Improving Spacecraft Design and Operability for Europa Clipper through High-Fidelity, Mission-Level Modeling and Simulation

Eric W. Ferguson[1], S.S. Wissler[2], B. K. Bradley[3], P. F. Maldague[4], J. M. Ludwinski[5], C. R. Lawler[6]
Jet Propulsion Laboratory, NASA/California Institute of Technology, Pasadena, CA, 91109, United States

NASA's planned Europa Clipper mission seeks to assess the habitability of Jupiter's moon Europa, which exhibits strong evidence of an ocean of liquid water underneath its icy crust. The sheer number of unique instruments, all of which require quiescent environments in order to operate, compounded with Jupiter's distance from the Sun and Earth, make this planned mission challenging and resource-constrained. High fidelity, mission-level simulations that model the spacecraft, ground, and environment from launch to end-of-mission with a given trajectory and mission plan have been employed early in the project lifecycle to better understand the interactions between various components of the mission and how design changes impact the entire system. These simulations have already resulted in tangible benefits to the project by providing vital input to key spacecraft trades, assessing impacts to operability, and quantifying how well the scientific objectives of the mission can be achieved. Improvements to simulation performance and to the process by which information defining the system is gathered and built into models used by simulations have the potential to further expand the scope of their use on Europa Clipper and future missions.

I. Introduction

MODELING and Simulation plays a vital role in the design and operation of most, if not all, space missions. Applications employing modeling and simulation techniques vary considerably from low-fidelity parameterized simulations trading different mission architectures during proposal development to highly complex hardware-in-the-loop simulations used for testing and command validation during operations. Mission-level software simulations integrated with high-fidelity spacecraft models have been used effectively for spacecraft activity planning on a number of missions [1, 2]. These types of simulations provide a means to predict the state of the spacecraft given an activity plan and to model the cumulative effect those activities have on mission resources (including, among many others, power, energy, pointing, data, and ground station tracking) over time.

As simulation frameworks and languages have evolved and the processing power of computers has improved, mission-level activity planning simulations typically used during operations have begun to appear much earlier in the project lifecycle. The Europa Clipper mission, one of the first flagship-class missions to adopt such simulations early in the NASA project lifecycle, began running simulations based on current mission concepts in Pre-Phase A. Developers of these simulations were able to leverage models from previous missions in order to quickly build a simulation infrastructure at an extremely low cost. Despite having a very small development team and limited budget, mission-level simulations have had immense benefits (many of which were unexpected) to the Europa Clipper project since they were established. Analyses based on simulation results have facilitated informed decisions during trade studies, hardware test plan development and resource allocation processes, and the development and verification of mission requirements. Even though the approach to run high-fidelity simulations early and often has been extremely successful on Europa Clipper, improvements could be made to further enhance their utility on this and future missions. With these improvements, it is even possible that these simulations will smoothly evolve into the simulations used to integrate and validate spacecraft plans during operations.

---

[1] Mission Planner, Europa Clipper Project
[2] Chief Engineer, Mission Systems and Operation Division
[3] Mission Planner, Europa Clipper Project
[4] Engineering Applications Software Engineer, Mission Systems and Operation Division
[5] Lead Mission Planner, Europa Clipper Project
[6] Engineering Applications Software Engineer, Mission Systems and Operation Division

Sections II and III introduce the Europa Clipper mission and present the Activity Plan Generator (APGen) simulation framework used to build and run the integrated mission-level simulations for the mission, respectively. Section IV describes the infrastructure wrapped around APGen that allows for quick ingestion of trajectory and design updates as well as comprehensive post-processing of simulation results. Section V provides concrete examples of how these simulations have been used for a diverse array of applications, along with some challenges faced by simulation developers thus far on the project. Section VI provides a brief exploration of recommendations for future improvements, including how this infrastructure could be directly transferred to operations.

II. Europa Clipper Overview

A. Mission Design

As one of the top candidates in our solar system with the potential to harbor present day extraterrestrial microbial life, Jupiter's moon Europa is an intriguing world to explore. Europa, roughly the size of Earth's moon, is hypothesized to possess the three main requirements for life: liquid water from a global-scale ocean, energy from tidal heating due to Europa's slightly elliptical orbit around Jupiter, and chemistry via interactions between materials on the surface and the ocean environment [3]. Currently in its preliminary design phase and scheduled to launch in 2022, the Europa Clipper mission's primary scientific objectives are to map the surface and subsurface structure, constrain the average thickness of the ice shell, characterize the composition of the surface and atmosphere, understand the formation and evolution of geologic features on the surface, and search for and characterize any current activity [4].

In order to achieve these objectives, the mission plans to send to Europa a robotic, solar-powered spacecraft carrying a suite of ten instruments: five remote sensing instruments covering the spectrum from thermal emissions to the ultraviolet, four in-situ fields and particles instruments, and a two-channel radar. Instead of orbiting directly around Europa where the spacecraft would be pummeled continuously by Jupiter's harsh radiation environment, the spacecraft would perform a series of repeated close flybys of Europa while in an elliptical orbit around Jupiter. This allows the spacecraft to quickly dive through the intense radiation zone near Europa while providing ample time during the remainder of the orbit to return scientific data collected during the flyby [5]. Nonetheless, the brief dips into the radiation environment pose a threat to sensitive electronics and cause the solar arrays to degrade and produce less power as the mission progresses. The simulations constructed for Clipper take this environmental effect into account when computing solar array power output.

The current mission baseline consists of launching on NASA's Space Launch System (SLS) rocket on a 2.5-year direct trajectory to Jupiter followed by a 3.5-year tour comprised of over 40 close flybys of Europa with closest-approach altitudes varying from several thousand kilometers to as low as 25 kilometers (Fig. 1) [6]. In the event the SLS is unavailable, an indirect trajectory requiring gravity assists from Venus and Earth and 4-5 years of additional cruise time would be used as no commercially available launch vehicles can provide sufficient energy to travel directly to Jupiter. Therefore, the spacecraft must be designed to withstand the lengthier cruise and lower minimum solar

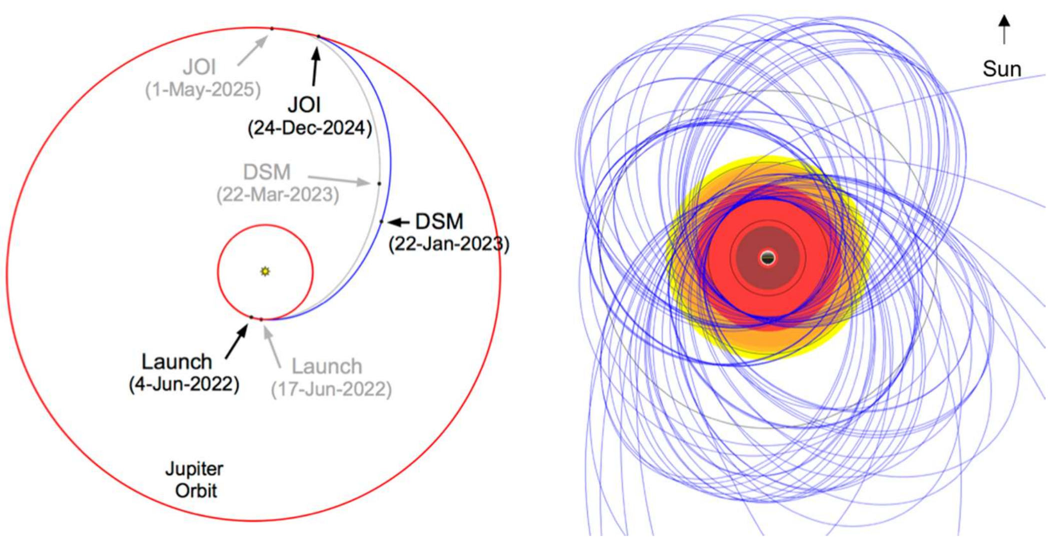

Fig. 1 Baseline mission design for Europa Clipper. Direct trajectory to Jupiter showing both early and late launch dates (left) and multi-flyby Prime Mission orbits around Jupiter (right).

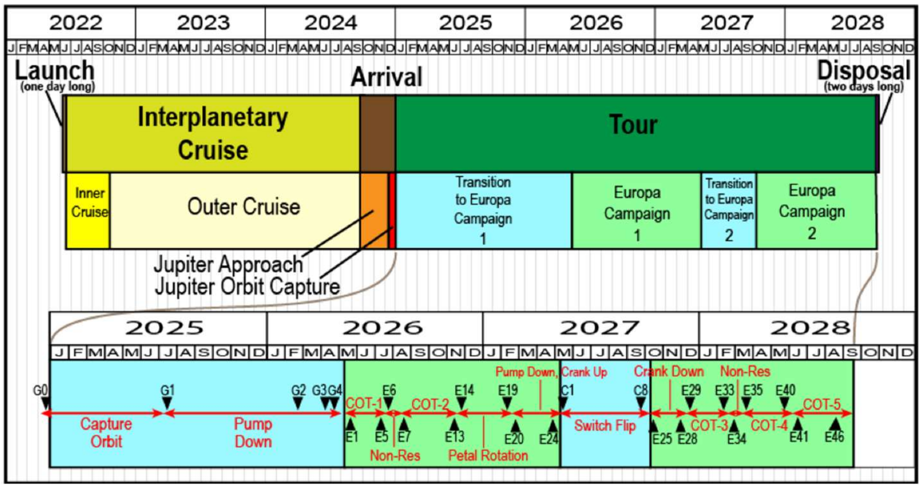

Fig. 2 Baseline mission timeline including mission phases along with a zoomed view of Tour.

distance seen on an indirect trajectory. Both the direct and indirect trajectories can be run through integrated simulations to help uncover operational issues due to geometry (e.g. preventing the solar arrays from pointing directly at the Sun during turns to and from the trajectory maneuver attitude near the Venus gravity assist).

The prime phase of the mission, known as Tour, consists of two distinct flyby campaigns, the first of which focuses primarily on Europa's anti-Jovian hemisphere (Europa is tidally locked to Jupiter and thus the same side always faces away from Jupiter) while the second campaign focuses on the sub-Jovian hemisphere (Fig. 2). Each campaign systemically takes regional-scale observations of its entire hemisphere, which will ultimately provide the global coverage needed to answer key questions about Europa's habitability by the end of the mission. Prior to the first campaign, a series of gravity-assist flybys of Ganymede reduce the orbit period to a duration of just over two weeks and set up lighting and velocity conditions for the first Europa flyby. This transition to the first Europa campaign also provides time to exercise the spacecraft and operations processes as well as calibrate instruments in preparation for operations at Europa. A transition from the first to second Europa campaigns using a series of Callisto flybys reshapes the spacecraft orbit so that it can observe the sub-Jovian side of Europa in the Sun.

During Tour, orbits containing a targeted Jovian moon flyby are divided into sub-phases (Fig. 3). The time between 2 days before and after closest approach, known as the flyby period, is further broken down into 3 parts: Approach, Nadir, and Departure. The Playback phase follows Departure, which in turn ends 2 days prior to closest approach of the next targeted flyby, and the cycle repeats. One such cycle is known as an encounter. Most of the mission's Jupiter orbits contain a targeted flyby; thus, the duration between flybys is typically just over 14 days. However, there are occasional "phasing orbits" which do not contain targeted flybys; in these cases, an encounter spans multiple orbits.

During Playback, the spacecraft spends most of its time pointing its high-gain antenna (HGA) at Earth to maximize data return and only occasionally departs from this orientation to perform activities that require specific pointing such as orbit trim maneuvers and instrument calibrations. During Approach, instruments prepare for and begin to make episodic observations of Europa, which requires the spacecraft to turn to and from a nadir orientation (where the remote sensing instruments are pointed at Europa). After a final communication pass roughly a half-day prior to closest approach, the spacecraft turns to point at or near nadir while twisting around the nadir axis to allow full Sun to fall on the solar arrays, and no longer turns back towards Earth until after closest approach. The majority of instrument observations and, consequently, data accumulation occur within a 4-hour period around closest approach. Just prior to this period, the spacecraft twists about the nadir axis to align the along-track fields of view (FOVs) of the remote sensing instruments and the boresights of the in-situ particle detection instruments with the spacecraft velocity vector present at closest

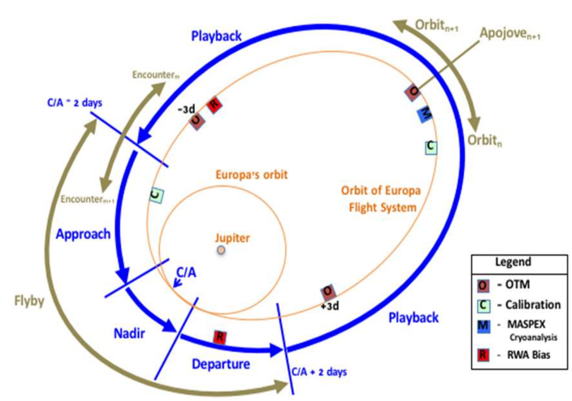

Fig. 3 Encounter phase definitions and typical timing of key activities during an encounter.

approach. Observations during departure generally mirror those found during approach with some differences driven by lighting conditions before and after closest approach, and also by the need to prepare for data playback. All spacecraft activities (including instrument observations) that occur during an encounter, and their effect on resources like spacecraft attitude, power, and data, are incorporated into the integrated Clipper simulations.

B. Flight System Design

The Europa Clipper Flight System is comprised of the spacecraft and the instrument payload (Fig. 4). Single-axis solar arrays provide power to the spacecraft during the majority of the mission while large batteries are used during long eclipses (lasting up to 9.2 hours) and periods when the spacecraft attitude is constrained such that the arrays cannot point directly at the Sun (e.g. maneuvers, flybys). Unfortunately, the periods of highest power demand are often the periods when the arrays cannot point towards the Sun. A unique aspect of the design is that the VHF and HF radar antennas of the REASON instrument are mounted directly to the solar array. During REASON radar sounding, which occurs below 1000 km altitude from Europa, the solar arrays must remain fixed in the position that points the VHF antennas towards nadir.

A suite of one HGA, one medium-gain antenna (MGA), two low-gain antennas (LGAs), and three fanbeam antennas are used for communication. All antennas support X-Band for uplink and downlink, but only the HGA supports Ka-Band downlink, which is the primary return path for acquired science data. Additionally, the LGAs and fanbeam antennas support the gravity science investigation conducted a few hours around Europa closest approach. Reaction wheels are used for precise attitude control while a bipropellant system is used both for coarse attitude control and trajectory maneuvers. Highly accurate attitude knowledge is maintained through a pair of non-co-aligned stellar reference units (SRUs) and an inertial measurement unit (IMU). Dual-redundant, non-volatile bulk data storage (BDS) units, each of which can hold at least 550 Gigabits, collect data from the instruments throughout the prime mission. The large capacity of these units is primarily driven by the fact that the Flight System may collect on the order of 100 Gigabits per flyby and the downlink capability may not always support returning all of that data before the next flyby. An active heat redistribution system (HRS) employing a pumped fluid loop, combined with passive systems such as a louvered radiator and multi-layer insulation (MLI) blankets, provide thermal control for the spacecraft in both the hot environment early in cruise and the cold, energy-starved environment at Jupiter. Each spacecraft subsystem and its respective hardware components are modeled as part of the mission-level simulations produced for Clipper, with updates to the models continually incorporated as the design matures.

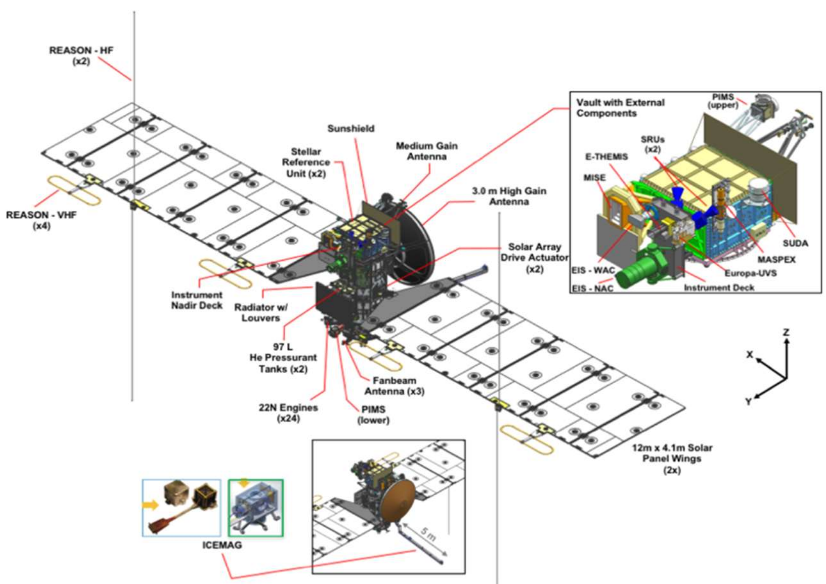

Fig. 4 Europa Clipper Flight System in its low-altitude science data acquisition configuration. Inset at upper-right depicts spacecraft vault and external components including co-boresighted remote sensing instruments (MISE, EIS, E-THEMIS, UVS) and co-boresighted particle detection instruments (MASPEX and SUDA).

III. Simulation Framework and Heritage

The central tool used to build and run simulations of the Europa Clipper Mission in its early stages was APGen, a modeling and simulation framework developed by NASA's Advanced Multi-Mission Operating System (AMMOS) organization [2]. APGen was originally developed to help mission planners create activity plans that did not oversubscribe critical resources such as data storage, electrical power and fuel during mission operations. Mission engineers specify ground and spacecraft activities, resources, and their interaction through a non-standard but highly effective domain-specific language (DSL) tailored to increase code clarity and coding efficiency when compared to conventional languages. The fundamental constructs within the DSL are parametrized activity types (much like classes in object-oriented programs) that can be instantiated within a plan, and resources which collectively describe the full state of the simulation. Code written in the DSL defining the activities and resources in a plan for a particular mission is often referred to as an APGen "adaptation." Activities and their effect on resources are modeled using APGen's discrete-event simulation engine, which only performs calculations at times specified by the simulation developer. This provides the capability to run simulations significantly faster than real-time (Europa Clipper end-to-end simulations run at over 8000x real-time).

One feature of APGen and its DSL that has played a fundamental role in recent mission adaptations including Europa Clipper is a complete framework for expressing activity hierarchies. High-level activities, which typically represent science or engineering goals, can be expanded into lower-level activities representing the commands that implement these goals. Detailed multi-year mission simulations like Clipper produce between 100,000 and a million activities; these massive simulations were not even possible until recent, significant upgrades were made to APGen.

Over the years the APGen DSL has been used to formulate a number of subsystem models for various missions, many of which have been completely or partially reusable by new missions [1]. Activities in the plan interact with these models via simple, well-defined interfaces. Additionally, APGen has the capability to interface with other model libraries written outside of the APGen DSL, which has allowed developers to leverage powerful tools like SPICE [7] or JPL-standard models like the Multi-Mission Power Analysis Tool (MMPAT) with ease. Table 1 describes the subsystem models used on Europa Clipper and their heritage (if any) from previous mission adaptations.

On missions that adopted APGen early in its development, planners would place activities in the plan manually or through scripts outside of the DSL. Later, the DSL was enhanced to allow mission engineers to automatically add science and engineering activities to the plan based on customizable scheduling criteria [9]. This provided planners with the ability to build algorithms that queried the state of the plan (via resources) to determine time windows where specified scheduling constraints were met. New activities were then placed into the plan at appropriate times based on these windows. The deterministic nature of the observing geometry of orbital missions like Europa Clipper compared with surface missions make these missions much more conducive to automated activity scheduling. Thus far, Europa Clipper simulations have relied almost entirely on automated APGen scheduling algorithms to build a notional plan of both science and engineering activities. As with the subsystem models, many of the scheduling algorithms, known colloquially as "schedulers," were based on heritage from previous missions.

While many previous uses of the APGen scheduling engine were focused on a single activity type or subsystem, Europa Clipper simulations are built from the ground up using automated schedulers. Complex logic based on various constraints, policies, and guidelines is frequently required when introducing activities to the Clipper mission plan, which is essentially empty at the beginning of a simulation run. Heuristics that had been introduced in early versions of the engine had to be removed and replaced with well-defined, predictably behaved algorithms that could be put to use with a minimum of difficulty by non-programmers. The resulting scheduling engine was perhaps the most significant ingredient to generating the Clipper simulations discussed in this paper. Without it, construction of realistic activity plans would have been significantly more laborious.

IV. Europa Clipper Mission Simulations

A. Simulation and Model Inputs

The foundation of all Clipper mission simulations is the spacecraft trajectory. Mission designers deliver new trajectories to the project periodically in the form of a SPICE SPK (ephemeris) kernel file with ancillary files for detailed information on maneuver times, magnitudes, and orbit determination error. Clipper simulations read the SPK directly along with other kernels defining the positions of natural bodies and ground stations on Earth to provide a basis for performing geometric calculations during activity scheduling and modeling. Additional kernels defining spacecraft and instrument coordinate systems and instrument FOVs (FKs and IKs respectively) were also created and then used by APGen for attitude scheduling as well as constraint checking. A new plan for the entire mission could

Table 1 Primary subsystem models used in the Europa Clipper APGen adaptation.

Subsystem	Model Description and Heritage
Geometry	NASA's Navigation and Ancillary Information Facility (NAIF) SPICE-based geometry calculations, including basic range and angular separation, as well as instrument field of view and constraint checking. Currently 50+ geometry resources are calculated and tracked. *Heritage: Deep Impact*
Ground Station	Models characteristics and usage of Earth-based ground stations. Composed primarily of the Deep Space Network (DSN), but also includes additional assets from the European Space Agency (ESA), the Japan Aerospace Exploration Agency (JAXA) and NASA's Near Earth Network. Includes the ability to generate view periods, transmitter limits, allocation patterns and ground station transmitter and receiver events. *Heritage: Mars Polar Lander, Deep Impact, Mars Odyssey, Mars Reconnaissance Orbiter, Mars Exploration Rovers, Mars Science Laboratory, and InSight.*
Telecommunications	Medium fidelity telecom link model to support the downlink modeling. Computes achievable data rates for X and Ka-band transmitters for all low-gain, fan-beam, medium gain and high-gain antennas. In the future, this model may be replaced by an interface to JPL's Telecom Forecaster Predictor (TFP) software, which is used by most JPL projects. *Heritage: New to Europa Clipper*
Data	High fidelity model of instrument and engineering data production, on-board data storage, and playback of data to the DSN. *Heritage: Mars Polar Lander, Deep Impact, and Phoenix*
Power	Models power loads from all spacecraft subsystems and instruments. Computes solar array output and battery state of charge. There is a medium fidelity model as well as an interface to JPL's high fidelity Multi-Mission Power Analysis Tool (MMPAT). *MMPAT Heritage: Mars Exploration Rovers, Deep Impact, Mars Science Laboratory, Phoenix, InSight, and JUNO*
GNC	Models commanded attitudes such as attitudes for communication, science observations, and trajectory maneuvers with high fidelity. Updated based on the algorithms and methods inherited from the Cassini Mission and adapted for the flight software implementation on Clipper. Important for input to other models, such as Power, Telecom Link and geometric constraint checking. *Heritage: Deep Impact*
Solar Array	Models solar array articulation for sun tracking and fixed modes, as well as hard-stop avoidance "flops." *Heritage: New to Europa Clipper*
Radiation	Models the radiation environment's effects on the solar arrays. The Jupiter radiation environment model is based on a gridded approximation of the Galileo Interim Radiation Environment (GIRE) model [8]. *Heritage: Galileo*
Propulsion	Medium fidelity fuel usage model. *Heritage: New to Europa Clipper*
Payload	Instrument models were implemented in the APGEN DSL based on specifications provided by the Europa science team. Models include instrument operating modes, power states and data production. *Heritage: New to Europa Clipper*
Mission Operations	Models mission operations process, timelines and shift schedules. *Heritage: Deep Impact*

be produced by APGen within a day or two after the delivery of a new trajectory. This provided the project with the power to quickly assess multiple trajectory candidates simultaneously and choose the trajectory that best met project requirements.

Many subsystem models used in the Clipper APGen adaptation were inherited from previous missions and already designed to be multi-mission in nature. Therefore, mission engineers needed only to modify configuration data for that model to make it effective for Clipper. For example, the setup for the on-board data collection, storage, and downlink model inherited from the Deep Impact mission only required changes to simple parameters such as on-board storage size, recorded and real-time engineering data rates, data collection bin names, and compression rates. Early in concept development, the spacecraft design was in flux and too, consequently, were certain model parameters. In fact, model parameters were often varied to understand how a given variation would affect the system and design margins.

The ease at which models could be reused from previous missions allowed a single engineer to build an integrated Clipper simulation modeling attitude, power, data collection, and data downlink within a few weeks. The main model development work was spent on instrument models, as these had to be built from scratch.

In order to accurately model energy demand and availability for the current flight system design, the simulation had to have knowledge about the hardware on-board and the various power states in which each piece of hardware could reside. Since Pre-Phase A, Clipper simulations have interfaced with the official project Power Equipment List (PEL), which provides the Current Best Estimate (CBE) and Maximum Expected Value (MEV) power draw for each mode of each hardware component that consumes power. Initially, the data in the PEL was manually scraped from a spreadsheet and added to the simulation configuration. Now, accessing the latest PEL data is much more efficient as it is pulled directly from the project's single source of truth "System Model" through the transformation process described in Ref. [10]. As the project progressed through concept development, there were major changes in the spacecraft power design–for example, changing from radioisotope thermoelectric generators (RTGs) to solar arrays–and the ability to provide rapid simulations of the mission was crucial in sizing the solar arrays and batteries as well as understanding the balance between data return strategies and battery depth of discharge.

B. Activity Scheduling

Mission simulations for Europa Clipper are built in APGen by layering activities into the plan through a series of automated schedulers. The order in which these activities are layered is a function of priority and/or dependency. For example, schedulers responsible for creating activities that specify key orbital events and ground station view-periods are only dependent on geometry and thus can first be run upon delivery of a trajectory. Similarly, critical engineering activities like orbit trim maneuvers are high-priority and inflexible in their timing, which means they are scheduled prior to other activities like science observations. Activities and their associated constraints are determined through conversations with systems engineers and subsystem experts. Based on activity constraints, each scheduler generates windows where it is valid to schedule each activity and adds one or more activities within these windows. Additional logic in the scheduler is often necessary to determine the number of activities to add, as well as the best place to schedule an activity within the available constraint windows. In rare cases, a subset of the activity scheduling is delegated to another tool, which then feeds back its results into APGen. For example, the scheduling of the EIS NAC instrument's imaging and gimbal motion activities are provided to APGen by SIMPLEX, a tool developed and maintained by the John Hopkins University Applied Physics Laboratory (APL). Input parameters for some schedulers also provided a means to enforce different constraints or alter scheduling behavior.

Each scheduler represents a separate instantiation of APGen core. Simple wrapper scripts driven by a configuration file were developed to change input parameters to schedulers and automatically trigger each scheduler when the previous scheduler had completed. The set of schedulers used for the Tour simulation of Europa Clipper generates approximately 165,000 activities, which would be impossible to build up manually. Activities scheduled often expand all the way down to the command level and include those developed to support attitude commanding, instrument data collection, data downlink, trajectory correction maneuvers, and engineering maintenance activities. A brief description of each scheduler, the key constraints it enforces, and its heritage is provided in Table 2.

C. Final Simulation Run

Once all the schedulers have generated activities and saved them in activity plan files (APFs), APGen reads in all of the activities and kicks off a final simulation run. The various subsystem models within the simulation are triggered by the scheduled activities, which in turn modify the state of the simulation over time through changes to resource values. Every change to the state is recorded, stored, and ultimately provided as an output once the simulation run completes. Not surprisingly, as the level of detail within the APGen Clipper adaptation has grown, so too has the time it takes to run them. For current end-to-end mission simulations covering a 6-year time frame, the activity scheduling process takes about 24 hours to complete; the final simulation takes 6-8 hours to run. There are a few potential opportunities to improve simulation run time discussed in more detail in Section VI. Fortunately, modifications to subsystem models often do not affect activity scheduling. In these instances, only the final simulation needs to be re-run, which results in significant savings of time.

D. Simulation Output and Visualization Products

The primary output of the APGen simulations for Clipper is a massive Extensible Markup Language (XML) time-ordered list (TOL) file on the order of 20 GB for end-to-end simulation runs. This file contains all of the state changes that occur within the simulation, every constraint violation, and all of the data associated with each scheduled activity (e.g. start time, duration, attributes). In addition to the XML TOL, a set of smaller files tailored to specific applications

Table 2 Activity schedulers used on Europa Clipper listed in the order in which they are run. Schedulers that can run in parallel with each other are noted in the description.

Activity Scheduler	Description and Key Scheduling Constraints
Ground Station View Periods	Complex combination of DSN station visibility and allocation, station elevation, S/C range, and computed S/N ratio
Geometric Events	Periapsis and apoapsis, occultation, eclipse, and transit times
Trajectory Correction Maneuvers	Given a table of maneuver times, magnitudes, and directions, add maneuvers to the plan.
Targeted Flyby Attitude Modes	Attitude changes between the Earth communications attitude and Europa science attitudes.
Gravity Science	X-Band communications activities in support of gravity science including RF operations and antenna selection and switching for the fan-beam and low-gain antennas
Flyby Science and Calibrations	All approach, nadir and departure Europa flyby science activities.
Ground Station Allocations	Since allocations are not available far in advance, the simulation mimics the allocation scheduling process based on station visibility and expected level of DSN commitment
Avionics Maintenance	Adds avionics maintenance activities, such as radiation scrubbing, copying science data from the prime to the backup avionics and pre-processing of stored science data into downlink products.
Instrument Calibrations (Outside Flyby Period)	Any instrument activities necessary for maintaining the calibrations of the instruments. These are scheduled away from the flyby science periods.
SUDA Surveys	Measurements of the Jupiter dust environment outside of the Europa close flybys
Downlinks	Playback of stored science and engineering data. Includes rate stepping for the Ka-Band downlink.

such as power and attitude reporting are produced. After each simulation run, the attitude report is automatically converted into an industry-standard SPICE C-kernel, which serves as an input to many other analysis tools on the project.

All of the raw data produced by simulations is useless unless it can be transformed into a product digestible by analysts and, ultimately, project management. A variety of data visualization tools have been utilized to post-process simulation results and produce valuable products that can be disseminated to a wide array of users including mission planners, subsystem teams, and instrument teams. The tools used to produce these products range from homegrown scripts that generate extensive PDF reports containing plots and tables to third-party or NASA-sponsored interactive timeline and 3D visualization tools. One commonality between the cores of these tools is that they are completely independent of the simulation itself so that other projects can also take advantage of their capabilities. What follows is a list of the primary visualization products currently generated from mission-level simulations along with a brief description of the product and the tool used to generate that product.

1. Interactive RAVEN Timelines

After a simulation completes, the output XML TOL file is automatically pushed to a server, divided into a plethora of resource and activity timelines, and stored in a database, which can be accessed by the web-based Resource and Activities Visualization Engine (RAVEN) tool. Created and managed by NASA's Multimission Ground System and Services (MGSS) program, RAVEN provides a way to view and interact with timeline data via a web browser. Each mission-level simulation generates thousands of individual timelines, which are organized into categories such as "Geometry," "DSN," and "Power." Users build a timeline view by first choosing the simulation data set/s of interest and adding individual timelines from that data applicable to their current analysis (Fig. 5). After one or more timelines have been added to the view, users can zoom in to the time periods of interest, customize how each timeline is displayed, reorder timelines, and save the view so it can be shared with others via a simple URL. Other useful features include overlaying timelines in a single "row" and saving a view layout so that it can be applied to other simulation data sets. On Europa Clipper, RAVEN views of mission simulations are used to answer questions in real-time that crop up during meetings, illustrate relationships between events in the mission plan, capture potential time periods of concern where constraints or guidelines are violated, and even build presentation material for reviews. RAVEN gives users the power to efficiently explore the vast quantities of data associated with these simulations and make connections between aspects of the mission that otherwise would have eluded them.

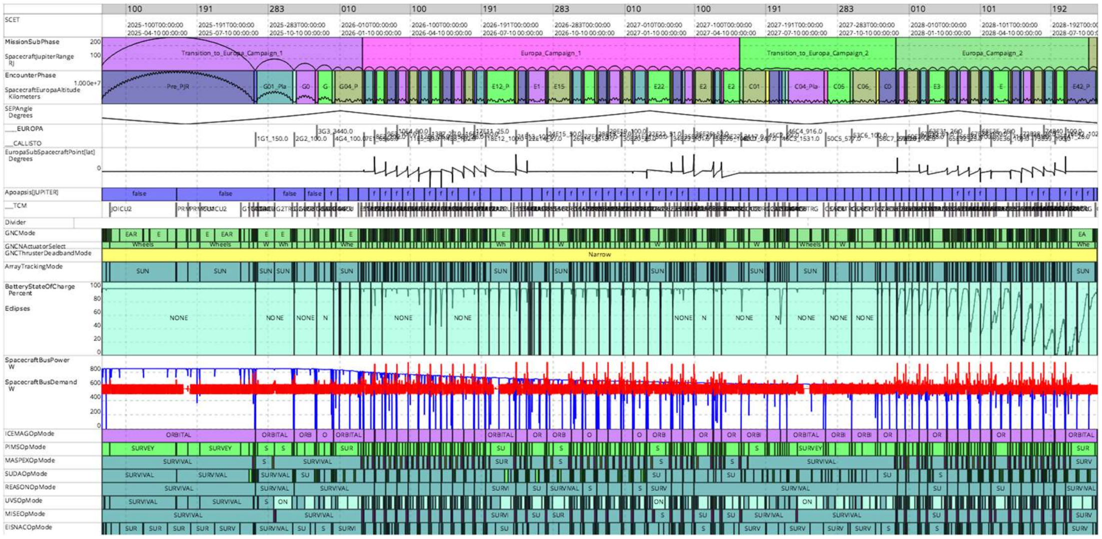

Fig. 5 Resource and Activities Visualization Engine (RAVEN) timeline view of the entire Tour phase of Europa Clipper. This only represents a small subset of the data available from each simulation. Users can zoom in to regions of interest and add additional timelines as desired.

2. *PDF APGen Simulation Report*

A PDF report consisting of hundreds of plots and tables summarizing the results of each simulation is produced automatically alongside every simulation run (Fig. 6). When a run has completed, the output TOL is split into a series of resource and activity timelines in an analogous manner to the process used prior to being viewed in RAVEN. These timelines are then read into MATLAB®, where they are used to perform the computations necessary to build various plots and tables. Finally, these plots and tables are organized and merged into a LaTeX-generated PDF report using a generic script that can be leveraged by other users. This report offers a one-stop shop for the project members to get high-level metrics on resources like data and power as well as detailed information about a particular encounter, subsystem, or instrument. The report is not intended to be read from start to finish; instead, users can jump to a

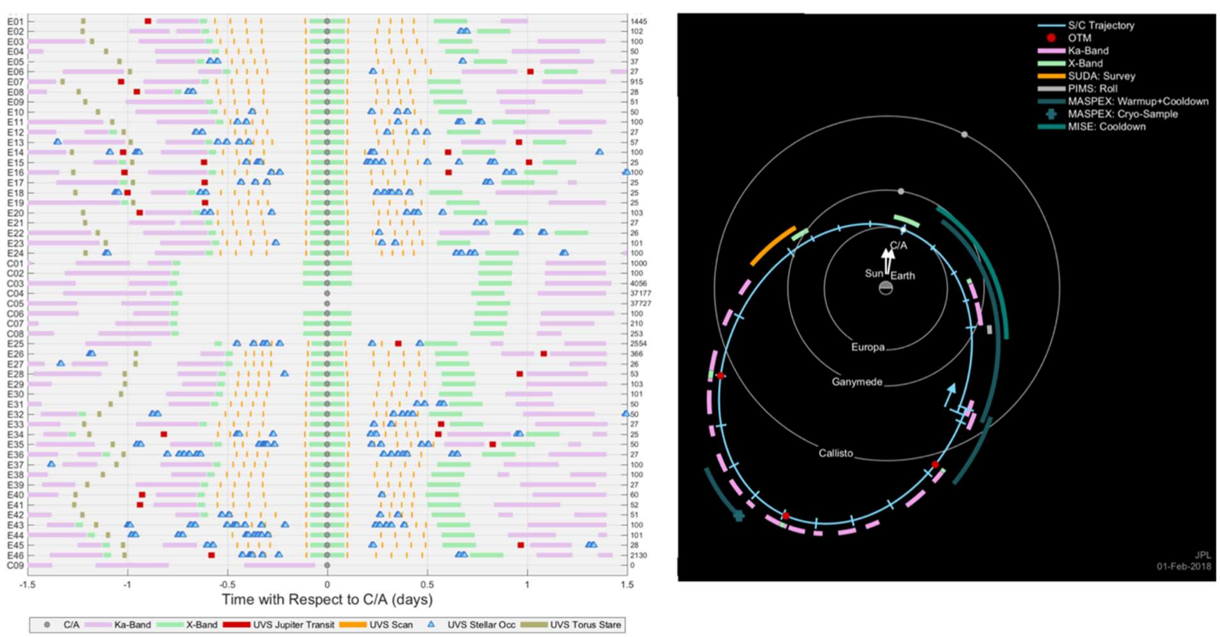

Fig. 6 Stacked timeline showing instrument activities around closest approach for every encounter (left) and diagram of Deep Space Network (DSN) passes in a single Jupiter orbit (right), two of the many plots included as part of the APGen Simulation Report.

particular section depending on their role or interest. Additional plots and tables can easily be added whenever a project member makes a new request. A completely separate PDF report has also been fashioned for the Clipper project to assess the satisfaction of science measurement requirements (see Section V.G). This report was built using the same infrastructure and generic scripts as the APGen Simulation Report.

3. Interactive Tableau® Workbooks

Data visualizations built using Tableau® have also proven effective for investigating simulation results, especially when cross-comparing results from multiple simulation runs (Fig. 7). Report files specifically structured to work effectively with Tableau® are output directly from each simulation. Once the data is connected to Tableau®, users can make beautiful, data-rich charts in seconds, interactively filter data, and perform statistics based on the data. Europa Clipper has constructed a series of workbooks containing a number of charts that track different spacecraft resources like data and power. When a new simulation run has completed, the data can be connected to the workbook and immediately compared to previous results.

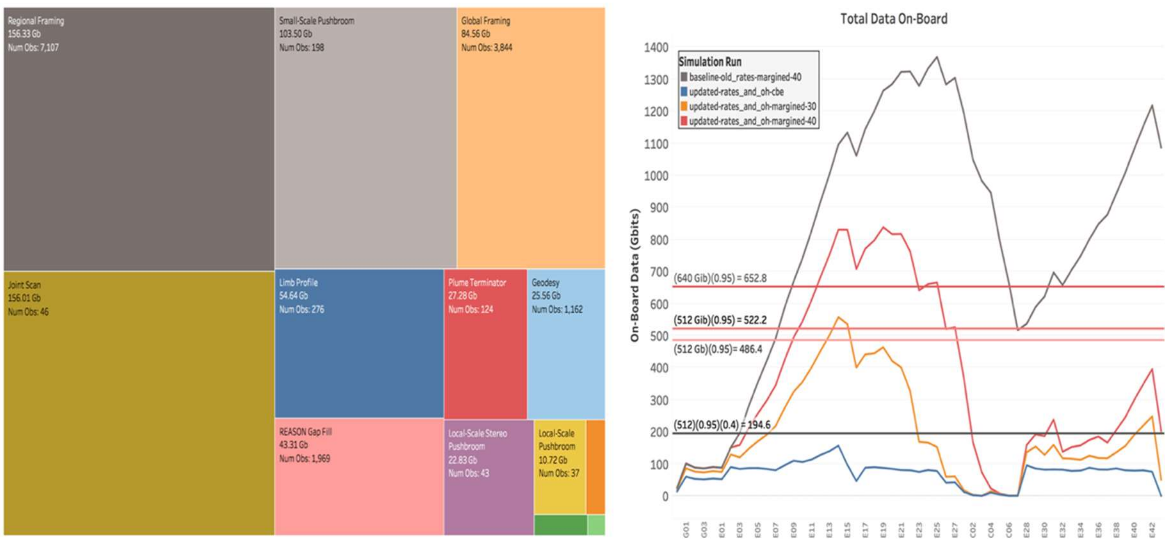

Fig. 7 Examples of Tableau visualizations used on Clipper. Breakdown of instrument data by type (left) and comparison of on-board data between multiple simulation runs (right).

4. Cosmographia 3D Visualizations

Understanding the relation of a spacecraft and its components to Earth, the Sun, and various other celestial objects as it turns, scans, or rolls to carry out various activities to meet mission objectives is extremely difficult without a way to visualize the geometry in 3 dimensions. 3D visualizations driven directly from mission simulations using Cosmographia, a publicly-available interactive tool that visualizes natural bodies within the solar system along with spacecraft trajectories, orientations, and observations, have been available to the project since Pre-Phase A (Fig. 8). Cosmographia was recently extended by NAIF so that SPICE files could be used directly to drive the positions and orientations of objects and display SPICE frames on objects within a visualization [11, 12].

After running an integrated Clipper simulation, the spacecraft attitude profile, solar array rotation profile, and instrument pointing profiles are automatically converted into SPICE C-kernels, which are then loaded into Cosmographia along with a 3D model of the spacecraft, its trajectory, and instrument observation times. When running Cosmographia, users have the freedom to jump to a specific time and place, zoom in on an object, and rotate the view to any desired angle. If desired, users can also script how the camera moves in space so the user doesn't have to change the view manually. This is especially useful when generating videos from the application.

3D visualizations using Cosmographia can be produced quite quickly once a simulation completes. Therefore, they have been used on a day-to-day basis during meetings as an engineering tool to ensure there is a mutual understanding amongst team members about the timing and geometries of specific spacecraft activities. For example, many project members didn't fully grasp that the solar arrays would have vastly different positions relative to the spacecraft while approaching Europa on different flybys due to the varying solar phase angle between those flybys until seeing the geometry in Cosmographia. These visualizations are also sleek enough that some have been used for public outreach and project reviews including (to date) the Mission Concept Review (MCR), Mission Definition Review (MDR), and Flight System Preliminary Design Review (PDR).

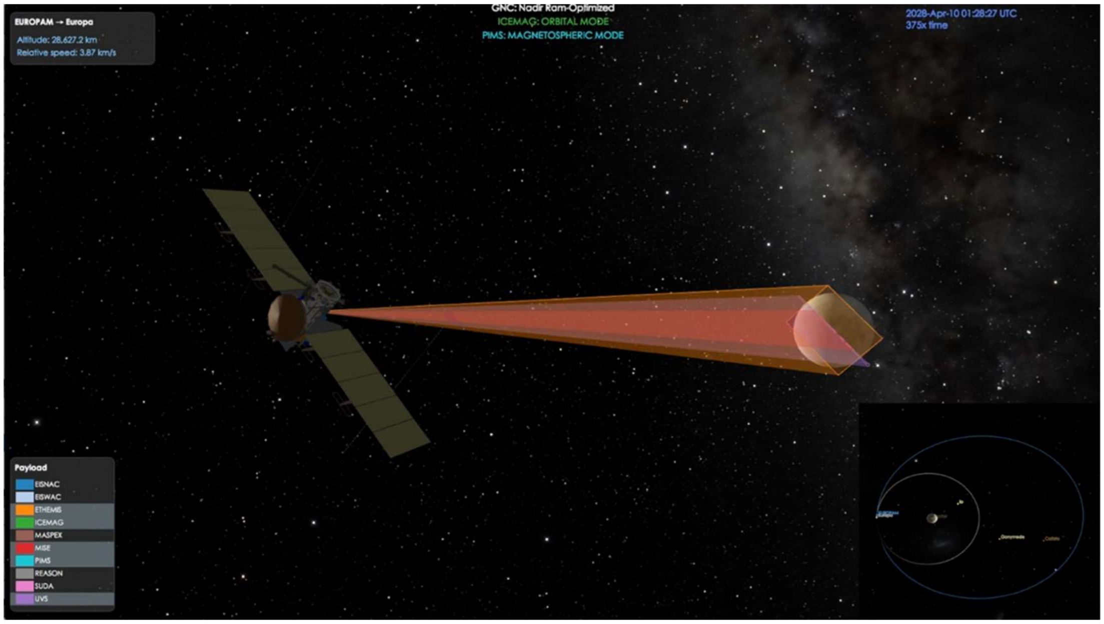

Fig. 8 Screenshot from a Cosmographia video directly driven from simulation data showing a scan from a typical flyby of Europa Clipper produced for the project's Mission Definition Review (MDR).

V. Applications of Mission Simulations During Project Formulation (Pre-Phase A to Phase B)

Since their inception in Pre-Phase A of the project lifecycle, mission-level integrated simulations using APGen have been applied to a diverse set of engineering challenges across the project. The applications of these simulations and their respective examples below are by no means an exhaustive list. However, they do provide a sense of the extensive role that mission-level simulations have played thus far on the project.

A. Mission Planning

The Mission Plan establishes the baseline strategy for successfully achieving mission objectives within the capabilities and constraints of the project systems, and subject to the project policies on the use of these systems. It does so by integrating the mission requirements from science, engineering, and navigation with the trajectory and ground station capabilities into a set of strategies and a time-ordered set of activities. For Clipper, the APGen mission simulations are the direct embodiment of the Mission Plan.

In earlier projects, a significant manual effort was necessary to produce timelines of these plans that were consumable by non-mission-planners, including decision makers in particular. The combination of APGen with RAVEN allowed for the direct production of timelines that can be used virtually as-is to communicate the Mission Plan. Furthermore, the relatively easy-to-use web interface of RAVEN allowed subsystem and instrument engineers to access the mission plan directly, to produce customized views of their activities, and to perform important tradeoffs that inform their design.

Simulations have provided an effective means of assessing impacts to resources brought about by an evolution in understanding of the flight system design and how that might alter strategies employed within the Mission Plan. For example, as the payload design matured, data volume estimates grew substantially and the project recognized an increased risk of exceeding the on-board storage capacity. Mission Planners were able to run simulations exercising different strategies such as arraying ground stations and adding additional data rate switches during downlink tracks to increase downlink capacity at times when the on-board storage was filling up. Although none of these strategies were employed at the time, project management could rest assured that there were ways to resolve data balance issues if data margin dipped below comfortable levels.

B. Project Trade Studies

Armed with high-fidelity simulations and tools to process and analyze results, the project could quickly assess whether a given trajectory and mission plan could meet mission objectives within the constraints imposed by the spacecraft design and environment. If mission objectives could not be met, the project sought out alternate solutions by modifying the spacecraft design, modifying the mission plan (with attendant impacts on operability), or reducing the scope of scientific objectives. The trade space of solutions could be explored by running multiple simulations to help quantify the cost and benefits of each solution.

On Clipper, this technique was used on several key trade studies, including one aimed at investigating potential solutions to curb energy demand and thus prevent solar array growth. The types of solutions explored fell into three broad categories: reduction of power demand via changes to spacecraft hardware, enhanced power generation capability through changes to solar cell cover glass thickness or revisions to the Jupiter environmental model, and operational changes such as lengthening each Jupiter orbit to provide additional time to recharge the battery after each flyby or reducing the number of downlinks late in the mission when power generation will be limited. Multiple APGen mission simulations were run with varying input parameters that reflected the proposed solutions. After cross-comparing simulation results using both RAVEN and Tableau®, the project quickly saw that the net energy gain of the operational solutions was limited when compared with other potential solutions. Moreover, the costs associated with longer orbits and less downlinks far outweighed the slight increase in available energy. Simulations also showed that using a new, more accurate Jupiter proton environment model compiled with the most recent data available lead to a large improvement in energy margin (Fig. 9). Ultimately the project decided to adopt the new model and add additional solar cells to the solar array yoke, which allowed the project to avoid having to add another solar panel to each wing.

Many other project-wide and subsystem specific trade studies were supported by the Clipper adaptation of APGen. All of these trades could be supported because analysis based on simulation results could be turned around relatively quickly (even more trades would likely have been supported if turn around could have been reduced further). A few additional examples of supported trades are listed and briefly described below.

1) Spacecraft Configuration Trade - an early project trade to compare various spacecraft hardware and instrument configurations. Simulations helped show that – for certain configurations – the restriction on solar array range of motion significantly reduced energy generation when it was most needed
2) Additional Instrument Accommodation - NASA requested the project examine the resource cost of adding a laser altimeter to the instrument suite. Numerous simulations were run with various assumptions on instrument operational concepts and resource usage. The results were used to help determine whether the additional science value of the instrument was worth the associated resource impacts.
3) MISE Radiator Position – the Mapping Imaging Spectrometer for Europa (MISE) near infrared instrument team needed to determine an optimal location for their passive radiator. APGen simulation of the spacecraft attitude throughout the mission showed the location on the spacecraft where bright bodies would pass over most frequently. Surprisingly, the radiator was designed to face the nadir direction towards Europa during flybys as that location provided the most protection from the brightest bodies like Jupiter and the Sun.

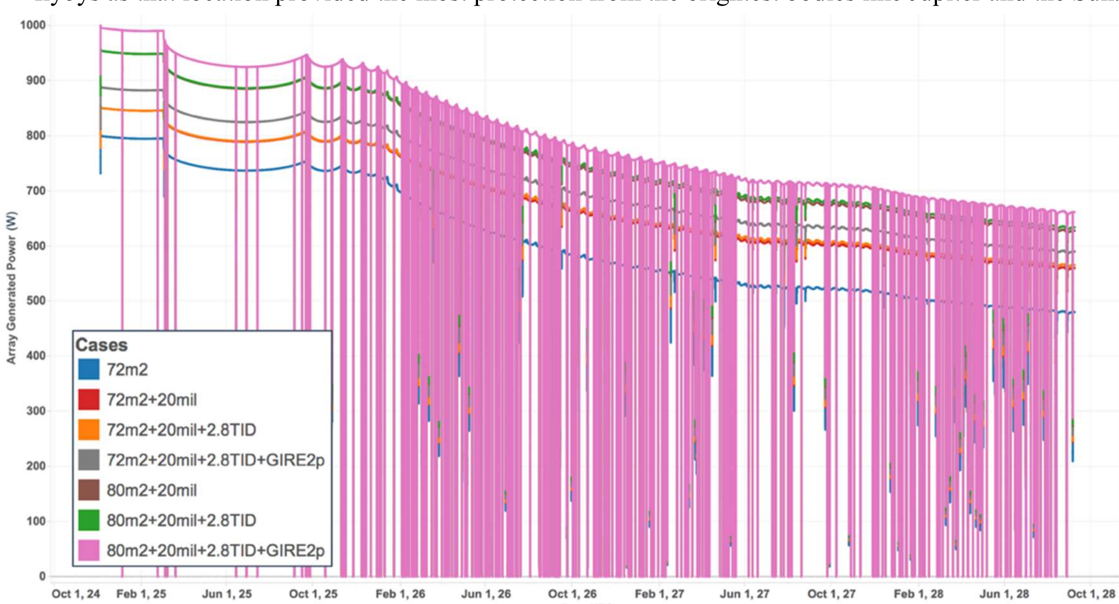

Fig. 9 Solar array power generation over Tour for multiple simulation runs varying array size, cover glass thickness, radiation dosage seen by the spacecraft, and Jupiter environment model.

C. Energy and Data Allocations

In order to determine the required energy generation and storage for the spacecraft, the project took a bottom-up approach by first deriving an energy allocation for the payload [13]. Setting energy allocations for a payload consisting of ten instruments, many of which have dynamic observation profiles that vary each encounter, was challenging. These allocations needed to be relatively robust to changes in tour design while also providing an easy means to sub-allocate to individual instruments. APGen simulations were run on recent tour designs using the Design Reference Mission (DRM) behaviors, a set of spacecraft and instrument activities that are designed to achieve the mission objectives; the DRM thereby provides a flight system design envelope. The DRM runs provide the CBE estimate for payload energy consumption. From these simulations, the project set the payload allocation based on the mean + 1σ energy consumption of the payload per orbit and per flyby (25% contingency plus 10% project reserve was used to compute the final values). Since the simulations also tracked the operational modes of each individual instrument, payload energy sub-allocations were computed based on the mean + 1σ durations for each mode [13].

Data volume allocations for the payload were also based on simulation runs using recent trajectory designs. Per direction of the project, allocations were based on current telecom capability and on-board storage size. While a per-encounter statistical approach similar to what was used for energy allocations initially looked promising, system engineers found it difficult to write useful requirements that would allow the payload to easily sub-allocate data to each instrument and report against their allocations. Instead, allocations were set based on two different portions of Tour: Transition to Europa Campaign 1 (TEC1) and the rest of Tour. Although there is excess downlink capacity during TEC1 due to its long duration and limited periods of high-volume instrument data acquisition, this capacity is not available for use later in the Tour. APGen simulations were run for three different tour designs, and the run with the limiting case for downlink capability for each segment of the Tour was used as the basis for setting allocations (holding 10% of the capability back as project reserve). Figure 10 shows the accumulated downlink capability for Europa Campaign 1 through the completion of Tour. Each simulation run also had CBE data estimates for the payload split out by instrument, so margin could be computed against the allocation and sub-allocations could be derived. In addition to ensuring allocations had sufficient margin between total data volume collected and data downlink capability, they had to be consistent with the spacecraft's on-board storage utilization. Simulations again proved that the current instrument operations plan during Tour using CBE data estimates did not violate the 60% project margin policy on maximum on-board storage utilization.

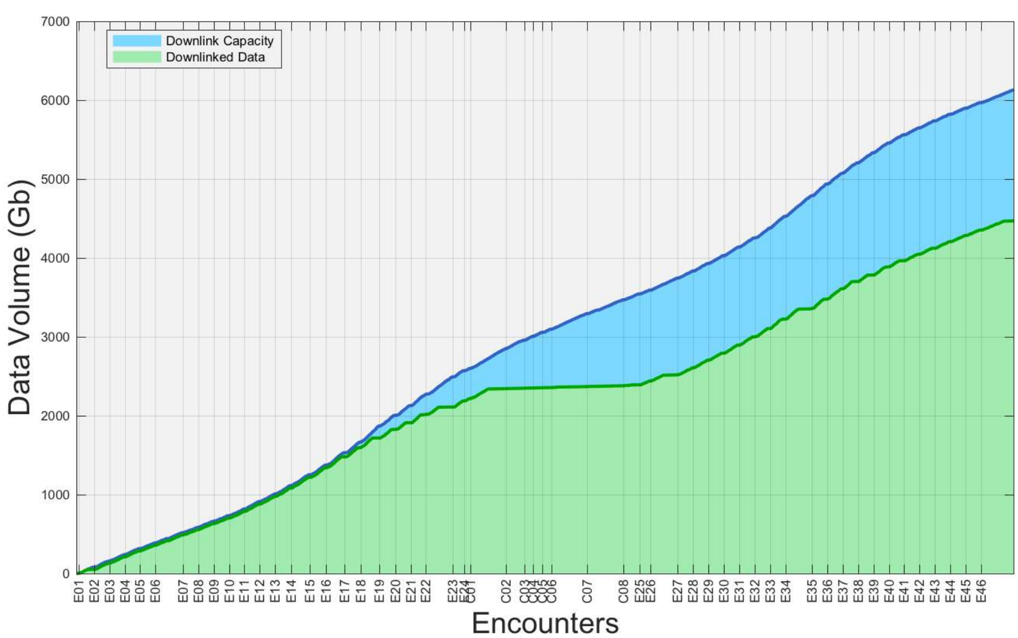

Fig. 10 Comparison of downlink capacity and total data downlinked during Tour starting at the beginning of Europa Campaign 1, which served as the basis for setting payload data volume allocation requirements.

D. Hardware Design

Clipper's solar arrays rotate about a single axis, but unfortunately the design could not support a full 360° range of motion. To protect the arrays from over-rotating and twisting cables running from the spacecraft, hard stops needed to be installed. Depending on the design, these hard stops would reduce the solar array range of motion between 10° and 40°. However, the location of these hard stops was rather flexible, which meant systems engineers could pick the best range of solar array rotation angles relative to the spacecraft to restrict based on operational considerations.

Two separate analyses were performed using APGen simulations to help determine the best location for the hard stops and the impact on energy generation of decreasing the array range of motion. First, a simulation of the entire Tour was run based on the mission plan and its resultant attitude and solar array strategies assuming a 360° range of motion. The solar array rotation angle over time, one of the outputs of the simulation, was processed and binned in 5° increments to determine how much time the arrays spent at different rotation angles (Fig. 11). Not surprisingly, for the majority of the Tour the solar arrays point in the same general direction as the HGA (defined as 0°) because the Earth and the Sun never stray too far apart when viewed from Jupiter. Moreover, the arrays spent very little time facing the opposite direction of the HGA (180°), which made this region ideal for the necessary array exclusion zone. Additional simulations were run that extended the size of the exclusion zone from 10° to 40° to see how this would impact battery state of charge (SOC). For certain flybys where solar array geometry was poor, SOC decreased with reduced range of motion. The impact was significant enough to require that the arrays be designed with at least a 350° range of motion. Ultimately, an exclusion zone between 165° and 175° was chosen because the arrays would need to be stowed in the 180° position for launch.

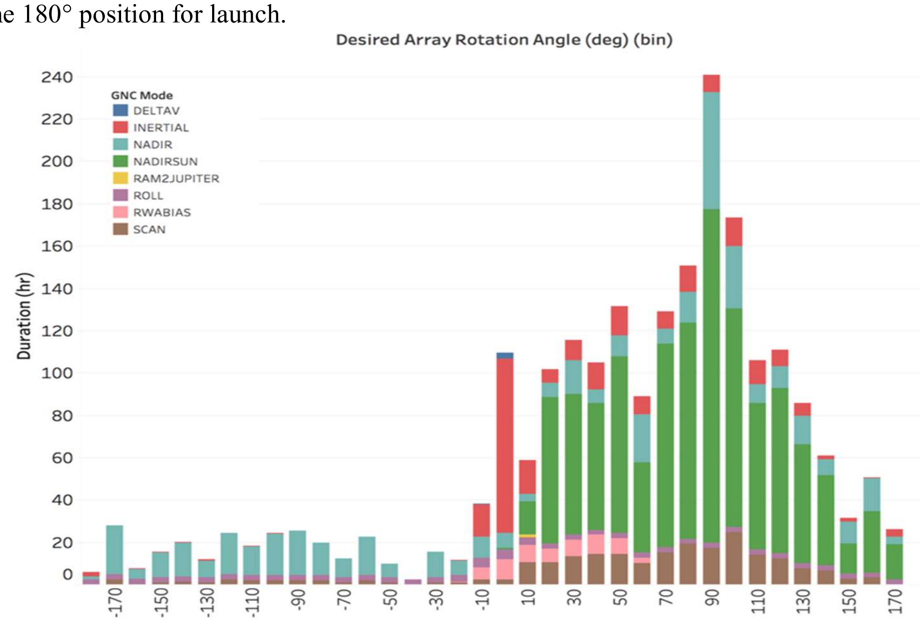

Fig. 11 Cumulative time solar array spends at different rotation angles over the Tour binned in 5° increments. The time spent turning and in the nominal attitude pointing the high-gain antenna (HGA) at Earth has been filtered out.

E. Hardware Test Plan Development

Given the criticality for the solar arrays to provide sufficient power to the spacecraft and the implication of their size to spacecraft mass and agility, detailed design work on the arrays had already begun early in Phase B. As part of that work, thermal engineers had to develop a comprehensive test profile for qualification testing of the solar arrays and solar array-mounted hardware like the REASON VHF and HF antennas that would not be temperature-controlled. The test profile had to define the number of thermal cycles to perform over different temperature ranges by determining the expected number of cycles that would occur during the mission and on the ground during pre-launch testing. In order to comply with design principles, the profile had to include at least three times the number of expected cycles.

Estimating the number of thermal cycles at different temperature ranges that occur during the mission is challenging as the temperature of the solar arrays during the mission is dependent on both the trajectory (solar distance

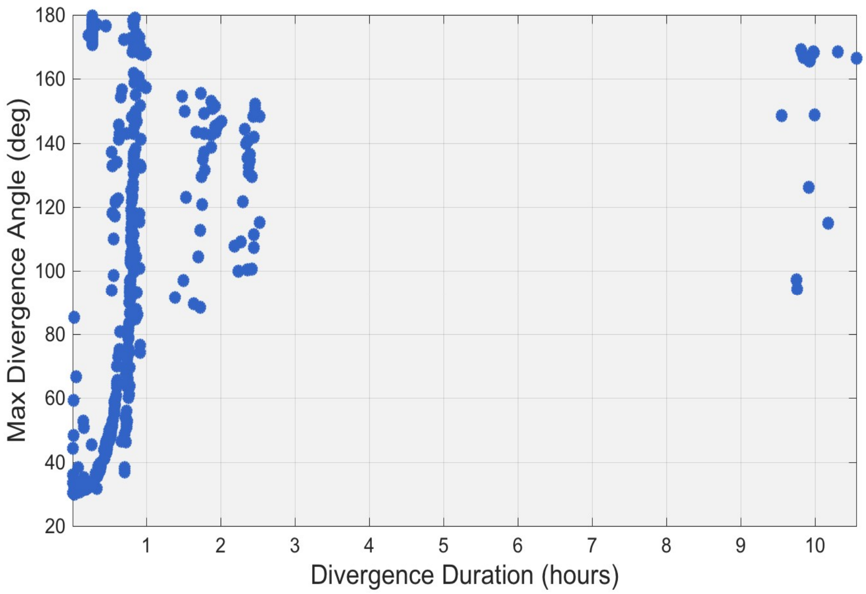

Fig. 12 Maximum off-point angle and duration of each solar array off-Sun-point event while in orbit around Jupiter.

and eclipses) and any planned spacecraft activities that may cause the solar arrays to point off the Sun. Given that the solar arrays can only rotate about one axis, activities that require a specific spacecraft attitude can cause the arrays to diverge from pointing directly at the Sun. Specific examples of such activities during the mission are flyby science activities near Europa closest approach and periodic instrument calibrations. Also, certain activities require the solar arrays to remain held still to prevent disturbances (e.g. EIS NAC imaging) or remain fixed at a specific solar array rotation angle with respect to the spacecraft body (e.g. REASON flyby operations, TCMs).

Fortunately, a notional set of planned spacecraft activities for the entire mission had already been incorporated into mission-level simulations using APGen. Since one of the outputs of these simulations was the solar array pointing profile, the time and duration of each event where the solar arrays pointed off the Sun was quickly determined. In order to assess the severity of each solar array off-point event from a thermal perspective, these events were binned by spacecraft-Sun range and maximum off-point angle. Figure Fig. 12 shows a plot of all the simulated off-point events while Clipper is in orbit around Jupiter. Note that most off-point events are under a few hours while there are some outlier events around 10 hours in duration. These longer instances of the solar array geometry remaining especially poor occur during flybys, when the spacecraft attitude is fixed such that the remote sensing instruments are pointed nadir and the in-situ particle detection instruments are pointed towards the spacecraft velocity vector.

Simulation results of the solar array off-point events were combined with eclipse time and duration data based on the reference trajectory to determine predicted thermal cycles for the mission. Additional margin was added to the results to account for uncertainties in the simulated mission profile. Thermal engineers then converted these results into their test profile as seen in Table 3, which consists of a series of discrete thermal cycle profiles with varying temperature ranges.

Table 3 Solar array qualification test profile accounting for predicted thermal cycles during the mission and on the ground prior to launch.

	Tmin (°C)	Tmax (°C)	Delta-T (°C)	No. of Cycles	3x life	Test Medium
PP Thermal Cycles	25	150	125	1	3	LN2
PF Hot Thermal Cycles	25	120	95	3	9	LN2
PF Cold Thermal Cycles	-240	-125	115	3	9	LHe
Thermal Cycles: Up to 2AU	-5	100	105	5	16	LN2
Thermal Cycles: 2-4AU	-80	-20	60	3	9	LN2
Thermal Cycles: 5AU, Tour, some Eclipse Thermal Cycles (short)	-180	-120	60	88	263	LN2
Eclipse Thermal Cycles (Long)	-240	-125	115	35	106	LHe

F. Operability

Operability is often overlooked early in the project lifecycle. However, designing a system that is difficult to operate can adversely affect mission performance, increase operations cost, and decrease reliability. Finding operational problems early is extremely beneficial, as late hardware design changes are expensive or may not even be possible. During early project formulation, the Europa Clipper project established a project policy to consider operability when making design decisions by establishing a project policy on operability and staffed an Operability Engineer [14]. Mission simulations built using APGen, a tool originally designed for operations, have helped enforce the operability policy by identifying operational concerns stemming from the trajectory, spacecraft, instrument, and initial operations design since Pre-Phase A.

For example, simulations of launch showed potential problems with the communication link between the spacecraft and Earth during post-launch critical events such as solar array deployment. Just after launch, the spacecraft will separate from the upper stage and continue its “barbeque roll” to keep the spacecraft thermally stable until it can obtain an inertial reference and transition to 3-axis control. While rolling, the best antenna for communicating with Earth will alternate between the two LGAs, and without inertial reference, the spacecraft would not know when to switch the signal from one antenna to another. Moreover, once the spacecraft establishes inertial reference and 3-axis control, the spacecraft must point the HGA towards the Sun to protect thermally-sensitive spacecraft components. This attitude places the Earth in a location that is nearly 90° away from the boresights of the LGAs. Operationally, having little to no capability to assess spacecraft health and command in the event of a problem during such a critical phase of the mission is extremely concerning. Since this operability issue was caught early, spacecraft engineers have examined different solutions including moving the locations of the LGAs on the spacecraft or adjusting the attitude strategy (as long as that strategy is also thermally safe.)

Some additional examples from the Clipper project where simulations and subsequent analyses have contributed to operability are provided in the following list:

1) Timelines produced by simulation runs showed the need for automation in the maneuver design process due to the criticality and short turnaround times of orbit trim maneuvers (OTMs).
2) Integrated plans scheduled by APGen showed that there are a limited number of conflicting activities during the flyby period, which should reduce the need for manual negotiation and activity planning.
3) Analysis of the positions of bright bodies like the Sun and Jupiter relative to the spacecraft given its orientation over the mission revealed there were no good locations to place instrument radiators. This lead to instrument design selections that included cryocoolers, thus avoiding operational mitigations requiring an adjustment of attitude to prevent over-heating the instrument (such mitigations were require on the Cassini mission.)
4) Simulations provided evidence that the post-flyby sorting of recorded science data required by the proposed BDS design could be accommodated with minimal impact to science opportunities and data playback [14].

G. Requirements Verification

Ensuring the concept of operations and the trajectory design can meet the science needs of the mission is one of the primary charters of mission designers, and mission-level simulations have proven to be an essential ingredient towards achieving this goal. After agreeing on the high-level scientific objectives of the mission, scientists quantified the amount and types of data each instrument would need to collect to meet these objectives and captured these needs as science measurement requirements. Evaluating achievement of over 100 science measurement requirements is not trivial, especially when coupled with 70+ geometric conditional requirements and a trajectory that is tweaked and redesigned annually. These factors make measurement requirement evaluation by inspection or by scripts spread across instrument teams infeasible.

A software tool called Verification of Europa Requirements Integrating Tour and Science (VERITaS) was built to aid mission planners and trajectory designers by quickly evaluating all science measurement requirements automatically based on simulations performed using APGen. VERITaS excels at assessing different trajectories and operations strategies to determine the level by which they meet measurement requirements. Written in MATLAB®, VERITaS analyzes requirement satisfaction and margin against those requirements based on a planned set of instrument observations and spacecraft attitude profile provided by APGen. With these inputs, VERITaS performs various computations necessary to check the requirements against the plan including complex geometrical and instrument surface coverage calculations. VERITaS makes widespread use of NAIF’s SPICE toolkit for many of its geometric calculations and frame transformations [7, 11]. The incorporation of SPICE allows for standard SPICE kernels containing trajectory, attitude, and frame information used across the project to be ingested. Instrument observations are supplied to VERITaS via a simple time-ordered-list CSV file, and the spacecraft attitude and solar

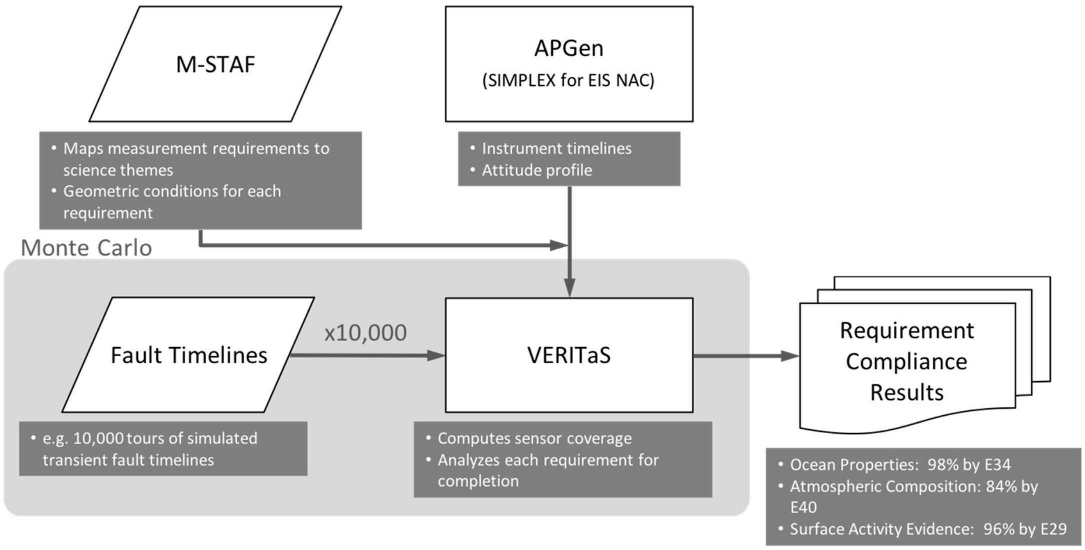

Fig. 13 Flowchart of inputs and outputs to the Verification of Europa Requirements Integrating Tour and Science (VERITaS) tool for science measurement requirement compliance assessment. The ability for VERITaS to ingest fault timelines in a Monte Carlo simulation is also shown, but is an optional feature [15].

array orientation, along with the independently pointable EIS NAC, are provided as SPICE C-kernels. A flowchart depicting the requirement assessment workflow is shown in Fig. 13.

In order to check an observation plan against requirements, VERITaS must also read in the requirements. Science measurement requirements are represented in the format of the Mission Science Traceability Alignment Framework (M-STAF), which is an Excel workbook that organizes Level-2 measurement requirements by instrument and relates them to corresponding geometry requirements (e.g., altitude, lighting, or ground speed) and science themes (e.g., ocean properties). While the M-STAF is useful as a stand-alone product on the project [16], the unique style of organizing requirement makes it indispensable for interpreting and assessing science measurement requirements. For ingestion into VERITaS, the Excel sheet is reduced to a simple text file with special syntax that VERITaS knows how to interpret.

Once VERITaS has a requirement set and observation plan, it begins by computing sensor coverage of Europa's surface for each remote sensing instrument using a specially built coverage tool. The coverage tool was built from scratch to accommodate input of SPICE kernels (including instrument attitude C-kernels), input of an observation timeline, and the generation and saving of coverage on a fine temporal scale. The fine temporal scale is necessary to take advantage of VERITaS's ability to ingest simulated fault timelines (see Sec. H for more information on the faulted Monte Carlo analysis). VERITaS then churns through each science measurement requirement, applying necessary geometry and lighting constraints, and determines if and when each requirement is met. All information coupled to each requirement is saved together for future reference. Data related to each requirement includes science theme, observation technique type, geometry constraints and their requirement IDs, pass/fail information, and the margin with which the requirement is met. Multiple plots are also automatically generated and saved in the process. Figure 14 contains an example visible imaging coverage map and the corresponding accumulation of surface coverage with each additional flyby of Europa over the course of the mission. Figure 15 highlights how this information can be used to provide insight into how much margin exists on the requirement given the current observation plan and trajectory.

Analyzing a VERITaS run can be done quickly by looking at an automatically generated Excel sheet and a PDF report full of both summary and detailed statistics and charts. The Excel sheet lists each requirement, provides a short description of the requirement, notes whether it passed and, if so, when it passed, and even color-codes the requirement based on whether it was achieved. The PDF report contains plots and tables for each requirement and is automatically compiled using LaTeX. The plots and tables provide background information and can quickly be incorporated into presentations and papers. As an example, Fig. 16 shows a histogram of how many requirements are met by encounter. Together these output products give the Europa project a quick assessment of which measurement requirements are failing or are more difficult to meet.

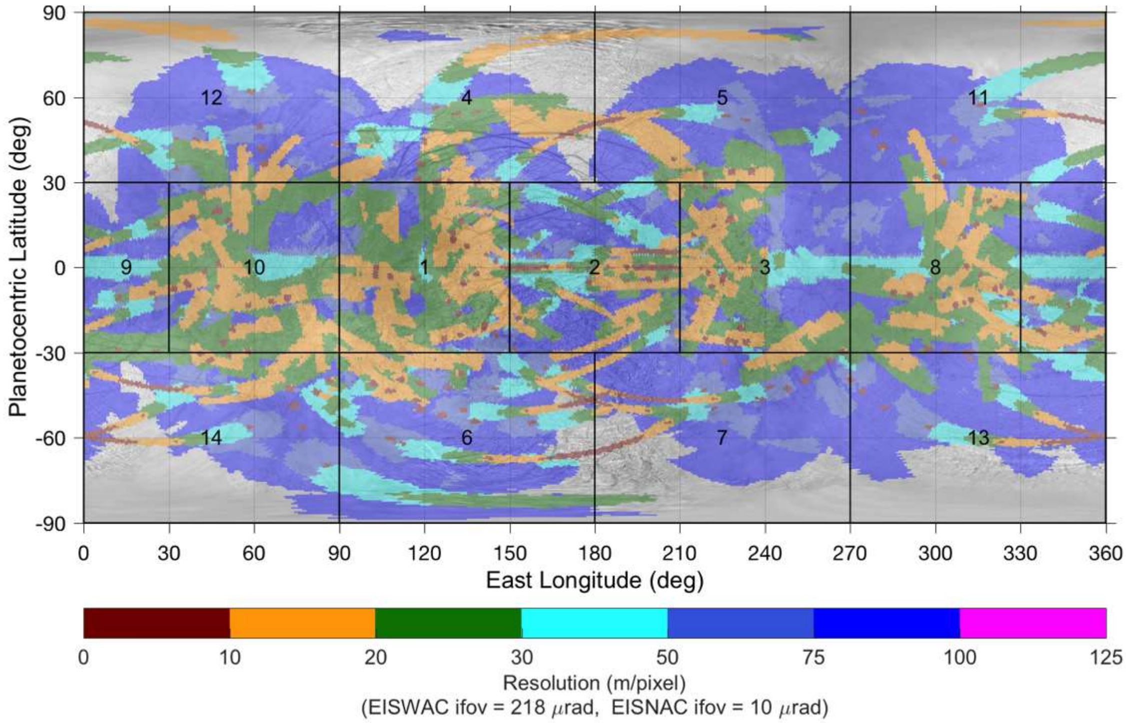

Fig. 14 Combined global-scale coverage of Europa for the EIS NAC and EIS WAC instruments colored by resolution. Constraints have been applied for altitude, pixel scale, and lighting.

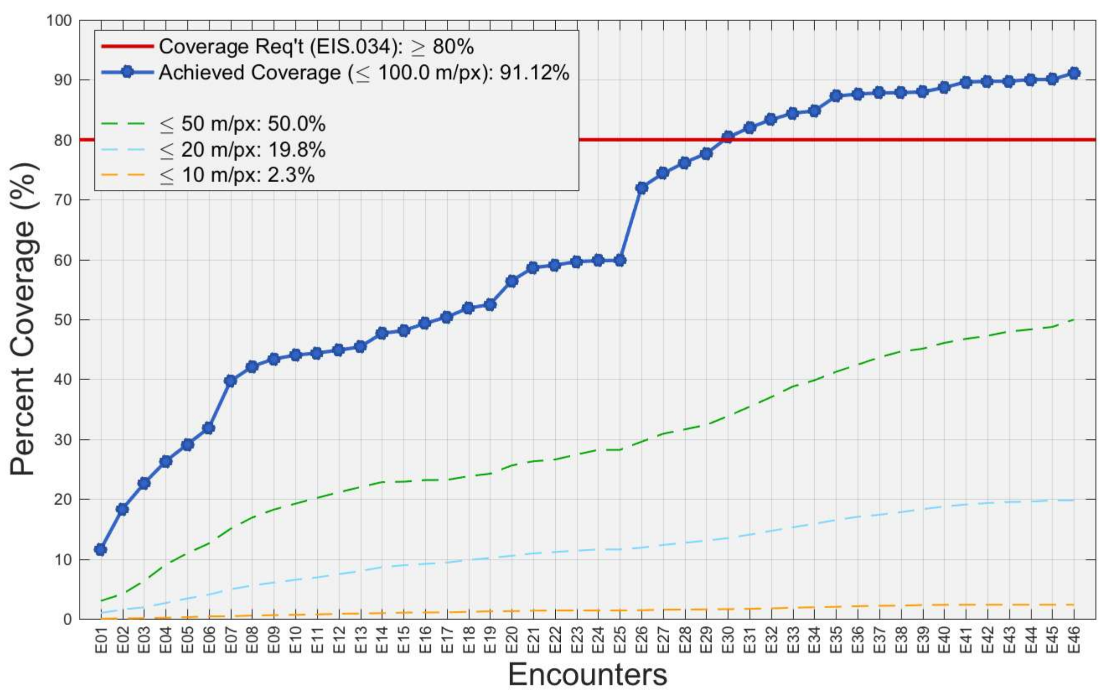

Fig. 15 The solid blue line shows the combined EIS NAC and EIS WAC visible imaging surface coverage of Europa over each encounter, cumulatively by encounter; the coverage requirement is met at E30. Coverage accumulations are also shown at different levels of resolution.

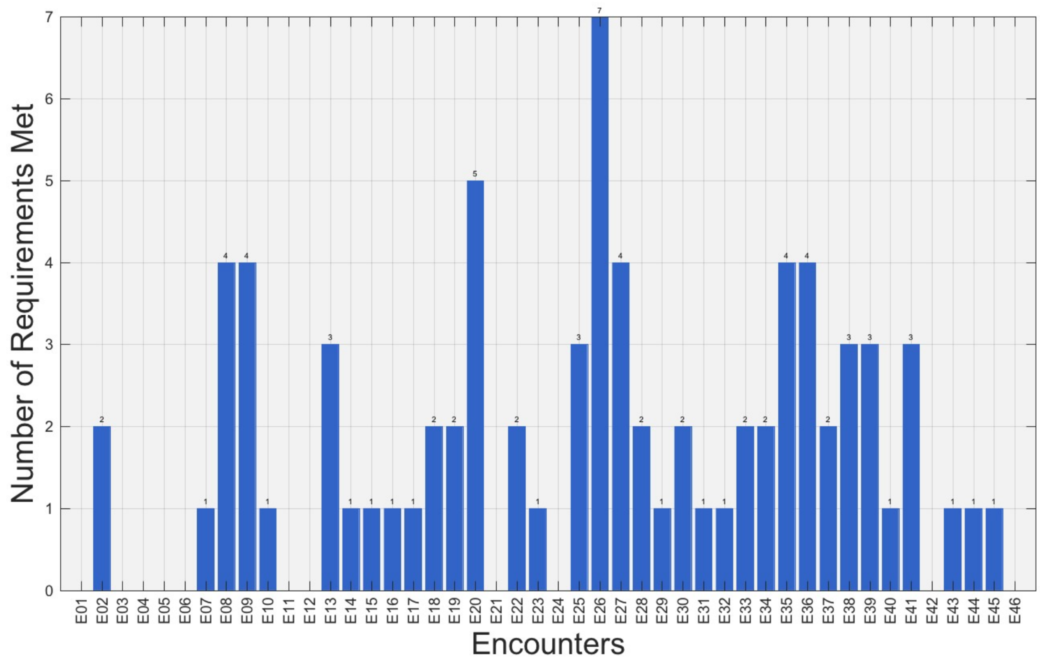

Fig. 16 Histogram of when each science measurement requirement is predicted to be met, providing insight into margin and robustness.

The use of VERITaS in conjunction with APGen simulations and their integration into the mission planning and trajectory design process has catalyzed updates to the mission concept of operations, initiated trajectory design tweaks, and even triggered updates to the measurement requirements themselves [15]. The tool has allowed for rapid requirement evaluation in the midst of an evolving flight system design and operations concept on a nearly weekly basis.

H. Fault Sensitivity Analysis

In addition to using VERITaS to assess a nominal mission, VERITaS can aid in analyzing the effect of fault rate and recovery time on requirement satisfaction and margin by ingesting simulated fault timelines [17]. As depicted in Fig. 13, thousands of simple fault timelines can be ingested by VERITaS and analyzed in a Monte Carlo process, effectively simulating thousands of "missions" with varying fault rates and recovery times. This capability is especially useful due to Europa Clipper's destination, Jupiter and its high-radiation environment.

Fault timelines are generated using historical spacecraft safing and transient fault rates along with models of the Jovian radiation environment based on data from Pioneer, Voyager, Galileo, and Juno. The probability of a fault occurring is a function of time and the flight system's distance to Jupiter. Still, there is uncertainty on the exact fault rate each instrument and the spacecraft will see in operations, so a range of fault rates have been assessed to date in order to bound the problem [17]. The goal of this type of analysis is threefold:

1) assess the robustness of the science requirements in the presence of transient faults and safings,
2) inform the design of the recovery time capability of the flight system, and
3) ensure the Tour duration is adequate to achieve the science objectives in the presence of potential disruptions.

Thousands of fault timelines have been run through VERITaS to simulate the loss of science data due to these putative transient faults. An example of resulting information is shown in Fig. 17, which illustrates how the probability of a particular requirement being met can be reported for each assumption about the fault rate and recovery time. This has become a powerful tool for steering the design of spacecraft recovery time, operations concepts, and even modifications to the measurement requirements themselves.

Baseline L1 RQ	Low transient fault rate *Current Design Recovery Time*	Low transient fault rate *No spacecraft recovery*	Low transient fault rate *No instrument recovery*
RQ 1	75.0%	49.0%	40.0%
RQ 2	95.0%	95.0%	90.0%
RQ 3	95.0%	90.0%	93.0%
RQ 4	100.0%	100.0%	95.0%
RQ 5	90.0%	81.0%	80.0%
RQ 6	100.0%	100.0%	95.0%
RQ 7	100.0%	100.0%	90.0%
RQ 8	100.0%	95.0%	98.0%
RQ 9	100.0%	100.0%	100.0%

Fig. 17 Example results of a Monte Carlo faulted VERITaS analysis showing the probability of each measurement requirement being met given the fault rate and recovery assumptions in each column [17].

I. Europa Lander Concept

The Clipper APGen adaptation was much quicker to set up than it would have been otherwise due to the multi-mission infrastructure that past projects had built up. Improvements made for Clipper are already proving very useful for other mission concepts being developed, especially Europa Lander. The Europa Lander simulation framework, as described in Ref. [18], branched off directly from the Clipper APGen adaptation. The Carrier and Relay Stage in the 2017 reference Lander mission, which would orbit Europa and provide relay support for the Lander, was directed to have as much heritage hardware as possible from Clipper, so it was easy to start with the Clipper adaptation as a starting point for the Carrier. Common components, which accounted for about half of the lines of code in the Clipper adaptation, were refactored into a 'Generic' directory from which both adaptations could import code, and that can form a more formal basis for all future APGen adaptations. The main generic components are: CDH, DSN, Geometry, GNC, Ground Station, MOS, Power, Propulsion, Radiation, Sequence, Solar Arrays, Telecom, Telemetry, and Thermal. The planned Psyche mission, whose adaptation will be starting up soon, should be able to reuse all of these components with minimal modifications, inheriting almost everything they would need run integrated simulations with the exception of the project-specific science instrument models.

VI. Improving Future Mission Simulations

A. Simulation Performance

For the most part, project customers of simulation results and analyses were pleased, if not impressed, by how quickly APGen adaptation developers could turn around requests. Nevertheless, there have been a few cases where lengthy simulation run times have impacted decision-making timelines. Full simulation runs including all activity scheduling can take up to 1.5 days to complete. Although multiple simulations can run in parallel, each simulation uses on the order of 40 Gigabytes of RAM. Even with a number of high-powered machines, the number of concurrent simulations has been limited. There are a number of enhancements, both to the APGen framework and adaptation, that together have the potential to significantly improve simulation performance by an order of magnitude or more.

After taking a hard look at the core APGen framework, developers determined one way to increase speed was to identify as much information as possible at ‘compile time’ (i.e. when APGen parses the adaptation). If information is known up front, parsed expressions from the adaptation can be organized very efficiently so that the CPU does not have to spend precious CPU cycles finding out what it needs to do. Work is currently ongoing to refactor APGen so that it can efficiently parse the adaptation and thoroughly analyze its content prior to execution. Test cases running a prototyped version of this “faster APGen” using a subset of the Clipper adaptation have already shown a boost in runtime performance by a factor of 3.4.

As the Clipper adaptation builds up its activity plan, it runs a series of activity schedulers, each of which invokes a separate instance of APGen. In order to correctly schedule the next set of activities based on resource constraints, many of these schedulers must first model the current state of the plan, which involves calculating the effect each activity in the plan has on various resources. In the current adaptation, many of these calculations are unnecessary because they focus on resources that turn out not to be needed in the evaluation of constraint windows suitable for adding new activities to the plan. Moreover, each scheduler often has to recalculate the same resource values even if their values do not change because each scheduler is a separate APGen process. Additional logic is being introduced into the adaptation so that resource calculations are only performed when necessary. Switches are now provided that turn subsystem models on and off based on whether resources within those models need to be calculated. Recent updates to the APGen framework also allow specific resource values to be saved off into 'partial' XML TOLs, which can then be read into subsequent instances of APGen. This functionality is used between many schedulers within the Clipper adaptation and has already reduced run time by over a factor of 2. There are still plenty of additional opportunities to modify the adaptation to take more advantage of these performance-enhancing capabilities.

B. Transformation of Models, Activities, and Constraints to Simulation Inputs

The vast majority of the time spent developing and refining the Clipper APGen adaption was devoted to understanding what activities encompass the mission plan, the constraints associated with each activity, and how those activities behave and interact with the rest of the system. Mission planners, who were the primary developers of the adaptation, would gather this information from each subsystem via meetings or documentation that was often vague or incomplete. Next, mission planners would manually create activity definitions and schedulers within APGen based on their interpretation of the gathered information and run a simulation. Simulation results would be reviewed by subsystem experts and inevitably the definitions and scheduling logic would be inaccurate. This would lead to one or more iterations between developers (planners) and customers (subsystems) before a satisfactory plan would be produced.

One way to prevent "lost in translation" issues would be to have subsystems develop their own models, activity definitions, and constraints within the simulation framework. This can be challenging as learning the APGen DSL is non-trivial and subsystem experts usually don't have the time or resources to devote to building additional system-level models for their subsystem. However, APGen does provide a way to interface with other model libraries outside of the APGen DSL, and this has been done successfully in the past for both power and telecom models.

Another solution would be to get project members to describe activities and model behavior in a more standard modeling language that could then be translated automatically into APGen for simulation execution. Since the inception of the project, Europa Clipper has emphasized and invested in Model Based System Engineering (MBSE) practices that employ common modeling languages such as SysML to describe various aspects of the system design [19]. Data captured in these languages serve as a Single-Source-of-Truth (SSoT) for the project that other tools can pull from and know they are receiving reliable information [20]. MBSE efforts on Clipper have had success in some areas, but little energy has been devoted to providing a simple way for subsystems to capture behavioral information that can then be easily queried by downstream simulation tools like APGen. Currently, the only information the APGen adaptation pulls directly from the SSoT Europa Clipper "System Model" is the PEL (see Section IV.A). Nonetheless, capturing information about activities, constraints, and behaviors using MBSE seems promising and should be pursued with more force as it has the potential to save hours of simulation development time and prevent misunderstandings between teams.

C. Parameterization and Automation

A single run of a mission simulation represents a single point design with a specific trajectory, spacecraft hardware configuration, and concept of operations. During preliminary design, engineers are challenged to explore the design space in order to find solutions that best meet mission objectives and constraints. One effective means of exploring this space is by running multiple simulations that each tweak different aspects of the mission design. Clipper found success using this method on targeted trade studies, but only a small number of design parameters were available to change and runs typically were manually kicked off. Building an architecture that easily exposes design parameters as simulation inputs would allow for a more thorough and more efficient exploration of the design space. Monte Carlo methods could be applied that automatically run simulations that vary design parameters to gauge which aspects of the design are the most driving. In order to effectively compare results from the multitude of simulations that would be run, a highly organized and efficient data storage infrastructure would be needed. The Clipper project has begun to tackle this "big data" problem through the use of industry standard unstructured databases as well as search and data analytics tools. However, there is still much work that could be done to improve the current infrastructure so it could support storing and analyzing multiple simulation runs on a much larger scale.

VII. Potential Use in Operations

It would be a shame to build up such a powerful simulation capability through project development just to throw it away when the mission begins operations. Fortunately, integrated simulations are expected to play a crucial role in operations for Europa Clipper and plans are already in the works to leverage as much capability as possible from the Clipper APGen adaptation. In fact, the operations process currently envisioned for Clipper revolves around planning and validating spacecraft behaviors at the activity level. Once planners agree upon a plan, the activities within the plan would automatically decompose into their respective spacecraft commands. Although the operations design is still rapidly evolving given the project is still in Phase B, some key aspects to the design include:

1) A centralized modeling and simulation environment used to perform all spacecraft and instrument planning and validation
2) Rapid, continuous integration and validation of plans
3) Highly automated processes to generate, validate, and approve final uplink products (e.g. commands and sequences)

A limitation of many previous mission efforts was that the evolution of the uplink products from the plan to sequences went through tool discontinuities as the level of detail increased. This lead to an iterative and often error prone planning process as the intent and behavior (such as resource usage) captured at the activity level regularly didn't match the delivered sequences. These inconsistencies would only get caught late in the planning process, which meant sequences had to be re-delivered and re-validated on extremely short time frames. With an integrated activity planning tool like APGen, however, increasing level of detail can be readily accommodated by capturing a more sophisticated algorithmic description within the definition of each activity. This means that the implemented plan is the direct ancestor of the sequence products that eventually get sent to the flight system. Such continuity ensures that the rationale for the planned strategies won't get lost as a result of switching between tools in the uplink process.

One such mission where the continuity between the plan and uplink products was maintained with success was the Deep Impact Mission. On Deep Impact, APGen was also used starting in Phase B to build integrated simulations of the comet Temple-1 encounter. The APGen adaptation was originally based on the preliminary design, but as actual flight system and flight software capabilities became available, the adaptation evolved from notional activities and design behaviors to actual spacecraft commands and measured behaviors. The commands produced by the adaptations were then validated by running integrated, system-level tests using hardware-in-the-loop test beds.

Even now in the early stages of the project, Clipper's APGen adaptation already contains an extraordinary level of detail about the planned spacecraft behavior for the entire mission. Most activities already decompose down to the command level and in some cases those commands are already being vetted and validated by subsystem experts. The current adaptation already has many of the desired qualities operators are looking for in a planning tool for Clipper. Simply evolving the Clipper adaptation and enhancing its capabilities where needed could significantly reduce cost both during development and operations. Whether or not the adaptation is used directly, the knowledge embedded in the adaptation about the system and how it behaves is readily available and will undoubtedly serve as the foundation for any future operational planning tool.

VIII. Conclusion

Mission level simulations have had a profound impact on the Europa Clipper project and the spacecraft and mission design. The unprecedented level of detail in the simulations have permitted managers and engineers alike to gain insights into the design and operations of Clipper that would have otherwise gone unnoticed until much later in the project lifecycle. Early in the design process, results from simulations gave mission operations engineers "a seat at the table" by providing quantitative evidence of operability concerns with the design. As the project progressed, the scope of simulations only expanded. Since their creation, end-to-end simulations have been used on Clipper on a daily basis to support mission planning, trades, hardware design and testing, requirements development, and even to help prove the mission can meet its bold scientific objectives. Ultimately, direct descendants of these simulations may even be used as the central engine behind planning during operations.

With high-fidelity simulations, the entire mission could be flown thousands, if not millions, of times before the spacecraft ever reaches the launch pad. If these simulations provide sufficient fidelity, run quickly, and simulation results are easily searchable and cross-compared, engineers could efficiently explore a large design space and produce a more optimal design. One could even imagine using automation and computer intelligence techniques in conjunction with simulations to discover the top design candidates for the engineers. Current simulations on Clipper represent a step in this direction, and future missions should consider adopting mission-level simulations early and advancing their capabilities towards this goal.

Acknowledgments

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The authors would like to thank Laureano Cangahuala, Nathan Strange, Kelli KcCoy, and Dave Mohr for taking the time to provide a thorough technical review of this paper. Also, a special thanks to Adam Roberts for providing additional feedback on grammar and style.

References

[1] Wissler, S., Maldague, P. F., Rocca, J., and Seybold, C., "Deep Impact Sequence Planning Using Multi-Mission Adaptable Planning Tools with Integrated Spacecraft Models," *AIAA 9th International Conference on Space Operations (SpaceOps)*, Rome, Italy, 2006.
doi: 10.2514/6.2006-5869

[2] Mitchell, A., et al, "MAPGEN: Mixed-Initiative Planning and Scheduling for the Mars Exploration Rover Mission," *IEEE Intelligent Systems* [online journal], Vol. 19, No. 1, pp. 8-12.

[3] McKinnon, W.B., Pappalardo, R. T., and Khurana, K.K., "Europa: Perspectives on an Ocean World," *Europa*, edited by R.T. Pappalardo, et al., Univ. Arizona Press, Tucson, pp. 697-710, 2009.

[4] Pappalardo, R. T., et al., "Science Objectives for the Europa Clipper Mission Concept: Investigating the Potential Habitability of Europa," *European Planetary Science Congress*, Vol. 10, EPSC2015-156, Nantes, France, Sep. 2015.

[5] Buffington, B., Campagnola, S., and Petropoulos, A., "Europa Multiple-Flyby Trajectory Design," *AIAA/AAS Astrodynamics Specialist Conference*, Minneapolis, MN, 13-16 Aug. 2012.
doi: 10.2514/6.2012-5069

[6] Lam T., Buffington B., and Campagnola, S., Scott, C., Ozimek, M., "A Robust Mission Tour for NASA's Planned Europa Clipper Mission," *2018 Space Flight Mechanics Meeting, AIAA SciTech Forum*, Kissimmee, FL, 8-12 Jan. 2018.
doi: 10.2514/6.2018-0202

[7] Acton, C.H., "Ancillary Data Services of NASA's Navigation and Ancillary Information Facility," *Planetary and Space Science*, Vol. 44, No. 1, pp. 65-70, 1996.

[8] Garrett, H. B., Martinez-Sierra, L. M., Evans, R., "Updating the Jovian Proton Radiation Environment," JPL Publication 15-9, Jet Propulsion Laboratory, National Aeronautics and Space Administration, Pasadena, CA, Oct. 2015.

[9] Maldague, P.F., Wissler, S., Lenda, M., and Finnerty, D., "APGEN Scheduling: 15 years of Experience in Planning Automation," *AIAA 13th International Conference on Space Operations (SpaceOps)*, Pasadena, CA, 5-9 May 2014.
doi: 10.2514/6.2014-1809

[10] Cole, B., Dinkel, K., "Multidisciplinary Model Transformation through Simplified Intermediate Representations," *IEEE Aerospace Conference*, Big Sky, MT, 5-12 Mar. 2016.
doi: 10.1109/AERO.2016.7500656

[11] Acton, C., Bachman N., Semenov, B., Wright, E., "A Look Toward the Future in the Handling of Space Science Mission Geometry," *Planetary and Space Science*, Vol. 150, Jan. 2018, pp. 9-12.
doi: 10.1016/j.pss.2017.02.013

[12] Semenov, B. V., "WebGeocalc and Cosmographia: Modern Tools to Access OPS SPICE Data," *AIAA 15th International Conference on Space Operations (SpaceOps)*, Marseille, France, 2018 (to be published).

[13] Oaida, B., Lewis K., Ferguson E., Day J., McCoy K., "A Statistical Approach to Payload Energy Management for NASA's Europa Clipper Mission," *IEEE Aerospace Conference*, Big Sky, MT, 3-10 Mar. 2018.

[14] Signorelli, J., Bindschadler, D. L., Schimmels, K. A., and Huh S. M., "Operability Engineering for the Europa Clipper Mission: Formulation Phase Results and Lessons," *AIAA 15th International Conference on Space Operations (SpaceOps)*, Marseille, France, 2018 (to be published).

[15] Buffington, B., et al., "Evolution of Trajectory Design Requirements on NASA's Planned Europa Clipper Mission," *68th International Astronautical Congress (IAC)*, IAC-17-C1.7.8, Adelaide, Australia, 25-29 Sep. 2017.

[16] Susca, S., Jones-Wilson, L. L., Oaida, B. V., "A Framework for Writing Measurement Requirements and its Application to the Planned Europa Mission," *IEEE Aerospace Conference*, Big Sky, MT, 4-11 Mar. 2017.
doi: 10.1109/AERO.2017.7943689

[17] McCoy, K., et al., "Assessing the Science Robustness of the Europa Clipper Mission: Science Sensitivity Model," *IEEE Aerospace Conference*, Big Sky, MT, 3-10 Mar. 2018.

[18] Lawler, C. R., Wissler, S. S., Kulkarni, T., Ferguson E. W., and Maldague P. F., "Europa Lander Concept: High fidelity system modeling informing flight system and concept of operations years before launch," *AIAA 15th International Conference on Space Operations (SpaceOps)*, Marseille, France, 2018 (to be published).

[19] Bayer, T. J., et al., "Model Based Systems Engineering on the Europa Mission Concept Study," *IEEE Aerospace Conference*, Big Sky, MT, 3-10 Mar. 2012.
doi: 10.1109/AERO.2012.6187337

[20] Dubos, G. F., Coren, D. P., Kerzhner A., Chung, S. H., Castet, J., "Modeling of the Flight System Design in the Early Formulation of the Europa Project," *IEEE Aerospace Conference*, Big Sky, MT, 5-12 Mar. 2016.