Chapter 21
Robotic Deployment and Installation of Payloads on Planetary Surface

Ashitey Trebi-Ollennu, Khaled Ali, Cristina Sorice, Won Kim, Steven Myint, Omair Khan, Philip Bailey, Hallie Abarca, Robert G. Deen, Jeng Yen, Justin N. Maki, Grace Lim, Nythi Udomkesmalee, and Jeffrey Umland

Abstract The InSight (Interior Exploration using Seismic Investigations, Geodesy and Heat Transport) mission is a Discovery Program lander to investigate the internal structure of Mars and the differentiation of the terrestrial planets (Banerdt et al. in Space Sci Rev 215:22 2018). The InSight flight system is a close copy of the Mars Phoenix Lander and comprises a lander, cruise stage, heatshield and backshell. The

A. Trebi-Ollennu (✉) · K. Ali · C. Sorice · W. Kim · S. Myint · O. Khan · P. Bailey · H. Abarca · R. G. Deen · J. Yen · J. N. Maki · G. Lim · N. Udomkesmalee · J. Umland
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA
e-mail: ashitey@jpl.nasa.gov

K. Ali
e-mail: Khaled.S.Ali@jpl.nasa.gov

C. Sorice
e-mail: Cristina.E.Sorice@jpl.nasa.gov

W. Kim
e-mail: Won.S.Kim@jpl.nasa.gov

S. Myint
e-mail: Steven.Myint@jpl.nasa.gov

O. Khan
e-mail: Mohammed.O.Khan@jpl.nasa.gov

P. Bailey
e-mail: Philip.Bailey@jpl.nasa.gov

H. Abarca
e-mail: Hallie.E.Gengl@jpl.nasa.gov

R. G. Deen
e-mail: bob.deen@jpl.nasa.gov

J. Yen
e-mail: Jeng.Yen@jpl.nasa.gov

J. N. Maki
e-mail: Justin.N.Maki@jpl.nasa.gov

G. Lim
e-mail: Grace.Lim@jpl.nasa.gov

V. Badescu et al. (eds.), *Handbook of Space Resources*,
https://doi.org/10.1007/978-3-030-97913-3_21

InSight science payload includes a seismometer, a wind and thermal shield, a heat flow probe and a precision tracking system to measure the size and state of the core, mantle and crust of Mars. InSight is NASA's first successful precision robotics instrument placement and release on another astronomical body since Apollo. This operations breakthrough enabled NASA's InSight lander to detect a 'marsquake', a faint trembling of Mars's surface on 6 April 2019, 128 Martian days after its landing on Mars. This is the first quake detected on an astronomical body other than Earth or the Moon. This chapter describes the operations of the robotics instrument deployment systems (IDS) that successfully deployed the InSight science payload to the surface of Mars, and the planning and command sequence generation process used for its successful deployment. Among its recommendations, the chapter identifies technology gaps in the operations of in-situ manipulators for planetary exploration.

21.1 Introduction

The InSight (Interior Exploration using Seismic Investigations, Geodesy and Heat Transport) mission is a Discovery Program lander to investigate the internal structure of Mars and the differentiation of the terrestrial planets (Banerdt et al. 2018). The InSight flight system is a close copy of the Mars Phoenix lander and comprises a lander, cruise stage, heatshield and backshell. The InSight science payload includes a seismometer (SEIS), (Longonné et al. 2019) and wind and thermal shield (WTS), a heat flow probe (Heat Flow and Physical Properties Package, HP3, Spohn et al. 2018) and a precision tracking system (RISE), (Folkner et al. 2018) to measure the size and state of the core, mantle and crust of Mars.

InSight is NASA's first successful precision robotics instrument placement and release on another astronomical body since Apollo. This operations breakthrough enabled NASA's InSight lander to detect the first known 'marsquake', a faint trembling of Mars's surface on 6 April 2019, 128 Martian days after landing on Mars on 26 November 2018. This is the first quake detected on an astronomical body other than Earth or the Moon. This chapter describes the operations of the Robotics Instrument Deployment Systems (IDS) that successfully deployed the InSight science payload to the surface of Mars. In addition, the chapter describes the IDS planning and command sequence generation process used for the successful deployment of SEIS, WTS and HP3 on the surface of Mars. The paper concludes with recommendations based on the experience gained from InSight IDS operations. This includes technology gaps identified in the operations of in-situ manipulators for planetary exploration.

N. Udomkesmalee
e-mail: nythi.udomkesmalee@jpl.nasa.gov

J. Umland
e-mail: Jeffrey.W.Umland@jpl.nasa.gov

NASA-JPL, 4800 Oak Grove Dr, M/S 157-205, Pasadena, CA 91109, USA

21.2 Robotic System

The InSight Instrument Deployment System (IDS) consists of the Instrument Deployment Arm (IDA), scoop, five-finger "claw" grapple, motor controller, arm-mounted Instrument Deployment Camera (IDC), lander-mounted Instrument Context Camera (ICC), and control software (Fig. 21.1). IDS is responsible for the first precision robotics instrument placement and release (seismometer and heat flow probe instruments) on a planetary surface. These instruments will enable scientists to perform the first comprehensive surface-based geophysical investigation of Mars. Table 21.1 list IDS driving requirements.

21.2.1 Instrument Deployment Arm (IDA)

The InSight IDA is a refurbished flight robotic arm from the Mars Surveyor 2001 lander mission (Bonitz et al. 2000). The IDA is a four degrees-of-freedom backhoe design manipulator with a 1.8 m reach that provides the following motion: yaw (shoulder azimuth, joint 1) and three pitch joints (shoulder elevation, elbow, and wrist, joints 2 through 4, respectively). The IDA links are made of titanium. During normal operations the IDA actuators are capable of generating 35, 120, 65, and 10.5 N-meters of torque at the joint output for joints 1 through 4, respectively. The

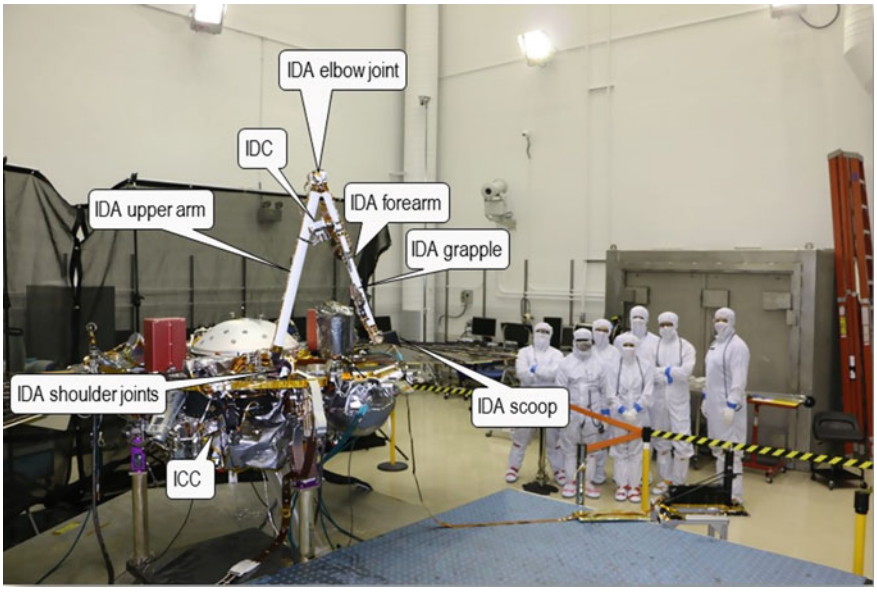

Fig. 21.1 InSight Mars lander with IDS elements labeled

Table 21.1 IDS driving requirements

	IDS driving requirements
1	The IDA shall operate for 4 h and for 111 Martian sols plus 10 h for ground testing
2	The IDS shall positively capture and retain deployable elements, including under loss of power, until placement on the surface is confirmed
3	The IDS shall deploy elements to the surface with the Lander deck tilted 15 degrees w.r.t. gravity
4	The IDS shall have a total mass of less than or equal to 9.41 kg
5	The IDS shall be able to lift a mass of up to 9.5 kg
6	The IDS shall acquire context images of the IDA workspace within the FOV of the InSight Context Camera (ICC)
7	The InSight Deployment Camera (IDC) field of view shall include the IDA end effector during capture and disengagement from deployable elements
8	The IDA shall have repeatability of 0.005 m
9	The IDA shall position the end effector with an absolute error of less than or equal to 0.015 m
10	The IDS shall determine the IDC imaging baseline to within 0.0028 m
11	The IDS shall acquire images of the Lander deck and Solar panels within the field of view of the IDC

IDA can lift and deploy a 9 kg payload on Mars (33 N) at 1.65 m distance. The force the IDA end effector can exert is configuration dependent, but the average force is typically about 80 N.

Each joint has a temperature sensor and heater and includes a dust seal to prevent contamination of the motor and gearbox. The IDA is designed to withstand expected environmental temperatures from $-110°$ to $+70$ °C, in a CO_2 atmosphere, with pressure as low as 5 Torr. Each of the IDA joints consists of a brushed DC motor with two-stage speed planetary gears and a harmonic drive at the output (except the wrist, which has a bevel gear at the output of the planetary gears). The IDA joints do not have mechanical braking systems but employ a dynamic braking system where actively shorting the motor leads slows the motor until magnetic detents capture the rotor. The magnetic detents are sized to provide the appropriate holding torque to assure no slippage while the IDA is powered off. Each joint has two position sensors: encoders on the joint input motor shaft and potentiometers at the joint output load shaft. Each joint is equipped with two mechanical hardstops at the end of their range of travel. The encoder counters are initialized based on potentiometer data or by running each joint up against their respective mechanical hardstops.

The IDA end effector consists of a five-finger "claw" grapple hanging on an umbilical cable, a scoop, and forearm-mounted camera IDC (closer to the elbow joint) facing the IDA end effector.

Two thermal characterization tests were performed on the IDS subsystem in a 13-foot sensor chamber at the Raytheon El Segundo Integrated Test Laboratory (ITL), California (Fig. 21.2). During the test the IDA heaters were characterized and

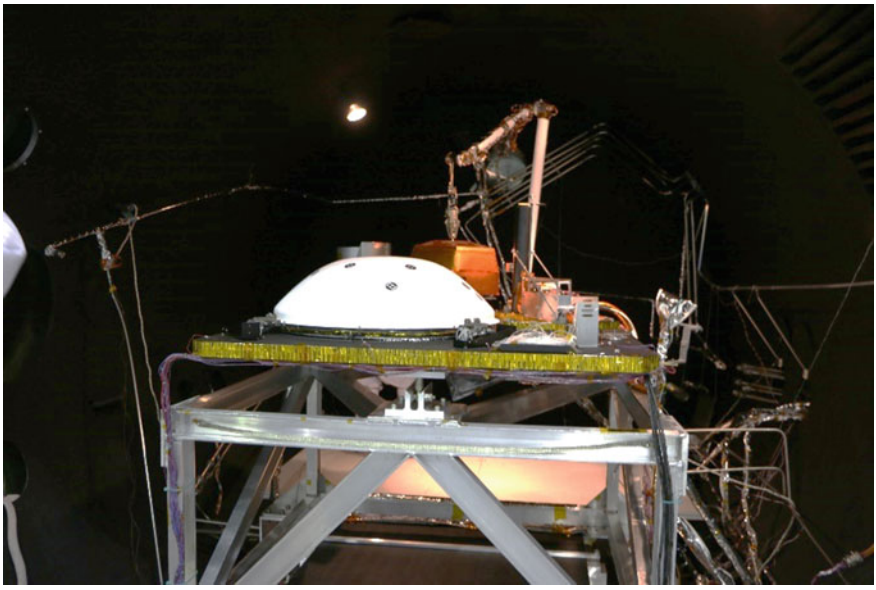

Fig. 21.2 IDS subsystem set-up in a 13-foot sensor chamber at the Raytheon El Segundo Integrated Test Laboratory (ITL), California for thermal characterization test

IDA functional qualification was successfully performed at proto-flight operational temperature. In addition, IDA stop-and-hold torques were characterized at various temperatures.

21.2.2 IDA End Effector Grapple

The grapple is a five-finger "claw" and hangs by an umbilical at the IDA end effector. The grapple is designed with five fingers to assure proper self-alignment and be position error tolerant while closing the grapple fingers around the spherical cap grapple hooks on the payloads. The grapple fingers are forced closed. The grapple fingers are opened by a single high-output paraffin (HOP) actuator that slowly heats up and melts the wax that pushes a rod out to open the fingers. When the fingers are fully open (as shown in Fig. 21.3), a limit contact switch trips and turns the grapple HOP heater power off. As the grapple HOP cools down in the ambient temperature, the grapple fingers slowly close passively without any actuation. The grapple design is robust against unexpected power loss because power is required to open the fingers. The grapple umbilical provides the necessary compliance (unactuated additional 2 DOF for the 4 DOF IDA) for engaging and deploying the payloads on tilted lander and uneven terrain. The grapple is stowed against the IDA forearm such that it does not obstruct the IDC FOV (shown in Fig. 21.3a). However, when the grapple is

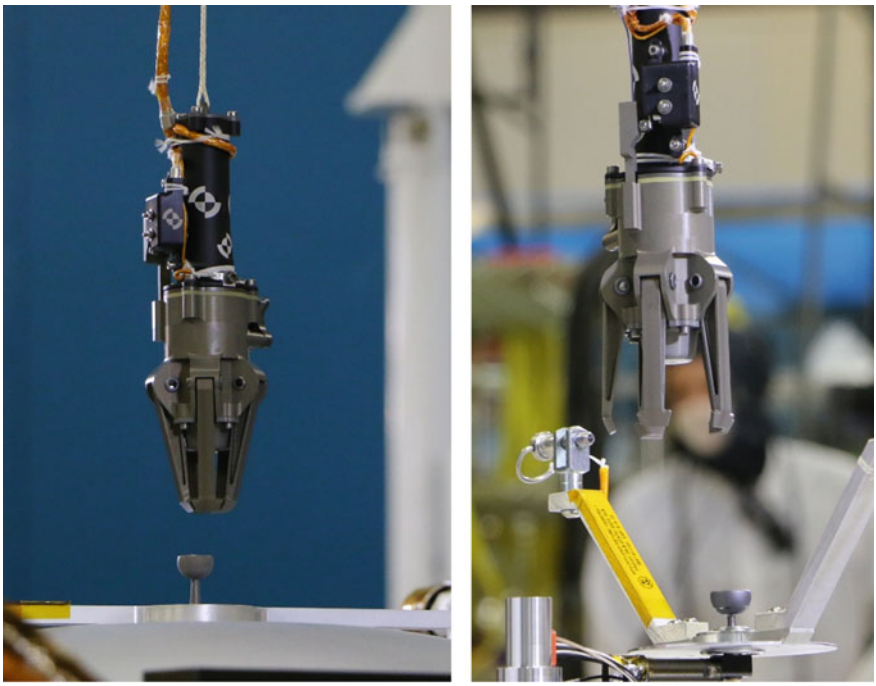

Fig. 21.3 Flight grapple with **a** fingers closed and **b** fingers opened

unstowed (shown in Fig. 21.4b) it hangs in the FOV of the IDC such that the IDC images can capture the opening of grapple fingers and the engagement of spherical cap grapple hooks on the payload.

During deployment the grapple is unstowed, hanging from the IDA end effector by an umbilical cable. The IDA can position the grapple to capture the payload's spherical cap grapple hook, lift, and place SEIS, WTS and HP3 on the Martian surface. The grapple can be stowed using the IDA in a "ball-and-cup" maneuver to the grapple restraint mechanism on the IDA forearm (shown in Fig. 21.4a).

21.2.3 IDA End-Effector Scoop

The scoop consists of a single chamber with a front blade and a secondary blade on the bottom side (as shown in Fig. 21.5). The scoop's front and secondary blades can be used to excavate materials (by digging or scraping) and collect materials excavated in the IDA workspace. The scoop will enable soil mechanics experiments for inferring mechanical properties of the Martian soil at the landing site using the IDA housekeeping data (motor currents, scoop position, etc.) to estimate the

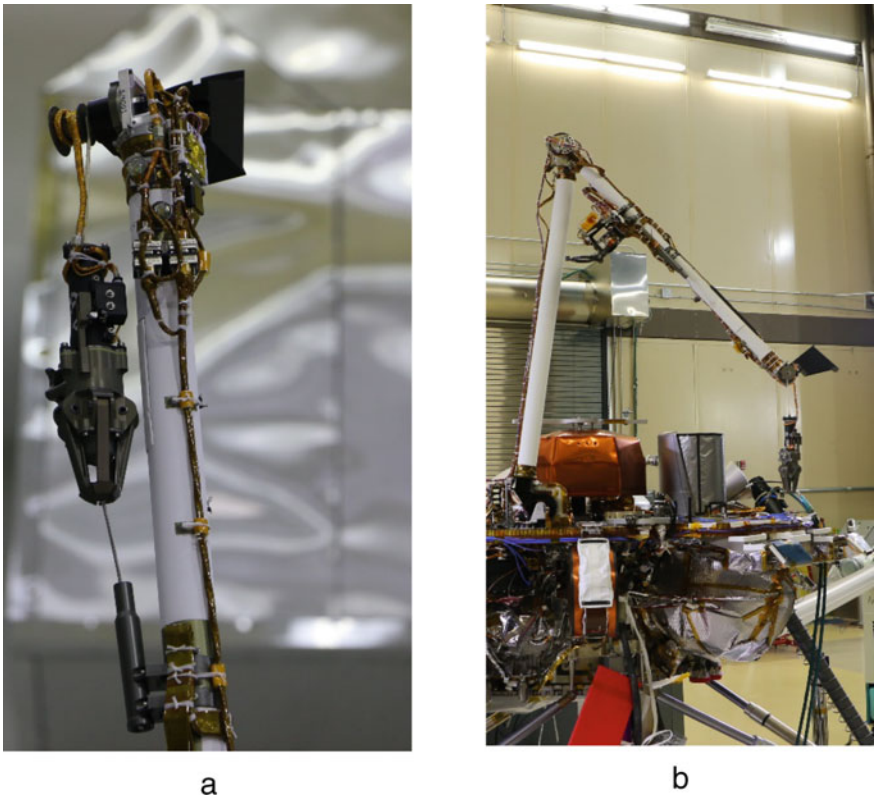

Fig. 21.4 **a** Stowed flight grapple and **b** Grapple unstowed from grapple pre-deployment restraint

scoop's applied force. The scoop is not required for nominal instrument deployment operations.

21.2.4 IDS Cameras

To assist in the deployment of the payloads, the robotic arm is equipped with two cameras: the Instrument Deployment Camera (IDC) mounted on the robotic arm and the Instrument Context Camera (ICC) mounted on the lander body underneath the top deck (Fig. 21.6) (Maki et al. 2018). The primary objectives of the IDC and ICC are to: (1) document the state of the lander and surrounding terrain; (2) support terrain assessment for the selection of the SEIS and HP3 instrument deployment locations; (3) facilitate and document the deployment activities; (4) monitor the location and state of the instruments post-deployment; and (5) measure and monitor atmospheric dust opacity (Banfield et al. 2020). The IDC has a FOV of 45° × 45° and an angular

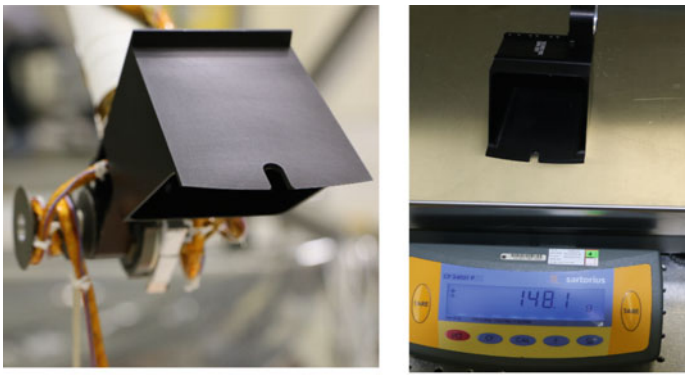

Fig. 21.5 **a** InSight scoop mounted on the IDA and **b** front view of the scoop before installation on the IDA

resolution of 0.82 mrad/pixel at the center of the image. The ICC is mounted to the lander and will acquire wide-angle views of the instrument deployment activities. The ICC has a FOV of 124° × 124° and an angular resolution of 2.1 mrad/pixel at the center of the image. The IDC and ICC cameras are flight spare engineering cameras from the Mars Science Laboratory (MSL) mission. The InSight project upgraded the inherited cameras from single-channel greyscale to red/green/blue (RGB) color by replacing the detector with a Bayer-pattern version of the same 1024 × 1024-pixel detector. Stereo IDC image pairs, acquired by moving the arm between images, are critical for characterizing the topography of the instrument deployment workspace, a 4 × 6-m area located in front of the lander. Images from the cameras are processed using software from previous Mars surface missions, with several new image products developed for InSight to support instrument placement activities.

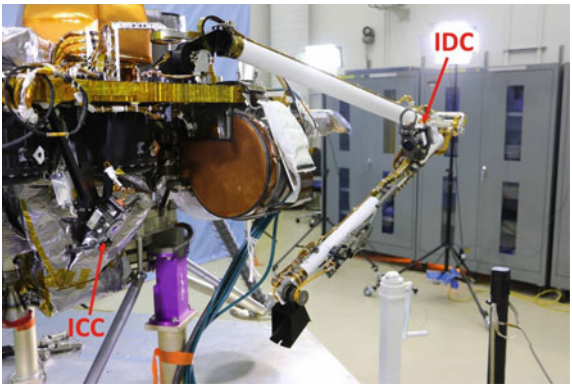

Fig. 21.6 Camera locations on the Insight lander. The ICC can be seen on the left, mounted to the lander, and the IDC is on the right, mounted to the forearm section of the robotic arm. The distance from the IDC to the scoop at the end of the arm is approximately 0.6 m

21.2.5 IDA Motor Controller

The IDA motor controller consists of two printed-circuit boards located in the lower payload electronics box (PEB) and provides power conditioning, motor voltage control and drivers, grapple heater drivers, joint encoder counting, and analog-to-digital conversion of potentiometer voltages, temperature sensor voltages, motor currents, and heater current. The PEB provides the interface to the lander command and data handling (C&DH) computer over a serial link. Firmware running on the IDA motor controller microprocessor provides for low-level motor command execution to move the joints to the specified positions, grapple heater command execution, analog-to-digital calibration, and sensor monitoring (Trebi-Ollennu et al. 2018).

21.3 Robotics Flight Software

The IDA flight software (FSW) provides both control of the IDA hardware and visibility of IDA hardware and software state. Running on the on-board flight computer, IDA FSW communicates with the IDA motor controller. It provides the interface for IDS ground operators to control the IDA through the motor controller. Telemetry data from the IDA FSW provides the hardware and software state to the ground operators. The IDA FSW inherits from and builds upon the Phoenix robotic arm FSW (Bonitz et al. 2008). The IDA FSW provides the following specific capabilities (Trebi-Ollennu et al. 2018):

- interface with external entities, including other spacecraft FSW components and the IDA PEB
- expansion of high-level IDA commands from the command sequencer into low-level IDA actions
- motion control of the IDA
- control of the grapple
- fault sensing, recovery, and safing
- collision prevention between the IDA, lander, and science instruments
- visibility of the IDS state in telemetry.

The lander sequencing engine sends sequenced IDA commands to the IDA FSW, one at a time. The IDA FSW responds to these commands, taking action as appropriate. While handling a command, the IDA FSW may communicate with external entities, such as the IDA motor controller or power-switching FSW.

Most IDA movement commands specify a single motion of the arm or action of the IDA FSW. This could be to move the IDA to a specific position, using the most direct path, to set a parameter, or to turn a heater on or off. After executing the activity, the IDA FSW sets a flag to let the command sequence in the sequencing engine know it is ready for the next command.

Other higher-level IDA FSW commands are more complicated, building on those low-level commands. These commands execute multiple actions, potentially including multiple IDA motions, before informing the sequencing engine of the command completion. Examples of these include commands to initialize the motor controller, calibrate the motors, dig a trench, and scrape the Martian surface.

21.3.1 IDA Motion Commands and Motor Control

There are multiple types of IDA commands to specify single motions. These different categories of motion commands allow the movement to be specified in terms of joint angles, Cartesian coordinates, and time durations for motion.

Joint commands specify the IDA motion in terms of the joint angles. These commands can specify the goal position as either absolute joint angles or relative offsets from the current joint angles. Joint commands can move either single joints or all IDA joints at the same time.

Cartesian commands specify the IDA motion in terms of Cartesian coordinates and an approach angle for the current end effector, or tool. These commands can specify the goal position using absolute or relative coordinates. Absolute coordinates can be in the IDA frame, which has its origin at the base of the IDA. Relative coordinates can be in the IDA frame or in one of many tool frames attached to the IDA end-effectors. As the IDA moves, the orientation of the tool frames changes relative to the IDA frame. Some of these tool frames are rigidly fixed to the IDA links, while others hang such that their z-axis is always the direction of gravity.

Timed motion commands are relative motion commands that specify a direction and duration of time to move each joint. The speed of the joint motion is set with another command prior to the motion command.

Both joint and Cartesian move commands can be commanded in a “guarded” manner. Guarded moves allow the IDA to safely contact other objects in its workspace. Normally, if the motor controller or IDA FSW detects excessive motor currents or joint torques, the FSW will “safe” itself, which means it will stop all motion, shut itself down, abort the currently executing command sequence, and wait for ground operators to tell it what to do next. If the motion is a guarded move command, however, the motion will stop after detecting high currents or torques, but the IDA FSW will not safe itself. Instead, it will inform the sequencing engine that the motion completed successfully, thus proceeding to the next command in the sequence.

For each motion command of any type, the IDA FSW breaks it into multiple via points, where each via point is a set of joint angles or Cartesian pose between the start and goal pose. One at a time, the FSW generates a joint velocity profile for each via point and passes the via-point encoder angles and a voltage that corresponds to the desired velocity to the motor controller. The motor controller closes the loop on the encoders to move the joints to the desired angles. During each control cycle, the IDA FSW monitors the motor state and uses a software-based PID controller to

compute a new voltage to send to the motor controller. All via points in a motion are checked for limit violations and potential collisions before beginning motion to the first via point.

The IDA FSW chooses via points based on whether the motion is intended to be joint interpolated or linear interpolated (straight line). The FSW uses joint-interpolated motion for joint move commands. Ground operators can choose between joint-interpolated and linear-interpolated motion for Cartesian move commands.

21.3.2 Grapple Control

To open the grapple (Fig. 21.7), the IDA FSW turns on one of a pair of redundant grapple heaters, pushing open the grapple fingers. The heat from the grapple's HOP actuator will damage the grapple if the actuator is left powered for too long. To avoid damage to the grapple, the IDA FSW monitors the grapple finger limit switches and the actuator temperature to determine when to power off the HOP actuator heaters. When the limit switches indicate the grapple fingers are fully open, or if the actuator becomes too hot, the FSW powers off the heaters. Additionally, the command to open the grapple has a timeout argument, and if the timeout is exceeded, the FSW will power off the grapple heater.

Ground operators use a spacecraft command to inform the IDA FSW of the state of the grapple. This grapple phase indicates whether the grapple is safe to open, stowed on the side of the IDA forearm, stowed for launch, or grappling one of the deployable payloads. If the grapple is grappling a payload, the grapple phase indicates which payload. Only spacecraft commands sent by ground operators can change the value of the grapple phase, and they are used to indicate both whether it is safe to open the grapple and which payload, if any, is grappled.

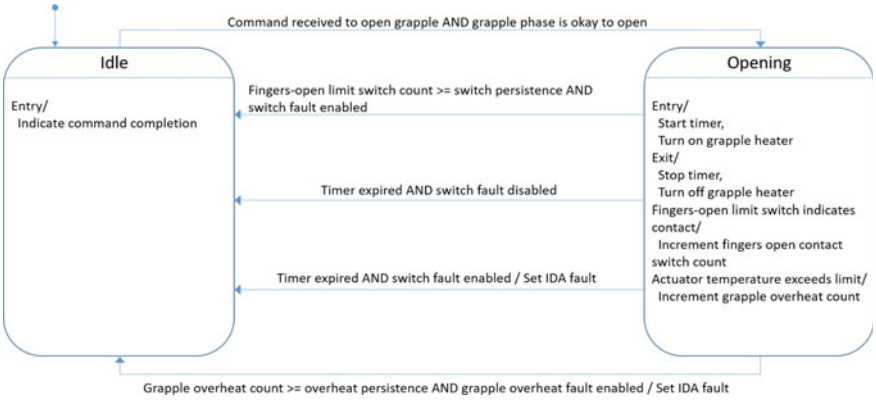

Fig. 21.7 IDA FSW grapple opening control

The grapple phase serves two purposes. First, it provides a level of protection against accidentally opening the grapple when it is not safe to do so, such as when a payload is suspended off the ground. The grapple phase must indicate that it is safe to open the grapple or the IDA FSW will reject commands to do so. Second, the grapple phase is used to indicate the grapple payload, and thus the suspended mass, when determining the deflection of the IDA due to gravity.

21.3.3 Kinematics and Deflection

For most IDA end-effectors, the IDA FSW computes the forward and inverse kinematics in the ordinary manner from the joint angles. The grapple, however, requires special treatment.

The grapple hangs on a flexible tether, suspended just prior to the fourth (wrist) joint of the IDA. The grapple and any suspended payload hang in the direction of gravity. The IDA FSW computes the forward kinematics for the grapple and suspended payloads by first computing the forward kinematics for the IDA at the grapple attachment point. Because the attachment point is prior to the fourth joint, this requires only the first three joint angles. Then it adds the length of the grapple, plus grapple cable, plus any suspended payload, times the unit gravity vector, to the grapple attachment point to determine the origin of the grapple frame or the frame for the suspended payload. The z-axis for the grapple or payload frame is set to point in the direction of gravity, and the x-axis points radially away from the base of the IDA.

To compute the inverse kinematics, the IDA FSW reverses the process. First, it determines the position of the grapple attachment point by subtracting the length of the grapple, cable, and suspended payload, times the gravity vector, from the position of the grapple or payload frame. Then, it computes the first three joint angles of the IDA using the standard inverse kinematics, ignoring the fourth joint.

Because the links of the IDA are not rigid and deflect under the weight of the IDA itself plus any suspended payload, the IDA FSW compensates for this deflection. When the FSW receives a command to move to an absolute Cartesian pose, it computes the deflection as a transformation of the end-effector frame. It uses a model of the stiffness and mass properties of each IDA link and the gravitational vector. It also uses the configuration of the grapple and mass of any suspended payload, as indicated by the grapple phase.

Once the IDA FSW has determined the deflection, it compensates by changing the commanded end-effector frame by the inverse of the deflection transformation. This compensated pose is substituted for the commanded goal pose prior to generating the via-point sequences to move the arm. Gravity deflects the IDA and payload back to the commanded goal pose.

21.3.4 Fault Protection

During IDA activity, the IDA FSW checks for off-nominal behavior of the hardware and software. The typical fault response is to stop all motors and heaters, including the grapple heater, announce a fault, generate a telemetry file with the recent history of the IDA hardware and software state, and then transition to a faulted state. In this state, the IDA FSW will not execute any further commands except commands to exit the fault state. Normal command sequences do not include these commands, so ground operators must send them after determining the problem.

21.4 Deployment Workspace Analysis

Figure 21.8 shows a top view of the deployment workspaces for SEIS and WTS. The coordinates are specified in the IDA frame as the reference frame to be used for surface operations on Mars. The IDA frame origin is at the IDA arm base, fixed on the lander deck with the x-axis towards workspace and the z-axis down perpendicular to the lander deck. The nominal height of the lander deck is 1.05 m, which makes the z-coordinate of the level surface 1.05 m in the IDA frame. In Fig. 21.8, the lander deck and its nearby footpads touching the surface are represented by circles of 97.8 cm and 14.5 cm radii, respectively. The white area including all inner overlapping regions is the kinematically reachable SEIS/WTS payload grapple-hook workspace, where the IDA grapple holding the payload grapple hook can reach and perform instrument placement on the level surface with nominal lander deck height. The SEIS/WTS grapple-hook workspace is bounded by (1) an outer circular boundary constrained by kinematic reachability of the arm for both SEIS and WTS placements, (2) an inner circular boundary constrained by collision prevention between WTS (larger than SEIS) and the lander structure, and (3) side boundaries constrained by collision prevention between the arm and the lander deck. The gray area including all inner overlapping regions is the SEIS footprint boundary (19.8-cm radius circle) workspace. The blue area including all inner overlapping regions is the WTS footprint boundary (50.8-cm radius circle) workspace.

The green zone is the nominal IDA grapple workspace for SEIS and WTS deployments. Its outer circular boundary is reduced from the kinematically reachable workspace by several constraints: (1) manipulability avoiding near singularity regions, (2) torque limits, (3) arm joints back-drive (IDA does not have mechanical brakes), (4) SEIS tether length, and (5) payload recapture for relocation contingency. The radius of the outer circular boundary of the green zone is 1.65 m from the arm base. Its side boundaries are confined by the yellow and pink zones. In the yellow zone, the ICC field of view is partially occluded. In the pink zone, WTS deployment at higher height over SEIS requires more maneuvering to handle collision prevention.

To minimize the effect of the noise contributions of the lander, scientists prefer to place the SEIS as far away from the lander footpads as possible. One such location is

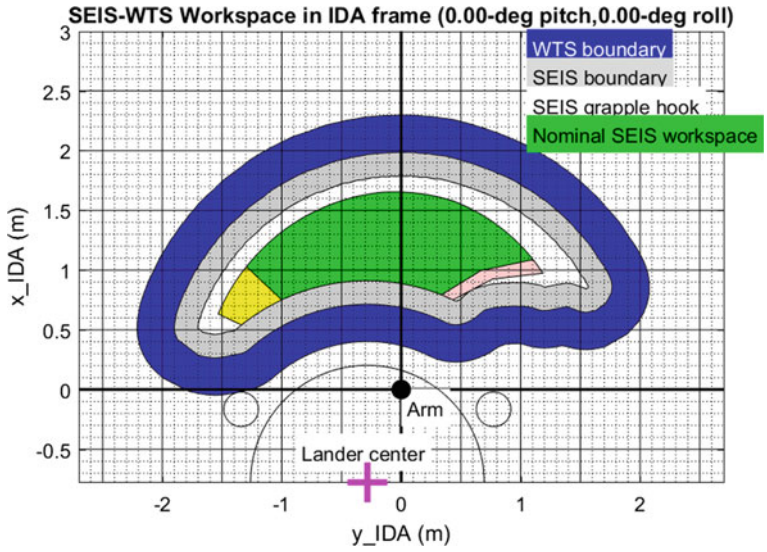

Fig. 21.8 SEIS-WTS deployment workspaces for level lander with level surface in IDA frame coordinates. The green zone is the nominal IDA grapple workspace for SEIS and WTS deployments. Its outer boundary's radial distance from the arm base is 1.65 m. In the yellow zone, ICC view partially occluded. In the pink zone, WTS deployment at higher height over SEIS requires more maneuvering to handle collision prevention

(1.65 m, 0 m) in xy coordinates at the intersection of the outer boundary of the green zone and the x-axis projected on the surface. Another location is (1.59 m, 0.44 m) along the tether peel direction.

Figure 21.9 shows the deployment workspaces for HP3. They are very similar to those for the SEIS and WTS deployments above. The blue area including all inner overlapping regions is the HP3 footprint boundary (31.2-cm radius circle) workspace. Scientists prefer to place HP3 as far away from the lander as possible, and more than 1 m away from SEIS.

Lander tilt has a significant effect on the payload deployment workspaces and must be considered. Figure 21.10 shows SEIS-WTS workspaces on the IDA frame xy plane for four different lander tilt cases with lander (IDA frame origin) height of 1.05 m from the level surface and lander footpads. Note that the workspaces and footpads on the surface are moving together, relative to the lander deck boundary. For positive lander pitch of Fig. 21.10 (a), the maximum x-coordinate in IDA frame for the nominal deployment was reduced to 1.45 m, versus 1.65 m for the level lander. For negative lander pitch of (b), it increased to 1.8 m, but the positive-y workspace zone shrank due to an arm collision issue. For positive lander roll of (c), the positive-y workspace zone increased while the negative-y workspace zone decreased. For negative lander roll of (d), the opposite trend happened.

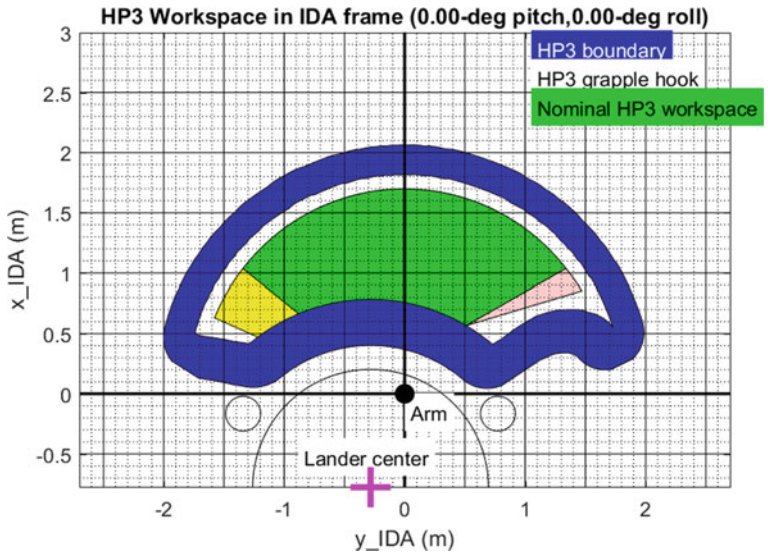

Fig. 21.9 HP3 deployment workspaces for level lander with level surface

Beyond the workspace reachability constraints, successful instrument deployments require knowledge of the 3-D workspace terrain to select the deployment site.

21.5 Workspace Imaging, Terrain Mosaic, and Site Selection

Prior to the instrument deployment, we need to know the workspace terrain in 3-D space, so we can model how the instruments interact with it. This requires a map, or digital elevation model (DEM) of the workspace, which is created by analyzing stereo images of the terrain acquired by the IDC. In order to minimize stereo baseline error, IDC stereo pairs are acquired by moving one arm joint only—the shoulder joint—while keeping all the other joints constant. The stereo overlap between left and right images is 80%, enabling generation of a workspace DEM comparable to those for the MER, PHX and MSL missions. IDC workspace imaging is done in event-driven mode with a sequence structure to capture images of the robotic arm's workspace in several tiers, starting with an inner tier close to the base of the lander and moving progressively outward. Only the IDA azimuth joint angle is changed within a tier. To move from one tier to the next, the IDA elbow joint angle is changed. Stereo images are identified using a unique 32-bit label, called the image ID, assigned to each image, which assigns each image as "left" or "right" in the stereo pair and specifies a number which is used to pair the images (Trebi-Ollennu et al. 2012).

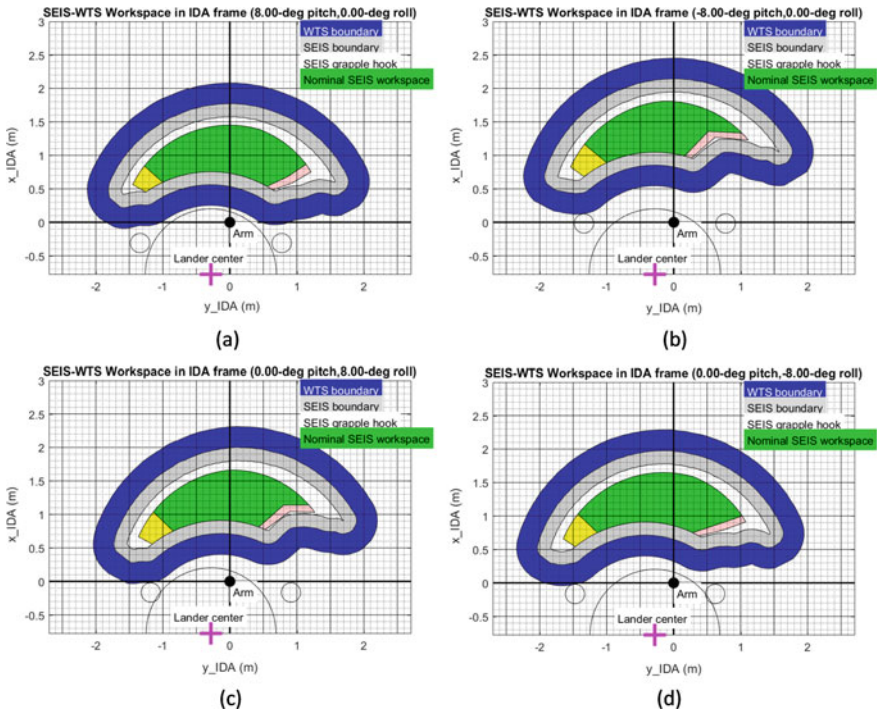

Fig. 21.10 SEIS-WTS deployment workspaces for tilted lander with level surface **a** 8-deg pitch and 0-deg roll lander **b** −8-deg pitch and 0-deg roll lander **c** 0-deg pitch and 8-deg roll lander, and **d** 0-deg pitch and −8-deg roll lander

The left eye and right eye image data products are used by a ground tool to perform stereo ranging. The range data from each pair of images is combined into a composite point cloud, which is then used to generate a 3-D model, also called a terrain mesh (see Fig. 21.11). Additional image data products are generated from the IDC stereo mosaic to support other science and operational goals (Abarca et al. 2018).

The workspace was imaged in three phases. First, an ICC context image was acquired on Sol 4. That provided an overview and was used for preliminary discussions. The primary workspace mosaic was acquired on Sol 12. This mosaic consisted of 56 images containing 26 stereo pairs in three tiers and four extra images between tiers for improved registration. This primary workspace mosaic was used for the bulk of the Instrument Site Selection Working Group (ISSWG). After preliminary deployment locations for the instruments were determined, high-resolution mosaics of just those two locations were acquired on Sol 16, consisting of 24 frames each. High resolution in this case means putting the camera closer to the terrain (1.2 m instead of 1.5 m). This could not be done earlier because of safety concerns; the workspace below the deck had to be proven to be free of obstacles (using the primary workspace mosaic) before the arm could go below the deck, as is required for the higher-resolution images.

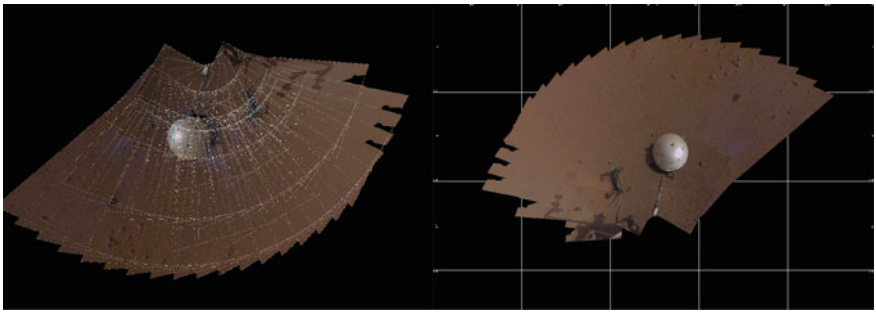

Fig. 21.11 Sol 110 post-instrument deployment vertically projected mosaic showing the location of the SEIS and HP3 instruments. Left: The white lines show the overlap of the 56 images that make up the workspace mosaic. Right: The resulting bundle-adjusted workspace mosaic

The pre-landing plan for workspace imaging and processing is fully described in (Abarca et al. 2018). Surface operations matched the pre-launch plan remarkably well, and the primary workspace mosaic was acquired and downlinked more quickly than anticipated. The actual landing site turned out to be much more benign than any of the testbed runs leading up to the mission.

As expected, the accuracy of telemetered arm positions was not sufficient to create reliable stereo products. The IDA requirement was to be able to determine stereo baseline accuracy to 2.8 mm. That amount of error in baseline knowledge translates to a theoretical ~2 cm of range error, which is too much to support accurate analysis of instrument placement. Therefore, the pointing knowledge of the arm must be corrected before the DEM can be computed. This is accomplished using a bundle adjustment procedure based on image tie points (Abarca et al. 2018). This process, developed before landing, worked very well, with mismatch of derived terrain in overlapping areas between frames measured at an average of 1.9 mm, up to 3.9 mm (compared to an average of 5.0 mm, up to 10.0 mm without bundle adjustment).

The high-resolution mosaics were similarly co-registered via bundle adjustment, additionally tying them to the "base map" created by the primary workspace mosaic as a control network. The process was repeated several times during instrument deployment with high-resolution mosaics co-registered to the base map on subsequent sols 35, 44, 58, 59, 62, 85, 182, 227, 230, and 240.

21.5.1 Deployment Image Products

A series of derived instrument placement products were created to show the instrument tilt, roughness of terrain under the instrument, and delta tilt between the SEIS and WTS at any point in the workspace, as well as an overall "goodness map" that combined them all, using thresholds for each instrument versus the instrument's requirements. These products were generated on the workspace and high-resolution

mosaics. The algorithms behind them are fully described in (Abarca et al. 2018). They were utilized by both the science and engineering teams to decide on the instrument deployment locations as part of the Instrument Site Selection Working Group (ISSWG) process. It was determined after landing that the delta-tilt product was not particularly useful operationally, but all others products were used as expected based on pre-landing testing.

The instrument tilt product analyzed, for each pixel, what the minimum and maximum tilt of the instrument would be if placed at that pixel (meaning, the grapple point was directly above that position in the map). This required modeling the interaction of the feet of the instrument with the terrain, taking into account the clocking angle as well as possible sinkage of feet into the terrain. Clocking angle means rotation of the instrument about its axis. Clocking of the InSight instruments is fairly well constrained by the tether, but there is no active control over clocking angle so a range of possibilities was analyzed. Also the feet have broad pads with a spike on the bottom; the foot could then sit on either the pad or the spike depending on the firmness of the material (e.g., rock vs. loose regolith). The minimum and maximum tilt angles across all combinations of clocking and sinkage were gathered for each pixel and compared to thresholds based on the instrument requirements. A "goodness" band indicated whether the tilt met the thresholds and is shown as green (good), yellow (marginal), or red (bad) in Fig. 21.12.

Instrument tilt should not be confused with terrain slope; terrain slope is an intrinsic feature of the terrain, while instrument tilt is a measure of how the instrument interacts with that terrain.

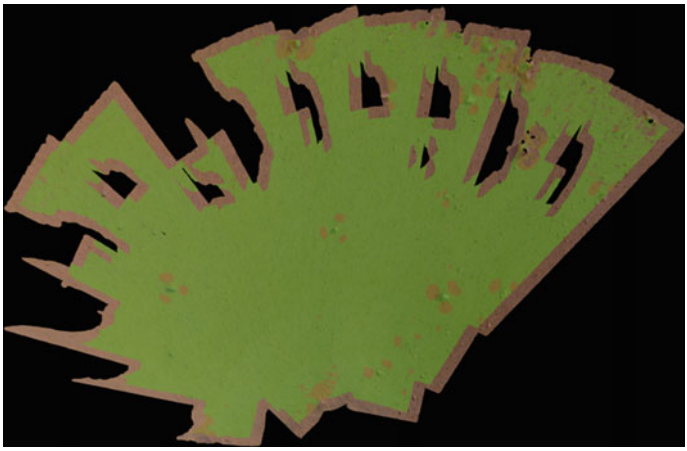

Fig. 21.12 Sol 12 orthorectified projection of the InSight Deployment Workspace with SEIS "instrument tilt" overlaid. Green pixels represent locations that pass all of the instrument placement criteria, orange represents one violated requirement, and red means at least two criteria are violated

Delta tilt looked at the difference in possible tilts between the SEIS and WTS; since the SEIS goes inside the WTS it is important to make sure they do not interact (touch).

Instrument roughness is a measure of the roughness of the terrain where it interacts with the instrument. Roughness generally is defined as the maximum peak-to-peak excursion of the terrain above or below a plane. There are two types: feet and body.

Feet roughness is a measure of roughness directly underneath the feet. A rough area for a foot means placing the foot there is unreliable; it could perch on a pebble or sink into a hole. As with tilt, the roughness was measured across a range of clock angles and sinkages.

Body roughness is more interesting; it is a measure of the terrain underneath the instrument. It measures only excursions above the plane of the feet (e.g., rocks); valleys or holes are ignored. This ensures the instrument is not "high-centered", i.e., placed over a rock where it sits on the belly rather than the feet. That would be particularly troublesome for the seismometer, which depends on good foot contact with the ground.

The roughness product also had a goodness band that indicates whether the requirements were met at that location. Green indicates both body and feet are within thresholds, yellow indicates one is out of spec, and red indicates both. See Fig. 21.12. (Fig. 21.13).

These maps—tilt, delta tilt, roughness, as well as the workspace reachability discussed in Sect. 21.4—were then combined into a single "goodness" map, which indicated where all criteria were met (in green). See Fig. 21.14.

The goodness map was the primary screening tool used by the ISSWG to determine instrument placement (Golombek et al. 2018). It provided a quick overview of

Fig. 21.13 Sol 12 orthorectified projection of the InSight Deployment Workspace with WTS "instrument roughness" overlaid. Green pixels represent locations that pass all of the instrument placement criteria, orange indicates one requirement has been violated, and red means at least two criteria are violated

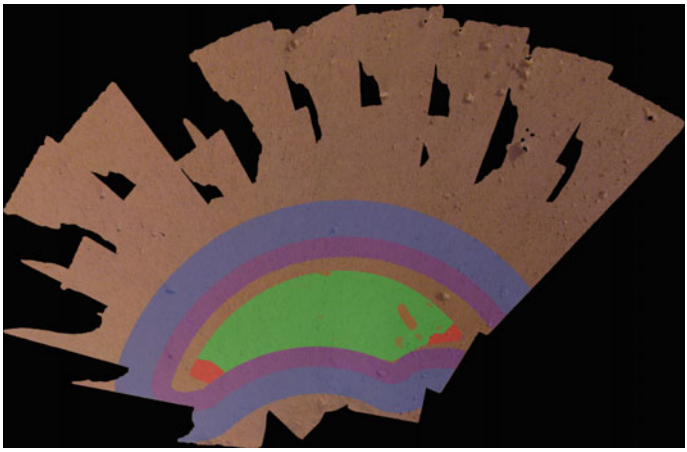

Fig. 21.14 Sol 12 orthorectified projection of the InSight Deployment Workspace with arm reachability and HP3 "goodness" overlaid. Green pixels represent locations that pass all of the instrument placement criteria, orange indicates one requirement has been violated, and red means at least two criteria are violated. Purple and blue colors indicate the extent of feet and other constraints

locations that were safe according to the analysis, with backup materials in all the other products to determine details for areas that might be marginal. As it turned out, the terrain was virtually all green across the entire workspace.

21.6 Payload Localization

On InSight, stereo imagery captured by moving the IDA was designed as the primary method for localization (Bailey et al. 2020). However, several constraints (both physical and planning related) on the system made capturing stereo data regularly difficult for day-to-day operations planning during deployment. For example, through the deployment of each instrument, the IDA was physically grappled to the instrument and thus could not capture the requisite stereo baseline. However, a measurement of the deployed location was desired as a check to know if it was safe to ungrapple. On the planning side, it was often difficult or impossible to schedule stereo imaging activities with proper lighting, or there was not enough downlink data volume available for stereo. As a result, robust monocular localization techniques were developed to allow the team to localize each instrument in the workspace with respect to the arm as well as trend their positions over time. The primary methods of monocular localization used were fiducial localization and grapple localization.

Fiducial markers were placed on each instrument and visible in the IDC field of view while the arm was grappled to the instrument; on the ground, we used this IDC image to calculate the pose of each instrument. The fiducial markers on the SEIS and WTS were "MSL-style" single-point fiducials, which provide a single measured

point on the instrument relative to the instrument origin. Unlike SEIS and WTS, the HP3 SSA visibility in the IDC field of view was much less favorable while grappled. As a result, many of the "MSL-style" fiducials on the HP3 were often not visible. To mitigate this, AprilTags were placed on the HP3. AprilTags consist of a rotationally asymmetrical pattern such that seeing a single tag in an image provided five points for position and orientation measurements. These methods of fiducial localization had an absolute error bound of about 1.5 cm and 1.5 degrees of accuracy. The sources of error on the localization were partly due to the relatively low resolution of the IDC, but had more to do with the kinematics of the IDA. For example, the computed IDC position error included propagated errors in robotic arm joint knowledge and in the robotic arm deflection model compensation. While this magnitude of error was acceptable with respect to mission requirements, the IDS team refined the measurements by calculating the bias locally to reduce risk on high-precision activities.

The method for calculating the bias for both stereo and monocular measurements was called "grapple localization." This activity was a qualitative check of an instrument position using the arm. After each deployment and un-grappling, the IDS team placed the grapple at a known offset of 4 cm above the measured position of the instrument's grapple hook. The team captured IDC images at this position and at 2 cm below this position. These images would give a clear qualitative indication of localization offset with respect to the arm frame location. The error in the direction perpendicular to the IDA would be very clear based on the alignment with the grapple hook and the grapple in the image. The error in the radial direction was visible in the alignment of the instrument and grapple shadows. Finally, the instrument height would be calculated using the distance from the grapple to the hook compared to the known baseline of the 2 cm move down between images. Thus, using grapple localization, the IDS team could calculate a local bias and refine future measurements. These methods became crucial for the WTS deployment in particular, where the margin for error on the deployment was very tight in order to prevent collision with the SEIS and LSA (Fig. 21.15).

21.7 Robotics Operation Tools

21.7.1 Robot Sequencing and Visualization Program

RSVP (Robot Sequencing and Visualization Program) is used to assist operators in planning InSight's arm motions and instrument deployments. Prior to uplinking commands to Mars, it is necessary to simulate and visualize the predicted motions of the arm for safety and correctness.

RSVP is composed of two main components. RoSE (Robot Sequence Editor) allows operators to write command sequences. HyperDrive simulates the sequences and visualizes the resulting arm motions.

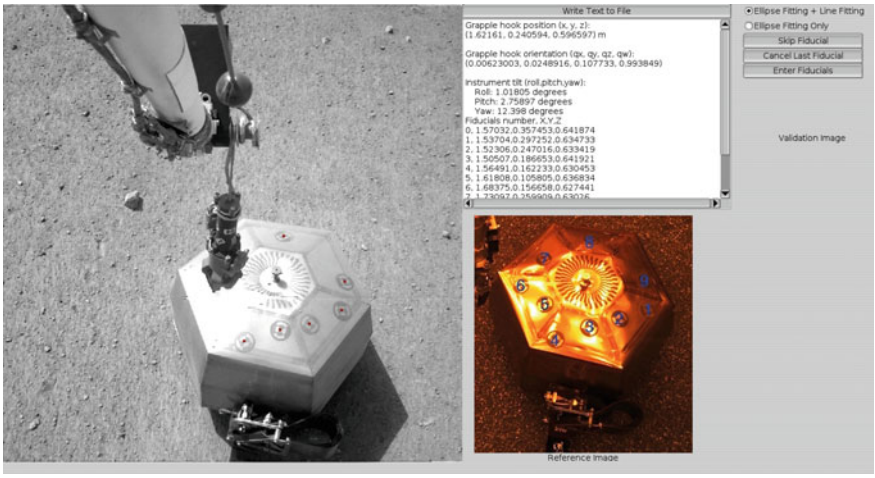

Fig. 21.15 Fiducial Localization Software with detected points remapped to the image

RSVP was first developed for driving the Sojourner rover. Newer versions of the software are used for many of the current and past Mars robotic missions. InSight's version of RSVP shares common multi-mission code with Mars Science Laboratory's version of RSVP. InSight's arm simulation is derived from the arm simulation used on the Phoenix Mars mission (Yen et al. 2004) (Fig. 21.16).

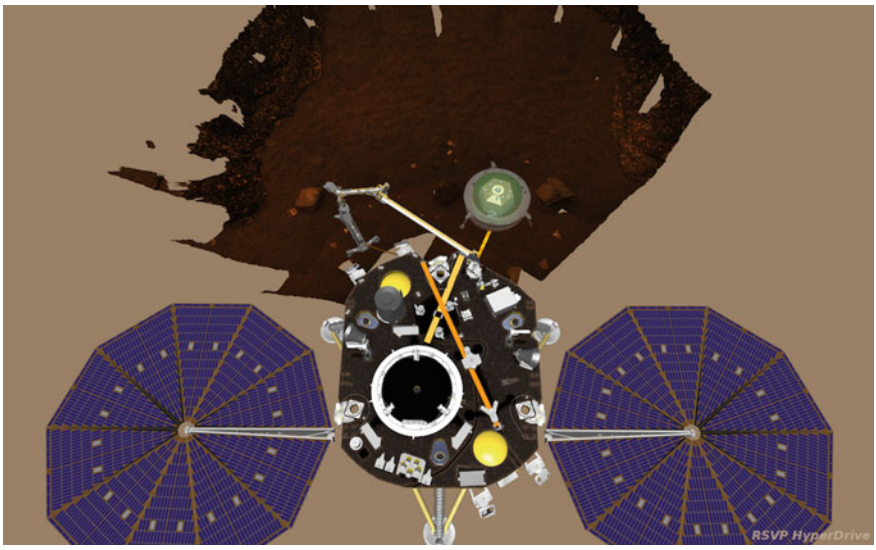

Fig. 21.16 InSight RSVP with instruments deployed

21.7.2 Simulation of Commands Using Flight Software in the Loop

RoSE is configured with knowledge of InSight's command dictionary, allowing operators to construct syntactically correct commands. These commands are stored in the Robot Markup Language (RML) format. But safely and correctly operating the arm requires more than mere syntactic correctness. This is where HyperDrive comes into play. HyperDrive is the companion program to RoSE. The two programs send messages to each other via inter-process communication (Fig. 21.17).

HyperDrive parses the arm commands from RoSE and runs them through its embedded InSight arm flight software. HyperDrive stubs out the lower-level hardware interfaces that would normally go to actuators on the arm. Instead, HyperDrive captures the kinematic motions of the arm, which it then visualizes in 3D for the user to evaluate and iterate on. Collision faults and other faults reported by flight software are reported to the user, allowing them to resolve these issues on the ground before uplinking the command sequences to Mars (Fig. 21.18).

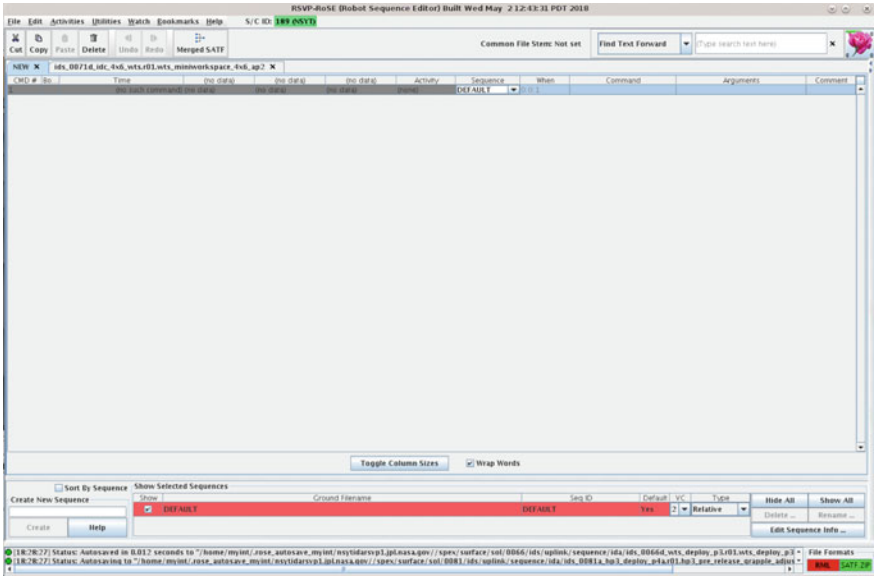

Fig. 21.17 RoSE user interface

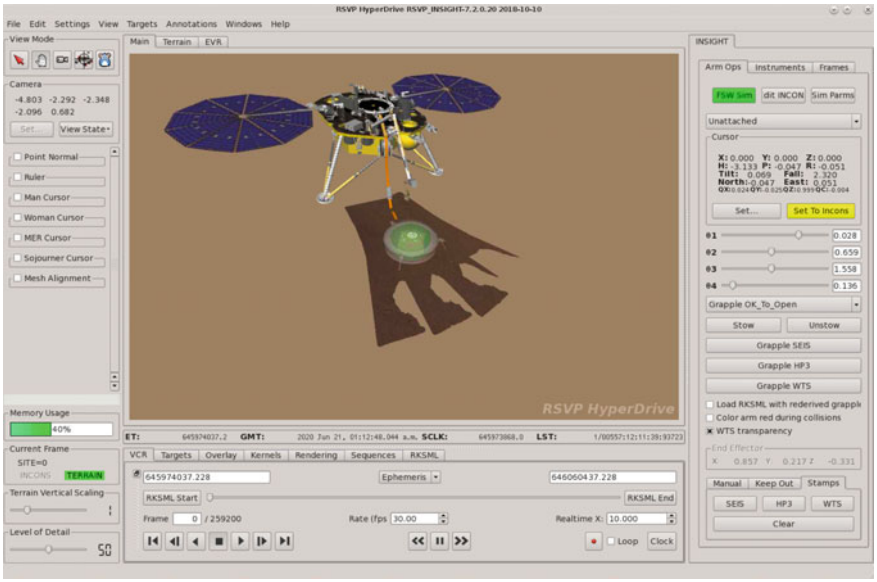

Fig. 21.18 HyperDrive user interface

21.7.3 Collision Volumes

In addition to the aforementioned collision checking during simulation of the command sequences, HyperDrive supports an interactive collision volume visualization. As the user moves the arm, HyperDrive can display the collision volumes for any given arm pose in real time. It does this by querying flight software for its collision volume information. HyperDrive also queries flight software for the collision state and displays objects that are in collision in red (Fig. 21.19).

21.7.4 Instrument Simulation

Building on top of the flight software-informed arm motions, HyperDrive simulates the motions of InSight's grapple, instruments, and tethers. Based on the grapple state in flight software, HyperDrive derives the position and the open/closed state of the grapple.

HyperDrive uses a catenary model for its simulation of the tethers connecting the SEIS and the HP3 to the lander. This gives the operators a rough estimate of where to expect the tether during each instrument deployment. To simulate the instrument/terrain interaction, HyperDrive fits the instrument plane to the local terrain patch (Fig. 21.20).

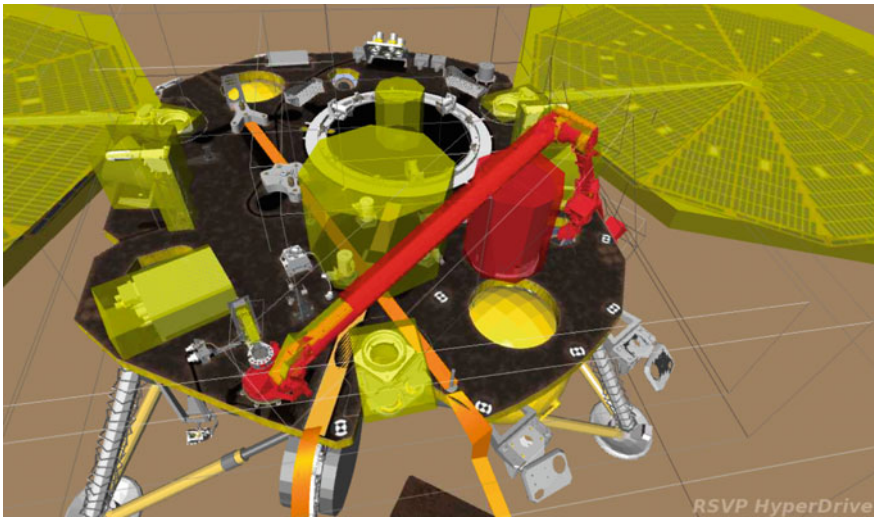

Fig. 21.19 Collision volumes in HyperDrive

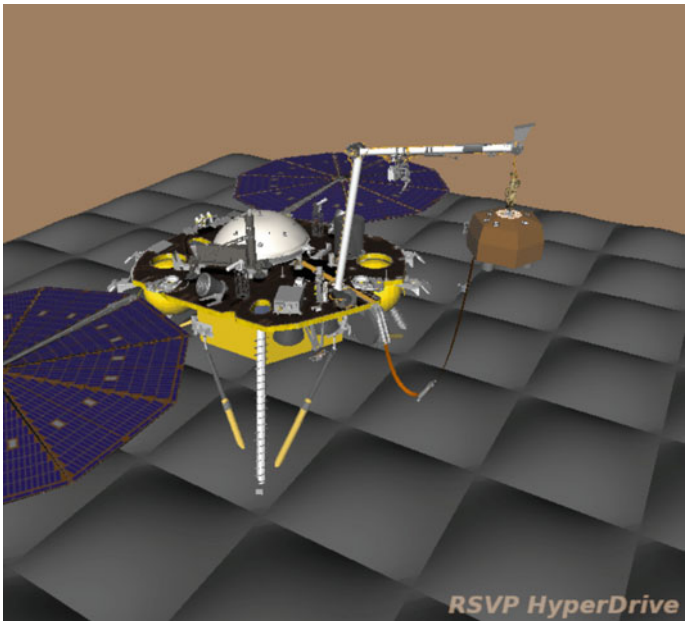

Fig. 21.20 Grapple in HyperDrive

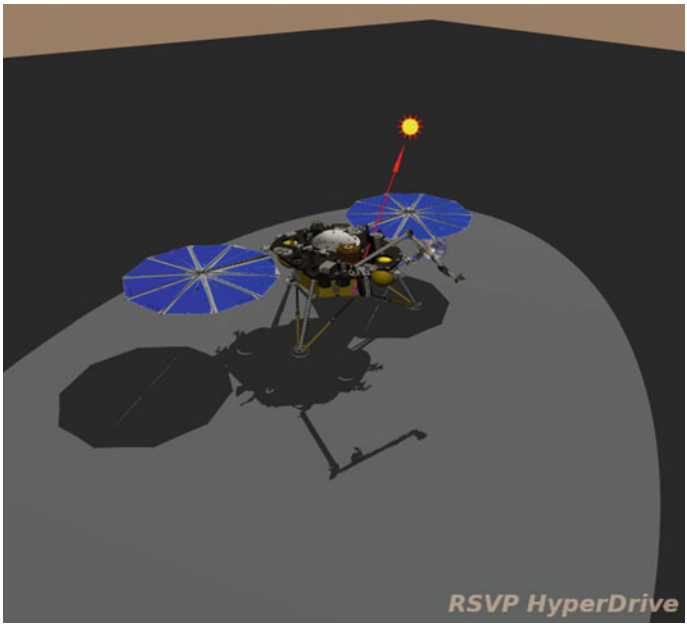

Fig. 21.21 Shadows in HyperDrive

21.7.5 Shadow Modeling

Using the InSight mission's SPICE (Spacecraft, Planet, Instrument, C-matrix, Events) kernels, HyperDrive supports visualizing the location of the Sun with respect to the lander. It renders the lander's shadow on the terrain to assist operators in constructing imaging commands (Fig. 21.21).

21.7.6 SEIS Tether Catenary Modeling

To successfully deploy the seismometer on the Martian surface, the team needed to consider the dynamics of the entire instrument system, which included its tether (Sorice et al. 2020). During SEIS deployment, Part 3 (as is described in Sect. 21.8), the robotic arm placed SEIS on the surface of Mars. In this deployment, key components of the tether were considered: the tether storage box (TSB), the field joint (FJ), the pinning mass (PM), and the load shunt assembly (LSA). Once the tether is peeled, the tether hangs from two points: one point at the exit of the TSB mounted on the lander, and the other point at the LSA loop where it attaches to the instrument. It is at this point that the tether takes on the catenary shape, and our quasi-static model

is interpreted to help us predict the position of the pinning mass and the field joint along the tether as we move the SEIS away from the lander and over the terrain.

When deploying SEIS in the deployment testbed and with the flight hardware, the team determined that the configuration of the tether had a significant impact on the SEIS instrument placement accuracy. After observing these effects, the team decided to develop a model to predict which part of the tether would touch the terrain first, given knowledge of the height of the terrain from the lander to the desired deployment target. By modeling the tether height throughout the deployment trajectory, we could control the deployment of the tether itself and utilize the tether–terrain interactions to our advantage. We chose to use a catenary equation rather than an arbitrary polynomial equation for the ease of computation, verification, and implementation in our simulation and visualization software. Additionally, the tether's stiffness and bending behavior made it unlikely that we could use a single polynomial equation across the entire deployment workspace.

In deploying our payloads, we commanded the robotic arm to a series of waypoints, each followed by a set of images with the IDC and ICC. During the deployment sequence, the arm (and the payload) remain still during the duration of the image acquisitions (roughly 1.5–2 min) before moving on to the next commanded position. By choosing to model the tether and solve the equation at these specific configurations along the trajectory of SEIS deployment, we took a quasi-static approach and ignored inertial effects. This simplified the problem while still remaining relevant for understanding the movement and position of the tether during deployment (Fig. 21.22).

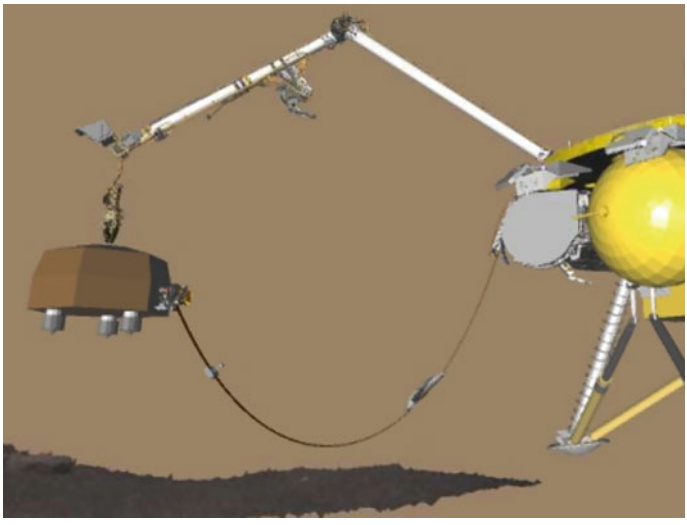

Fig. 21.22 RSVP simulation with suspended SEIS tether over terrain mesh

21.8 Surface Operations Results

21.8.1 SEIS Deployment

Each payload deployment (lift from the lander deck to placement on surface) consists of four parts (Part 1, Part 2, Part 3, and Part 4). After completion of each part there is a ground-in-the-loop GO/NO GO decision to execute the following part. In addition, there can be no intervening motion of the IDA between the four parts. On Sol 18 SEIS deployment Part 1 was successfully executed on Mars. It entailed moving the unstowed grapple to the SEIS teach point, leaving the grapple with about 4 cm between the bottom of the grapple fingers and the top of the SEIS grapple hook (see Fig. 21.23). Four IDC images were acquired at this IDA pose to document and verify successful alignment needed for the GO/NO GO decision for SEIS deployment Part 2.

SEIS deployment Part 2 (SEIS capture) was successfully executed on Sol 20. The SEIS deployment Part 2 sequence opened the grapple fingers, moved the IDA grapple end effector 4 cm toward the SEIS grapple hook, and then waited for the grapple fingers to close around the SEIS grapple hook. Grapple telemetry and IDC images were used to confirm successful SEIS capture (see Fig. 21.23).

On Sol 22 SEIS deployment Part 3 was successfully executed on Mars. First the SEIS launch restraint frangibolts were actuated successfully and SEIS was lifted 30 cm above the SEIS launch restraint pedal stools. Then, the SEIS tether, which was held to the lander deck and TSB by hook-and-loop fasteners, was peeled off. This was followed by the IDA extracting additional SEIS tether from the TSB up to the SEIS tether chock. The IDA then successfully placed the SEIS on the surface of Mars.

SEIS deployment Part 4 (SEIS release) was successfully executed on Sols 24 and 25. On Sol 24 the grapple cable slack was reduced by moving the IDA grapple end effector 1 cm up from the SEIS grapple hook. On Sol 25 the grapple was opened, and moved 10 cm up from the SEIS grapple hook to release the SEIS, marking the first ever successful precision robotics instrument placement and release on another astronomical body.

From Sols 26 to 36 the SEIS science operations team successfully performed SEIS health checkouts. On Sol 37 the remaining SEIS tether in the TSB was released by opening the TSB door. This was followed by IDC documentation of the SEIS tether configuration on the surface of Mars. On Sol 40 the SEIS load shunt assembly (LSA) frangibolt was actuated successfully but the LSA failed to meet the minimum separation gap that is intended to dampen the effects of tether thermoelastic noise on the seismic measurements. A successful minimum separation of the LSA is required for SEIS to meet its performance requirements.

From Sols 41 to 50 we acquired additional IDC images of the SEIS tether pinning mass (PM), LSA, and SEIS tether configuration on the surface on Mars in preparation for using the scoop to separate the LSA to meet the minimum separation gap required for SEIS to meet its performance requirements. On Sols 56 to 59 we successfully used the IDA scoop to move the PM via the grapple hook to assure successful LSA

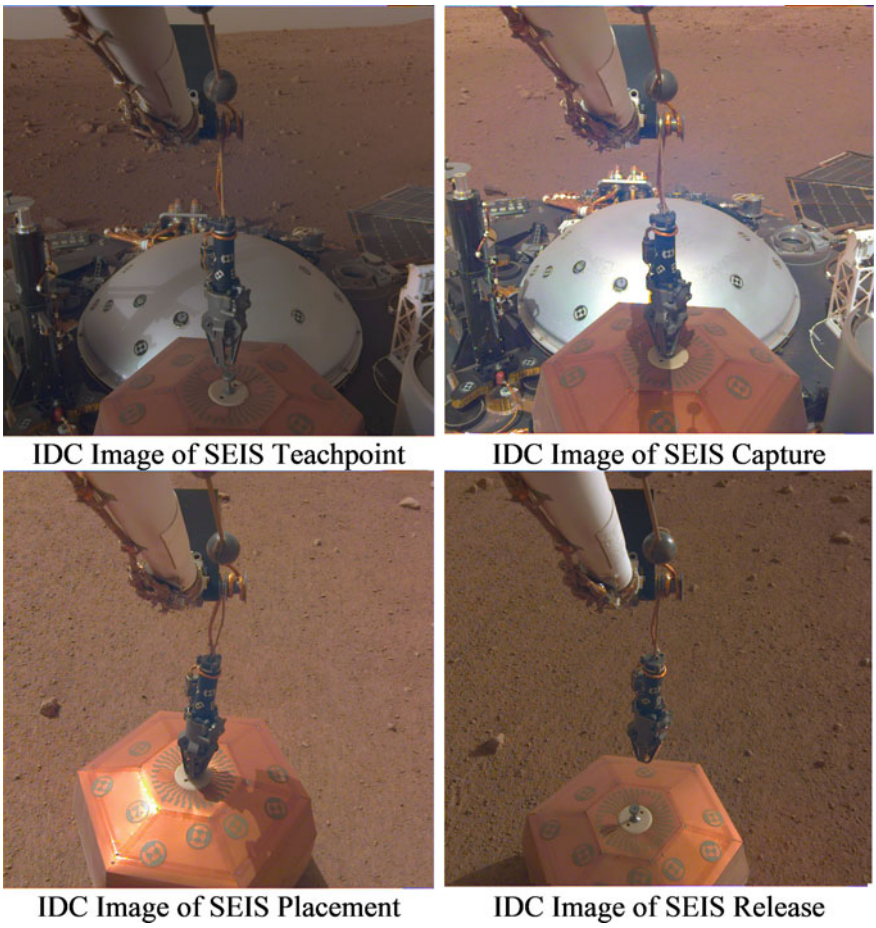

Fig. 21.23 IDC images showing successfully SEIS Deployment Parts 1–4 on the surface of Mars

minimum separation (see Fig. 21.27). From Sols 60 to 62 the SEIS science operations team successfully performed SEIS health checkouts in preparation for placing WTS over the SEIS.

21.8.2 WTS Deployment

The wind and thermal shield (WTS) is a cover that goes over the SEIS (Fig. 21.24). It has a rounded hollow aluminum cap supported by three feet. The feet were folded up beneath the cap when the WTS was on the lander deck (Fig. 21.25), and springs forced

them to unfold when the WTS was lifted from the deck. There is a compliant bellows-like aluminized-Kapton skirt, weighted down by steel scale mail and chainmail, which descends from the hard cap. The mail pulls the skirt down and makes it conform to the undulations of the Martian terrain and SEIS tether. The purpose of the WTS is to reduce noise in the seismic data collected by the SEIS. It does this by shielding the SEIS from the Martian wind and by reducing the magnitude of temperature swings throughout the Martian day.

The deployment of the WTS required accurate and precise placement. The WTS dome is not much larger than the combined SEIS remote warm electronics box (RWEB) (Fig. 21.26), which is the outer body of the SEIS, and the SEIS load shunt array (Fig. 21.27), where the tether attaches. To preserve the integrity of the SEIS data, the WTS must not contact the SEIS when placed over it. Therefore, there is not much margin for error in WTS placement. Error in WTS placement is a combination of IDA positioning error, error in SEIS location knowledge, error in surface slope knowledge, and error introduced by WTS movement by wind.

A WTS do-not-exceed (DNE) envelope was defined to account for up to 4 cm of placement uncertainty. If the chosen WTS placement location does not cause the SEIS to penetrate the WTS DNE, then the WTS should be guaranteed not to contact the SEIS. Potential differences in instrument tilt, due to undulations in the terrain, can shrink the range of placement locations that do not cause a WTS DNE violation.

The choice of WTS placement location was based on knowledge of the SEIS location and knowledge of the Martian surface slope in the area surrounding the SEIS. A relative offset, along the SEIS tether in the direction of the lander, was added to accommodate the opening of the load shunt array (LSA) (Fig. 21.27). Based on the maximum expected LSA opening distance and a maximum expected WTS placement error of 3 cm, a WTS/SEIS offset of 5 cm was chosen. This provides

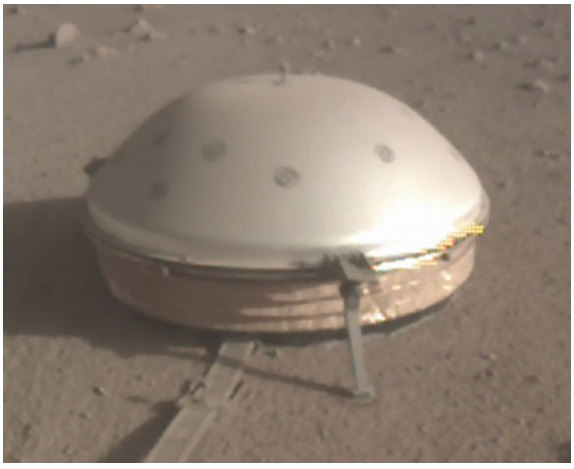

Fig. 21.24 The wind and thermal shield (WTS) on the Martian surface on Sol 394. The SEIS is underneath the WTS, and the SEIS tether can be seen protruding from underneath the WTS skirt

Fig. 21.25 The WTS on the lander deck before launch. The feet were folded to the side when the WTS was on the deck. They were pushed open by springs when the WTS was lifted

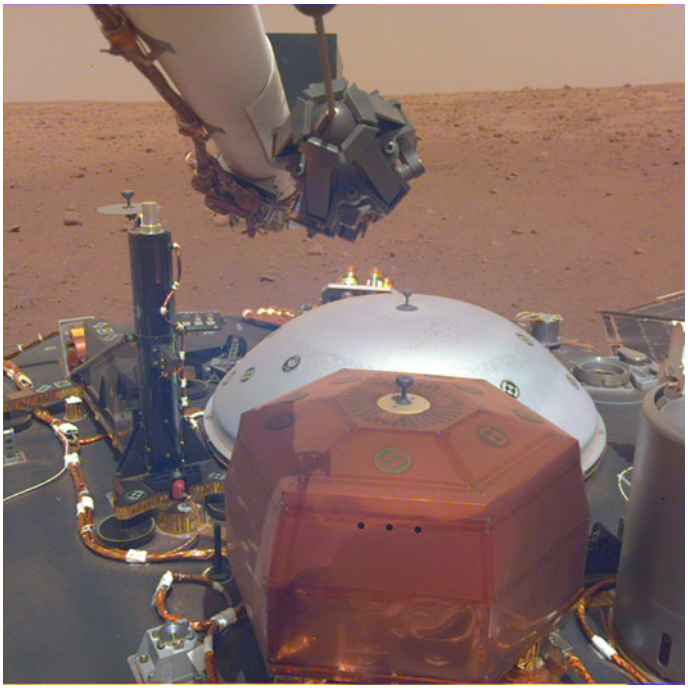

Fig. 21.26 The SEIS (foreground), WTS (behind the SEIS), and HP3 (to the left of the WTS) on the lander deck. The WTS is in its folded-leg configuration. The relative sizes of the SEIS and WTS can be seen

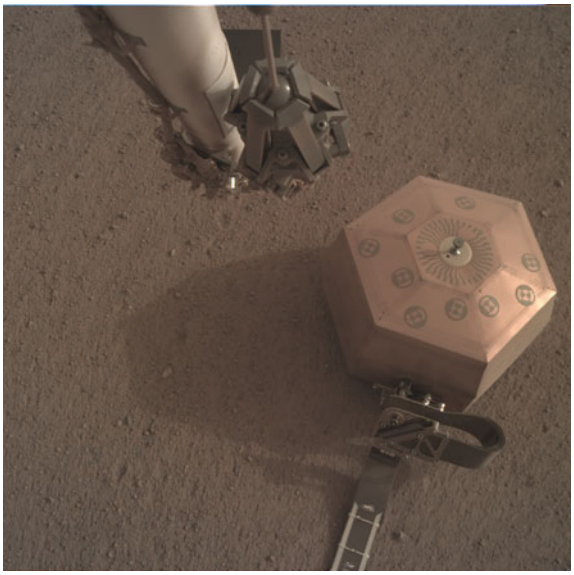

Fig. 21.27 The SEIS on the Martian surface with the LSA opened. Both the body of the SEIS and the LSA must fit beneath the WTS without contacting it

a margin of 3.5 cm of space between the far side of the SEIS and the WTS and a margin of 0.5 cm between the LSA and the WTS. The WTS/SEIS offset was chosen to allow more margin between the RWEB and the WTS than between the LSA and the WTS, because contact with the RWEB would cause a greater degradation to the quality of the scientific SEIS data.

The surface slope can cause an offset of the WTS as it is set down, where one foot contacts the ground before the others and then the WTS pivots around this foot. There was no control over the clocking, or rotation, of the WTS during deployment. The grapple and WTS grapple hook are designed to minimize clocking, but the compliant grapple cable still allows some clocking. Therefore, there was no control over the exact positions of the WTS feet on the surface. Because the positions of the feet on the Martian surface determine the effective surface slope in the WTS deployment location, this uncertainty in clocking added to the uncertainty in deployed WTS tilt, which in turn, affects the relative SEIS/WTS separation.

The temperature and wind sensors (TWINS) instrument, which is part of the auxiliary payload sensor system (APSS) (Banfield et al. 2020), provides wind speed data in the environment around the lander. This data was used to help choose a time of day for WTS deployment that was not excessively windy. Because the grapple cable is compliant, excessive winds could cause undesired movement of the WTS during deployment, which would affect the accuracy of the placement.

The IDA actuators were not designed to hold heavy loads such as the WTS, which has a mass of 9.5 kg, in an outstretched pose while not powered. In between each individual motion of the IDA, the motors are powered off. This means that the motors may back-drive while under a heavy load such as the WTS. Back-drive during WTS deployment contributes to WTS placement error. To minimize the amount of IDA back-drive, the WTS was deployed when the motors were cold, and therefore there was more friction in the joints. The IDA joints needed to be less than or equal to -20° Celsius to prevent back-drive, with colder temperatures preferred. However, the IDA has minimum temperature constraints during use to prevent damage to the motors. Additionally, there must be enough sunlight to allow useful images to be taken with the cameras during deployment. These restrictions created two time windows in which the WTS could be deployed. One was in the morning and one in the late afternoon. After consideration of the wind speeds and the timing of the available communications windows, the morning deployment window was chosen.

Like the SEIS and HP3 deployments, the WTS deployment was divided into four main parts, with an additional grapple positioning adjustment that had to be performed prior to executing the fourth part. Each of these parts, plus the grapple positioning adjustment, required ground-in-the-loop confirmation of success before proceeding to the next part. Therefore, each part had to be executed in a planning cycle of its own. There could be no intervening motion of the IDA between these deployment parts.

WTS deployment Part 1 was executed on Sol 63 of the mission. On this sol, before starting the WTS deployment, we commanded a move to put the grapple above the SEIS grapple hook to confirm our localization of the SEIS. To start the deployment, we moved the IDA to our "deck-ready-in" pose, a neutral pose from which the IDA can easily move the grapple to above any of the three instrument grapple hooks. Then we moved the grapple to a position such that the grapple frame was 5 cm above the WTS grapple hook frame. This leaves about 4 cm between the bottom of the grapple and the top of the grapple hook. We took an IDC image here and disabled the collision checking in the IDA FSW between the grapple fingers and the WTS grapple hook. Then we moved the grapple down 2 cm closer to the grapple hook and took another IDC image. We moved up 2 cm and then down 2 cm again, taking images (see Fig. 21.28) at each of the poses. We used the IDC images and the IDA joint angles to confirm that the grapple was in the correct location for the capture in the next deployment part. The images at different heights above the grapple hook assist in determining whether the grapple is sufficiently aligned over the grapple hook, because we do not have stereo vision and cannot move side to side for this determination. On Sol 64, in preparation for deployment Part 3, we monitored the IDA joint temperatures at the chosen deployment time, to verify that the temperatures were still cold enough.

WTS deployment Part 2 was executed on Sol 65. We opened the grapple, moved the grapple down 4 cm, and then let it close around the WTS grapple hook. We used IDC images, grapple telemetry showing correct grapple operation, and IDA joint angles to confirm that the IDA was in the correct pose and had successfully captured the WTS grapple hook.

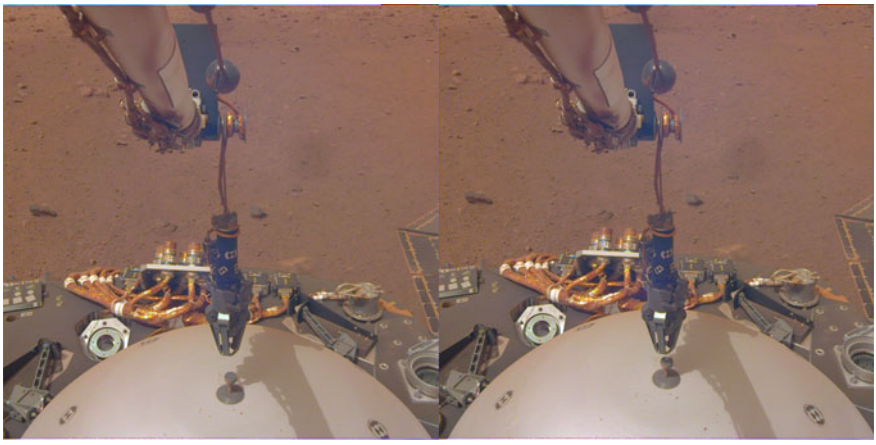

Fig. 21.28 The grapple above the WTS grapple hook during WTS deployment Part 1. In the left image, the bottom of the grapple fingers are approximately 4 cm above the top of the WTS grapple hook. In the right image, they are 2 cm above the top of the grapple hook, and this is the final position after the completion of WTS deployment Part 1

WTS deployment Part 3 was executed on Sol 66. In this part, we moved the WTS from the deck to the Martian surface. First, we actuated frangibolts which held the WTS to the deck, freeing the WTS. Then, we lifted the WTS 18 cm in the direction opposite the pull of gravity. This lifted it to its standup position, such that the legs extended to their upright position supporting the weight of the WTS. Then, we lifted it another 43 cm in the anti-gravity direction, such that it was high enough for us to swing it over the rest of the deck without hitting anything. We then moved it a few centimeters outward (straight ahead in Fig. 21.29), so that it would not fully eclipse the UHF antenna when we moved it around to the IDA workspace. By rotating only the IDA azimuth joint, we moved the WTS to the front of the lander over the workspace. In a series of stair-step motions, we moved the WTS back in towards the base of the arm and downward towards the surface. Once the WTS was lower to the surface, but not so low that it would strike the SEIS, we moved the WTS to the high standoff position above the SEIS, directly over the chosen WTS deployment location. After lowering the WTS in the anti-gravity direction to the low standoff position, just above the height of the SEIS, we disabled the IDA FSW collision checking between the WTS and SEIS. Finally, we lowered the WTS to the Martian surface in multiple steps, taking IDC images along the way. Because there was some uncertainty in the height of the ground, we overdrove the final movement, meaning that we moved the grapple lower than necessary to place the WTS on the surface. This overdrive guaranteed that the WTS would not be partially suspended at the end of the deployment, even with uncertainty in the terrain height. We used IDC and ICC images, and IDA joint angles, to determine that the WTS was in the correct location and fully supported by the surface. We looked for slack in the grapple cable in the images to determine that the WTS was not still partially suspended (Fig. 21.30).

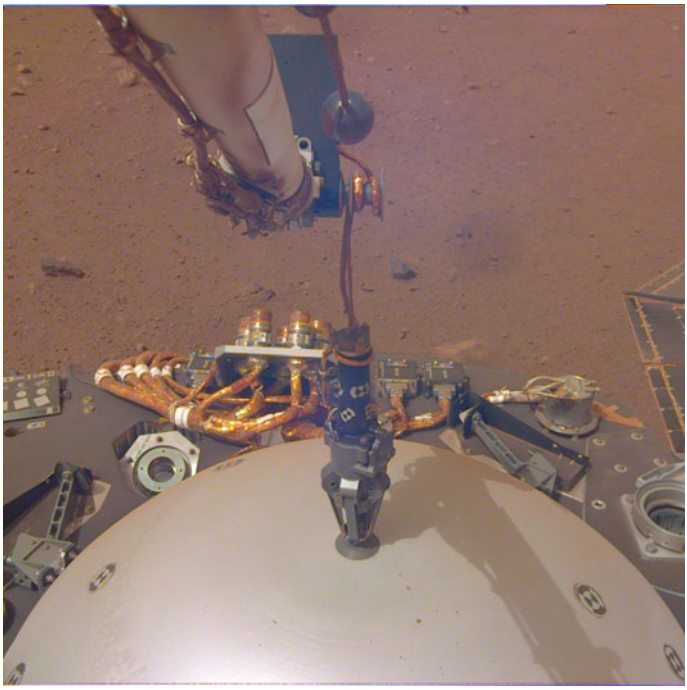

Fig. 21.29 The grapple closed around the WTS grapple hook at the completion of WTS deployment Part 2

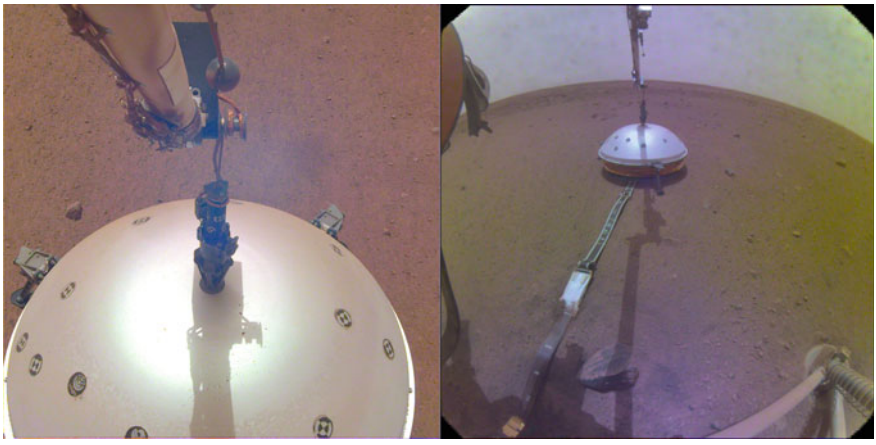

Fig. 21.30 The WTS and grapple after completion of WTS deployment Part 4a. The grapple has been moved up and over a little to reduce the slack in the grapple cable and center it over the WTS grapple hook

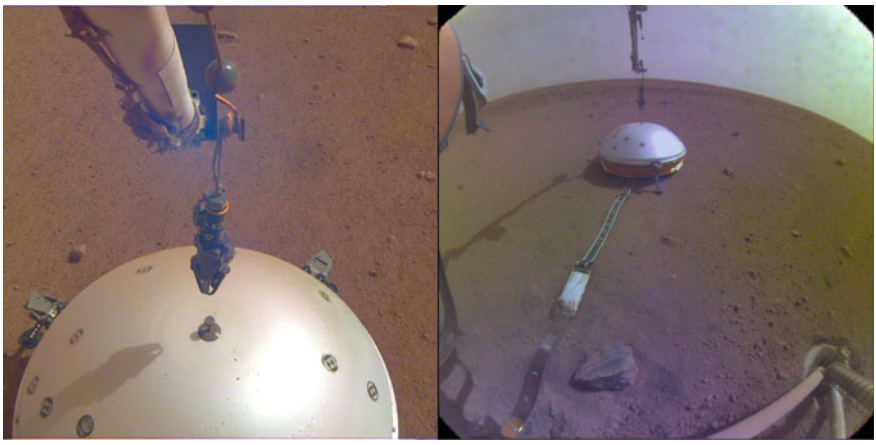

Fig. 21.31 The WTS and grapple after completion of WTS deployment Part 4. The grapple has released the WTS grapple hook and the IDA has moved upwards

WTS deployment Part 4a, or the grapple positioning adjustment, was executed on Sol 67. This part included moving the IDA slightly in order to pull out slack in the grapple cable and align the grapple over the WTS. This prepares the grapple for opening during WTS deployment Part 4. Before moving the IDA, we set the collision parameters in the IDA FSW to position the WTS collision model at our current best estimate for the WTS position. We moved the grapple up and over to center it over the grapple hook. While we wanted to lift the grapple up to remove the slack in the cable, we biased the motion to the lower portion of the grapple hook to ensure there was no upward force on the top of the grapple hook. To confirm the correct positioning, we looked at both IDC and ICC images (see Fig. 21.30), plus IDA joint angles.

WTS deployment Part 4 was executed on Sol 70. During this final part of WTS deployment, we moved the grapple up one more cm to reduce the slack in the grapple cable more, opened the grapple, and then moved the grapple up 10 cm to release the grapple hook. Then we restored the collision parameters in the IDA FSW so that it would again consider collisions between the grapple and the WTS. We used IDC and ICC images (see Fig. 21.31), grapple telemetry showing correct grapple operation, and IDA joint angles to confirm that the IDA was in the correct pose and had successfully released the WTS grapple hook.

The subsequent localization of the WTS indicated a ground position offset by 6.5 cm in the desired direction from the SEIS center point. Because we were trying to achieve a 5 cm offset, the deployed position was within 1.5 cm of the desired placement location, which is very good accuracy for this first-of-its-kind robotic stacking deployment on another planet. The SEIS science data showed an immediate improvement after the WTS was placed over the SEIS in WTS deployment Part 3.

Based on the SEIS data, the SEIS team was able to confirm that there is no contact between the WTS and SEIS.

21.8.3 HP3 Deployment

The Heat Flow and Physical Properties Package, or HP3, consists of a support structure assembly (SSA), a "mole" intended to penetrate into the Martian regolith, and a fixed radiometer. The mole is attached to the SSA via a science tether, and the SSA is attached to the lander via an engineering tether. The HP3 experiment is attempting to understand Mars' subsurface heat flow and physical properties (Spohn et al. 2018). During the deployment phase, the HP3 was deployed last, on Sol 76. As described in Sect. 21.5, the HP3 placement location was selected earlier in the mission by the Instrument Site Selection Working Group. The final decision was based on the necessary engineering and science criteria as well as a desire to be far away from SEIS/WTS.

The surface operations team determined the deployment time of day by evaluating constraints such as robotic arm motor temperature requirements, engineering tether temperature requirements, and sunlight in the IDC field of view. Ultimately, 10.45 a.m. local mean solar time (LMST) proved to satisfy all of the constraints.

Like SEIS and WTS deployments before it, HP3 deployment and release was divided across five operational sols and five parts. Part 1, which consisted of moving the IDA's grapple above the HP3 grapple hook, known as a teach point, took place on Sol 73. Part 2, or capturing the HP3, took place on Sol 74. Part 3, or lifting the HP3 from the lander deck and placing it down on the Martian surface, took place on Sol 76. Part 4a, or adjusting the grapple above the HP3 to be vertical, took place on Sol 79. Part 4, or releasing the grapple from the HP3, took place on Sol 83.

The HP3 engineering tether that connects the SSA to the lander was stored inside the SSA while it was bolted to the lander deck. During deployment, the tether unfurled and was pulled out of the SSA as the SSA was pulled further and further from its position on the deck. The robotic arm lifted the SSA and moved it across the lander deck and over the workspace (Fig. 21.32). At that point, we commanded the arm to an outstretched position over the lander deck which allowed us to further extract the engineering tether before bringing the SSA towards its designated placement site.

On Sol 87, the mole was released from the SSA via a frangibolt firing. On Sol 92, the first mole hammering cycle was commanded. The mole did not reach its target depth of 70 cm during this hammering cycle. The subsequent hammering tests were unsuccessful; the mole remained partially above ground and inside the SSA (Fig. 21.33). These unexpected hammering failures resulted in the creation of an Anomaly Resolution Team to resolve this issue. Mole recovery efforts are still ongoing.

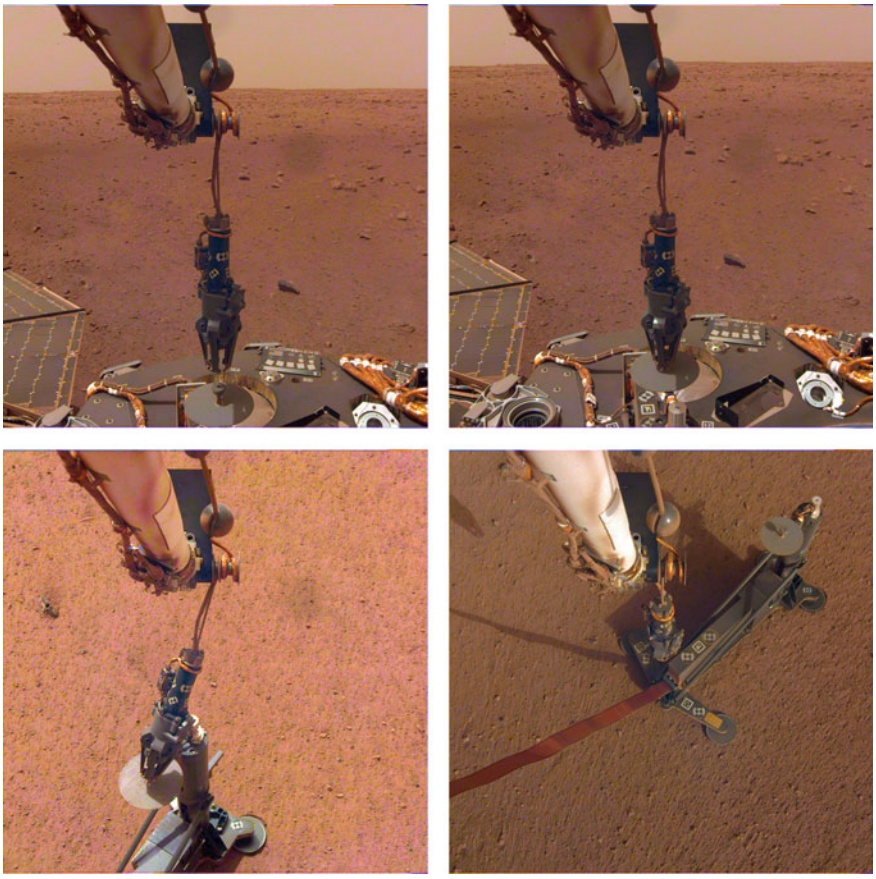

Fig. 21.32 (upper left) Grapple at HP3 Teach Point; (upper right) HP3 grappled; (bottom left) HP3 touchdown on Martian surface; (bottom right) HP3 SSA and engineering tether visible after grapple release

21.9 Summary and Conclusions

This chapter presented flight operations results for the InSight Robotics Instrument Deployment Systems (IDS) that successfully deployed SEIS, WTS and HP3 on the surface of Mars and enabled scientists to perform the first comprehensive surface-based geophysical investigation of Mars' interior structure. NASA's first successful precision robotics instrument placement and release on another astronomical body since Apollo has paved the way for future human precursor robotics planetary construction missions. In addition, the success of the IDS has paved the way for the development of autonomous manipulation as a key enabling technology for successful in-situ payload installation and geophysical science investigations in planetary exploration missions.

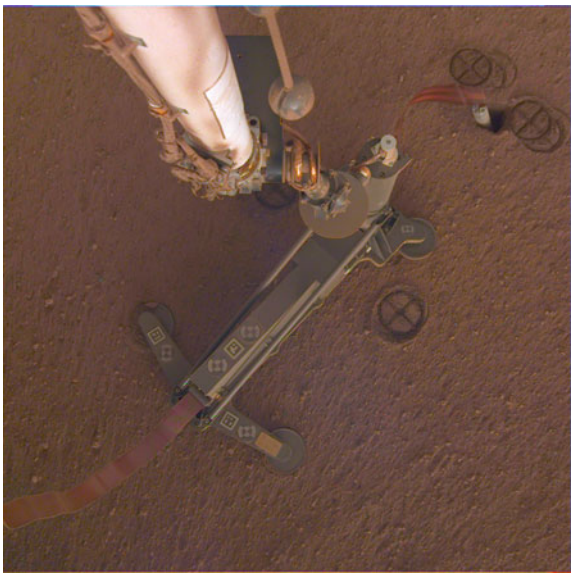

Fig. 21.33 IDC view on Sol 209, after the HP3 SSA was lifted and re-placed down behind the mole by the Anomaly Resolution Team to investigate the configuration of the mole in the Martian surface

Acknowledgements The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (NASA). © 2020. California Institute of Technology. Government sponsorship acknowledged.

References

Abarca, H., et al. 2018. Image and data processing for InSight lander operations and science. *Space Science Reviews* 215: 22 (2019).

Bailey, P., et al. 2020. Deployed instrument monocular localization on the InSight Mars lander. In *IEEE aerospace conference*, pp. 1–13. Big Sky, MT, USA 2020. https://doi.org/10.1109/AERO47225.2020.9172343

Banerdt, W.B., et al. 2018. The InSight mission, space science reviews. *Space Science Reviews* 215: 22 (2018).

Banfield, D., A. Spiga, C. Newman, et al. 2020. The atmosphere of Mars as observed by InSight. *Nature Geoscience* 13: 190–198. https://doi.org/10.1038/s41561-020-0534-0.

Bonitz, R.G., T.T. Nguyen, and W. S. Kim. 2000. The Mars surveyor '01 rover and robotic arm. *IEEE aerospace conference. Proceedings* (Cat. No.00TH8484), vol. 7, pp. 235–246

Bonitz, R.G., et al. 2008. NASA Mars 2007 Phoenix lander robotic arm and icy soil acquisition device. *Journal of Geophysical Research* 113: E00A01. https://doi.org/10.1029/2007JE003030.

Folkner, W., et al. 2018. The rotation and interior structure experiment on the InSight mission to Mars. *Space Science Reviews*. https://doi.org/10.1007/s11214-018-0530-5. (2018 this issue).

Golombek, M., et al. 2018. Geology and physical properties investigations by the InSight lander. *Space Science Reviews* 214: 84.

Lognonné, P., et al. 2019. SEIS: Insight's seismic experiment for internal structure of Mars. *Space Science Reviews* 215: 12.

Maki, J.N., et al. 2018. The color cameras on the InSight lander. *Space Science Reviews* 214: 105.

Sorice, C., P. Bailey, A. Trebi-Ollennu, W.S. Kim, and S. Myint. 2020. Catenary model of InSight SEIS tether for instrument deployment. In *2020 IEEE aerospace conference*, Big Sky, MT, USA, pp. 1–8. https://doi.org/10.1109/AERO47225.2020.9172783.

Spohn, T., M. Grott, et al. 2018. The heat flow and physical properties package (HP3) for the InSight mission. *Space Science Reviews*, this issue.

Trebi-Ollennu, A., et al. 2012. Lunar surface operation testbed (LSOT). In 2012 *IEEE aerospace conference*, Big Sky, MT, pp. 1– 16. https://doi.org/10.1109/AERO.2012.6187057.

Trebi-Ollennu, A., et al. 2018. InSight Mars lander robotics instrument deployment system. *Space Science Reviews* 214: 93.

Yen, J., B. Cooper, F. Hartman, S. Maxwell, and J. Wright. 2004. Sequence rehearsal and validation on surface operations of the Mars exploration rovers. In *Proceedings of SpaceOps 2004*, Montreal, Canada.