, 3)$, $[3, 6)$, $[6, 9)$, and $[9, 12)$. We then calculated the frequency

at which intersections for each time bin were identified, expressed as a percentage of the total number of days analyzed. Additionally, we found the conditional mean per day for each time bin.

By now, we are privy to the temporal agreement between self-intersections for either CubeSat, and as such we interpret them interchangeably going forward. The results for a single CubeSat are represented in Table 5. Note that the shortest time bins, namely $[0, 0.5)$ and $[0.5, 1)$ are not applicable to self-intersections since a single orbit never resamples itself. It therefore follows that we have no self-intersections in the first two bins. Neglecting them accordingly, we find that 100% of the time we have self-intersections in the remaining four bins. Interpreting the mean per day, we find an increasing trend with longer Δt , yielding a maximum of 91 intersections per day in $[9, 12)$, up from the minimum of 30 per day in $(1, 3]$.

| Time bin (hours) | $[0, 0.5)$ | $[0.5, 1)$ | $[1, 3)$ | $[3, 6)$ | $[6, 9)$ | $[9, 12)$ |
|------------------------------|------------|------------|----------|----------|----------|-----------|
| Frequency (% of days) | 0 | 0 | 100 | 100 | 100 | 100 |
| Mean per day (counts) | 0 | 0 | 30 | 60 | 60 | 91 |

Table 5: The frequency and mean number of self-intersections per day for six hourly time bins. Results shown are for a single CubeSat. For each bin, the hourly range shown is inclusive of the lower bound but excludes the upper bound.

Results for SAT1-SAT2 intersections are shown in Table 6. The dynamic coverage we previously highlighted for this intersection type is reflected in the first two bins, where we now see nonzero frequencies less than unity. For 67% of the days analyzed, we see SAT1-SAT2 intersections with Δt less than a half an hour, representing relatively rapid intersections. Frequency drops to 40% when we consider the next bin, $[0.5, 1)$. This arithmetic reveals that intersections belonging to each of the first two bins can and commonly do exist concurrently. Moreover, on days when intersections in $[0, 0.5)$ and $[0.5, 1)$ occur, we see an average of 29 and 50 daily

intersections, respectively. The remaining bins contrast with the first two in that each of the four longer-term bins achieve 100% frequency. Unlike self-intersection coverage, the maximum number of SAT1-SAT2 intersections occurs for intersections [3, 6), suggesting that intersections between CubeSats tend to be shorter on average.

| Time bin (hours) | [0, 0.5) | [0.5, 1) | [1, 3) | [3, 6) | [6, 9) | [9, 12) |
|------------------------------|----------|----------|--------|--------|--------|---------|
| Frequency (% of days) | 67 | 40 | 100 | 100 | 100 | 100 |
| Mean per day (counts) | 29 | 50 | 81 | 135 | 103 | 115 |

Table 6: The same as Table 5 but for SAT1-SAT2 intersections. For time bins with frequency less than 100%, the mean shown is conditional and specifies the average number of intersections on days during which intersections within the respective time bin occur.

3.4 Simulated short-term revisits

Suppose the PREFIRE CubeSats register different spectral radiances, denoted I_λ , for a mutual region. To interpret their spectral differences as true indicators of variable conditions, the mission must achieve reliable inter-calibration between instruments. Intersections with arbitrarily short Δt will be integral to inter-calibration and are therefore characterized separately in this section. Although such intersections may be familiar to some readers as near-simultaneous nadir overpasses, this thesis hereafter refers to them more generally as “short-term revisits.”

To begin, we must clarify what qualifies an intersection as a short-term revisit. Theoretically, an intersection is useful for calibration if conditions remain practically static from the first crossover to the second. In reality, the upper bound on Δt is not fixed and will vary pending local conditions; however, we demonstrate this concept by applying an arbitrary threshold of ± 30 minutes. The centers of intersections that satisfy this threshold are emphasized in Figure 27.

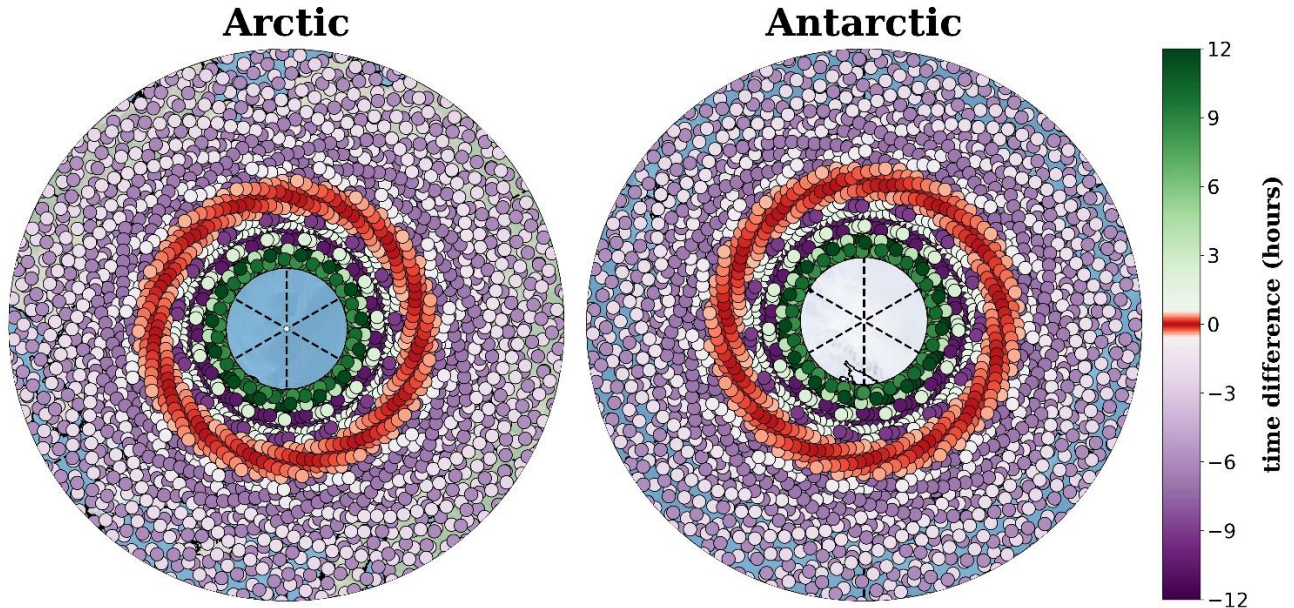


Figure 27: The same as Figure 26 (bottom) but with emphasis on intersections with Δt less than 30 minutes, shown in red. These intersections will be ideal candidates for inter-calibration.

In the resulting mosaic, we note meandering bands of short-term revisits that encircle the highest latitudes, extending from approximately 72° to 78° in either hemisphere. It follows that this latitude range would be appropriate to target for inter-calibration, with the important caveat that the precise range may differ operationally. The temporal pattern moreover implies that we achieve progressively shorter revisits as we move radially inward from the outermost latitudes of the range specified. Going a step further, we note that short-term revisits exhibit a minimum Δt near the center of the corresponding latitude range, approaching near-instantaneous timescales near $\pm 75^\circ$.

Alternatively, we can reconsider our arbitrary threshold of ± 30 minutes in favor of a more inclusive analysis. Figure 28 shows the Δt of the shortest SAT1-SAT2 intersection for every successive CubeSat 1 orbit, including those with Δt greater than ± 30 minutes. Our resulting plot confirms that short-term revisits exhibit periodicity with respect to latitude, gradually shifting equatorward over the course of about 324 orbits. This implies a hypothetical cycle of about three weeks. Periodicity is also observed with respect to Δt over this duration, with our shortest

intersections transitioning from a little over 45 minutes to near-instantaneous elapsed time during the first half of the cycle, followed by the reverse transition in the second half.

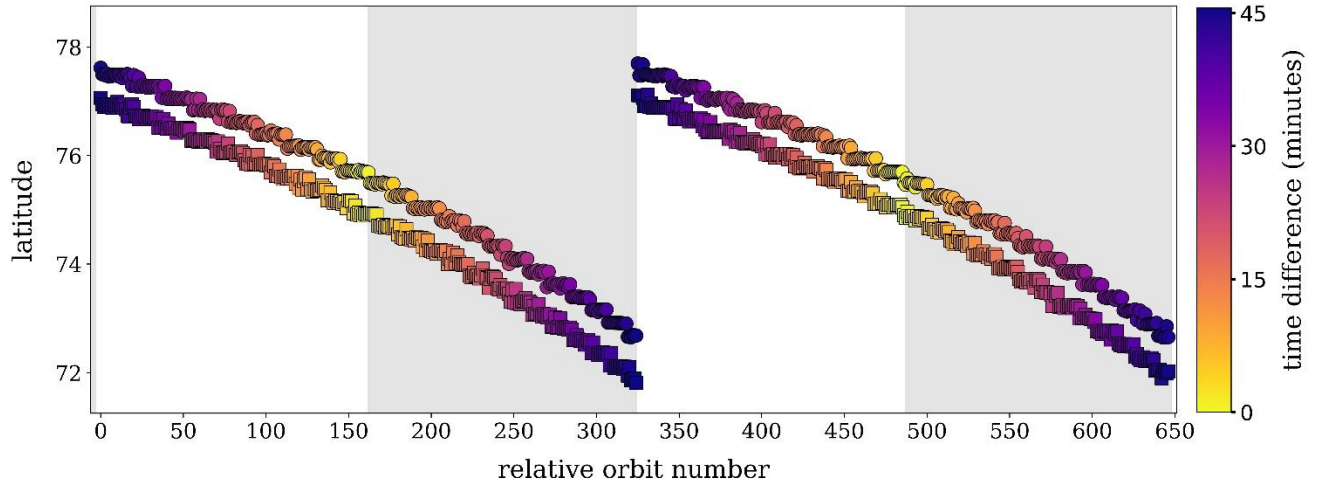


Figure 28: The central latitude of the shortest SAT1-SAT2 intersection per relative CubeSat 1 orbit. Circles denote Northern Hemisphere intersections, whereas squares denote Southern Hemisphere intersections. Unshaded (shaded) regions indicate intervals over which CubeSat 1 (CubeSat 2) was the leading satellite.

4. Conceptual cases

4.1 Arctic heat wave case results

The first case lends insight into the spectral signature of Arctic warming across three distinct surface types. Neglecting land, the selected intersection observes both ocean and sea ice during the first crossover. Although we are fundamentally interested in the OLR responses to sea ice melt, if any, this thesis first addresses whether a PREFIRE CubeSat can distinguish sea ice from ocean.

Figure 29 (a) shows I_λ for two adjacent offshore pixels during the first crossover, representing one ocean and one sea ice pixel. Since we are targeting surface type in this exercise, we are considering only wavelengths in or near our window channels. We note that ocean has slightly greater I_λ across most of the wavelengths shown. This is explicitly plotted in Figure 29 (b), where we show the difference in I_λ between the ocean and sea ice pixel. Our attention is drawn

to a maximum difference of about $0.2 \text{ W m}^{-2} \text{ sr}^{-1} \mu\text{m}^{-1}$ near $11.8 \mu\text{m}$, with a secondary maximum near $12.6 \mu\text{m}$. It should be noted that the brightness temperature (T_B) sensitivity threshold for either TIRS is approximately 1 K. This implies that either instrument can discern between the two surfaces if ΔT_B is at least 1 K. Figure 29 (c) confirms that we achieve sufficiently large ΔT_B from the channels centered near 11.8 and $12.6 \mu\text{m}$, where we previously noted a primary and secondary maximum in ΔI_λ , respectively. This thesis therefore verifies that sea ice and ocean exhibit spectral differences on the order of TIRS sensitivity.




Figure 29: (a) Simulated spectral radiances for a single pixel containing 100% sea ice and the same pixel containing 100% ocean, shown with respect to TIRS2 window channels. Central wavelengths of each channel are represented via squares. (b) Differences in radiances between sea ice and ocean for the pixel in (a). (c) Differences in brightness temperatures between sea ice and ocean for the pixel in (a) and (b). The shaded region represents brightness temperature differences less than the approximate TIRS sensitivity of 1 K.

As detailed in section 2.3.1, this vicinity coincides with sensible heating and variable moisture fluxes that are moreover accompanied by idealized sea ice melt, where applicable. Collectively, such scene changes constitute an arbitrary Arctic warming event, and we now consider how such an event modifies OLR in the context of our hypothetical intersection.

Figure 30 shows I_λ for $6.73 \mu\text{m}$ during the first and second crossover (top), followed by the percent change in I_λ (bottom). With regard to high water vapor sensitivity at this wavelength, we see a sharp increase in I_λ for the onshore pixels that incurred the greatest loss in TCWV, as expected, with a peak increase of just over 11% during the second sampling. For the regions to which this increase applies, the corresponding weighting function peaks closer to the surface as a result of drier conditions. Elsewhere, most pixels registered a decrease in I_λ during the second crossover, decreasing by as much as nearly 15%. This is consistent with an increase in TCWV, which results in our weighting function peaking higher in the atmosphere where temperatures are presumably colder relative to conditions closer to the surface. (It was outside of the scope of this analysis to consider the vertical profile for every pixel, and so we rest our interpretation on universal assumptions.) It is noteworthy that a modest increase in TCWV is associated with a sharp decrease in I_λ . This reinforces the strong sensitivity of the $6.73 \mu\text{m}$ channel to the vertical moisture profile.

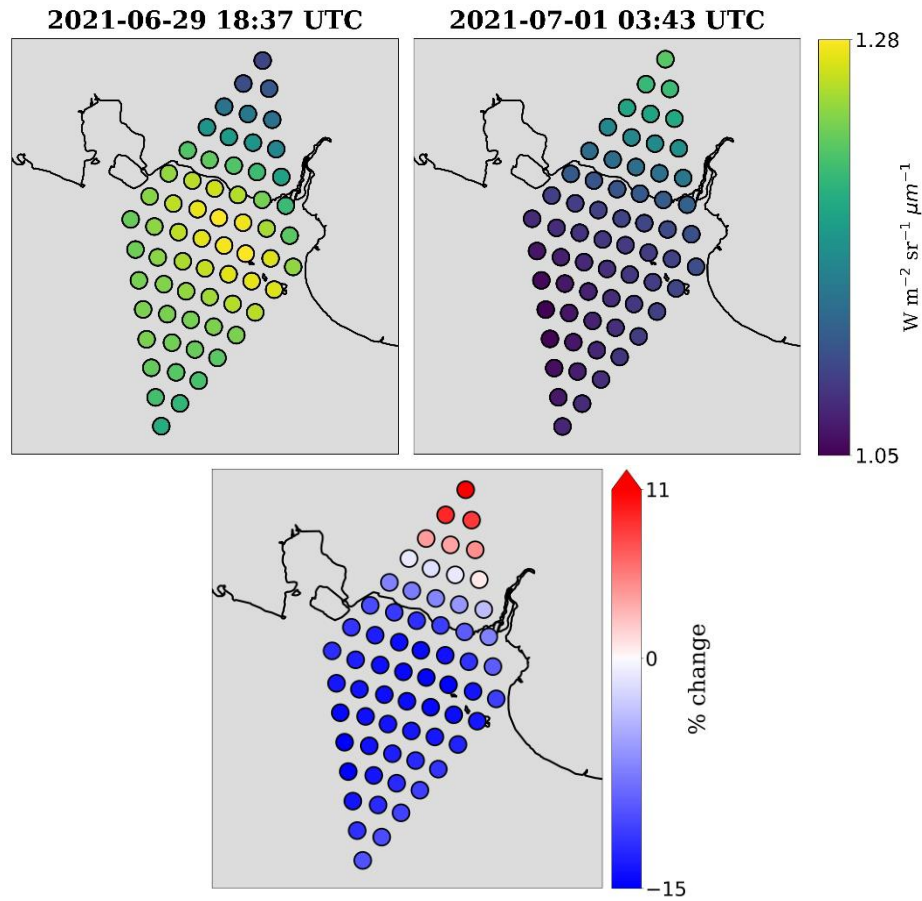


Figure 30: Simulated spectral radiances for the TIRS2 channel centered on $6.73 \mu\text{m}$ during the first and second crossover (top), followed by the percent change (bottom). Pixels shown represent the common pixels of this intersection. Additionally, the upper and lower bounds selected were based on the spectral range for the common pixels at either crossover.

We now consider two representative window channels, starting with $10.95 \mu\text{m}$. This is shown in Figure 31. Comparing crossovers, we see subtle markers of sea ice pixels in the first sampling, where we note a relatively low I_λ compared to surrounding ocean pixels. By the second sampling, during which we have imposed total melt of sea ice, we see a more uniform distribution of I_λ for offshore pixels relative to the offshore distribution observed during the first sampling. Notably, the difference between the two crossovers is everywhere greater than zero. In other words, no pixels registered a decrease in I_λ at this wavelength. Due to variable rates of warming, the percent increase also varies across the resampling region, with peak increases of greater than 12% noted for onshore pixels where the greatest sensible warming was noted. Moderate increases

in I_λ were also observed in offshore pixels that experienced appreciable warming, notably in pixels corresponding to sea ice melt.

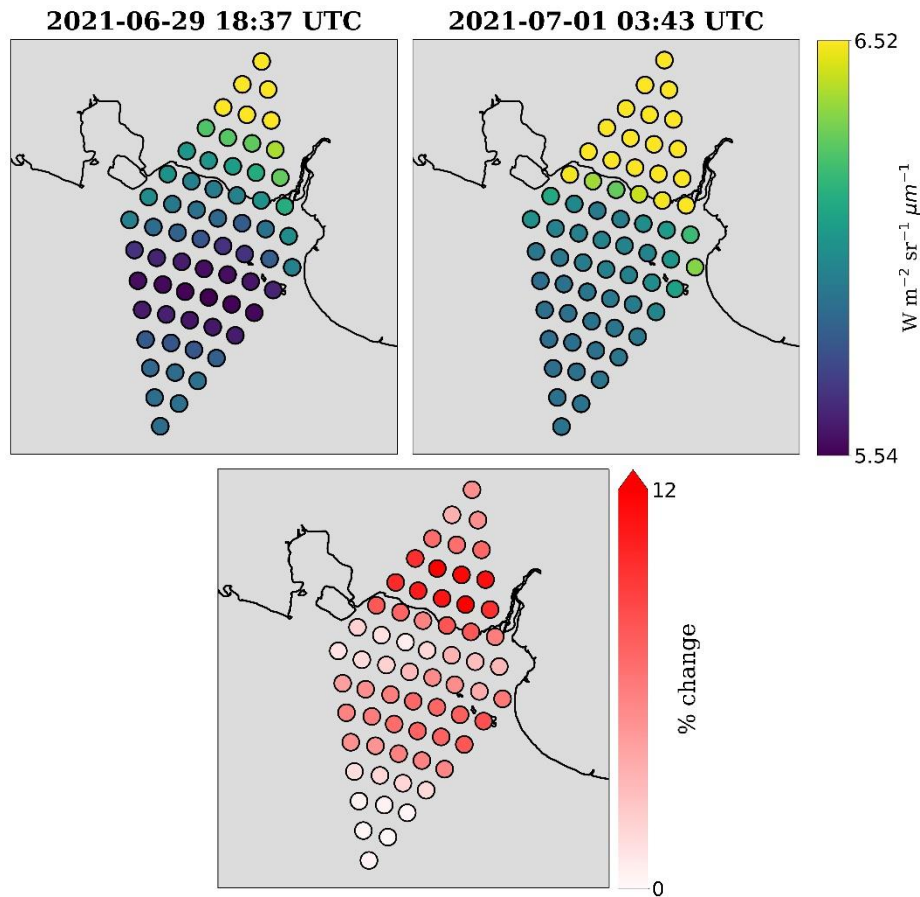


Figure 31: The same as Figure 30 but for the TIRS2 channel centered on $10.95 \mu m$.

The second window channel we consider is $11.8 \mu m$, the channel at which our previous assessment confirmed maximum ΔI_λ between ocean and sea ice. This is shown in Figure 32. As with the previous window channel, we again note the relatively low I_λ corresponding to sea ice pixels in the first crossover, which contrasts with the higher I_λ of the surrounding ocean pixels. This difference between sea ice and ocean is sharper for $11.8 \mu m$ compared to 10.95 , resulting in an enhanced distinction between the two surface types. Reviewing the later sampling time, we again find that, following total sea ice melt, most offshore pixels converge with respect to I_λ during the second crossover. When the differences between observations are plotted, our attention is again

drawn to notable increases in I_λ along the coastal land, co-located with maximum local warming. The peak increase in I_λ is slightly lower than that for $10.95 \mu\text{m}$, with some pixels exceeding +11%. Unlike the previous window channel, we encounter a more well-defined response in pixels corresponding to sea ice melt, where we note magnitudes approaching our onshore peak at approximately +9%. Our findings imply that, while changes in our window channels are strongly driven by sensible warming, sea ice melt yields measurable differences in the OLR response, particularly for the TIRS2 channel centered on $11.8 \mu\text{m}$.

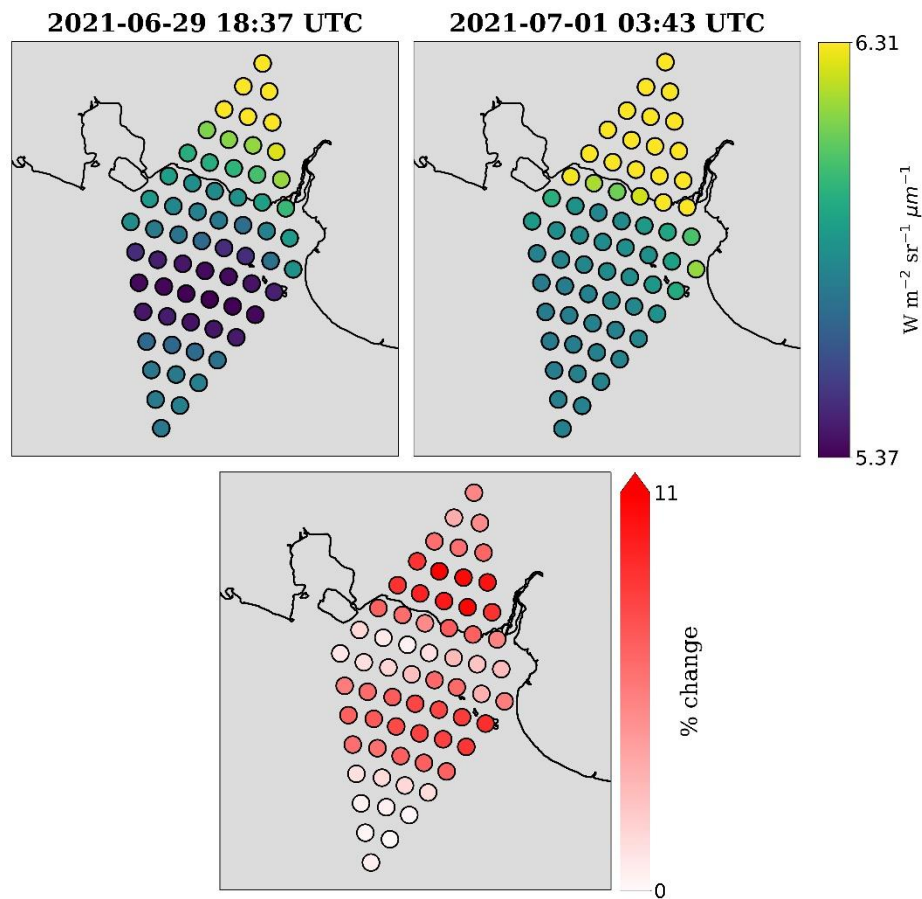


Figure 32: The same as Figure 30 but for the TIRS2 channel centered on $11.8 \mu\text{m}$.

The next three channels we consider are 19.39 , 26.99 and $39.64 \mu\text{m}$, which represent a partially transparent, mostly opaque and strongly opaque far-infrared channel, respectively. Beginning with $19.39 \mu\text{m}$, shown in Figure 33, we note the emerging signs of decoupling between

outgoing FIR and surface conditions. For instance, the boundary between sea ice and ocean is no longer detectable in the first crossover. Similarly, I_λ observed during the second crossover are consistent with the TCWV profile, such that the pixels with the highest (lowest) I_λ coincide with the lowest (highest) TCWV values. Either implies that, as water vapor sensitivity increases relative to the MIR channels, our weighting functions are peaking lower in the atmosphere when TCWV is relatively low and higher in the atmosphere when TCWV values surge. Taking ΔI_λ into account, we verify that a small fraction of onshore pixels registered an increase of greater than 4% during the second sampling compared to the first. Again, this is consistent with a shift from relatively high to low TCWV, as observed in the vicinity of the noted pixels. Interestingly, we see practically equal and opposite ΔI_λ over a much larger fraction of pixels, primarily offshore. That is, where we observed modest increases in TCWV, we now see a decrease in I_λ , peaking at just over -4%. This finding strongly implies that our partially transparent channels still exhibit sensitivity to enhanced moisture, evident by the disproportionately large OLR response to TCWV increases compared to the OLR response to TCWV decreases.

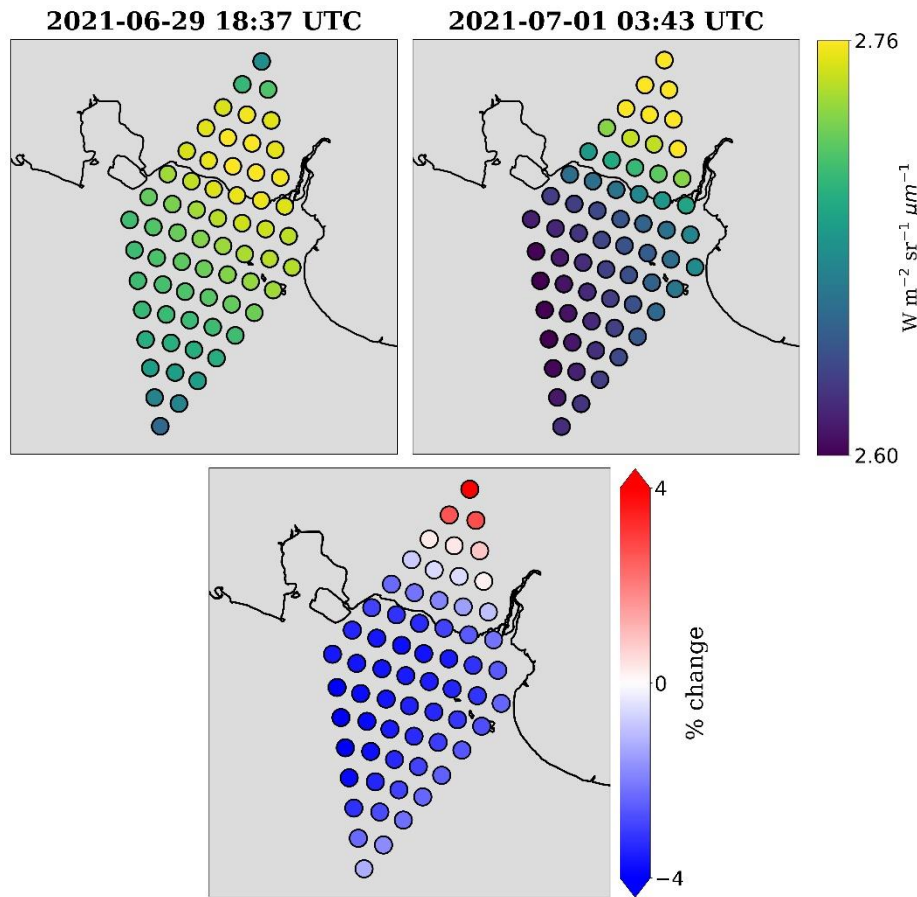


Figure 33: The same as Figure 30 but for the TIRS2 channel centered on $19.39 \mu m$.

Figure 34 illustrates I_{λ} for our TIRS channel centered at $26.99 \mu m$. We again remark that either crossover reinforces a clear decoupling of outgoing FIR from underlying surface conditions. Many of the comments issued for the preceding channel apply to $26.99 \mu m$ as well, including the absence of a clear sea ice-ocean boundary in the offshore population of pixels during the first sampling or the tendency for regions with relatively low (high) TCWV to exhibit higher (lower) I_{λ} during either the first or second sampling. It is worth citing that the fidelity between TCWV and I_{λ} is even more apparent for this channel than for $19.39 \mu m$, reaffirming our expectations for mostly opaque channels. With reference to ΔI_{λ} , we recognize the general pattern from the previous channel, such that we again note positive ΔI_{λ} for onshore pixels with corresponding decreases in TCWV and negative ΔI_{λ} for the majority of pixels where modest increases in TCWV were

previously cited. Unlike 19.39 μm , however, the percent change is not nearly equal in magnitude between the positive and negative changes and this time favors a percent decrease. For example, we see a peak increase in I_λ of just over 3% in the regions farthest inland where sharp moisture decreases were registered; conversely, we see many pixels associated with modest positive moisture changes decrease by over 5% between crossovers. Among other interpretations, it might strike the reader that this suggests a greater sensitivity to enhanced moisture than that which was noted for the partially transparent channel. This satisfies our expectations.

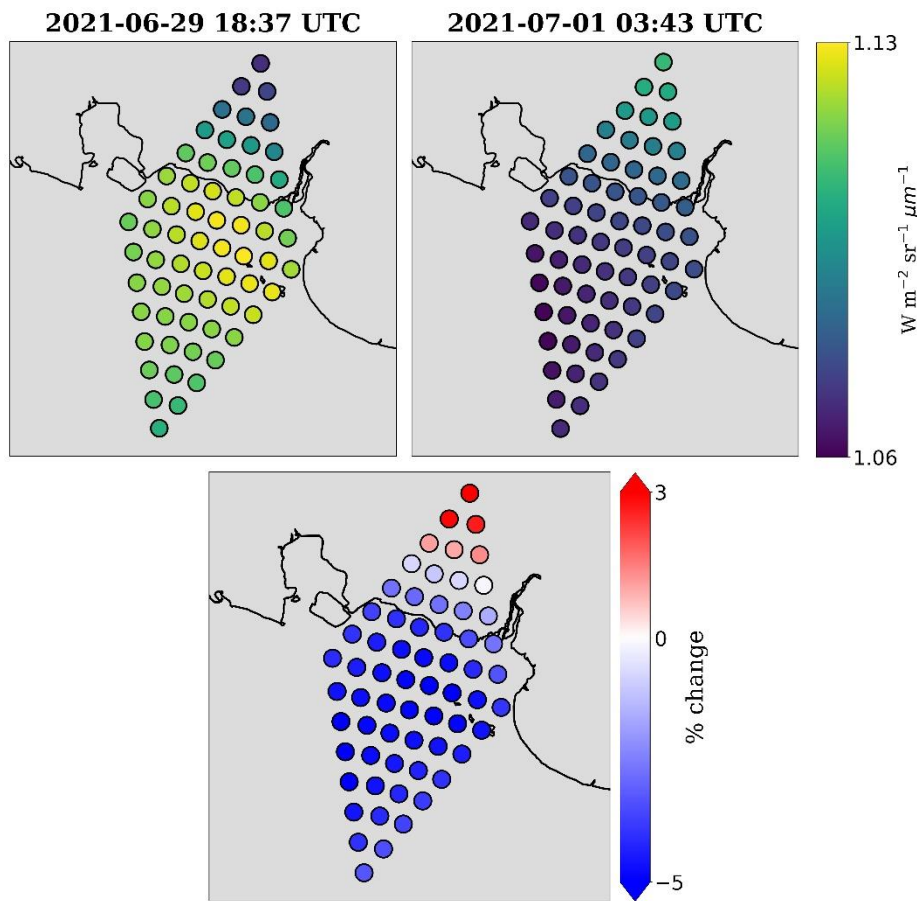


Figure 34: The same as Figure 30 but for the TIRS2 channel centered on 26.99 μm .

For completeness, we also consider a strongly opaque channel centered at 39.64 μm , shown in Figure 35. Two observations greet us at the outset. First, the magnitude of I_λ is quite small, confirming relatively weak emissions measured by the lowest frequency infrared channels. This

contributes to the second observation: Variability in I_λ is limited at this wavelength, particularly during the second crossover when a fairly uniform distribution emerges. Taking these observations into account, we emphasize the clear influence of vertical moisture on I_λ at either sampling time. During the first crossover, we see minimum (maximum) I_λ where TCWV values were found to be highest (lowest). Similarly, we see a reversal in the location of minimum and maximum I_λ during the second crossover, such that our minimum and maximum I_λ are found offshore and onshore, respectively. This corresponds to the approximate reversal in TCWV noted earlier. In either case, lower radiances imply that our weighting function is peaking higher in the atmosphere, while the opposite is generally valid for higher radiances. It is important to note, however, that the small range of I_λ suggests that vertical shifts in the weighting functions may be modest, highlighting the opacity of the 39.64 μm channel. For ΔI_λ , we see a pattern that diverges somewhat from the pattern observed for 26.99 μm . That is, the second crossover registers a slightly greater response to reduced vertical moisture compared to the first crossover, with regions of maximum moisture loss corresponding to increases in I_λ by up to nearly 4%. Conversely, in regions exhibiting a peak reduction in TCWV, we note a decrease in I_λ by up to approximately 3%. Overall, shifts in vertical moisture dominate ΔI_λ , resulting in an incidental division between onshore and offshore regions of the intersection.

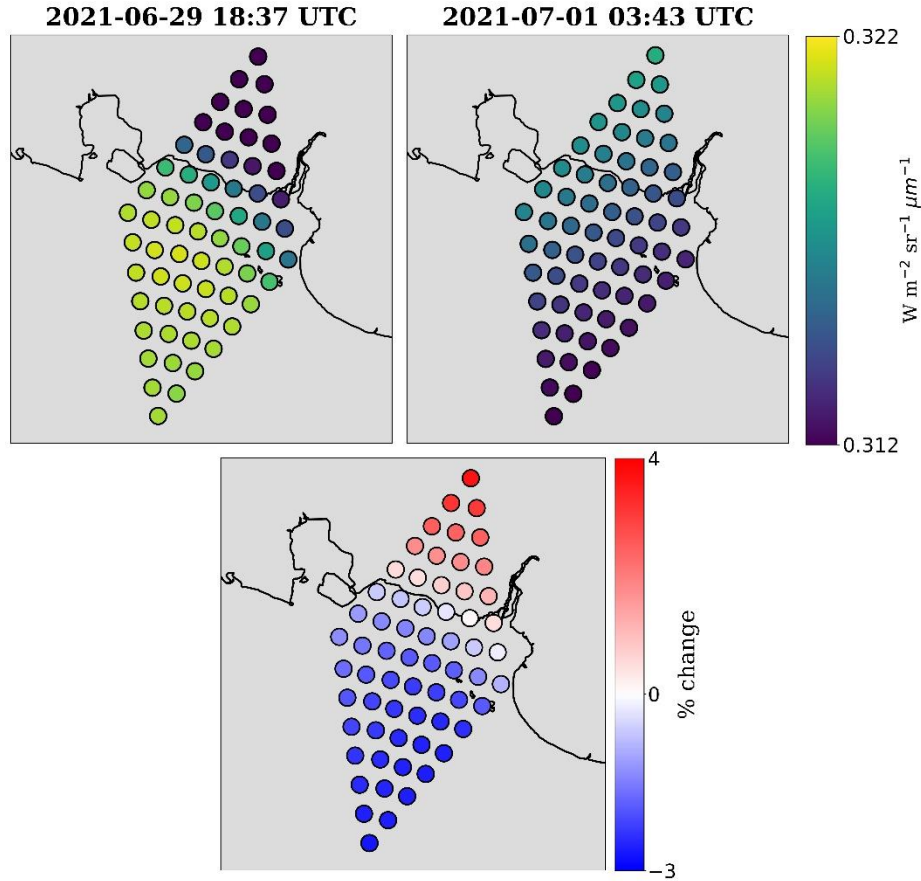


Figure 35: The same as Figure 30 but for the TIRS2 channel centered on $39.64 \mu m$.

Finally, we conclude our first case by plotting the mean difference in I_{λ} for every channel as captured by our hypothetical intersection. This is shown in Figure 36. Each surface type has been plotted separately to facilitate a simple comparison between land, ocean and sea ice melt pixels. Our attention initially falls to the window channels from roughly 9 to $13 \mu m$, which yield the greatest absolute differences in I_{λ} . On average, land pixels registered the greatest difference in I_{λ} , with a maximum difference approaching $0.6 W m^{-2} sr^{-1} \mu m^{-1}$ at the channel centered on $10.95 \mu m$. Reviewing the mean differences for offshore surfaces, we find a greater spectral response, on average, for the sea ice melt pixels compared to nearby ocean pixels. For pixels that underwent sea ice melt, we verify a maximum difference of about $0.4 W m^{-2} sr^{-1} \mu m^{-1}$ at the channel centered on $11.8 \mu m$, consistent with our earlier findings. Ocean pixels registered a more modest spectral

response in our window channels, peaking at under $0.2 \text{ W m}^{-2} \text{ sr}^{-1} \mu\text{m}^{-1}$ at the channel centered on $10.95 \mu\text{m}$. This represents about half the maximum average spectral difference noted for sea ice melt pixels. Additionally, we see similar mean decreases in I_λ near $6.73 \mu\text{m}$ for either offshore surface type, equaling about $-0.2 \text{ W m}^{-2} \text{ sr}^{-1} \mu\text{m}^{-1}$. On average, we see negligible difference in the land pixels at this wavelength; we caution, however, that this is the result of competing differences observed onshore, which offset each other when averaged together. This same caution applies to FIR wavelengths, where negligible differences shown in Figure 36 must be acknowledged in the same context. Either sea ice melt or ocean pixels exhibit modest decreases across the partially transparent channels on average, with the former surface type registering a slightly greater mean decrease. This pattern holds albeit to marginal extent in the mostly opaque channels. Very little absolute difference in I_λ is observed across the low-frequency water vapor continuum channels. On average, either offshore surface type registered a slight decrease in I_λ , whereas little mean difference is detected in land pixels.




Figure 36: Differences in simulated mean spectral radiances between crossovers, averaged over all 64 common pixels and taken with respect to each of the three surfaces represented in this intersection. Squares denote the central wavelengths of the TIRS2 channels and missing wavelengths correspond to TIRS2 spectral gaps.

Critically, our analysis cannot quantify the contribution of sea ice melt to the spectral responses observed. In addition to surface emissivity, variable rates of warming and uneven moisture fluxes undoubtedly impart an influence on OLR. It is thus impossible, absent a more advanced analysis, to isolate the radiative contributions of sea ice melt. Constraining OLR sensitivity to such an extent would require an inclusive examination of local changes across pixels, an endeavor that exceeds the scope of this research. This author speculates that, for an intersection like the one profiled here, a future study might characterize the vertical temperature and moisture profiles at each pixel; snow cover over the land and sea ice pixels; snow grain size and type, if applicable; and spectral emissivities based on surface type adjustments if snow is determined present. Additionally, although we have removed clouds from the simulated radiances, a more realistic study would also need to consider cloud cover, including whether and what type of clouds are present in the vicinity of each pixel. It may alternatively be more practical to target intersections that capture a less ambiguous change in the underlying scene to help isolate spectral processes. Ultimately, this case demonstrates not only the rich information content we anticipate within highly variable, multi-surface polar scenes, but also the considerable challenges of interpreting them in the context of PREFIRE resampling.

4.2 Cloud dissipation case results

As a matter of course, we expect PREFIRE intersections to capture cloud cover changes in high latitudes, where underlying polar surfaces favor unique cloud radiative effects. In the service of illustrating cloud radiative forcing in a frigid polar environment, we now discuss the results of a simulated local transition from ice clouds to clear skies under the three optical depth scenarios described in section 2.3.2.

We begin by considering the radiative impact of cloud dissipation for a window channel, focusing on the TIRS1 channel centered on $11.39\ \mu\text{m}$. Figure 37 shows the spectral OLR for the common pixels of our Antarctic intersection. Although the simulation was structured so that clear sky conditions were sampled during the second crossover, we begin with the clear sky results to foster a more sequential comparison. Thus, Figure 37 shows spectral OLR for clear sky in the first panel (a), followed by spectral OLR for the same region given a uniform ice cloud with an optical depth equal to 1 (b), 2 (c) and 4 (d). It strikes us that the cold wintertime surface of interior East Antarctica (as described in section 2.3.2) contributes to relatively low emissions when clouds are absent. Moreover, spectral OLR increases everywhere when a uniform ice cloud is present, with the magnitude of increase scaling with optical depth. In other words, we verify that cloud dissipation corresponds to a decrease in OLR for $11.39\ \mu\text{m}$, with the difference between cloudy and clear conditions increasing in step with cloud optical depth. This implies that the cloud top level is, despite its relatively high elevation near 400 hPa, warmer than the wintertime polar surface, resulting in a positive cloud radiative effect. In all scenarios, the highest radiances coincide with the warmest regions of the resampling area.

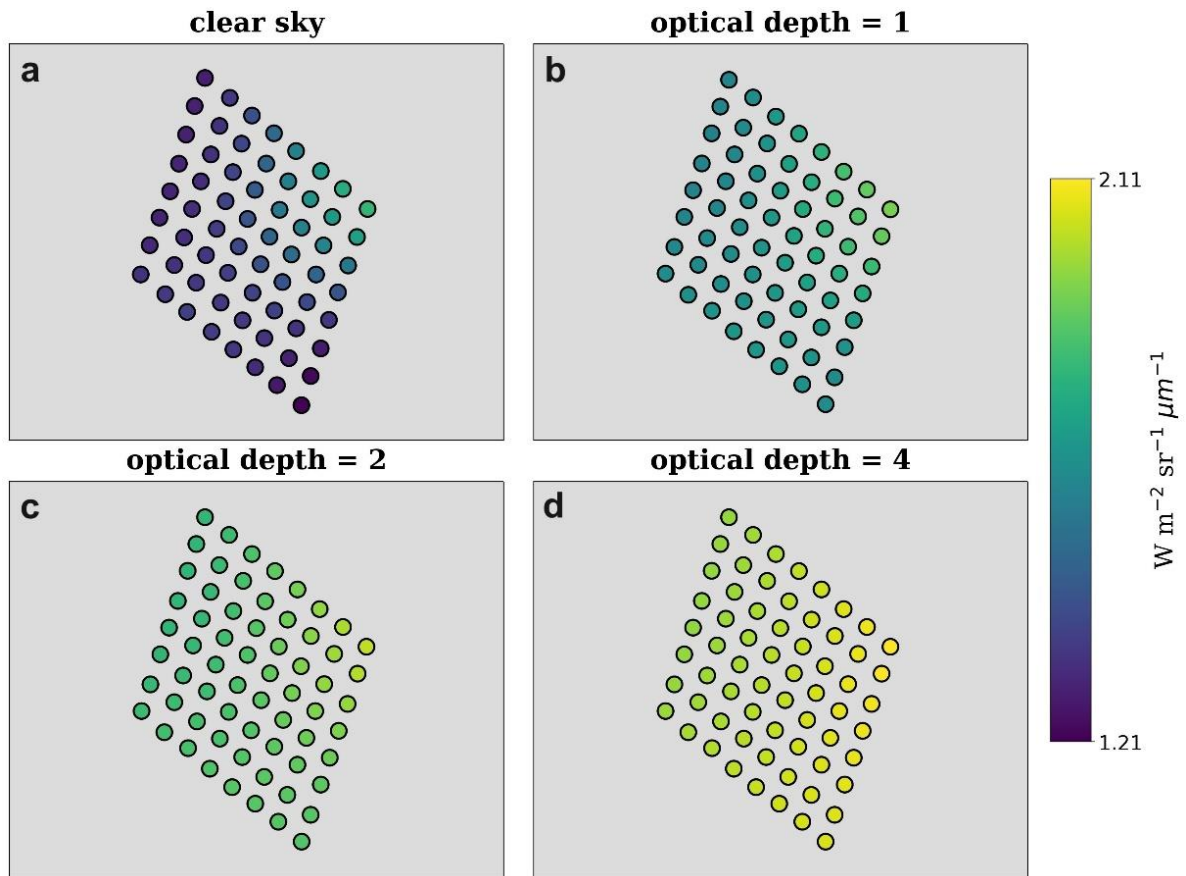


Figure 37: Simulated spectral radiances at $11.39 \mu\text{m}$ for (a) clear sky and for a uniform cloud layer with cloud optical depths (b) 1, (c) 2, and (d) 4.

Figure 38 is the same as Figure 37 but for the TIRS1 channel centered on $18.99 \mu\text{m}$. At this wavelength, the clear sky crossover (a) looks nearly as it does for the window channel, save for a marginally greater magnitude compared to $11.39 \mu\text{m}$. This implies our weighting function is peaking at a slightly warmer level just above the surface. The inverse can be seen in the three uniform cloud scenarios (b-d), where we see I_λ values are somewhat less than those seen in the previous channel, but our cloudy results are otherwise practically identical compared to $11.39 \mu\text{m}$. Taken as a whole, this reveals that our simulated ice cloud has a positive radiative effect at $18.99 \mu\text{m}$ that is similar to the transparent channels but somewhat damped.

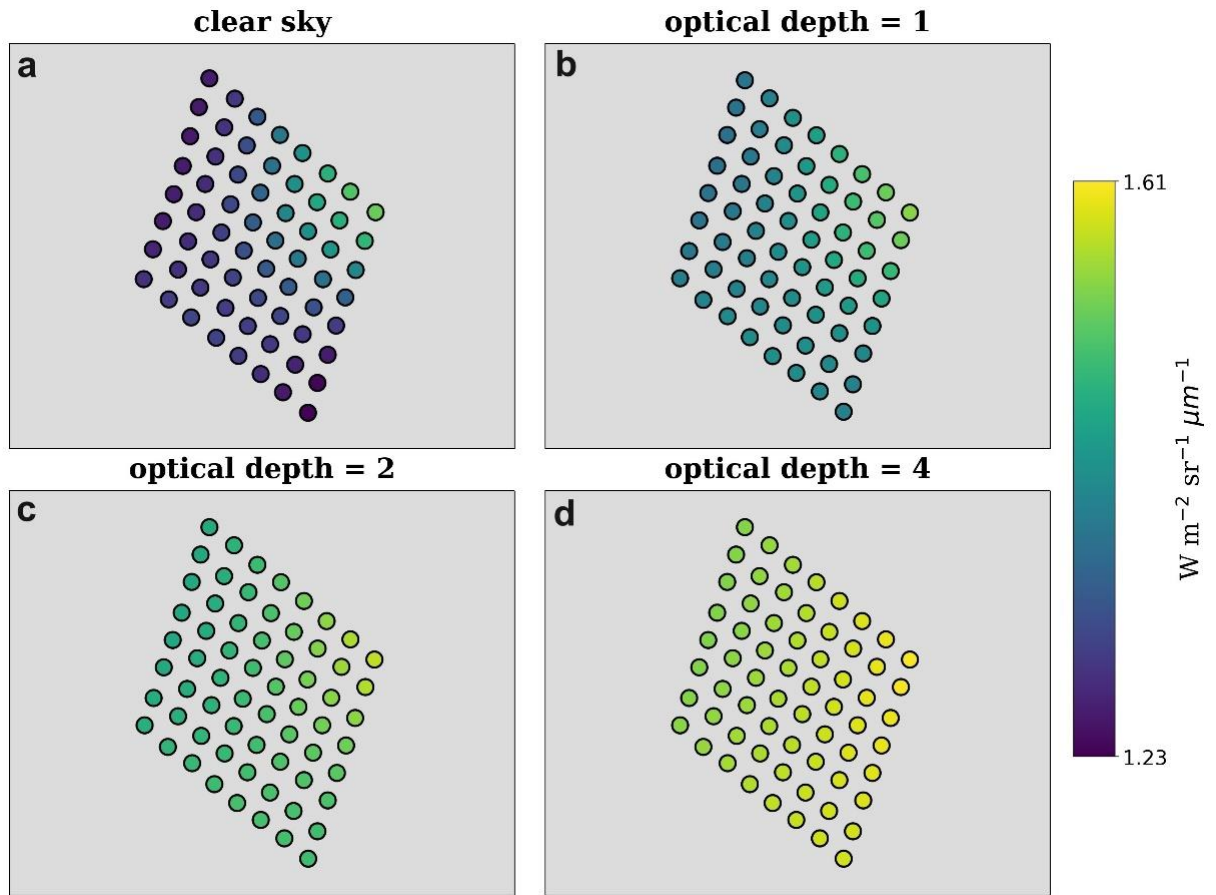


Figure 38: The same as Figure 37 but for 18.99 μm.

Going further, we consider the cloud radiative effect for the TIRS1 channel centered on 26.58 μm, shown in Figure 39. Compared to the preceding clean and dirty window channels, we now encounter a more restricted range of OLR values as well as lower output in general, peaking at less than $0.85 \text{ W m}^{-2} \text{ sr}^{-1} \mu\text{m}^{-1}$. A more important point of contrast is found near the right vertex, where a fraction of pixels registered greater spectral OLR during the clear sky crossover (a) compared to cloudy conditions (b-d). This vicinity coincides with the local maximum for skin temperature and TCWV, as described in section 2.3.2. Note that this positive radiative effect with cloud dissipation becomes somewhat neutralized with increasing optical depth. In the remaining regions of the intersection area, where skin temperature and vertical moisture remain lower, we see a response to cloud dissipation that mirrors the radiative responses for the more transparent

channels. That is, pixels excluding those near the right vertex generally exhibit a subtle decrease in spectral OLR at $26.58 \mu\text{m}$ when a uniform ice cloud dissipates. We also see that the distribution of spectral OLR becomes much more uniform with increasing optical depth, more so than was observed at smaller wavelengths.

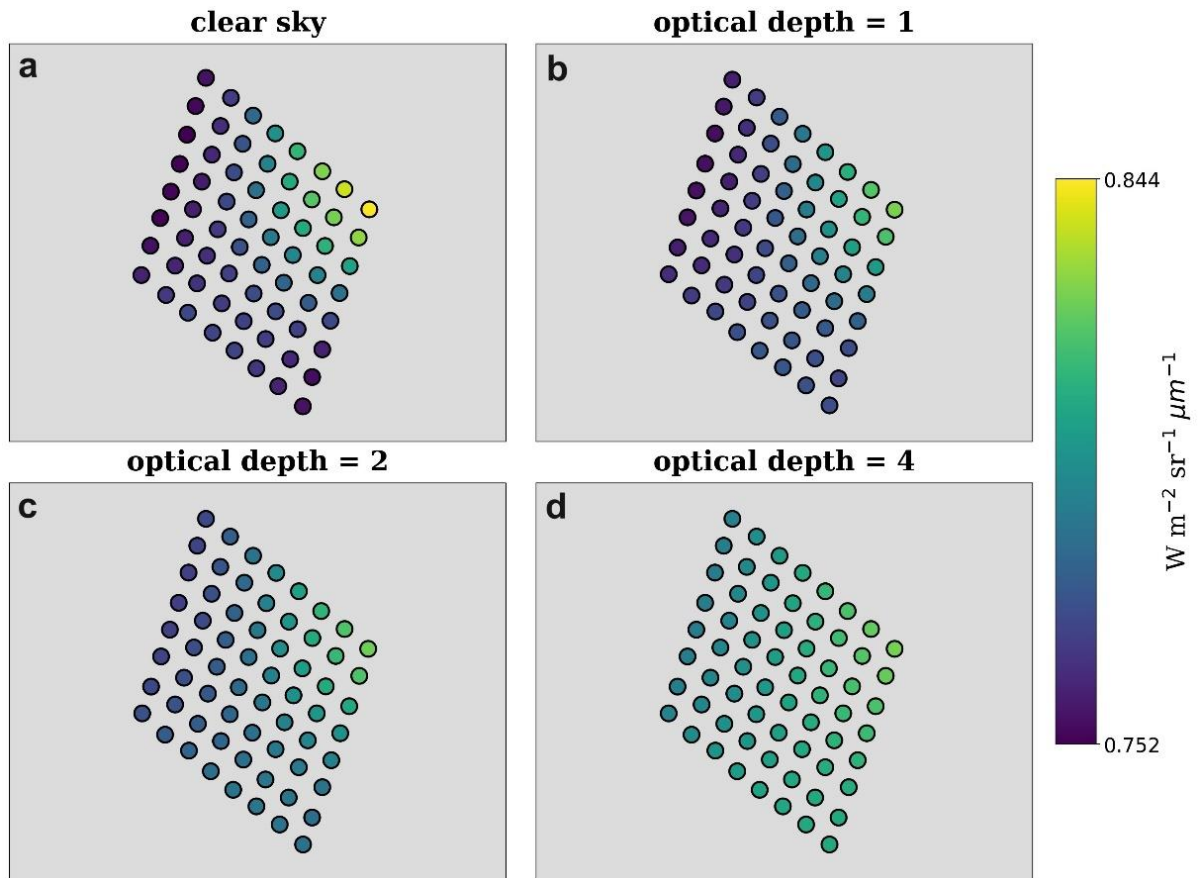


Figure 39: The same as Figure 37 but for $26.58 \mu\text{m}$.

The last wavelength we profile is the TIRS1 channel centered on $40.08 \mu\text{m}$, a wavelength representing a strongly opaque far-infrared channel. As shown in Figure 40, spectral OLR is relatively weak at this wavelength and exhibits low variability with respect to absolute output. It also becomes apparent that the cloud radiative effect at $40.08 \mu\text{m}$ is generally negative, such that a transition to clear sky (a) from an ice cloud (b-d) yields an increase in spectral OLR. This constitutes a reversal in the radiative response noted earlier. Up until this point, we observed a decrease in spectral OLR with cloud dissipation (excepting the sector in which this did not hold

for the previous channel), and noted that cloud optical depth enhanced the response. We again see an enhanced radiative response with cloud optical depth, but this time in the opposite direction. Additionally, maximum spectral OLR occurs during the clear sky sampling near pixels that are less frigid and correspond to lower TCWV, taken relative to other pixels. As cloud optical depth increases, spatial variability decreases and the signal of relatively high temperature and low moisture content becomes extinguished for clouds that are optically thickest. It implicitly follows that, absent clouds, our weighting function peaks at an elevated level where it is warmer than the surface and cloud top height.

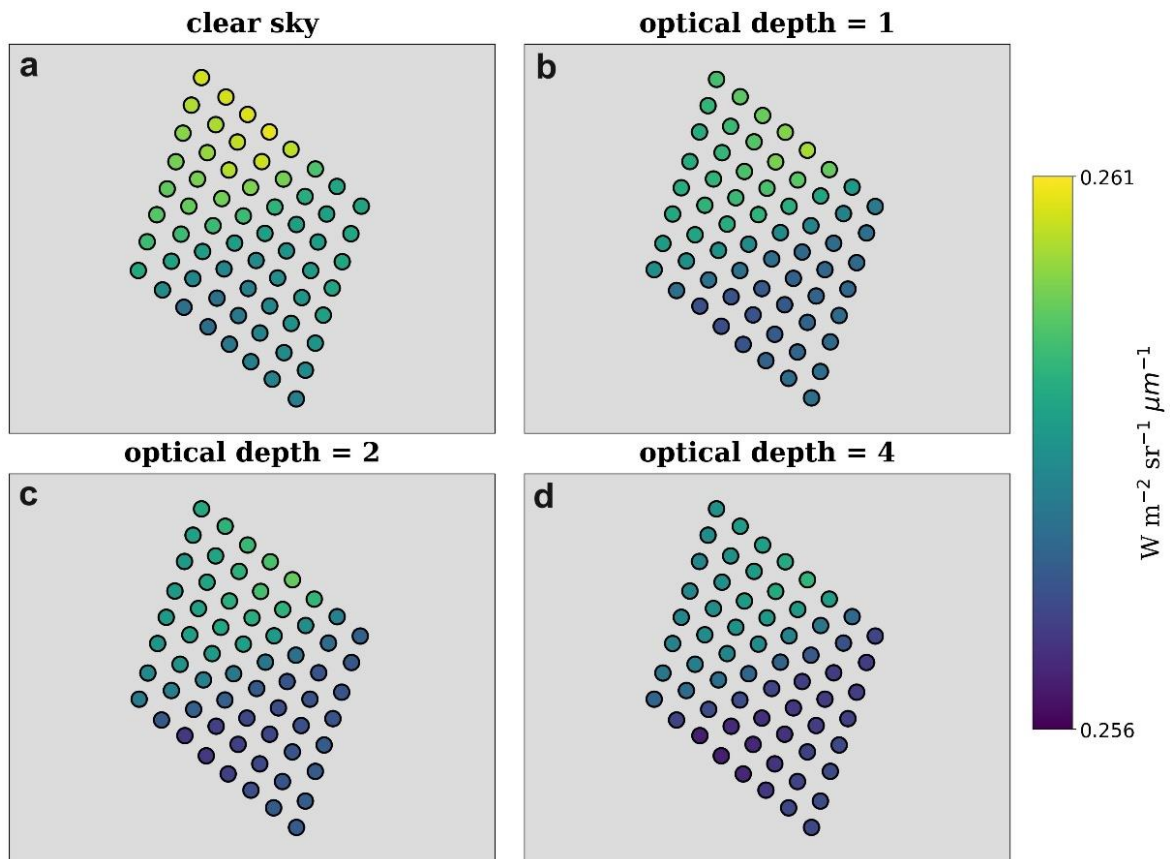


Figure 40: The same as Figure 37 but for 40.08 μm.

Figure 41 quantifies the radiative impact of cloud dissipation for our selected channels, represented as a percent change in spectral OLR following a transition from cloudy to clear sky. Each of the three cloud optical depths is shown per column. Differences in absolute I_{λ} at 11.39 μm

translate to percent decreases ranging between approximately -6% to -26%, -11% to -35%, and -14% to -41% for cloud optical depths 1, 2 and 4, respectively (a-c). To a lesser extent, we encounter the same trend at 18.99 μm , with percent decreases ranging from approximately -1% to -12%, -3% to -17%, and -5% to -22% for cloud optical depths 1, 2 and 4, respectively (d-f). Recall that we see local differences in the radiative response at 26.58 μm , where pixels near the right vertex exhibited an increase in OLR when clouds dissipated. For optical depth 1 (g), we see a corresponding percent increase for pixels near the upper right sector, peaking at just over 2%. Other pixels, primarily near the lower vertex, register roughly equal and opposite changes, decreasing by about the same amount (-2%) following cloud dissipation. Excepting these regions, we see negligible changes elsewhere. When we review the same wavelength for cloud optical depth 2 (h) and 4 (i), we find that the share of pixels for which the percent change is positive decreases as optical depth increases. In all scenarios, the percent increase peaks around 2%. As the number of pixels registering increases in spectral OLR shrinks with thickness, the opposite is also true. In other words, the number of pixels for which cloud dissipation reduces spectral OLR increases in line with optical depth. For optical depth 2 and 4, we see peak decreases of about -4% and -6.5%, respectively. For 40.08 μm , we generally observe marginal increases in spectral OLR, with a maximum increase of about 0.5%, 0.65% and 0.8%, for cloud optical depth 1, 2 and 4, respectively (j-l).

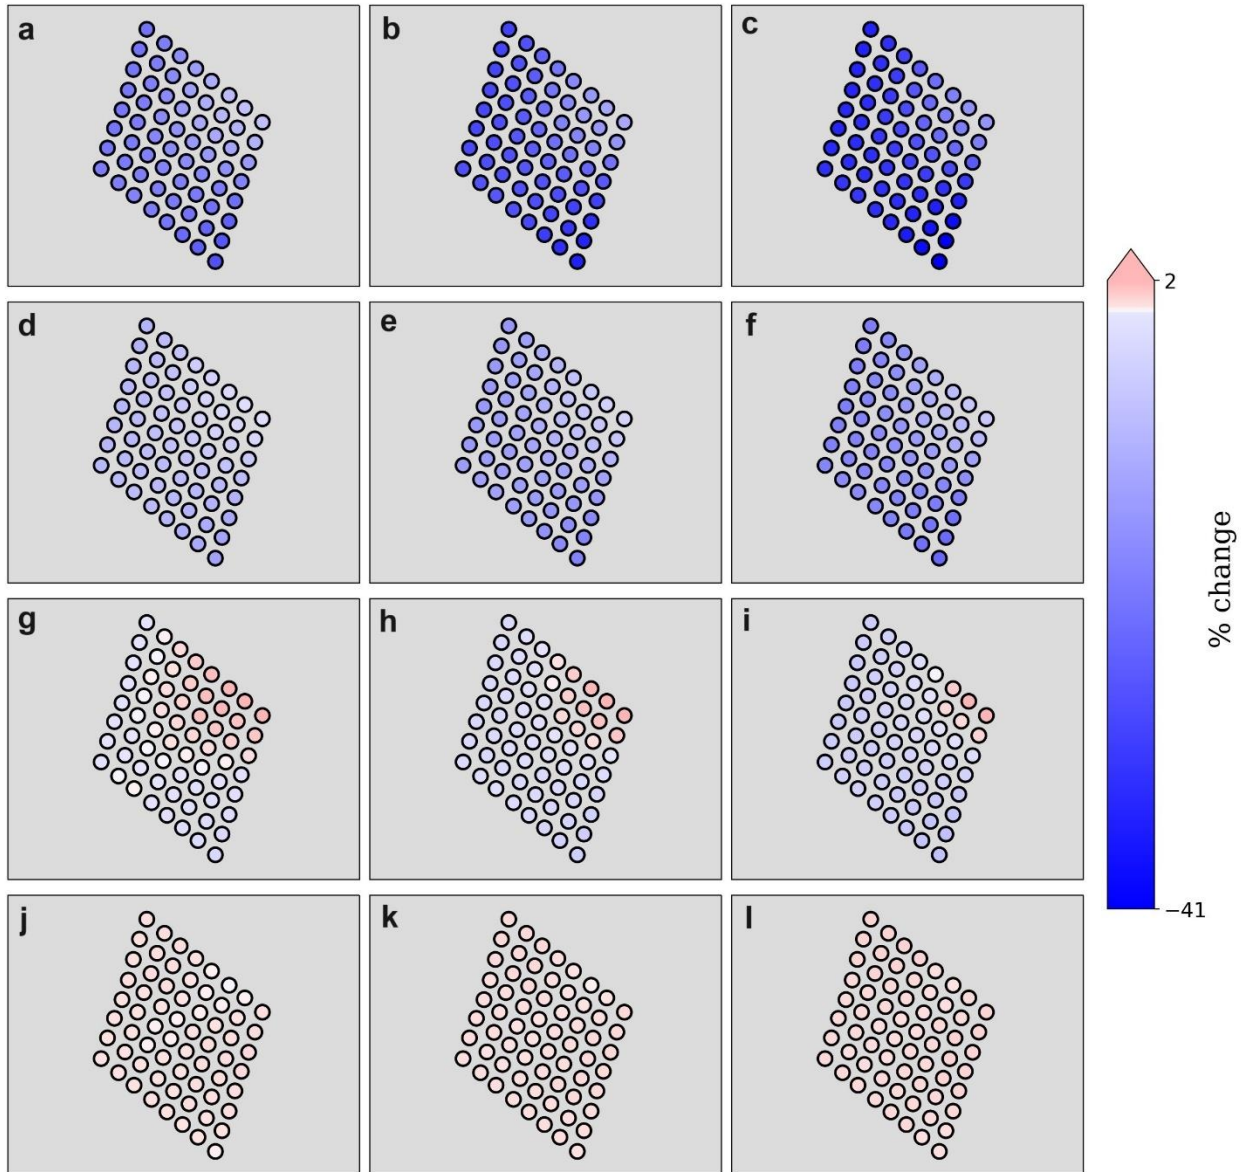


Figure 41: Percent change in spectral radiances given a transition from cloudy to clear skies, where columns correspond to cloud optical depths 1, 2, and 4, respectively. Each row corresponds to the four channels profiled, namely (a-c) 11.39 μm ; (d-f) 18.99 μm ; (g-i) 26.58 μm ; and (j-l) 40.08 μm .

In closing, Figure 42 presents mean I_λ during each cloud scenario and the subsequent clear sky sampling. On average, cloud dissipation results in reductions in spectral OLR that scale with cloud optical depth, particularly for transparent and partially transparent channels, as detailed earlier. For transparent channels, changes in OLR do not scale linearly, such that each doubling of optical thickness does not yield a twofold increase in mean emissions. Instead, the most

pronounced difference in mean OLR across our window channels corresponds to the difference between clear sky and an initial cloud of optical depth 1. This difference is greater, that is, than the differences between each cloud scenario. This holds to a more modest degree for partially transparent windows. Figure 42 also reveals that the peak in our window channels shifts to longer wavelengths for clear sky versus cloudy conditions. Conversely, the opposite is true for partially transparent channels, where peak radiances shift to shorter wavelengths for clear sky.




Figure 42: Mean spectral radiances (averaged over the 64 common pixels) for clear sky and each cloud scenario presented (optical depths 1, 2, and 4). Squares denote TIRS1 channel centers and missing wavelengths correspond to TIRS1 spectral gaps.

Alternatively, we can look at the mean radiative response to cloud dissipation by plotting the spectral differences explicitly, as shown in Figure 43. For every cloud optical depth, we represent the mean spectral OLR for the clear sky crossover minus mean spectral OLR for the cloudy sampling. Figure 43 verifies that the cloud radiative effect increases with optical depth. For our transparent channels, peak differences range from about 0.3 for optical depth 1 to over 0.6 $W m^{-2} sr^{-1} \mu m^{-1}$ for optical depth 4. Peak differences for partially transparent wavelengths are on the order of about 0.1 for optical depth 1 to 0.3 $W m^{-2} sr^{-1} \mu m^{-1}$ for optical depth 4.




Figure 43: Explicit differences in mean spectral radiances, computed by subtracting mean radiances during the clear sky crossover from the mean radiances when clouds of optical depths 1, 2 and 4 are present. As before, squares denote TIRS1 channel centers and missing wavelengths indicate spectral gaps.

4.3 Calibration case results

Our third case serves as a conceptual model for inter-calibration. During the mission, the PREFIRE CubeSats will achieve intersections on near-rapid timescales, forming the so-called short-term revisits characterized in section 3.4. By virtue of their tendency to sample static conditions, this thesis proposes a general strategy for using short-term revisits to identify instrument drift. We now illustrate this strategy for the hypothetical intersection described in section 2.3.3.

Figure 44 shows ΔI_λ for each scene of TIRS1 compared to the eight scenes of TIRS2, taken with respect to select wavelengths between 6 and 40 μm . Recall that underlying conditions are effectively identical for either crossover, such that ΔI_λ results solely from the combined differences in channel centers and spectral response functions. What emerges is therefore the baseline spectral offset across instruments per scene, referenced to discrete wavelengths. For this exercise, we neglect specific values of ΔI_λ and instead focus more broadly on the degree of consistency across

scenes. Granted the static and homogeneous scene observed, as described in section 2.3.3, we encounter a specific sequence of ΔI_λ that roughly recurs for each scene.

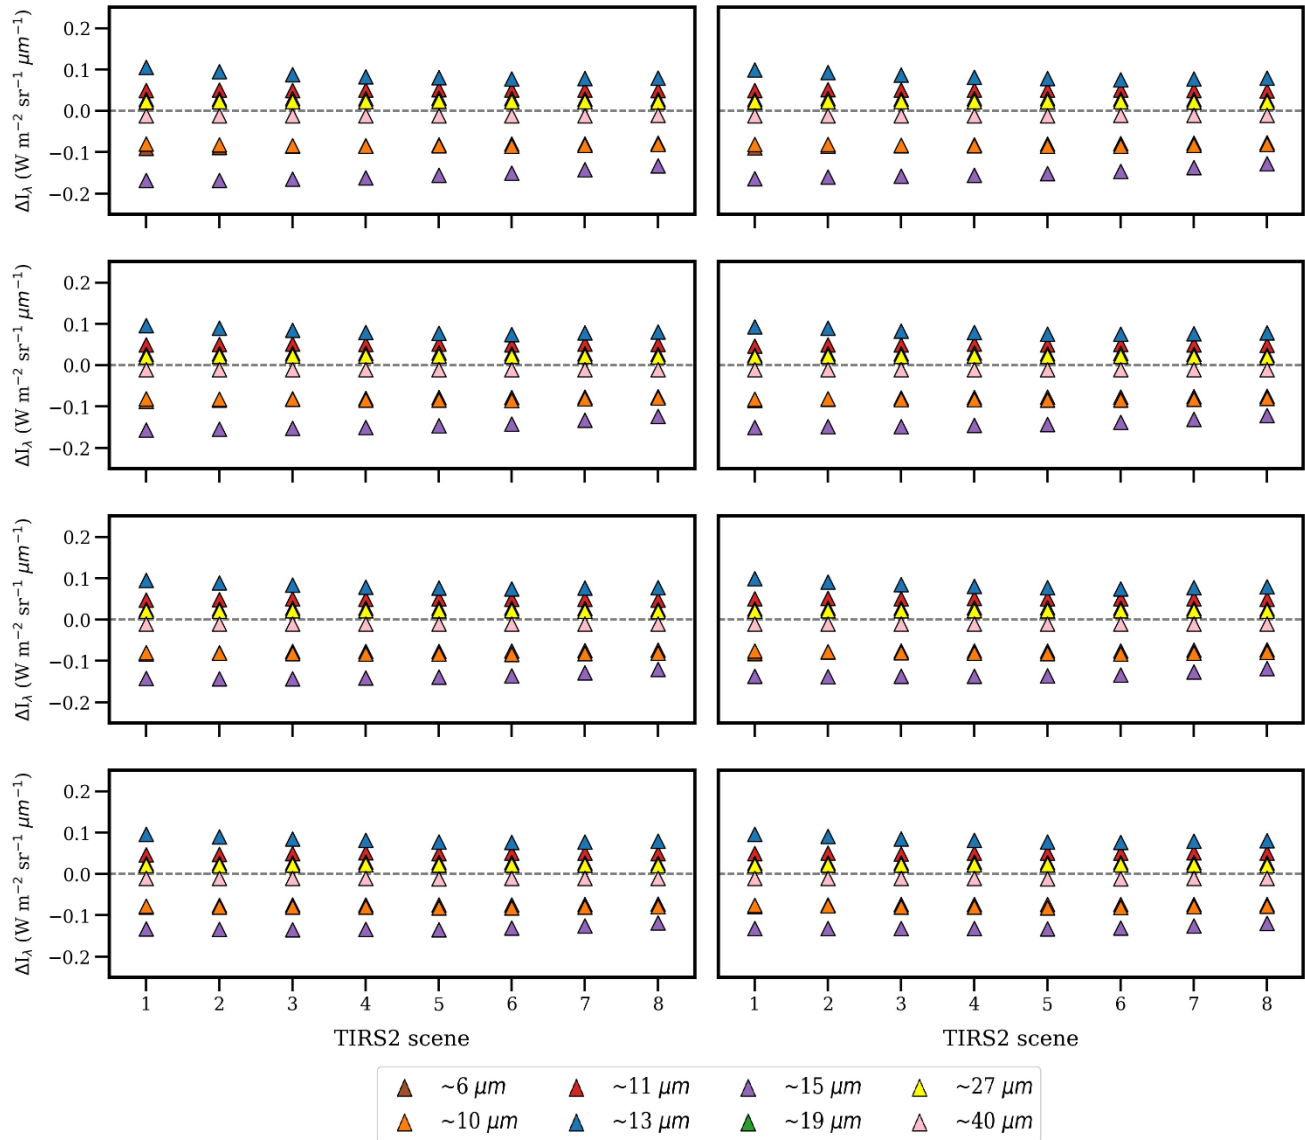


Figure 44: Nominal differences in spectral radiances between TIRS1 and TIRS2 for our calibration case. Each panel represents the spectral differences between a single TIRS1 scene and each of the eight TIRS2 scenes. Triangles indicate the target wavelengths for which differences are computed. Note that the actual wavelengths for either sensor are not strictly identical to those targeted, such that the differences are calculated between the TIRS1 channels and the TIRS2 channels nearest the target wavelengths, which are not equal across sensors (hence the nonzero differences shown).

We can extend this analysis to every integer wavelength and plot the differences accordingly. This is shown for a single TIRS1 scene in Figure 45. It is obvious that ΔI_λ is greatest

near missing channels, denoted by the gold and gray shading. Variability across TIRS2 scenes is also greatest surrounding spectral gaps, where we note the largest disagreement near the second band of missing channels. Excepting these sensitive wavelengths, we generally find strict consistency in ΔI_λ for all TIRS2 scenes. In the context of our hypothetical intersection, we regard the ΔI_λ profile shown in Figure 45 as a reasonable representation of the baseline spectral offset between the corresponding TIRS1 scene and each scene of TIRS2. Similar results are found for the remaining TIRS1 scenes (not shown).




Figure 45: Nominal spectral differences between TIRS1 scene 1 and the eight scenes of TIRS2, taken with respect to every integer wavelength between 5 and 54 μm . As noted previously, differences are computed between the TIRS1 channels nearest the targeted wavelengths (i.e., integer wavelengths) and the TIRS2 channels nearest the targeted wavelengths. Shaded gold and gray regions indicate spectral gaps for TIRS1 and TIRS2, respectively.

Having established baseline conditions, this thesis proceeds with a demonstration of idealized drift. To do so, we apply a 10% increase in I_λ to the eighth scene of TIRS2 and then recalculate ΔI_λ , as outlined in section 2.3.3. Although it is not yet known how drift will present in each PREFIRE sensor, it is reasonable to assume that drift may narrowly impact a single scene given that each scene is physically independent. Figure 46 is the same as Figure 44 but with our applied enhancement. Clearly, when we restrict drift to a single scene, the affected scene is distinguished by its unique ΔI_λ sequence. In this case, ΔI_λ is negative for every wavelength, a

distinct feature of the impacted scene vis-à-vis the remaining seven. This "drift signature" holds for all TIRS1 scenes.




Figure 46: Differences in spectral radiances between TIRS1 and TIRS2 when idealized constant drift (10%) is applied to all channels of the eighth scene of TIRS2. As before, each panel represents differences between a single TIRS1 scene and every scene of TIRS2. Note that the y-axis has been modified compared to Figure 44 to adequately show the spectral differences in the drift-affected scene.

As before, we generate ΔI_λ for every integer wavelength but now include the idealized drift described previously. Figure 47 shows the associated profile for the same TIRS1 scene as Figure 45. Our results confirm that the "drift signature" emerges for most wavelengths, with a clear

distinction in ΔI_λ between the impacted scene and the nominal scenes. This distinction is most prominent for wavelengths falling between the first and third set of missing channels, with a subtle drift signal persisting at longer wavelengths.




Figure 47: The same as Figure 45 but with constant 10% drift applied to all channels of the eighth scene of TIRS2. As before, shaded gold and gray regions denote spectral gaps in TIRS1 and TIRS2, respectively.

Similarly, we can translate our preceding results into a more direct expression of drift. Recall that our proposed method infers changes in instrument calibration from the differences of the differences, where differences in the latter context refer to ΔI_λ . Having previously established baseline ΔI_λ for each scene, we then compute the differences between nominal ΔI_λ for TIRS2 scene 8 and ΔI_λ for the same scene after imposing drift. The resulting differences of the differences, denoted $\Delta(\Delta I_\lambda)$, are plotted in Figure 48. We ultimately find that $\Delta(\Delta I_\lambda)$ is everywhere nonzero, peaking near 15 μm with a maximum absolute value of about $0.15 \text{ W m}^{-2} \text{ sr}^{-1} \mu\text{m}^{-1}$. In reality, the relationship between $\Delta(\Delta I_\lambda)$ and drift is likely far more complex than our idealized results imply. This thesis therefore does not quantify this relationship. Instead, we sought to broadly demonstrate how this relationship can be measured via short-term revisits and leveraged accordingly.




Figure 48: Differences in spectral differences after constant 10% drift is applied to all channels of the eighth scene of TIRS2. Differences are taken with respect to TIRS1 scene 1. As before, shaded gold and gray regions denote spectral gaps in TIRS1 and TIRS2, respectively.

For reference, we conclude this case with mean I_λ for either crossover, excluding missing channels. This is plotted in Figure 49. We note that, courtesy of the shaded regions denoting the maximum and minimum I_λ observed for either sampling, this intersection exhibits very little spectral variability, thereby confirming the relative homogeneity of the calibration region. More importantly, we observe that the I_λ profile for either crossover converges to roughly the same (incomplete) Planck function. Although we do not interpolate between spectral gaps in this thesis, doing so would yield a full Planck function peaking near 15 μm . This peak coincides with the maximum $\Delta(\Delta I_\lambda)$ noted earlier. In addition to inferring drift from $\Delta(\Delta I_\lambda)$, as previously detailed, the mission may derive value from generating mean I_λ for select short-term revisits and confirming whether their implicit Planck functions remain reasonably congruent.




Figure 49: Mean spectral radiances for all common pixels from either crossover. As before, squares denote the central wavelengths for TIRS1 (orange) and TIRS2 (purple). Minimum and maximum spectral radiances are indicated for TIRS1 and TIRS2 pixels by the shaded orange and purple regions, respectively. Missing channels correspond to spectral gaps in either sensor. We note that the resulting spectra form (incomplete) Planck functions.

5. Discussion

This thesis underscores the novelty of PREFIRE, which will rely in part on the frequent resampling the mission's CubeSats afford. By simulating representative orbits and approximating where and when their ground tracks intersect, we establish a preliminary record of PREFIRE resampling, referred to as intersections. Our efforts motivate the so-called "fourth-scene method," a technique that approximates the location of intersections by strategically applying geospatial distance calculations. This straightforward approach grants an initial framework for compiling resampling statistics sans robust computations. Additionally, the underlying premise informs the process with which we characterize and extract co-located pixels from intersection regions, a process exercised in our conceptual cases. The procedures prescribed in this thesis therefore empower us to estimate mission resampling before launch and may help guide future CubeSat missions tasked with similar objectives.

Recall that PREFIRE resampling takes three forms: SAT1 self-intersections, SAT2 self-intersections and SAT1-SAT2 intersections. Having compiled numerical, spatial and temporal statistics for each, we now draw on our results to expound the conclusions that follow.

For the orbit parameters prescribed in this thesis, we would reasonably expect hundreds of intersections per day for each intersection type, rounding out to nearly 1,000 combined daily intersections. Approximately one-quarter constitute self-intersections for a single CubeSat, while SAT1-SAT2 intersections account for the remaining half. These ratios are preserved when we restrict our sample of intersections to include only the polar subset. Taken together, we conclude that intersection frequency nearly quadruples with the addition of a second CubeSat, which alternatively suggests that we would forfeit about 75% of intersections in the event that one satellite becomes inoperable. Note that the number of self-intersections is practically identical regardless of CubeSat, such that any statistical difference between their numerical means is unremarkable. Our results therefore imply that an altitude difference on the order of 10 km has little impact on self-intersection frequency, particularly for sub-daily timescales.

We encounter broad agreement not only in the numerical distribution of the self-intersections of each CubeSat but also in their spatial and temporal profiles. For either satellite, we anticipate about three-quarters of their respective self-intersections will be centered in high latitudes, with the remaining quarter evenly split between low- and mid-latitude sectors. Two critical findings emerge for this intersection type: First, self-intersections mutually coalesce in discrete, time-invariant latitude bands, forming eight bands in either hemisphere. We therefore expect latitudinal gaps in coverage. Second, the elapsed time between crossovers decreases monotonically as intersections move poleward, ranging from roughly 12 hours near the equator to the theoretical lower bound of 1.5 hours near the maximum sampling latitude. We note that

latitudinal coverage is not strictly identical between CubeSats, with self-intersections of the lower CubeSat exhibiting marginal poleward displacement relative to those of the higher CubeSat. This implies that altitude exerts a modest impact on precise intersection location and may prompt future missions to consider a lower (higher) altitude to shift intersections poleward (equatorward). The noted differences remain slight, however, and the underlying distribution prevails regardless of altitude, suggesting that self-resampling will tend toward the spatial and temporal patterns shown.

Ultimately, we expect to achieve greater spatial and temporal coverage with intersections between CubeSats compared to self-intersections. Under the simulated conditions imposed in this thesis, the CubeSats were shown to resample one another in a cyclic, time-varying fashion. We emphasize that the simulated SAT1-SAT2 intersections exhibit staggered modes, such that they oscillate meridionally over time. This sharply contrasts with simulated self-intersections. Similarly, the elapsed time between SAT1-SAT2 intersections did not decrease monotonically with increasing latitude. Instead, the simulated SAT1-SAT2 intersections oscillated between various time differences, attaining the greatest range near the maximum sampling latitude, where intersections on sub-orbital timescales approach intersections with time differences up to our imposed upper bound of 12 hours. This oscillatory behavior is a distinguishing feature of the simulated SAT1-SAT2 intersections and a fundamental conclusion of this research. Additionally, between roughly 72° and 78° in either hemisphere, the synthetic orbits yielded "short-term revisits," which denote SAT1-SAT2 intersections with rapid elapsed time. Short-term revisits are ideal candidates for inter-calibration between CubeSats and will be used accordingly during the mission. Excepting where simulated self-intersections and SAT1-SAT2 intersections are different, we also note where they are similar: Just as for self-intersections, we anticipate that a majority of SAT1-SAT2 intersections will be centered in high latitudes. In this case, more than 75% of

intersections are expected at or poleward of 60° . The remaining latitudes receive variable coverage, with the resulting distribution exhibiting a monotonic increase in frequency with poleward displacement.

Our discussion has as yet neglected the implicit contributions of altitude on the distribution of resampling between CubeSats. We note that cyclic resampling corresponds to greater spatial coverage and thus makes time-varying coverage an appealing feature for future CubeSat missions. Previous analyses undertaken by the author suggest that cyclic coverage depends on a nonzero difference in altitude. During highly idealized assessments in which both CubeSats were prescribed identical altitudes of 525 km, resampling between CubeSats notably did not yield cyclic coverage and instead produced discrete latitude banding. Results from this previous assessment are shown in Figure 50. To the best of the author's knowledge, this discovery is unique and remains paramount to ongoing discussions concerning CubeSat utility, informing us that resampling may be optimized by staggering the altitudes of two or more CubeSats. We may therefore take the case of identical altitudes as representing the lower bound on cross-satellite resampling. Even in this context, however, we find greater spatial and temporal coverage for SAT1-SAT2 intersections compared to self-intersections. This is apparent by more abundant latitude banding for SAT1-SAT2 resampling as well as the familiar staggering of time differences, particularly in high latitudes. Future work may endeavor a robust comparison of SAT1-SAT2 intersections for different altitude scenarios.

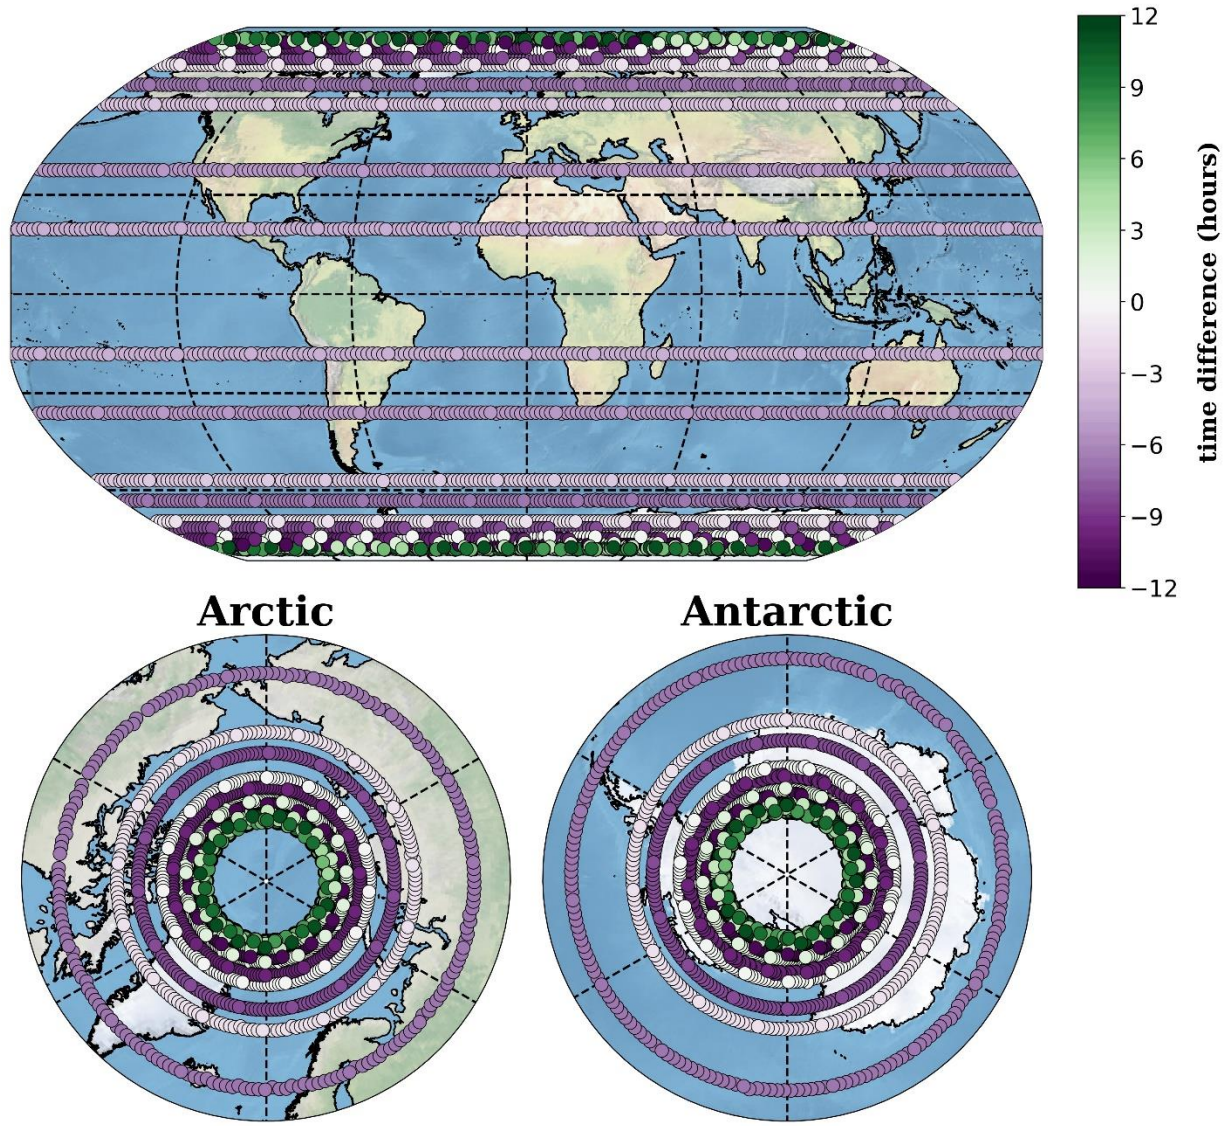


Figure 50: Centers of all SAT1-SAT2 intersections for a previous (unpublished) analysis conducted by the author in which both CubeSats were prescribed identical altitudes equal to 525 km.

Our conceptual cases further underscore the observational potential of PREFIRE intersections as well as their foreseeable challenges. The first case, an idealized Arctic heat wave sampled by a simulated coastal SAT2 self-intersection, features three distinct polar surfaces and constitutes the longest and most complex of the three simulated intersections presented. Observed changes between crossovers include variable rates of warming, moisture fluxes in either direction, and melt of snow-free ice. Our results confirm that TIRS will have sufficient sensitivity to

distinguish sea ice from ocean, with maximum sensitivity corresponding to channels between 11 and 12 μm . Accordingly, we conclude that well-defined melt or freeze events will generally exhibit distinct spectral signatures when sampled by PREFIRE. Uncertainty remains, however, concerning the spectral response for partial melt or freeze events. Additionally, since we imposed a transition from snow-free ice to ocean, the skill with which TIRS can distinguish ocean from snow-covered ice remains undefined. For transparent channels, which are naturally associated with maximum TIRS sensitivity to surface type changes, emissivities of medium and fine snow are nearly identical to ocean emissivities. Consequently, the spectral signature of ice may become ambiguous when snow is present, indicating that snow-covered ice could present unique challenges in the way of observing melt or freeze events via PREFIRE intersections. Neglecting transparent channels, we must also contend with observational challenges in the FIR, where spectral OLR becomes partially or fully decoupled from surface processes, pending scene conditions. In the Arctic heat wave case presented, sea ice melt is accompanied by sensible heating and variable changes in integrated moisture across the vicinity, making it impossible absent more advanced analyses to quantify OLR sensitivity for a single variable. Such analyses, while exceeding the scope of this research, may be necessary to fully constrain AGHE and will likely motivate future work. For now, we conclude that care is required when interpreting highly variable, heterogeneous scenes using PREFIRE resampling. Future users of PREFIRE data may find utility in targeting isolated pixels and analyzing them apart from the broader resampling region (i.e., the intersection polygon).

In our second case, we consider a relatively short-term and fairly homogeneous simulated SAT1 self-intersection over the backdrop of wintertime East Antarctica. Recall that we simulated three cloud dissipation scenarios in the resampling vicinity, which we achieved by imposing and

then removing a uniform ice cloud during the first and second crossover, respectively. Cloud optical depth was doubled for each iteration, resulting in optical depths 1, 2 and 4. Percent differences confirm that, given the extreme cold, clouds generally have a positive radiative effect in transparent channels, and their impact increases with optical depth. Despite a cloud top height near about 400 hPa, the surface of interior Antarctica during austral winter remains colder by contrast, resulting in reduced OLR when a cloud is absent. We also note a positive cloud radiative effect for partially transparent channels, implying that weighting functions still peak near the surface for some FIR frequencies. This is consistent with the arid conditions commonly observed during polar winter. And yet this holds only to a point: A small but opposite effect was observed for opaque channels in the FIR. For example, cloud dissipation induced a marginal *increase* in spectral OLR compared to cloudy conditions for the channel centered near 40 μm . This suggests that, compared to mid-infrared spectra, cloud radiative effects may differ for far-infrared PREFIRE channels, for which weighting functions peak some distance from the surface. In particular, we might expect cloud dissipation to enhance spectral OLR in at least some FIR channels under a range of common polar conditions, an expectation that will remain pending until real PREFIRE data are examined.

Finally, the third case introduces and conceptually demonstrates a technique for inter-calibration. Such a protocol remains necessary to identify instrument drift and is therefore integral to the mission. Our proposed method leverages the unique feature of cross-satellite resampling we have herein referred to as short-term revisits, or intersections on timescales sufficiently small so as to ensure conditions remain virtually static between crossovers. Owing to unique channel centers and spectral response functions between instruments, PREFIRE TIRS will inevitably register different spectra even in the event that they sample identical conditions. Further, since the

spectral response functions not only vary by sensor but also by scene, measurements will also differ across each of the eight scenes. Thus, we claim that a reasonable first step in inter-calibration includes computing simple differences in OLR between the eight scenes of TIRS1 vis-à-vis the eight scenes of TIRS2. As a demonstrative exercise, we illustrate this procedure by first selecting a simulated rapid SAT1-SAT2 intersection for which practically identical conditions have been confirmed and then calculating OLR differences for the channels nearest every integer wavelength for either TIRS. After establishing the baseline differences in spectral OLR, we urge that drift may be identified when differences in the differences are nonzero. When we impose idealized drift by artificially increasing spectral radiances for one scene of TIRS2, as undertaken in our conceptual exercise, we observe a corresponding nonzero difference in the differences taken with respect to the affected scene. In this highly idealized scenario, a constant 10% drift produces differences in the differences with peak absolute values of approximately $0.2 \text{ W m}^{-2} \text{ sr}^{-1} \text{ um}^{-1}$. Alternatively, mean spectral radiances for the "common pixels" of each crossover can be shown to yield partial Planck functions, one for each TIRS, and we would expect them to converge under most circumstances if the two sensors are inter-calibrated. This may supplement the differences of the differences calculations as a form of validating one TIRS against the other. We caution that, in reality, inter-calibration will assuredly present greater challenges than those accounted for in this thesis. Nonetheless, the fundamental premise described earlier – that is, inferring sensor drift from differences in the baseline differences over time – will help guide PREFIRE calibration and, we hope, future generations of CubeSat missions.

6. References

- Barkstrom, B. R., E. F. Harrison, G. L. Smith, and R. D. Cess, 1989: Results from the Earth Radiation Budget Experiment (ERBE). *Adv. Space. Res.*, 9, 7, 75-82, doi:10.1016/0273-1177(89)90146-4.
- Blackwell, W. J., and Coauthors, 2018: An overview of the TROPICS NASA Earth venture mission. *Quart. J. Roy. Meteorological Soc.*, 144, 1, 16-26, <https://doi.org/10.1002/qj.3290>
- Buckhard, C., and S. Weston, 2021: The Evolution of CubeSat Spacecraft Platforms. AVT-336 Specialists' Meeting, virtual, NATO Science and Technology Organization, https://ntrs.nasa.gov/api/citations/20210021348/downloads/Platforms_paper_final.pdf
- Cao, C., M. Weinreb, and H. Xu, 2004: Predicting Simultaneous Nadir Overpasses among Polar-Orbiting Meteorological Satellites for the Intersatellite Calibration of Radiometers. *J. Atmos. Oceanic Technol.*, 21, 4, 537-542, [https://doi.org/10.1175/1520-0426\(2004\)021<0537:PSNOAP>2.0.CO;2](https://doi.org/10.1175/1520-0426(2004)021<0537:PSNOAP>2.0.CO;2)
- Chen, X., X. Huang, and M. G. Flanner, 2014: Sensitivity of modeled far-IR radiation budgets in polar continents to treatments of snow surface and ice cloud radiative properties, *Geophys. Res. Lett.*, 41, 6530–6537, doi:10.1002/2014GL061216.
- Clough, S. A., F. X. Kneizys, and R. W. Davies, 1989: Line shape and the water vapor continuum. *Atmospheric Research*, 23, 3-4, 229–241. doi:10.1016/0169-8095(89)90020-3
- Davis, G. R., 1993: The far infrared continuum absorption of water vapour. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 50, 6, 673-694, [https://doi.org/10.1016/0022-4073\(93\)90034-F](https://doi.org/10.1016/0022-4073(93)90034-F)

- Drouin, B. J., B. Kahn, B. Lim, A. Merrelli, E. Nelson, G. Quinn, F. Nagle, and T. L'Ecuyer, 2022: Orbital Trade Study for the PREFIRE Mission, 2022 IEEE Aerospace Conference, Big Sky, MT, USA, 1-7, doi: 10.1109/AERO53065.2022.9843312.
- Feldman, D. R., W.D. Collins, R. Pincus, X. Huang, and X. Chen, 2014: Far-infrared surface emissivity and climate. *Proceedings of the National Academy of Sciences*, 111, 46, 16297–16302. <https://doi.org/10.1073/pnas.1413640111>
- Gill, E., P. Sundaramoorthy, J. Bouwmeester, B. Zandbergen, and R. Reinhard, 2010: Formation flying within a constellation of nano-satellites: The QB50 mission. 6th International Workshop on Satellite Constellation and Formation Flying, Taipei, Taiwan, International Astronautical Federation, <http://resolver.tudelft.nl/uuid:59a8ba1d-070d-4879-a048-eee208d04983>
- Glumb, R., J. Predina, and J. Fennelly, 2003: Overview of the cross-track infrared sounder (CrIS). *Fourier Transform Spectroscopy*, 84, FMC4-1, <https://doi.org/10.1364/FTS.2003.FMC4>.
- Hanel, R. A., B.J. Conrath, V. G. Kunde, C. Prabhakara, I. Revah, V. V. Salomonson, and G. Woford, 1972: The Nimbus 4 infrared spectroscopy experiment: 1. Calibrated thermal emission spectra. *J. Geophys. Res.*, 77, 15, 2629– 2641, doi:10.1029/JC077i015p02629.
- Harries, J., B. Carli, R. Rizzi, C. Serio, M. Mlynchak, L. Palchetti, T. Maestri, H. Brindley, and G. Masiello, 2008: The far-infrared Earth. *Rev. Geophys.*, 46, RG4004, doi:10.1029/2007RG000233.
- Hersbach, H., and Coauthors, 2020: The ERA5 global reanalysis. *Quart. J. Roy. Meteor. Soc.*, 146, 1999-2049, <https://doi.org/10.1002/qj.3803>
- Huang, X., X. Chen, G. L. Potter, L. Oreopoulos, J. N. S. Cole, D. Lee, and N. G. Loeb, 2014: A Global Climatology of Outgoing Longwave Spectral Cloud Radiative Effect and

- Associated Effective Cloud Properties. *J. Climate*, 27, 9, 7475-7492, <https://doi.org/10.1175/JCLI-D-13-00663.1>
- Huang, X., X. Chen, D.K. Zhou, and X. Liu, 2016: An Observationally Based Global Band-by-Band Surface Emissivity Dataset for Climate and Weather Simulations. *J. Atmos. Sci.*, 73, 9, 3541-3555, <https://doi.org/10.1175/JAS-D-15-0355.1>.
- Jordan, P., V. Blaes, L. Roszman, and J. Cooley, 1993: Analysis of the Effects of Mean Local Node-Crossing Time on the Evolution of Sun-Synchronous Orbits. *Flight Mechanics/Estimation Theory Symposium 1992*, Greenbelt, MD, NASA Goddard Space Flight Center, 177-190, <https://ntrs.nasa.gov/api/citations/19930015505/downloads/19930015505.pdf>
- Justice, C. O., and Coauthors, 1998: The Moderate Resolution Imaging Spectroradiometer (MODIS): land remote sensing for global change research. *IEEE Transactions on Geoscience and Remote Sensing*, 36, 4, 1228-1249, doi: 10.1109/36.701075
- Kempe, V., D. Oertel, R. Schuster, H. Becker-Ross, and H. Jahn, 1980: Absolute IR-spectra from the measurement of Fourier-spectrometers aboard Meteor 25 and d28. *Acta Astronaut.*, 7, 1403-1416, [https://doi.org/10.1016/0096-5765\(80\)90015-6](https://doi.org/10.1016/0096-5765(80)90015-6)
- L'Ecuyer, T. S., and Coauthors, 2021: The Polar Radiant Energy in the Far Infrared Experiment: A New Perspective on Polar Longwave Energy Exchanges. *Bull. Amer. Meteor. Soc.*, 102, 7, E1431-E1449, <https://doi.org/10.1175/BAMS-D-20-0155.1>.
- Lee, Z. T., 2017: CubeSat constellation implementation and management using differential drag. M.S. thesis, Dept. of Aeronautics and Astronautics, Massachusetts Institute of Technology, 117 pp., <http://hdl.handle.net/1721.1/112471>

- Leslie, R. V., M. DiLiberto, and S. Hazard, 2023: User Guide for TROPICS Data Products, https://ntrs.nasa.gov/api/citations/20210018318/downloads/TROPICS_UserGuide_base_Sept2023.pdf
- Liu, X., W. L. Smith, D. K. Zhou, and A. Larar, 2006: Principal component-based radiative transfer model for hyperspectral sensors: Theoretical concept. *Appl. Opt.*, 45, 201–209, <https://doi.org/10.1364/AO.45.000201>.
- Malan, D. F., K. Wiid, H. Burger, L. Visagie, and W. H. Steyn, 2017: The Development of "nSight-1" - Earth Observation and Science in 2U. 31st Annual AIAA/USU Conference on Small Satellites, Logan, UT, American Institute of Aeronautics and Astronautics/Utah State University, <https://digitalcommons.usu.edu/cgi/viewcontent.cgi?article=3675&context=smallsat>
- Miller, N. B., A. Merrelli, T. S. L'Ecuyer, and B. J. Drouin, 2023: Simulated Clear sky Water Vapor and Temperature Retrievals from PREFIRE Measurements. *J. Atmos. Oceanic Technol.*, 40, 6, 645-659, <https://doi.org/10.1175/jtech-d-22-0128.1>
- Pagano, T. S., M. T. Chahine, and E. J. Fetzer, 2010: The Atmospheric Infrared Sounder (AIRS) on the NASA Aqua Spacecraft: a general remote sensing tool for understanding atmospheric structure, dynamics, and composition. *Proc. SPIE 7827, Remote Sensing of Clouds and the Atmosphere XV*, 78270B, <https://doi.org/10.1117/12.865335>.
- Palchetti, L., and Coauthors, 2020: FORUM: Unique far-infrared satellite observations to better understand how earth radiates energy to space. *Bulletin of the American Meteorological Society*, 101, 12, E2030–E2046. <https://doi.org/10.1175/bams-d-19-0322.1>
- Puig-Suari, J., C. Turner, and W. Ahlgren, 2001: Development of the standard CubeSat deployer and a CubeSat class PicoSatellite. 2001 IEEE Aerospace Conference Proceedings, *Cat.*

- No.01TH8542*, Big Sky, MT, Institute of Electrical and Electronics Engineers,
<https://doi.org/10.1109/AERO.2001.931726>
- Rees, W. G., 1992: Orbital subcycles for Earth remote sensing satellites. *International Journal of Remote Sensing*, 13, 5, 825-933, doi:10.1080/01431169208904156
- Rew, R. K., and G. P. Davis, 1990: NetCDF: An Interface for Scientific Data Access. *IEEE Computer Graphics and Applications*, 10, 4, 76-82.
- Saeed, N., A. Elzanaty, H. Almorad, H. Dahrouj, T. Y. Al-Naffouri, and M.-S. Alouini, 2020: CubeSat Communications: Recent Advances and Future Challenges. *IEEE Communications Surveys & Tutorials*, 22, 3, 1839-1862, doi:10.1109/COMST.2020.2990499
- Schmugge, T., A. French, J. C. Ritchie, A. Rango, and H. Pelgrum, 2002: Temperature and emissivity separation from multispectral thermal infrared observations. *Remote Sens. Environ.*, 79, 2-3, 189-198, [https://doi.org/10.1016/S0034-4257\(01\)00272-3](https://doi.org/10.1016/S0034-4257(01)00272-3)
- Shiroma, W. A., L. K. Martin, J. M. Akagi, J. T. Akagi, B. L. Wolfe, B. A. Fewell, and A. T. Ohta, 2011: CubeSats: A bright future for nanosatellites. *Open Engineering*, 1, 1, 9-15, <https://doi.org/10.2478/s13531-011-0007-8>
- Silber, I., and Coauthors, 2019: Persistent Supercooled Drizzle at Temperatures Below -25 C Observed at McMurdo Station, Antarctica. *JGR Atmosphere*, 124, 20, 10878-10895, <https://doi.org/10.1029/2019JD030882>
- Smith, G. L., K. J. Priestley, N. G. Loeb, B. A. Wielicki, T. P. Charlock, P. Minnis, D. R. Doelling, and D. A. Rutan, 2011: Clouds and Earth Radiant Energy System (CERES), a review: Past, present and future. *Adv. in Space Res.*, 48, 2, 254-263, doi:10.1016/j.asr.2011.03.009.

Wielicki, B. A., B. R. Barkstrom, E. F. Harrison, R. B. Lee III, G. L. Smith, and J. E. Cooper,
1996: Clouds and the Earth's Radiant Energy System (CERES): An Earth Observing
System Experiment. Bull. Amer. Meteor. Soc., 77, 5, 853-868,
[https://doi.org/10.1175/1520-0477\(1996\)077<0853:CATERE>2.0.CO;2](https://doi.org/10.1175/1520-0477(1996)077<0853:CATERE>2.0.CO;2)