Distributed Operations as Applied in a Large Multi-Instrument Space Mission: Lessons Learned from the Cassini-Huygens Program

L.Y. Cheng, B. Larsen,

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, U.S.A

ABSTRACT

Launched in October 1997, the Cassini-Huygens Mission sent the largest interplanetary spacecraft ever built in the service of science. Carrying a suite of 12 scientific instruments and an atmospheric entry probe, this complex spacecraft to explore the Saturn system may not have gotten off the ground without undergoing significant design changes and cost reductions. As a means to control operations cost, a novel concept, called "Distributed Operations", was chosen. This concept utilized advances in information technology and distributed computation to decentralize the mission control room and science operations.

Although other interplanetary missions have decentralized science operations, none of them can match Cassini-Huygens in the scale and complexity of the distributed system. With 12 science teams and one Huygens Probe team distributed worldwide, the coordination effort to design and integrate command sequences posed a significant challenge in mission operations. Furthermore, the spacecraft design of mounting remote sensing instruments and in-situ instruments on a fixed spacecraft bus introduced operational constraints for even the simplest of maneuvers.

As Cassini-Huygens approaches its tour of Saturn, we reflect on the broad implications of the Distributed Operations (DO) concept, especially to large scale, multi-instrument spacecraft. We find the roles of the Scientist and the Engineer under the DO environment to be blurred. We find the act of decentralization necessitated the Scientist to redefine their roles in the mission operations system. In particular, the Cassini spacecraft's lack of a scan platform for remote sensing instruments, some of whom have complex pointing constraints, blurred instrument health and safety with spacecraft bus health and safety. In other words, to safely point his or her instrument, the Cassini Remote Sensing Scientist has a vested interest in spacecraft attitude simulation and constraint checking. This expands the Scientist role from a spacecraft user to include spacecraft safety assurance, a responsibility traditionally assigned to the Engineers.

We conclude by arguing the case for Distributed Operations must be made carefully. We propose some criteria to assess if a particular mission design is suited for a Distributed Operations environment. The criteria are: a) Number of science instrument payloads; b) The scope of shared resources among science

instruments (i.e. non-inclusion of a scan platform means spacecraft attitude is a vital shared resource); c) The inclusion of a mission director to resolve Science vs. Engineering conflicts.

BACKGROUND: The Cassini-Huygens Mission to Saturn is a complex mission with a large multi-instrument spacecraft.

The Cassini-Huygens Mission is a joint NASA and ESA mission to study Saturn. Set to arrive at Saturn on July 1, 2004, the spacecraft carries on board a sophisticated array of cameras, spectrometers, and sensors, covering much of the electromagnetic spectrum. This suite of 12 science instruments and one planetary entry probe were designed to address the following varied set of mission objectives: a) to investigate the chemical composition and physical state of Saturn's atmosphere; b) to investigate the chemical composition and physical state of Titan's atmosphere and surface; c) to investigate the chemical composition and physical state of icy Saturnian satellites; d) to investigate the chemical composition and physical state of Saturn's rings; e) to investigate the structure and physical dynamics of the Saturnian magnetosphere.

Because we are an international mission, the Cassini-Huygens flight operations teams are located in 8 states in the United States as well as Germany, Great Britain, France and the Netherlands. This poses major challenges in space mission operations. New concepts that go beyond the traditional "Mission Control" room paradigm must be utilized in order to effectively perform the mission objectives, while controlling operations cost.

The Need To Decentralize Cassini-Huygens Mission Operations

Voyager and Galileo: A Comparison

Cassini-Huygens is similar to two prior missions to explore the outer planets in the solar system: Voyager and Galileo. All three missions sent a spacecraft (in Voyager's case two spacecraft) with a large package of science experiments. The instrument packages all included an array of remote sensing cameras and spectrometers and in-situ fields and particles sensors. Galileo shares with Cassini-Huygens a passenger, in the form of a planetary entry probe, and its associated probe telecommunications relay system. Finally, all three spacecraft spent many years in cruise before arriving at its destination.

There are also major differences between the three missions. First, the Voyager missions were not designed to orbit a planet. Planetary and satellite flybys and encounters were the name of the game for the Voyager flight operations team. Secondly, both Voyager and Galileo carried a scan platform, which allowed the remote sensing instruments to point independent of the rest of the spacecraft. This allowed the spacecraft to point to scientifically interesting targets, while having its high gain antenna point towards the Earth. As we shall see, the lack of a scan platform for Cassini-Huygens created challenges to mission operations. Thirdly, both Voyager and Galileo mission operations were modeled after a traditional centralized operations scheme. Cassini-Huygens, on the other hand, choose to decentralize, or "distribute" mission operations. Although some aspects of distributed operations has been implemented on other JPL missions (e.g. Mars Global Surveyor), we believe the Cassini-Huygens implementation, with its much more varied instrument suite and international scope, has truly pushed distributed operations to its limits.

The Advantages of Distributed Operations for Cassini-Huygens

For Cassini-Huygens, the choice of distributed operations held many advantages. First, the non-real time nature of the mission allows this to be done. Cassini-Huygens will be nominally utilizing only one Deep Space Network (DSN) tracking per day. There is no need for a continuously staffed flight team to support round the clock station coverage, seven days a week. And the long one-way light time (1.5 hours on average) precludes much real-time commanding. Therefore, the spacecraft command process is a relegated to the development of stored command sequences. Secondly, the data architecture on-board the spacecraft is distributed. Each science instrument has its own processor, independent of the spacecraft central processor. The spacecraft central processor (called CDS), serves as a router, relaying science instrument commands. All science instrument commanding can be pre-planned outside of JPL, and "merged" into a main stored command sequence, by a central JPL team, sometime before uplink. Thirdly, operations cost can be controlled by "outsourcing" science operations to non-JPL sites. This along with the international nature of the mission, lends itself to decentralization. Unlike Galileo and Voyager, where there are "liaison" science teams at JPL representing the Principal Investigator (P.I.) teams located outside JPL, Cassini-Huygens directly coordinates with the P.I. teams. Those P.I. teams interact with JPL just as any other spacecraft subsystem team. Utilizing distributed computing networks, the P.I. teams would be able to command the spacecraft and process instrument telemetry from their home institutions. Given the multiple time zones in which the P.I. teams reside, distributed operations makes sense. Finally, distributed operations has the added effect of bringing in educational and research institutions to the space exploration effort.

Distributed Operations Works Best When Subsystems And Constraints Are Decoupled

The geographic separation of the P.I. teams, the non-real time nature of command sequence development, and the independent processors internal to each science instrument on-board the spacecraft have one thing in common: the decoupling of subsystems, responsibilities and constraints.

Case study 1: Issue a Command Internal to a Science Instrument - Successful application of Distributed Operations

As a case study in how distributed operations works best, we will look at how a Cassini-Huygens distributed science team issues a command internal to the instrument itself. When a P.I. team wants to issue a command to the team's own instrument, there is only one spacecraft constraint they have to satisfy: make sure command does not exceed power allocation. Since Cassini-Huygens manages power through the use of pre-defined power modes, as long as the team's command does not allow the instrument to draw more power than it's allocated, the command can be issued at virtually any time in the command sequence. Several weeks before the uplink of the command sequence, the P.I. team can submit the instrument command to JPL, to be merged with the rest of the sequence commands as part of the sequence development process, and radiated to the spacecraft.

The ease and simplicity of this process can be attributed to the decoupling of the science subsystem, on-board the spacecraft via independent processors, and the clear definition of responsibilities for checking spacecraft constraints through the use of pre-defined power modes. As we shall see in the next case study, distributed operations will not be so easy if the subsystems are coupled to each other to share major resources, and constraint check responsibilities are blurred.

Distributed Operations Will Face Many Challenges When Subsystems Share Resources And Responsibility

The Cassini-Huygens spacecraft does not have a scan platform for remote sensing instruments. Both the remote sensing instruments and the in-situ fields and particles instruments are mounted on the spacecraft bus. Two fields and particles instruments can articulate, but its range of motion is not enough to be considered to point independently from the main bus. In order to point a remote sensing instrument to observe a target, the entire Cassini spacecraft must be moved. This can expose science instruments to thermal radiation.

As with all spacecraft with remote sensing instruments, there are certain orientations which can expose the sensors and cooling elements (also called radiators) to solar and other planetary thermal radiation. Those viewing constraints are a necessary evil in the world of spacecraft operations. However, if a spacecraft has a scan platform, an additional degree of freedom from an articulating platform can provide a margin of safety.

Case Study 2: Performing Remote Sensing in the Distributed Operations Environment, Without a Scan Platform

In the distributed operations environment of Cassini-Huygens, spacecraft pointing commands originate from the P.I. teams. They have the responsibility to design spacecraft attitude control changes, which do not violate the multitude of viewing constraint flight rules. For example, if the imaging team wants to image a particular portion of the Rings of Saturn, a team member must use ground software to model the spacecraft turn from the initial attitude to the target attitude, and perform the necessary slews in order for the camera's field of view to cover the region of space. In the Cassini-Huygens mission, this software also checks for violations of viewing constraints and generate the actual spacecraft pointing commands. It is up to the team member to not only design an observation that meets its science objectives, but also not to violate any viewing constraints. Furthermore, it is up to the P.I. team member to use as the initial attitude, the end attitude of the previous observation. Not doing this will could cause the spacecraft to attempt a dynamically impossible attitude change, trigger fault protection, thereby jeopardizing the safety of the entire spacecraft. The lines have been blurred between instrument health and safety and spacecraft health and safety. The Scientist now have the additional responsibility of assessing the health and safety of not just his or her instrument, but the entire spacecraft. Before we address the implication of this in the distributed operations environment, we must understand the traditional roles the Scientist and the Engineer play in all mission operations environments.

Engineering vs. Science: Checks and Balances

Although all members of the flight team are responsible for ensuring safe spacecraft operations, traditionally the Engineering team performs an independent assessment of spacecraft health and safety. In our imaging example above, an independent assessment of the pointing commands generated by the imaging team must be made. This is done when the imaging team delivers their command file to the appropriate lead engineer (also called sequence lead) responsible for integrating the sequence, in the sequence development process.

Once the sequence lead receives the imaging team file, he or she will create a master sequence file by merging all command files received from subsystem teams. Then, the AACS (Attitude and Articulation Control Subsystem) team, an arm of the Engineering team, will perform the independent check of the attitude control commands submitted by all teams. They will model the pointing profile to determine if any geometric (viewing constraint) or dynamic violations (excessive turn rates and/or accelerations) has occurred.

Because of the checks and balances set up between the Scientist and Engineer, a healthy “creative tension” exists in all JPL missions. The Scientist may want to push the operating envelope to collect some unique science data. The Engineer, on the other hand, may push the envelope back because it’s too risky. Each player knows their role; each team knows their responsibility.

Blurring of the Lines

Returning to the example of the imaging team’s Ring observations, if the P.I. team member designs a turn without any viewing constraint violations, using the approved ground software, he or she then delivers the command file to be checked by AACS. However, AACS using its own tool to check viewing constraints finds a violation from the imaging team’s observation. The question then becomes who’s assessment of the pointing profile is correct? If we look at this from the point of view of the centralized operations environment, the answer would most likely to use the AACS result. AACS is the expert when it comes to assessing pointing profiles.

In the distributed operations environment, the situation is trickier. Lets assume that the violation is excessive heating of the infrared spectrometer’s radiator element. With this in mind, the imaging team asks the infrared spectrometer team for an assessment. Given that it’s the infrared team’s flight rule, they should give the definitive answer as to whether the violation is real or not. If the infrared team responds with a “no, this isn’t a real violation”, then the case for AACS is weakened substantially. Even if the infrared team's findings are consistent with AACS, the lines have been blurred. Under distributed operations, the engineering team supporting the P.I team are often managed by the P.I., or have greater authority than previous centralized operations environments.

This is not a hypothetical scenario, but one which has played out many times during the development of Saturn tour sequences. Although the discrepancies between various models of pointing profiles have been resolved, they often involved many work hours among the various parties, therefore, creating operations complexity. Additionally, the blurring of the roles and responsibilities created un-healthy tension between Science and Engineering. One possible remedy to the increased workload and stress for the model discrepancy issue is to have the Mission Director resolve it. The Mission Director is the appropriate person to mediate conflicts between Science and Engineering.

Conclusion

In conclusion, we would like to present a method to assess mission operations concepts based on the Cassini-Huygens experience. The decoupling of instrument internal commanding lead to a successful implementation of distributed operations.

We should also point out the many challenges faced by the Cassini-Huygens because it lacked a scan platform. Figure 1 shows the relationship between a scan platform and the impact to centralized or distributed operations. Any savings from adopting distributed operations could be mitigated by the increase in mission complexity resulting from more viewing constraint violations, and the blurring of the lines between Science and Engineering.

	Scan Platform Exists	**Scan Platform Does Not Exist**
Centralized Operations	Less constraint violations No blurring of the lines between Science and Engineering	More constraint violations No blurring of the lines between Science and Engineering
Distributed Operations	Less constraint violations Blurring of the lines between Science and Engineering	More constraint violations Blurring of the lines between Science and Engineering ▪ Recommend Mission Director to mediate conflict

Figure 1: The relationship between the existence of a spacecraft with a scan platform and centralized vs. distributed operations.