

The Orion Microsatellite Mission: A Testbed for Command, Control, and Communications for Formation Fleets

Brian Engberg

Research Asst., Dept. of Aeronautics and Astronautics, Stanford Univ.
(650)-723-5450, briguy@stanford.edu

Robert Twiggs

Professor, Dept. of Aeronautics and Astronautics, Stanford Univ.
(650)-723-8651, btwiggs@stanford.edu

Abstract

The Orion microsatellite, under development at Stanford University, will fly along with two other Stanford satellites ("Emeralds") as part of a NASA-funded project. The primary objective is to demonstrate, for the first time, the use of carrier-phase differential GPS (CDGPS) for the relative sensing, navigation and coordinated control of satellites to form a *virtual spacecraft bus*. Launch of this mission has been tentatively scheduled for late 2001. Formation flying offers an exciting new approach to conducting space science missions. Instead of employing a single, large satellite, a fleet of similar, smaller spacecraft is coordinated to perform mission-related tasks. While formation flying architectures have a significant amount of operational flexibility, the internal system complexity increases with the number of satellites in the fleet. In addition, constraints on satellite resources play a particularly key role. This paper is a summary of work conducted at Stanford to investigate the influence of resource constraints on mission and current-task planning. By making efficient use of knowledge associated with mission goals and operations, optimal strategies can be used to increase fleet life-cycle performance. In addition to discussing this topic, the role of the Orion mission as a testbed for these concepts is included.

I. Introduction

Formation flying technologies will enable and enhance the performance of a variety of new space observation missions. Some of these missions aim to see further into space and with greater detail. Other missions strive to observe events on Earth with improved precision. In either case, the sensing instruments usually require very long baselines (up to 1 km or more) in order to achieve the desired image resolution. Conventional wisdom indicates that a single, large structure would be extremely expensive to design, construct, and launch for this purpose. Another way to achieve the desired functionality would be to construct a fleet of smaller, but well-coordinated spacecraft. However, this introduces a new set of technological challenges associated with fleet control, resource management, and inter-satellite communications. Addressing these topics is at the forefront of a number of investigations.

There is a great deal of current support for formation flying missions. Interferometry

missions such as Space Technology 3 (ST3) and Terrestrial Planet Finder (TPF) hope to utilize the benefits of a distributed satellite array in order to see the surfaces of stars and detect extra-solar planets [Refs 1,2]. TechSat21 will employ several satellites to form a synthetic aperture radar [Ref 3]. This will allow objects moving over the Earth's surface to be detected and tracked with a high level of detail. Other concept missions include LISA, a mission to assemble three satellites over extremely long baselines (millions of km) in order to detect relativistic phenomena such as gravity waves [Ref 4]. Finally, there is Orion, a mission under development at Stanford University [Ref 5]. The purpose of Orion is to identify and demonstrate technologies that will benefit these and future formation flying missions. Orion is currently scheduled to launch aboard the Space Shuttle, along with the Stanford *Emerald* mission [Ref 10], in November 2001.

This paper summarizes the Orion satellite design and its use as a testbed for formation control and fleet operations. Section II begins

with an overview of some of the issues associated with formation flying, and the research conducted at Stanford to address them. This indicates the general class of desirable experiments to be performed on Orion. Sections III and IV provide an overview of the Orion mission and briefly describe the design of the Orion spacecraft itself. This highlights the ability of the satellite to serve as a platform for testing hardware and theory. Section V then describes the actual operational experiments currently planned for the Orion mission, and the expected results.

II. Fleet Management and Mission Planning

There are some obvious benefits to a formation flying mission architecture. When compared to a single large satellite, a distributed satellite system will, in general, have a lower total mass and be easier to launch. In addition, a distributed platform inherently has a higher degree of operational flexibility, as it can assume a variety of “shapes”. But the key benefits of smaller satellites are that they cost less to reproduce and generally require fewer resources to construct. This makes a distributed, formation-type architecture most attractive from an economic standpoint. Moreover, less money and fewer resources will be required for extra satellites. This is attractive if coverage or reliability is a performance issue, and it also positively impacts the cost for replacing and upgrading the fleet.

An architecture that uses many small satellites has many advantages, but the fact that the constituent satellites are small is itself a weakness. The main reason is that the limited size of small satellites (such as the micro-satellite class) makes them inherently deficient in certain key resources. Structural and surface area limitations inhibit the ability to provide power through solar cells; volume limitations bound the amount of maneuvering fuel that can be carried; furthermore, the reduced mass properties make individual satellites more susceptible to environmental disturbances. All of these problems make the fleet control-and-coordination task that much more difficult. In addition, mission operations must be planned so that fleet performance requirements are

achieved, and valuable resources (such as fuel) are not squandered.

Independent of vehicle size, there are other design challenges inherent to a formation fleet architecture. The chief concern is that the relative positions of the fleet vehicles must not only be determinate, but such information must also be effectively communicated and distributed to the fleet. Depending on the chosen control architecture, this could have a significant impact on system performance.

Research conducted at Stanford that addresses these issues includes:

- CDGPS hardware for position sensing
- Performance of various control architectures
- Path planning and optimal trajectories
- Life-cycle mission planning

CDGPS Hardware

In order for a formation flying architecture to work, the system must have some knowledge of the relative positions of the different vehicles. Carrier-Phase Differential GPS (CDGPS) performs very effectively for this purpose [Ref 8]. CDGPS is capable of providing precise relative position (~1 cm) and velocity (~1 mm/s) knowledge. This is accomplished by performing differential calculations on the carrier phase signal between GPS receivers. The process works best if the receiver units run off the same clock. This is a difficulty encountered in formation flying, as the all vehicles must constantly and consistently synchronize their clocks through some time transfer protocol. For the Orion mission, a low-power GPS receiver has been developed to provide CDGPS measurements as a position and attitude sensor.

Control Architectures

The control strategy is a key design choice for a formation flying mission. Three main control architectures have been identified [Ref 6,7]:

- Centralized – vehicle movements based on relative states of whole formation
- Leader-referenced – vehicle movements based on position relative to a designated “leader”; the leader controls only its absolute position
- Absolute – vehicle movements based on its own absolute position (relative positions not known or controlled)

Each of these schemes can be used as part of a larger overall control strategy that translates high-level ground commands into a coherent, organized fleet deployment plan. Each of these control schemes was tried on a ground-based 2-D formation flying test bed at Stanford. Performance was compared for a variety of necessary fleet maneuvers, such as resizing, retargeting, and initialization. The Orion mission will be the first opportunity to try and compare each of these architectures on-orbit in a 3-D experiment.

Optimal Trajectories

The most important resource for a space-borne vehicle is fuel, as it is generally non-replenishable. As a result, on-orbit maneuvers must execute with a high degree of fuel efficiency. In conjunction with the work on control strategies, a linear programming technique was employed to determine fuel-optimal paths for a fleet of three vehicles [Ref 6]. The objective function was to minimize a weighted sum of fleet fuel, with constraints related to the system on-orbit dynamics, and operational restrictions. The Orion satellite will employ a similar algorithm as part of its overall formation maneuver control. Again it will be the first opportunity to verify results with an on-orbit test.

Optimal Mission Planning and Coordination

The overall control architecture can be used to plan and execute a single experiment in some optimal fashion. However, such techniques can be improved to take into account a more global and long-term view of the fleet mission. For instance:

- *Fleet combinations* – not all spacecraft may be required to perform an experiment, or spare satellites may be present within a fleet. In either case, a choice must be made about which satellites to use. Work initiated recently at Stanford aims to elaborate on this issue, starting with a way to choose the “weighting” associated with the optimal fuel consumption method mentioned above.
- *Performance quality trades* – minimize the amount of fuel expended to acquire a particular quantity of data, or conversely, commit to expending a particular amount of fuel and maximize the amount of data acquired. This may involve using a stochastic method of determining the probability that enough data has been

collected, or the locations in which the “best” data can be found.

Ultimately, the fleet can become more autonomous by acquiring a greater ability to make intelligent decisions without intervention from an operator. The Orion mission may employ such algorithms in order to demonstrate these concepts and thereby extend the total number and quality of experiments.

III. The Orion Mission

The main function of the Orion project is to identify and evaluate technologies required for formation flying missions. It is crucial to understand what types of components will be needed, and to what degree spacecraft resources are critical for this particular class of multi-satellite missions. In conjunction with this effort, some basic proof-of-concept formation flying experiments will need to be executed in order to build a knowledge base of experience from which future, more sophisticated missions can benefit. If successful, Orion will accomplish each of these tasks.

The Orion mission has two primary mission goals [Ref 9]:

- (1) Demonstrate the use and operation of a low-cost, low-power, multi-channel GPS receiver for real-time determination of the attitude and position of a small satellite.
- (2) Demonstrate the ability to organize a group of small vehicles into a pre-determined formation on orbit.

Orion will be launched aboard the Space Shuttle. It will be deployed side-by-side with the Emerald nanosatellites on the Multiple Satellite Deployment System (MSDS) platform, designed and constructed by engineers at the Air Force Research Laboratory. A diagram of the launch configuration appears in Figure 1. The entire MSDS-Emerald-Orion system will be ejected from the Shuttle; afterwards, there will be a time window for start-up operations before a timer on the MSDS causes Orion and the Emeralds to be released. At this point, nominal operations and experiments may begin. The target orbit is 325-350 km altitude and 28.5 degrees inclination; however, 400 km and 50 degrees are preferred parameters, since atmospheric drag severely limits mission lifetime at lower altitudes and

higher inclinations increase ground contact visibility times.

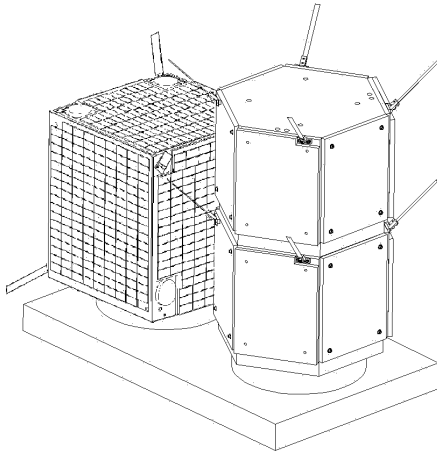


Figure 1: Orion-Emerald-MSDS launch configuration

In order to achieve the goals mentioned above, GPS data will be exchanged between satellites. Orion, and to some degree the Emerald satellites, will then use that data to execute pre-planned, organized maneuvers. The maneuvering process will be governed by real-time autonomous control software that will be directed at a high level from the ground. An operations plan will be designed so that mission resources are adequately conserved.

The minimum success criteria for the Orion mission are as follows:

- (1) The GPS receiver payload must be able to calculate absolute orbital position in real time to within 50 meters. Attitude must be calculated to within 2° .
- (2) The attitude of each spacecraft in the formation must be controlled to within 10° .
- (3) At least two satellites must be arranged on command in an in-track formation. The satellite spacing must be even over a range of 1 kilometer, and the formation must be held for at least 30 minutes. Relative position between satellites must be known to within 5 meters so that a 20-meter precision of control may be enforced. The process must be repeatable two times over the period of one week.

Orion will achieve the above criteria by maneuvering relative to one of the Emerald satellites. To demonstrate repeatability, the

process will be attempted 5 times over a period of two weeks. When complete, formation experiments involving both Emeralds and Orion will be performed.

IV. Orion Design

Now that some of the mission details have been described, the resulting design is discussed. It is important to first reflect upon the key design requirements and how they flow down to the subsystem level. Then, a brief description of each Orion subsystem is given.

Key Design Requirements

The mission requirements and goals demand that the final integrated system achieve several key functions. First, the GPS receiver must function properly at all times. Without precise knowledge of relative position and attitude, the relative position control will not meet specs, or may even be unstable. Second, the communications system must allow data to be exchanged between satellites during all phases. Without a data cross-link, position information cannot be distributed and again the position control loop cannot be closed properly. The third key requirement is that the flight control software must be able to control the satellites to the required degree of accuracy. While accuracy may not be as important for some missions, it is nonetheless required by the design team to meet minimum success for this project. Lastly, the design must incorporate adequate system resources to perform the experiments. The Orion design team has assembled a complete mission requirements document that outlines these key criteria and the subsequent subsystem requirements.

The primary requirements for the GPS receiver are similar to those for a ground-based GPS unit. To obtain an initial position lock, at least one antenna must view four GPS satellites simultaneously for a few minutes. At least 3 antennas must then track a minimum of 4 common GPS satellites to maintain position and attitude solutions. The accompanying electronics must be able to calculate the appropriate position and attitude solutions from the acquired signals, store that data, and share it with the control software. The GPS receiver also places requirements on other subsystems. Since it must be active all the time, it is a constant draw on the power subsystem. In order to

maintain a lock on the GPS signal, the spin rate and pointing accuracy of each satellite must be kept below a particular threshold. Finally, the size and accumulation rate of GPS data influences the required memory capacity and data bus speed for the command and data handling (CDH) subsystem.

Subsystems

Structure

A blow-out diagram of Orion is shown in Figure 2. The Orion satellite is a cube 17.5” per side. The mass target for the flight spacecraft is 35 kg. The primary structural material is aluminum honeycomb. The main load-bearing structure is a pair of square plates (top and bottom) connected by a set of plates that form an X along the diagonal. Each of these structural members is 1/2” honeycomb. The X-pattern thereby creates four chambers into which the remaining components and electronics are placed. Three of these chambers house nitrogen tanks for the propulsion system. The outer panels are 1/4” honeycomb plates. These will bear relatively little load and will be mainly used for mounting solar cells and torque coils. Plate connections are made via aluminum L-channels, brackets, screws, and honeycomb inserts. Some plate corners are cut away to allow for the routing of propulsion plumbing and bus wiring. A slot on the top plate allows access to the propulsion tanks’ fill/drain valve. An MSDS platform adapter occupies most of the bottom panel. One of the challenges in laying out the subsystem components is to minimize the moments of inertia (to increase rotational maneuverability), and place the fuel tanks such that the center of mass moves very little over time. A summary of the Orion mass distribution appears in Table 1.




Figure 2 - Orion structure

Table 1 - Orion mass budget

Subsystem	Mass (g)
GPS Payload	1072
Structure	12403.68
CDH	900
Comm	696
Torquer Coils	2473
Propulsion System	10648
Power System	7140
<hr/>	
Total	35332.68
Budget/Allocation	35000
Margin	-1.0%

Power

The Orion power consists of the following components:

- Spectrolab dual-junction GaAs solar cells
- One pack of ten Sanyo NiCd KR-10000M (10 A-hr) batteries
- Lambda PM30-12S05 power regulator
- Ten magnetic latching relays (operational power inhibits required for Shuttle safety)

The Orion power system delivers the following performance:

- 5V regulated power up to 6A
- 12-14V unregulated power up to battery current limit
- Expected solar input ~18W time-averaged

The power system design is fairly straightforward. Solar panels are sized so as to make battery over-charge almost impossible. Solar input feeds directly to the batteries and to the main power bus. The batteries act as a buffer. A total of four basic flight modes exist. These modes include the initial start-up mode, a “cruise” mode (essentially a stand-by mode), the experiment mode, and a data downlink mode. A summary of the power budget appears below in Table 2. The power system is capable of sustaining a (power-hungry) experiment for up to 3 orbits (~4.5 hours).

Table 2 - Orion power budget

Subsystem	Startup (mW)	Cruise (mW)	Contact (mW)	Exper (mW)
GPS Payload	1450	1450	1450	5225
Structure	0	0	0	0
CDH	448	448	448	6948
Comm	1426	1426	5890	2376
Torquer Coils	3650	250	250	250
Propulsion System	300	300	300	22380
Power System	0	0	0	0
<hr/>				
Total	7274	3874	8338	37179

CDH

The CDH and communication subsystems are unique, as these two subsystems are virtually identical between the Orion and Emerald satellites. The aspiration is to allow both teams to integrate efforts and to mutually develop the end product. Since the satellites must communicate mission-critical data in order to achieve success, this development strategy is key. The CDH subsystem consists of a SpaceQuest CPU and an I²C bus monitor. The SpaceQuest Rev C motherboard is a flight-ready CPU; it was purchased as an off-the-shelf solution so as to conserve manpower and time. The BekTek operating system, previously used on earlier revisions of the board has been purchased in order to expedite software development.

The spacecraft data bus is based on the I²C specification. As such, all subsystems on the data bus have a built-in I²C adapter. PIC16C74A microcontrollers are thereby used to do the low-level control for the attitude control systems, which offloads some burden from the main CPU. The microcontrollers are manufactured by Microchip, and have 4 KB ROM, digital I/O lines, PWM output and A/D channels in addition to RS-232 and I²C interfaces. Another advantage is the reduced amount of wiring required. Instead of having to run separate data lines to each subsystem, the subsystems can now be daisy-chained on the I²C data bus. In addition, all system telemetry is conveyed over a Dallas 1-wire bus. This eliminates the need for routing every single telemetry sensor wire back to the CPU. All sensors are wired locally within subsystems and telemetry values are addressed

and read over the Dallas bus. A bus monitor integrated with the CPU has full control over I²C-controlled systems and the flow of telemetry. The architecture for the CDH system appears in Figure 3.

Communications

Orion’s communication subsystem consists of two transmitters, two receivers, and a SpaceQuest modem. Each of the receivers was constructed from a Hamtronics amateur developers kit. Only slight modifications will be required to make these components space-worthy. The transmitters operate on the standard 70-cm amateur band as does one receiver. This provides adequate hardware to cover the cross-linking duties. The remaining backup receiver operates on the 2-m band and can be used to upload commands during contingency operations. All RF components are routed through the SpaceQuest modem; this is a flight-ready modem purchased as part of a CPU-communications package solution, and is specially designed both in hardware and software to integrate smoothly with the SpaceQuest CPU.

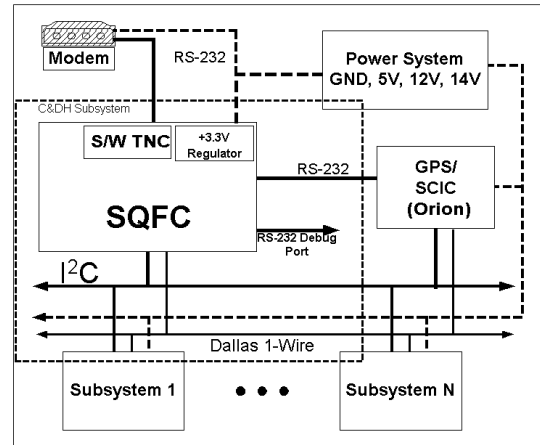


Figure 3 - Orion CDH architecture

Attitude Control

Orion uses a cold-gas propulsion system as its primary means of maneuvering and attitude control. Four 3-axis clusters of thrusters are located at four corners of the satellite. The flow through each thruster unit is regulated by a solenoid valve, which in turn is controlled by a PIC16C74A microprocessor. A diagram of the propulsion system is shown in Figure 4. Each thruster is capable of generating 50 mN of thrust

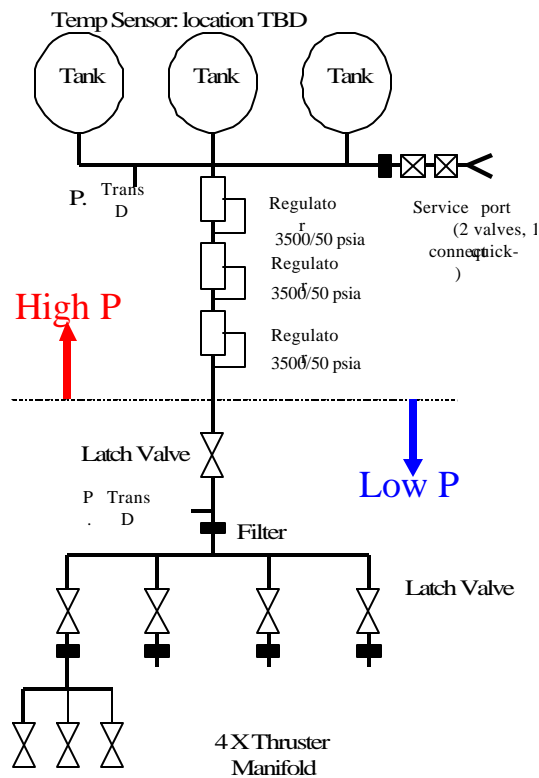


Figure 4 - Propulsion system schematic

at a specific impulse of ~ 70 sec. A fair amount of the system design and redundancy was driven by NASA Shuttle safety requirements.

Without GPS data, however, the propulsion system controller would essentially be running open-loop. Thus, in the absence of a reliable GPS signal (such as during start-up or contingency operations), an auxiliary control method is needed. Torquer coils are used for this purpose. Three coils consisting of 300 turns of magnet-wire on an aluminum frames are mounted on the inside surface of three side-panels. This system is capable of generating a magnetic moment of 5 Am^2 , which is equivalent to $1.25 \cdot 10^{-4} \text{ Nm}$ at 500 km altitude. The amount of current through the coils is controlled by a PIC16C74A microcontroller through power MOSFETs. This subsystem also includes the necessary 3-axis magnetometer as a sensor for the feedback loop. The coil system slowly de-tumbles the spacecraft until a GPS signal lock is obtained. Initial simulation results show that Orion can be de-tumbled in less than 3 orbits for appreciably large initial Euler angle rates

(several degrees/sec). A picture of Orion's prototype coils appears in Figure 5.

GPS Receiver

The Orion GPS receiver consists of a single 6-antenna attitude and relative navigation receiver using carrier-differential GPS. The GPS receiver design is based on the Mitel Plessey GPS chipset, using the GP2015 RF front end and the GP2021 12-channel correlator. A diagram of the GPS receiver system is shown in Figure 6.




Figure 5 - Prototype torquer coils

Modifications to the original receiver design include a second RF front end and an external clock input. This receiver has two RF front ends, a correlator (six channels per RF front end), an ARM60 processor, and the required EPROM and RAM memory. Another board provides the 5V regulated power input and RS232 serial input and output. The attitude receiver uses two of these modified cards with a common clock. Integrated Carrier Phase data is shared between the two cards over the serial ports. The ARM60 closes the low level code and carrier tracking loops on both cards. Furthermore, on one ARM60 the absolute position solution is determined while the other determines the attitude. This process is currently run at 5 Hz. The relative navigation uses a single receiver card, with the Integrated Carrier Phase data being sent from a second receiver though the serial port. The processor computes both the absolute and relative position solutions. Because of the greater computational load, this process is performed at only 1 Hz. Current tests show relative position accuracy on the order of 2 cm.

Several hardware and software changes have been completed or are in progress. Hardware upgrades tie together the single six-antenna receiver with a science computer capable of performing all solutions (attitude, absolute position, and relative navigation). Software upgrades include:

- Improved bias initialization algorithm – CDGPS techniques require knowledge of the number of integer cycles between antennas. This bias must be initialized and calculated in the software
- Orbit estimator – the relative velocities between the user and the GPS spacecraft is much larger than in terrestrial applications. A much larger Doppler space must be searched for the GPS signal. Without estimating the expected relative velocity, a signal lock may never be acquired.
- Non-aligned antenna compensation – phase differences between non-aligned antenna bore-sights will certainly occur between multiple spacecraft
- Low power mode – a single antenna will be employed to maintain a signal lock

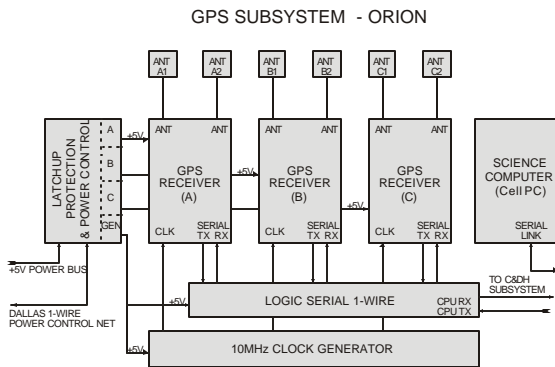


Figure 6 - GPS receiver design

V. Orion as a Formation Flying Testbed

As outlined in Section II, a strong effort has been made to develop the applicative theory for coordinating and controlling a fleet of cooperative vehicles in space. Orion will aid in these efforts by providing a suitable space platform from which to test these principles.

For CDGPS hardware, the Orion mission will provide a means by which to evaluate the performance of the receiver system in space,

including the chip set, the ARM60 processors, and the RF front-end equipment. A successful flight will grant these items much-valued flight heritage, while also achieving one of the mission's primary goals. It is expected that the hardware modifications will perform as required for the intended mission life.

Orion and the Emerald satellites together will provide the necessary environment in which to try the proposed formation control architectures. The first phase of formation flying experiments will see Orion employing more of a leader-following control strategy; one Emerald will broadcast its own GPS solution over the cross-link, which Orion will use to calculate relative positioning and act accordingly. System performance in this respect is expected to be good. The next phase will involve Orion and an Emerald actually having more of a conversation, such as in the distributed control architecture. The Emerald satellites each have a very limited ability to control position with a set of drag panels. Orion will be burdened with performing most of the calculations, however both Orion and Emerald will nevertheless be able to exchange relative positioning information and take action accordingly. The final phases will attempt both leader and distributed control experiments with both Emerald satellites instead of just one. The results from these experiments are expected to show benefits for distributed control (as it did at the Stanford formation flying testbed). However, since the Emerald's actuation performance is so weak and limited, some of the data might end up appearing subjective.

In conjunction with the control architecture experiments, optimal path planning software will be included on Orion's number-crunching science computer. During the cycle of an experiment, relative positioning data, thruster on-off times, and propulsion tank pressure will be recorded. During analysis, thruster firings can be matched against drops in tank pressure to characterize how efficient the propulsion system behaves. With that done, the noted fuel consumption can be collated with relative position measurements in order to realize the true path taken by the Orion spacecraft, and the corresponding quantity of fuel spent. These results can then be compared with ground simulations to determine how well the path-planning algorithm performs. These results are expected to be nominal; however, uncertainties

in the plant model (Orion) may cause the path planning loop to be less efficient than expected.

A major concern addressed during the design phase was the balance between the quantity and quality of experiments. More orbits of data necessarily mean fewer opportunities to run experiments. In addition, extremely long experiments could lead to power shortages. These problems highlight the need for attempting to invoke mission planning techniques that take into account more than a myopic view of fuel usage. The objective will be to incorporate mission planning algorithms to gather good data and support the mission goals while extending the total number of missions (mission life) as much as possible.

VI. Conclusions

This paper has summarized the relationship between the formation flying research being conducted at Stanford University and its relationship to the Orion mission. The main thrust areas for research included developing CDGPS hardware, investigating various architectures for fleet control, optimized vehicle trajectories, and optimized long-term mission planning. The Orion mission provides an appealing formation flying test platform, and is the primary candidate for which to develop and demonstrate these concepts. As a result, a great deal of knowledge and experience will be gained that will benefit future formation flying missions. It is expected that valuable insights will be gained into the performance of CDGPS hardware on-orbit, as well as information about the practical application of various control architectures and optimal planning methods in such an environment.

VII. References

1. Space Technology 3 homepage, <http://jpl.nasa.gov/st3/>
2. C. A. Beichman, "The Terrestrial Planet Finder: The Search for Life-Bearing Planets Around Other Stars", in Proc of the SPIE Conference on Astronomical Interferometry, March 1998.
3. AFRL SV Directorate, Techsat21 homepage,

- <http://www.vs.af.mil/factsheets/TechSat21.html>
4. LISA homepage, <http://lisa.jpl.nasa.gov>
5. J. How, R. Twiggs, D. Weidow, K. Hartman, and F. Bauer, "Orion: A Low-Cost Demonstration of Formation Flying in Space Using GPS," in AIAA Astrodynamics Specialists Conference, August 1998.
6. A. Robertson, G. Inalhan, and J. P. How, "Formation Control Strategies for a Separated Spacecraft Interferometer," in Proceedings of 1999 ACC, June 1999.
7. A. Robertson, G. Inalhan, and J. P. How, "Spacecraft Formation Flying Control Design for the Orion Mission," AIAA-99-4266, 1999.
8. Olsen, E. "CDGPS Sensing for Formation Flying Vehicles," PhD Thesis, Dept of Aeronautics & Astronautics, Stanford University. Dec 1999.
9. Orion design team. "Orion Mission Requirements Document" Unpublished internal document.
10. Emerald homepage, <http://ssdl.stanford.edu/emerald>