Automated fleet commissioning workflows at Planet

Jose Angel Gutierrez Ahumada , Kelsey Doerksen , Stefan Zeller
Planet Labs Inc., 645 Harrison St floor 4, San Francisco, CA 94107; (415) 829-3313
jose.gutierrez@planet.com

ABSTRACT

In order to maintain the world's largest Earth Observation satellite constellation and its goal of providing daily Earth coverage, Planet procures regular launch opportunities, each including up to several dozens of new satellites. Commissioning these batches of satellites, known at Planet as "flocks", from deployment until imaging is done by a small team of operators and depends heavily on a number of automated workflows developed at Planet. Bringing all satellites to production in a safe and timely manner has a direct link to the productivity of the entire fleet. The main bottleneck during commissioning is ground station access time with the satellites, which are concentrated close to one another on-orbit after launch. To circumvent this, it is key to develop a solution that optimizes the number of contacts needed per satellite and distributes the available accesses effectively. Through the power of automation, Planet has been able to develop a commissioning process that can support a large number of satellites at any given time progressing in parallel through the different phases of commissioning, creating a fast, efficient, and safe commissioning workflow. This paper will present an overview of how Planet commissions new SuperDove satellites, the automated solutions and improvements implemented over time, and on-orbit results and lessons learned from the most recent commissioning campaigns: Vega VV16 (Flock 4V), Electron "In Focus" (Flock 4EP) and SpaceX Transporter-1 (Flock 4S).

LIST OF ABBREVIATIONS

ADCS	Attitude Determination and Control System
CPU	Central Processing Unit
EMI	Electro-Magnetic Interference
FPGA	Field-Programmable Gate Array
HSD	High Speed Downlink
HSU	High Speed Uplink
LST	Low Speed Transceiver
RFI	Radio-Frequency Interference
SSD	Solid-State Drive
TT&C	Telemetry, Tracking and Command

INTRODUCTION

Founded in 2010 by a team of ex-NASA scientists, Planet's mission is dedicated to imaging the entire Earth everyday, thereby enabling global change to be visible, accessible, and actionable. Since its inception, Planet has successfully launched and operated over 400 "Dove" Earth-imaging satellites, resulting in Planet's success of operating the world's largest Earth Observation satellite constellation.

This paper is structured as follows: first, it gives an overview of the latest Dove design and a brief history of Planet's launch and commissioning campaigns, followed by an overview of commissioning itself and the phases within the process. Key challenges to commissioning will be highlighted, and the Automation in Commissioning section will detail the tools that have been developed to automate the process. A comparative analysis between some of the most recent launches will help to quantify the impact of automation on the performance of commissioning and key takeaways from the work will be highlighted.

SuperDove Design

Planet's latest Dove design is the medium-resolution SuperDove, which is a significant step forward from the original Dove model with improved multispectral imaging capability that enables new markets, downstream applications and value-added services. Figure 1 shows a rendering of the SuperDove satellite in its deployed configuration.

The SuperDove design is based on the 3U+ CubeSat form factor. Its payload is a telescope and CCD camera which are used for capturing imagery over Earth's landmass. Two deployable solar arrays generate the power needed by the spacecraft. There are also three types of communication radios used by the SuperDove; the UHF Low-Speed Transceiver (LST), S-band High Speed Uplink (HSU) receiver and an X-Band High Speed Downlink (HSD) transmitter.

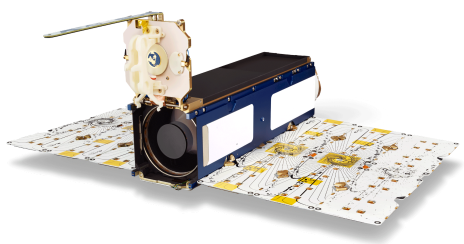

Figure 1: Planet SuperDove

The SuperDove satellites also have a few different on-board processors: the always-on main on-board computer, called the "superconductor", in charge of attitude control, power management and thermal control; and the more powerful Central Processing Unit (CPU) and Field-Programmable Gate Array (FPGA), needed for commanding the high-speed X-band radio and the optical payload.

Finally, the main components of the Attitude Determination and Control System (ADCS) are its sensors (magnetometers, gyroscopes, horizon sensor, photodiodes and star camera) and actuators (reaction wheels and magnetorquers).

First launched in 2019, the main focus of this paper will be the SuperDove commissioning process, although some mention of the Dove campaigns is included in the results section to complement the comparative analysis.

Planet Dove Launch and Commissioning History

To date, Planet has had 30 successful launches over 8 years, utilizing commercial launch services from a broad range of launch providers across the globe. The unique feature of Planet's strategy in the procurement of regular launches allows Planet to maintain the on-orbit capacity and coverage requirements to successfully image the Earth's entire land mass on a daily cadence. Launch size is dependent on a number of factors, and with the rise of rideshare programs, available capacity for Doves on-board can dictate the relative flock size. Planet has had over 450 satellite deployments from 10 different rocket types, in 10 different sites, in 7 countries across the globe; a testament to the agile aerospace techniques used at Planet. Table 1 summarizes the history of Planet's successful Dove launches since 2013. Note that this timeline includes only Dove launches, and does not encompass the SkySat or RapidEye satellites, as the focus of this paper is specific to Dove commissioning.

Table 1: Dove Launch History

Name	Vehicle	Launch Date	# of Sats
Dove 2	Soyuz 2.1b	4/19/2013	1
Dove 1	Antares	4/21/2013	1
Dove 3/4	Dnepr	11/21/2013	2
Flock 1	Antares	1/9/2014	28
Flock 1c	Dnepr	6/19/2014	11
Flock 1b	Antares	7/13/2014	28
Flock 1dp	Falcon 9	1/10/2015	2
Flock 1e	Falcon 9	4/13/2015	14
Flock 2b	H-11B	8/19/2015	14
Flock 2e	Atlas V	12/6/2015	14
Flock 2ep	Atlas V	3/23/2016	20
Flock 2p	PSLV	6/22/2016	20
Flock 3p	PSLV	2/15/2017	88
Flock 2k	Soyuz 2.1a	7/14/2017	48
Flock 3m	Minotaur-C	10/31/2017	4
Flock 3pp	PSLV	1/12/2018	4
Flock 3r	PSLV	11/29/2018	16
Flock 3s	Falcon 9	2/3/2018	3
Flock 3k	Soyuz-2	12/27/2018	12
Flock 4a	PSLV	4/1/2019	20
Flock 4p	PSLV	11/27/2019	12
Flock 4v	Vega	9/3/2020	26
Flock 4ep	Electron	10/28/2020	9
Flock 4s	Falcon 9	1/24/2021	48

Planet's most recent and largest SuperDove launch was on January 24th, 2021, onboard the historic Transporter-1 launch. Transporter-1 was SpaceX's first dedicated rideshare mission, with a record-setting 143 spacecraft carried to orbit, the most spacecraft ever deployed on a single mission. Planet procured capacity on Transporter-1 to launch the 48 SuperDove satellites, forming Flock 4S. The deployment phase of the launch was broken into two parts; first, an immediate post-launch deployment of 40 SuperDoves, and second, a delayed deployment of the remaining 8 SuperDoves onboard D-Orbit's ION Deployer, wherein the satellites were deployed at T-plus 2 weeks after launch, with one satellite deployed every two days to assist with satellite phasing. Following the successful launch, the Missions

Operations team was tasked with commissioning the flock as quickly and safely as possible. The targeted commissioning length was 13 days for the initial 40 deployed satellites, with a total planned commissioning duration of one month.

COMMISSIONING AT PLANET

The main goal of commissioning is to bring the satellites from their initial state after deployment from the launcher all the way to nominal operations. In practice this includes:

- Establishing first contact with all satellites and performing initial basic health checks,
- Obtaining enough ranging points for initial orbit determination of all satellites,
- Using the ADCS system first for detumbling the satellites and then for 3-axis pointing,
- Deploying the solar arrays,
- Booting the on-board CPU and establishing high-speed X-band communications,
- Calibrating all sensors and actuators via specific calibration activities,
- Performing on-board software updates, and
- Enabling imaging and all relevant maintenance activities regularly scheduled on the production fleet.

For more clarity, within the context of this text the term “bus commissioning” is used to encompass the commissioning phase from first contact with the satellites until deployment of the solar arrays, including a checkout of the most basic components of the satellite bus. The gateway to the next phase of commissioning is the checkout of the on-board CPU and the X-band HSD radio. Following this, the ADCS calibration occurs next, which encompasses a pair of on-orbit maneuvers used to calibrate ADCS sensors and actuators before enabling imaging. If these different phases are successful, the satellites move on to "first light", or the first acquisition and downlink of Earth-facing imagery after launch. From a satellite operations point of view, the extent of commissioning reaches up to this point. The satellites are now regularly acquiring and downlinking imagery and can be treated as operational and in line with the rest of the production fleet. The calibration and validation of the actual payload (telescope and camera) is another necessary step before imagery from these new satellites is published to customers, however that process is outside of the scope of this paper.

Bus Commissioning

As mentioned, Bus Commissioning is the first active state of commissioning a satellite. The satellites are loaded into the launcher with all deployables (antennas, star camera, solar arrays) stowed, and with their power system and all electrical components inhibited. Upon ejection from the launcher, the power system switches on and allows power to flow from the solar arrays to the bus and batteries. The spacecraft's main on-board computer (superconductor) and the TT&C radios (UHF and S-band) power on within seconds. At this point, communications between ground stations and the satellites become possible.

The CubeSat deployers used for ejection typically introduce some rotation when deploying the satellites. The release of the antenna flap and star camera, which are deployed automatically, add to this rotation. The rotational speed reported by the gyroscopes at first contact with the ground can be up to 40 degrees per second.

Before establishing communications with ground, the satellites remain in a passive safe mode, with no active attitude control or on-board activity execution. In most instances the spacecraft is power positive in this configuration and could safely remain in this state indefinitely.

The first contact with a satellite involves a basic subsystem-level check-out and the collection of some initial telemetry data. The first action taken on-board is to start the ADCS system to stop the aforementioned rotation rate. Detumbling is achieved with the satellite’s magnetorquers following a simple “bdot” algorithm using the Earth’s magnetic field. Before activating the magnetorquers indefinitely, a sanity check is performed to make sure both sensors and actuators are functioning as expected. This is done through a so-called “Detumble-Toe-Dip” activity, wherein the satellite detumbles for a few minutes and the rotation rates before and after the activity are compared to confirm the satellite is indeed detumbling and not spinning up.

After the magnetorquers are checked out, they can remain active indefinitely. Once the gyroscope telemetry from the satellite shows it is fully detumbled, the reaction wheels are activated next to perform 3-axis control. At the same time, a “panel deploy” activity is scheduled. During this activity the deploy mechanism of the solar arrays is activated and the wings extend, finally exposing all the solar cells to the Sun. At this point, all deployables have been released and the satellite is power-positive and capable of TT&C communications and 3-axis attitude control, marking the end of the first major phase of commissioning.

CPU and X-band radio check-out

Now that the satellite is stable, the CPU, FPGA, and X-band radio can be checked out. This is done during the first high-speed contact with the satellite, using S-band for uplink communications and X-band for downlink. The on-board X-band radio is not omnidirectional, so for a successful high-speed downlink the spacecraft needs to orient itself towards the ground station and power on all the necessary subsystems: CPU, FPGA, and X-band transmitter.

Once all these components are confirmed to be working, any essential upgrades of the on-board CPU and FPGA are carried out if needed, and the spacecraft can move on to the next phase of commissioning: ADCS calibration.

ADCS Calibration

Before enabling Earth-facing imaging on the spacecraft, the sensors comprising the Attitude Determination and Control System are calibrated with data from two different on-orbit maneuvers. A brief overview of these activities is given next.

Calibration Maneuver 1 (C1)

The first calibration maneuver, C1, is used to calibrate the photodiodes and magnetometers. It spans two orbits or approximately three and a half hours, and it is used to collect high-rate telemetry from both a slow controlled tumble and a number of orientations with respect to the Sun. The main electronic components of the satellite, namely the CPU, FPGA, X-band radio, and camera, are also powered on to determine the effect of electro-magnetic interference (EMI) on the attitude sensors. High-frequency telemetry data from the ADCS sensors as well as component voltages and currents are collected by the CPU and afterwards downlinked over X-band.

Calibration Maneuver 2 (C2)

The second calibration maneuver, C2, generates the data required to calibrate the star camera, horizon sensor, and gyroscopes. C2 takes place during the eclipse period of one orbit, which is roughly thirty minutes long, and uses the spacecraft main camera along with the star and horizon cameras simultaneously to capture a known constellation in the sky. The images are downlinked and compared to determine the offset of the ADCS sensors and the star camera.

Imaging and Nominal Operations

After ADCS calibration is completed, the acquisition and downlink of Earth-facing imagery for the first time, referred to as "first light", is the main milestone left before considering a satellite to be in production. At this point, the satellite joins the rest of imagery-producing satellites that make up the fleet. A centralized fleet-wide scheduler optimizes the allocation of imaging opportunities and ground station contact time for all spacecraft including these new production satellites, with the goal of maximizing global acquired and downlinked coverage. The details of this scheduler are out of the scope of this paper, but it suffices to note that at this point the satellites are instructed to continuously acquire and downlink new imagery.

A number of other on-orbit activities that are also part of nominal operations are enabled at this point as well. These include periodic monitoring and maintenance of certain subsystems (solid state drives (SSD’s), reaction wheels) or image acquisition over calibration sites used for payload radiometric calibration and validation. The satellites also start following a differential drag schedule, alternating between high-drag and low-drag flying configurations with the goal of spreading around across the orbital plane to maximize imaging and downlinking opportunities.[1]

CHALLENGES

Competition for ground station contact time

Planet maintains an extensive ground station network with antenna sites in many parts of the world. Whenever there is a launch, contact time with these antennas is shared between the production fleet and the new commissioning satellites. Some ground stations are used solely for TT&C over UHF while others are capable of high-speed payload downlink over X-band. However, they are built uniformly and can be used to contact any Dove satellite.[2]

It is important to note that, within our concept of operations, one ground antenna can only be in contact with one on-orbit satellite at any given time. This makes ground station-satellite contact time a limited commodity that should be allocated as optimally as possible across the fleet. Optimizing the commissioning workflow becomes a matter of limiting the number of passes needed to progress through commissioning, and allocating those passes to satellites that need them most at any given time.

Newly launched satellites typically concentrate rather close to each other along the orbit they are injected in. This poses an additional challenge in terms of contact allocation because all of the satellites in a new flock are effectively competing for the same accesses over ground stations, and concurrent passes with different satellites can result in radio-frequency interference

(RFI). With no propulsion on-board, it is only later in the orbital lifetime that the satellites achieve full separation via differential drag.[1]

In addition, the new commissioning satellites compete for ground station contact time not only among themselves but also with the rest of the production fleet. It is the goal of the concept of operations during commissioning that downlinking imagery from production satellites for publication should not be affected by contact allocation on the newly launched satellites.

Small Operations Team in need of automation

The satellite operations team at Planet is responsible for maintaining the uptime and productivity of the PlanetScope constellation, commissioning new satellites, resolving on-orbit anomalies, executing on-orbit experiments and testing and rolling out on-board software upgrades. In addition, the team is greatly focused on developing and improving automation tools and better processes to make all the aforementioned activities more efficient and maximize the total productivity of the fleet.

The team is formed by just a handful of engineers located mostly in Germany and the United States. With a satellite-to-operator ratio a couple orders of magnitude higher than in a typical space mission, the need for automation becomes apparent. Commissioning new flocks of up to several dozen new spacecraft while also maintaining the productivity of over a hundred imaging satellites requires a high level of automation for all on-orbit tasks, in order to free up time for operators to focus on specific satellites and issues requiring their attention.

AUTOMATION IN COMMISSIONING

Introduction to Automation at Planet

As outlined in the previous section, commissioning a satellite requires a variety of steps and procedures. For the given scale of Planet's launch campaigns and the relatively small number of operators responsible for satellite operations, this makes a high level of automation paramount. Automated commissioning procedures reduce not only the necessary operator time but also the time after deployment until a satellite can contribute to the fleet's nominal operations, thus taking part in Planet's mission to image Earth at an unprecedented cadence. In addition to that, automations play a vital role in continuous and thorough health monitoring and consequently swift and effective anomaly response.

With each Dove and SuperDove launch, new automations were put into place and the functionality of the existing automations was iteratively extended. This follows Planet's approach of Agile Aerospace, where each design iteration builds upon and improves the previous one. With Planet's latest commissioning campaign for 48 SuperDove satellites onboard SpaceX's Transporter-1 rideshare mission, the commissioning process has reached a level of automation where operator input is only required in off-nominal situations. The commissioning procedure is started once per satellite and then carried out in its entirety by automated systems. The operators only need to monitor high-level metrics reflecting the overall progress of the campaign or respond to anomalies, which are flagged by the system.

The automation system has six main components, each taking care of a different part of the commissioning process. These systems are namely: the Bus Commissioning State Machine, Sequencer, the Post-Firstlight-Engine, MissionMetrics, Autobot and the Contact Allocator.

The former three are responsible for leading a satellite through the different commissioning phases that were outlined in the previous Section. The latter three agents are designed for health monitoring, resource optimization and automated anomaly response. The following sections will provide details for each agent's primary functionality as well as their main design characteristics.

Bus Commissioning State Machine

The Bus Commissiong State Machine schedules and triggers all the necessary operations the satellite has to perform up to the point where its solar panels are deployed. This is the most critical phase of the entire commissioning process, as the satellite's capacity for recharging its batteries is limited. Although SuperDoves are designed to be power positive and therefore safe with their panels being stowed, it is desirable to detumble the satellite and deploy its panels quickly after launch, since this represents a more stable configuration.

The Bus Commissiong State Machine script itself executes during each contact of a satellite with a ground station antenna and is integrated into the main software framework running directly on the ground stations, which is also used to perform nominal in-pass actions for regular operations on the fleet. This makes the Bus Commissioning State Machine easy to integrate into Planet's Mission Control system and makes best use of existing software infrastructure.

In contrast to that, for the early Dove launch campaigns, the commissioning workflow was integrated into the satellite's onboard software, meaning the individual steps were triggered and executed not from the ground, but directly on the satellite.[3] This had the major caveat that the process was much less transparent from an operator's perspective and reacting to problems along the way could pose a significant challenge. Especially for cases of unforeseen behavior in the commissioning scripts, the ground-based solution for managing the workflows provides much more flexibility for swift adaption of the automation and satellite-specific anomaly responses.

Since the commissioning procedure requires the satellite to perform several sequential steps that cannot all be performed during a single contact, the system needs state-machine capability. Depending on connectivity, the satellite's state upon deployment, and availability of access time, this process can take multiple orbits or even days to complete. Each contact is used to schedule activities on the satellite, e.g. initially stabilizing it using the magnetorquers or deploying the solar panels. These activities are then executed by the satellite during its orbit. On the next ground station pass the Bus Commissioning State Machine queries onboard sensor data to assess whether the commissioning activity was executed successfully. If successful execution of the current step can be verified, the script automatically proceeds and schedules the following activity on the satellite.

The process is repeated until a secure radio link to the satellite can be established, and the satellite has detumbled and reached a stable state with the solar panels deployed and fully charged batteries. Upon successful completion, the Bus Commissioning State Machine activates Sequencer, which is another automated agent responsible for managing the following commissioning phase, namely the ADCS calibration.

Sequencer - ADCS Calibration

Sequencer's underlying functionality is rather similar to the Bus Commissioning State Machine's, as it is an automated agent with state-machine capability. Sequencer is used in regular production for software updates on-orbit as well as for scheduling, processing and error handling the ADCS calibration that is part of the commissioning procedure.

First, Sequencer schedules a calibration activity to execute onboard a satellite. After execution, the satellite downlinks sensor data and experimental imagery from the calibration activity. Sequencer then processes this data to calculate the ADCS calibration parameters, which are uploaded to the satellite on the following contact. This entire process is executed twice, once for each of the ADCS calibration activities C1 and C2 described in the previous section. Should any anomalous behavior be identified during the process, e.g. a satellite not executing a calibration activity or some of the telemetry data not meeting quality thresholds, Sequencer repeats the process until it executes successfully.

In contrast to the Bus Commissioning State Machine, Sequencer runs not on the ground stations during contacts, but on separate servers in a periodic fashion. This has the main advantage that it is not limited by ground station access time for the processing of telemetry data and the calculation of ADCS calibration parameters. As soon as both calibration activities are executed and processed successfully, Sequencer changes the satellite's scheduling mode in Mission Control, so that it gets its first imaging jobs assigned by the fleet's scheduling optimizer.

Post First Light Engine

The last step of the commissioning process is achieved when a satellite reaches first light, i.e. when it takes nominal and calibrated, Earth-facing imagery for the first time. A separate agent, the post first-light engine, runs periodically in the background and monitors all commissioning satellites, checking whether first light pictures have been downlinked successfully.

The script doesn't trigger any actions on the satellite directly but rather changes all necessary and still outstanding configuration parameters in the Mission Control system in order to transition it from commissioning to regular production. This includes enabling additional health monitoring for the satellite as it joins the production fleet, enabling frequent maintenance activities and also sending automated notifications to the operator team about the satellite's successful start of nominal production.

Dynamic Contact Allocator

In addition to the automated systems that guide the satellite through the individual steps of the commissioning process, there are several other agents running in the background that are used for satellite health monitoring, resource optimization, and anomaly response. One of these agents is the so-called Contact Allocator. The Contact Allocator's main task is to optimize the use of ground station access availability and to re-distribute contact time among satellites accordingly. As previously mentioned, ground station access time is a limited commodity and especially for large launches this can be a bottleneck for commissioning speed. Satellites that already have their panels deployed, are running through regular orbital charging cycles and are using HSD/HSU passes for TT&C do not have the need of additional LST contacts to maintain nominal operations. These passes can and should consequently be reallocated to satellites that need them for scheduling onboard activities in order to progress within the commissioning procedure.

The task of redistributing contacts is taken care of by the Contact Allocator, which runs in addition to the fleet's regular schedule optimizer. The schedule optimizer uses a sophisticated optimization algorithm which only allows for an updated schedule every few hours. Thus, by the time a satellite is about to execute a contact several hours later, its commissioning state might already have changed or an unexpected anomaly might have occurred. This previous latency issue in scheduling is resolved by the Contact Allocator, which runs on a much higher frequency (every few minutes). Because of its high cadence, it is capable of reacting to short-term changes in a satellite's status, its health telemetry or progress within the commissioning process. It therefore dynamically assesses for upcoming ground station access slots, which satellite can make the best use of them for advancing the commissioning campaign. The Contact Allocator can do this by assigning each satellite a unique priority that reflects its short-term necessity for additional contact time. It then looks for opportunities where a scheduled contact can be passed on from a satellite with a low priority to a satellite with a high scheduling priority. Thus, the Contact Allocator makes sure that the commissioning satellites make use of available ground station accesses in the most efficient way and drastically reduces the amount of operator input for rescheduling accordingly.

The constant reassessment and rescheduling of ground station contacts becomes even more important when considering satellites that are showing critical power levels or other anomalous behavior. By assigning extra contacts to those satellites, the Contact Allocator makes sure that such satellites transmit up-to-date health telemetry and the extra contacts can be used for debugging their anomalies.

Mission Metrics

The more than one hundred operational on-orbit satellites that are part of Planet's fleet generate a constant stream of telemetry from hundreds of sensors each, as well as on-board CPU logs describing every event occurring on-orbit. The vast amount of data is infeasible to parse manually, so a number of health indicators and metrics are defined and computed automatically based on the large volume of data.

The data processing procedures are also applied to commissioning satellites. Telemetry is downlinked and ingested on each contact, and immediately triggers alerts if certain thresholds are violated. A number of derived metrics can then be visually analyzed to assess the progress of a whole flock through the commissioning workflow at a glance. Any off-nominal behavior can easily be spotted, identified, and triaged.

Anomaly Detection (Autobot)

The automation of satellite operations at Planet is not exclusive to our commissioning process. Autobot is the main automation agent in charge of several aspects of fleet management but mostly automated on-orbit anomaly detection and reaction.[4] Anomalies whose symptom signature and possible mitigation are well understood can have an automated response in this manner.

There are many automated reactions implemented by Autobot, mostly based on an analysis of telemetry data from the satellites. Examples of these include reactions to low bus voltage or unexpected tumbling.[5] Constant health monitoring for the entire fleet is essential for maximizing satellite uptime and productivity for production satellites, but also serves a purpose during commissioning by keeping the newly launched spacecraft safe. Especially during the more time-sensitive bus commissioning phase, swift detection and response to any unexpected anomalies like a declining bus voltage is vital, and, for the number of satellites that Planet typically commissions, once more only really achievable through the extensive use of automation.

RESULTS

In this section, an overview of the performance of the last three Planet commissioning campaigns to date will be discussed, followed by a comparative analysis between all SuperDove launches. The most recent launches, Flock 4V, Flock 4EP, and Flock 4S are specifically highlighted as these campaigns integrated a significant and increased amount of automation into their concept of operations from early campaign. The launches will first be highlighted for their differences in automation features, followed by a section quantifying their relative performance with respect to one another.

Flock 4V

Flock 4V launched on September 3rd, 2020, onboard the Vega VV16 mission, adding 26 SuperDoves to the Planet fleet. The commissioning campaign utilized the Autobot anomaly detection, Mission Metrics, and the Bus Commissioning State Machine, but of the three featured campaigns, had the least amount of automation present.

Figure 2 shows the time in each phase for the satellites moving through the commissioning process. The satellites are ordered from the top down in the order of commissioning kick off, with the y-axis denoting the satellite identification number (hardware ID) and the x-axis the time since launch on September 3rd. Note that the first 14 SuperDoves in this flock were injected directly into orbit by the launcher, while the next 12 were deployed every two to three days by D-Orbit's ION deployer, starting three weeks after launch. It can be noted that 2274 spent the most time in commissioning due to anomalous behaviour which is reflected in it's time spent in phases C1 and C2, and would be considered as an outlier in the relative performance of the flock moving through the commissioning process.

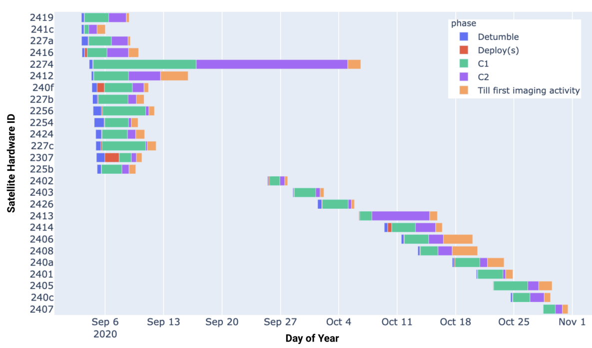

Figure 2: F4V Commissioning Phases

Flock 4EP

Flock 4EP launched on October 28th, 2020, onboard Rocket Lab's Electron as part of the "In Focus" mission, adding 9 SuperDoves to the fleet. Less than two months from the previous Flock 4V launch, the 4EP commissioning campaign introduced several new automation features to improve the process: a streamlined detumble and panel deploy procedure no longer requiring operator intervention to confirm success, an automatic transition from the bus commissioning state machine to ADCS calibration via Sequencer, and the use of the post-first light script to move the satellite from commissioning into production quicker. Through the implementation of these new features, the commissioning process was streamlined, reducing the amount of contacts needed to move through phases and decreasing the overall time to complete commissioning.

Figure 3 similarly shows the time spent in each phase by each of the satellites during this commissioning campaign. As was in the F4V campaign, anomalies can occur during commissioning that affect the overall time a satellite spends in specific phases of the process. Satellite 225a a part of F4EP clearly showcases an example of this; the satellite spent the largest amount of time in commissioning phase C1, but was resolved and moved into production by the end of November 2020.

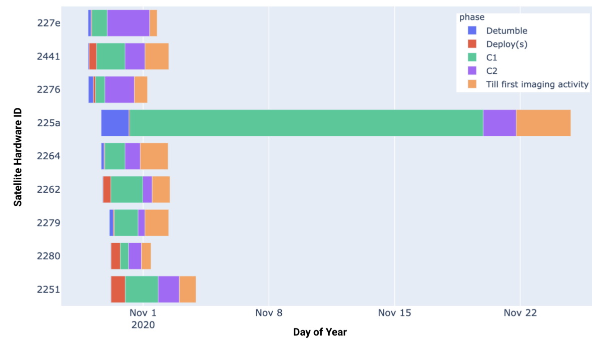

Figure 3: F4EP Commissioning Phases

Flock 4S

Planet has a 100% record of contacting satellites that have been successfully deployed to orbit. Flock 4S was no exception; less than 24 hours after launch, operators had made contact with all 40 of the first set of deployed satellites. Flock 4S complements the existing constellation of over 100 satellites and delivers unmatched medium-resolution global coverage. The targeted commissioning campaign timeline for the first 40 deployed satellites was 13 days, which was achieved by the Missions Operations team. Figures 4 and 5 show their progress through commissioning.

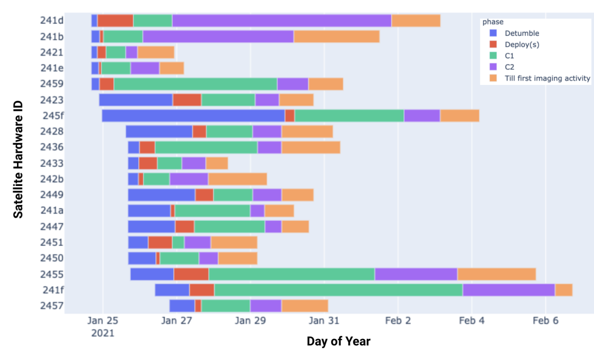

Figure 4: F4S Commissioning Phases Pt. 1

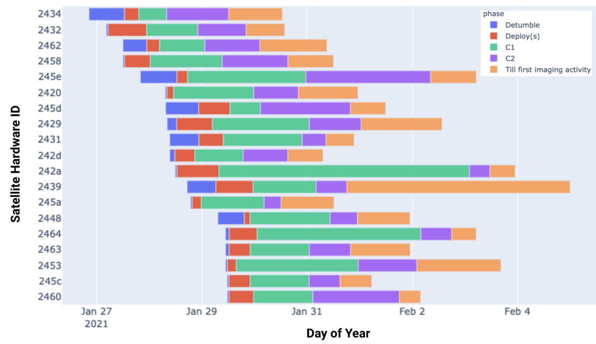

Figure 5: F4S Commissioning Phases Pt. 2

Flock 4S was the first time the Dynamic Contact Allocator automation was implemented into Planet's commissioning workflow. Due to the large number of satellites of this launch (48), it was critical in effectively allocating contacts to those satellites that needed them to quickly and efficiently move through the commissioning process. The following Section, Quantifying the impact of Automation, will detail the relative impact of this automation on the commissioning process.

Quantifying the Impact of Automation

The integration of automation tooling into Planet's SuperDove commissioning workflow has resulted in a streamlined process where little to no operator intervention is necessary for a satellite to move between phases. Quantifying the impact and improvements of the automation tools developed between launches is key in highlighting the gaps that have been filled and where there is room for further improvement. In this case, the authors use the latest launch, Flock 4S, as the standard of commissioning to compare previous SuperDove campaigns to. There are two distinct ways that will be used to compare the performance of commissioning campaigns: the number of contacts a satellite requires to complete each commissioning phase and the total time needed to progress through these phases. The reason why we use the number of contacts as well as the total time in commissioning as metrics of performance is because no two satellites will have the exact same ephemeris describing their orbit and therefore will not have access to the same number of contacts over the same time period, i.e. a twelve hour window of time will not have the same accesses for any two satellites a part of the constellation. Therefore, using the number of contacts a satellite requires in phases of commissioning is an unbiased evaluation to quantify the efficiency of phases that require communication with the satellites.

Quantifying Improvements: Number of Contacts

There are three distinct phases of commissioning that the average number of contacts per satellite per flock will be compared for:

A. Bus Commissioning
B. The time between Bus Commissioning completion and beginning of ADCS calibration
C. ADCS Calibration

The reasoning as to why phase B is included in this comparison is because this process was initially done manually by operators in past commissioning campaigns and is now fully automated, which directly captures a before and after comparison of introducing automation into this portion of commissioning. As can be seen in Figures 7 and 8, there is a significant difference between the average number of contacts an average satellite required for phases B and C from Flock 4A onwards. Looking at the average number of contacts for C, Flock 4A took an average of ~180 contacts to move through ACS calibration, whereas Flock 4S averaged ~20, a reduction of almost 90%. The difference in contacts required to complete commissioning is a massive improvement and not only speeds up the time a satellite takes during the ACS calibration commissioning phase, but also relieves the load on the ground station network and therefore reduces the impact of the new commissioning satellites on the contact availability for the remaining production satellites that are actively imaging.

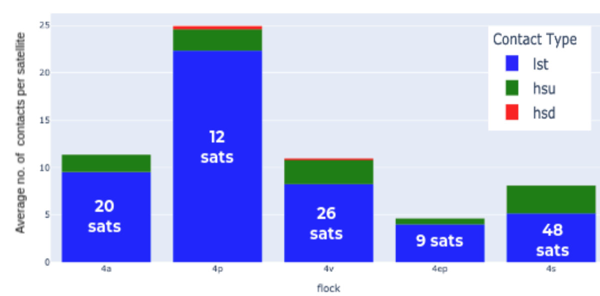

Figure 6: Avg. number of contacts needed during Bus Commissioning

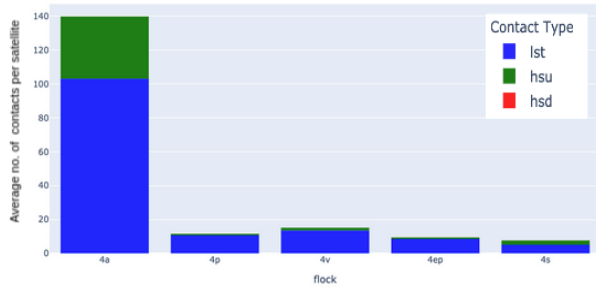

Figure 7: Avg. number of contacts before the first High Speed Downlink

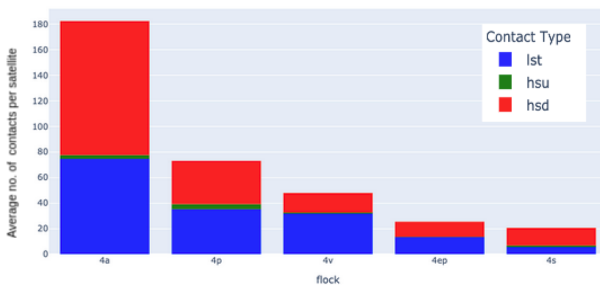

Figure 8: Avg. number of contacts needed during ACS Calibration

Quantifying Improvements: Total Time in Commissioning

Table 2 synthesizes the average time it took for a satellite from each flock to complete commissioning. The relative improvement from the earliest launch, F4A, to the latest launches, F4EP and F4S is paramount; moving from an average time for a satellite to complete commissioning of almost 39 days to between 3-4 days. Note that it is not unexpected that F4S took slightly longer than F4EP in its average satellite time through commissioning, because of the large difference in number of satellites a part of the launch and therefore greater strain on ground station resources during the F4S campaign.

Table 2: Average Time between First Contact and First Light

Flock	# of Sats	Avg. Time to Completion	Fastest Sat Time
4A	20	38 days 22:56:02	13 days 17:22:36
4P	12	17 days 00:10:38	10 days 16:38:00
4V (incl. ION)	26	5 days 12:07:13	1 day 13:06:16
4EP	9	3 days 16:16:35	2 days 06:11:21
4S (incl. ION)	48	4 days 05:35:50	1 day 07:36:38

Table 3 shows the average time it took per flock from the detumbling phase (i.e. beginning of bus commissioning) to first light. This table is included because the actual active portion of the Planet commissioning process begins with the detumble activity, which is first scheduled after an operator begins the commissioning workflow. It is possible that a satellite had a period of time between its first contact and beginning its commissioning process, and so this table encompasses the time from

beginning bus commissioning to completion of commissioning. What is particularly commendable about these results is that it took only 15.5 hours for a satellite a part of the F4S campaign to complete the phases of detumbling to downlinking its first imagery, a massive improvement and decrease from the ~2 days it took for F4EP.

Table 3: Average Time between Starting Detumble and First Light

Flock	# of Sats	Avg. Time to Completion	Fastest Sat Time
4A	20	38 days 07:56:31	13 days 07:07:59
4P	12	16 days 04:56:35	10 days 13:33
4V (incl. ION)	26	5 days 07:31:11	1 day 12:47:27
4EP	9	3 days 13:52:14	2 days 05:51:59
4S (incl. ION)	48	3 days 15:40:25	15:28:32

Going farther back in time, Table 4 is a comparison between Flock 3P's commissioning campaign (and largest launch to date) which launched February 15th, 2017 and Planet's latest launch Flock 4S which launched January 24th, 2021. Although the launches encompassed two different generations of satellites, Dove and SuperDove, the comparison of the time it took to move through the commissioning process of one of Planet's earlier launches to its most recent is an excellent way to highlight the significant improvements to the process.

Table 4: Average Time between First Contact and First Light 3P vs 4S

Flock	Launch Date	# of Sats	Avg. Time to completion	Fastest Sat Time	Commissioning Length
3P	February 15th 2017	88	35 days 03:44:51	3 days 19:42:23	125 days
4S (No ION)	January 24 2021	40	4 days 13:33:34	1 day 07:36:38	13 days

Note that for the above comparison, the ION satellites in Flock 4S were removed as they would artificially increase the time through commissioning, since the spacecraft onboard the deployer were only deployed starting T-plus 2 weeks from launch, one satellite every 2 days. It is clear across all metrics that the commissioning process has greatly improved. The average time it took for a satellite to complete commissioning dramatically decreased from F3P to F4S by over a month; the large discrepancy between the two launches compounded by the fact that F3P experienced power issues as well as the development of automation features over time. The fastest satellite through commissioning was also achieved by F4S, which had a quicker time by approximately 2.5 days as compared to F3P. Finally, although F3P's launch size was only roughly twice the size of F4S, it took over 9.5 times as long to commission F3P than it did F4S; a testament to the improvements of Planet's commissioning process and, in particular, its ability to rapidly commission large numbers of satellites.

CONCLUSION

Planet is able to operate the world's largest Earth Observation satellite constellation, with hundreds of satellites a part of their fleet, through the development of automation tools and techniques. Throughout its history, Planet has iterated its commissioning process to be able to seamlessly add new flocks of spacecraft to the existing fleet within a few weeks of launch. Flock 4S has been the largest SuperDove-generation launch to date with a total of 48 satellites, 40 immediately deployed and another 8 deployed via D-orbit's ION Deployer. As shown through a comparative analysis between past launches, it was one of the most successful. The commissioning campaign took less than two weeks to complete for the first 40 satellites with 100% first contact rate and a record 15.5 hours from starting its detumble to downlinking its first image in the case of the fastest satellite.

With each launch campaign, Planet has incrementally implemented a variety of automated systems that have improved the safety, speed, and reliability of the commissioning process, while simultaneously drastically reducing the amount of required operator oversight and ground station access time. Through the power of automation, Planet's small team of operators can quickly, effectively, and successfully commission a SuperDove flock with little to no human intervention necessary, while concurrently minimizing the impact of commissioning efforts on the regular fleet in production.

Acknowledgements

The authors would like to thank the entire Space Systems team, and the entire Planet team, for their contributions and support throughout every launch, deployment and commissioning process. Although this paper focuses on the results and challenges of commissioning from a satellite operator's perspective, it is undoubtedly a team effort across the entire Planet community that makes these campaigns a success.

References

1. Cyrus Foster, James Mason, Vivek Vittaldev, Lawrence Leung, Vincent Beukelaers, Leon Stepan, and Rob Zimmerman. Constellation phasing with differential drag on Planet Labs satellites. *Journal of Spacecraft and Rockets*, 55(2):1-11, 2017.
2. Kyle Colton and Bryan Klofas. Supporting the flock: Building a new ground station network for autonomy and reliability. In *30th Annual AIAA/USU Conference on Small Satellites,* Logan, Utah, USA, 2016.
3. Rob Zimmerman, Deanna Doan, Lawrence Leung, James Mason, Nate Parsons and Kam Shahid. Commissioning the World's Largest Satellite Constellation. In *31st Annual AIAA/USU Conference on Small Satellites*, Logan, Utah, USA, 2017.
4. Mirue Choi. Closing the satellite operations loops with automation via Autobot. In *70th International Astronautical Conference*, Washington D.C., USA, 2019.
5. Pablo Bernal-Mencia, Cruz Peregrina, Jose Angel Gutierrez Ahumada, Kattia Flores Pozo, and Lisa Mcgill. Automated anomaly response for Planet constellations. In *16th International Conference on Space Operations*, Cape Town, South Africa, 2021.