The Coral Reef Satellite Mission

Robert F. Goeke, Member SPIE
MIT Center for Space Research
Room NE80-6099, Cambridge, MA 02139
Email: goeke@space.mit.edu

ABSTRACT

The Planetary Coral Reef Foundation is proposing to privately fund a Coral Reef Satellite Mission to perform a global mapping survey of shallow water coral reefs and monitor their health over a period of at least 5 years. Using available technology, but restricting the goals of the Mission to the limited objectives of finding and monitoring coral reefs, it is possible to design a coherent and cost effective flight system. We describe the general design of the mission and its components, giving particular emphasis to the science payload.

KEYWORDS

Coral reefs, remote sensing, oceanography, environmental sciences, satellites, PCRF

1. INTRODUCTION

That the world's oceans in general, and its coral reefs in particular, are in a state of cataclysmic decline is neither news nor an event lacking in documentation[1]. There is, as well, much of the "government will have to act" literature in circulation[2], though to date governments have seen it far easier to decry the current – not to mention potential – state of affairs than to act upon them. The Planetary Coral Reef Foundation (PCRF) <www.pcrf.org> has taken an unusually active approach to the problem by gathering together a diverse team of individuals and organizations to define and execute a privately funded Coral Reef Satellite Mission (CRSM). The Mission will establish a high quality set of baseline data for the global shallow water coral reef population and track temporal changes of these reefs on monthly time scales. Mission-funded education, conservation, and public outreach programs will use this publicly-accessible data set to drive the world's governments into action.

This is not a low cost mission. One could, for $10M or so, fly an instrument which would take color pictures of reefs as opportunity arose -- launched as a co-manifest to some larger mission -- publish a few news articles and declare victory. PCRF wanted to do more, much more. They wanted to do science, building on their years of *in situ* coral reef observations with the *RV Heraclitus*, and to make publicly available the basic remote sensing data both to researchers and to those who earn their living from the oceans. NASA or NOAA could do this, of course, and so could ESA. They have chosen not to, and that is less a criticism than a statement of fact; no organization has the resources to do every researcher's bidding. So PCRF has decided to launch a private fund-raising effort, but not for a minimal pathfinder mission. PCRF wants a high-probability-of-success, high-quality-science, long-operational-life mission. This leads to costs which, compared to a LANDSAT model, might be viewed as inexpensive, but are, ne'er the less, not for the faint of heart. Our estimate for the cost of the flight hardware plus launch is $76M (FY03). Five years of operations, data analysis, education, conservation, and public outreach will not quite double that cost, but it is money that must be paid if the mission is to fulfill its purpose.

PCRF has designated the Massachusetts Institute of Technology Center for Space Research (CSR) as the manager for the flight hardware development, launch, and mission operations. This paper describes our current state of flight hardware design.

2. OVERVIEW

2.1 Science Requirements

While past and presently orbiting satellites can and are being used to generate crude maps of coral reefs, none of these simultaneously possess both the spectral and spatial resolutions required to monitor reef health at the scale of the reef communities. This task requires a dedicated reef satellite with a unique suite of optical sensors to provide the information necessary for a radiometric solution to estimating coral reef community composition. Such a science payload must have a pixel footprint of 5 to 10 meters and a set of carefully chosen narrow bandwidths to resolve functional algal groups into their major taxonomic groupings based on photosynthetic pigments. In Table 1 we show the spectral coverage of two contemporary satellites which, unfortunately for coral reef studies, are limited to 1 km spatial resolution; the CRSM requirements are also shown.

Wavelength (nm)	SeaWiFS	MODIS	CRSM	Algal Type	Photo synthetic Pigment	Animal Fluorescence	Interest
410	412	412	**X**				
443	443	443	**X**	All	Chlorophyll a		Absorption max
450							
480-490	490	488	**X**	Brown	Fucoxanthin, Peridinin	Blue-green	Absorption max Type 1 peak
500						Green	Type 2 peak
510	510		**X**			Blue-green	Type 1 shoulder
515						Green	Type 3 peak
530		531	**X**				
540				Red, Blue-green	Phycobilins		
550	550	551	**X**				CZCS hinge
570			**X**	Red	Phycobilins		Absorption max
600							
630				Brown	Chlorophyll c		Absorption max
675	670	667	**X**	All	Chlorophyll a		Absorption max
683		678	**X**	All	Chlorophyll a		Natural Fluorescence
690							
700			**X**	All	Chlorophyll minimum		Chlorophyll reflectance
765	765	765	**X**				Atmosphere
865	865	865	**X**				Atmosphere

Table 1: Wavelengths of Interest for Remote Sensing of Reef Organisms.

Complete hyperspectral imagery is not required to monitor reef health, but we do need to measure the relative intensities of the wavelengths of interest to a percent or two. This, in turn, requires additional measurement channels (included in this list) to take out the effects of both atmospheric and water column.

2.2 Observation Strategy

The mission is to be flown as a sun-synchronous 97.8 degree inclination orbit with a 10:30 am local time ground track at a nominal altitude of 600 km. Depending upon cloud cover, sea state, *etc*., a target list is uploaded to the spacecraft once per day for reefs that will be located along the ground track. By trimming the altitude after launch, the ground track will drift 450 km per day so that the entire surface of the earth will be visible to the spacecraft once in every 6 day interval.

It is not necessary for a reef to be directly below the track of the satellite (*i.e.*, at nadir). With the science payload fixed to the nadir face of the spacecraft bus, the bus can roll up to 30 degrees cross-track, so that any reef within plus or minus 300 km of the ground track can be observed.

When no known reef targets are available for a particular ocean region – and cloud cover and sea state allow – a global survey will be conducted with a single camera. After a period of time, the entire ocean system will be surveyed and the complete set of (shallow) reefs identified and then added to the regular target list.

2.3 Science Payload

The science payload is composed of twelve individual, co-aligned cameras; each tuned to a particular wavelength. Eight of these cameras cover wavelengths diagnostic to the health of the coral reefs. The other four cameras provide data that allows the removal of both atmospheric and water attenuation artifacts in the primary data. Additionally a single, wide spectral band camera, co-aligned with the primary twelve, provides 2x higher spatial resolution data over the central field. Finally, a single camera provides detailed spectral information on a narrow central strip co-aligned with the primary cameras. These last two cameras are primarily used for diagnostic purposes during the data analysis process.

2.4 Spacecraft Bus

The spacecraft bus is provided by EADS Astrium, Friedrichshafen, based on an established design optimized for earth resource missions. As is common for such an application, this spacecraft generates a fair amount of data (though orders of magnitude less than a regular mapping mission – the benefit of focusing only on reefs). Aside from normal attitude control, power, command and data handling, and telemetry functions, it is necessary to carry a substantial amount of fuel for periodic orbit adjustment. It is beneficial to keep the altitude around 600 km to reduce radiation effects in the science payload. Unfortunately, aerodynamic drag becomes an issue, and below about 525 km it would be intolerable.

We specify to the spacecraft computer (via commands from the Science Operations Center) the longitude and latitude of the center of the pictures to be taken during the next few orbits. The on-board computer must calculate both the times to take the pictures and the appropriate satellite orientations (cross-track offsets from nadir). This requires that the spacecraft calculate its ephemeris on-board (from Global Positioning System data) and be able to determine the satellite orientation to about 10 arc seconds during the daylight portion of the orbit. While the second requirement has been done several times in the past, it is not trivial and close attention must be paid to star tracker positioning and orientation, and the design of the sunshades. On the other hand, satellite-borne GPS systems have been available for only the last few years. Therefore, this mission could not have been implemented in the same way as little as five years ago.

2.5 Launch

The spacecraft is designed to be compatible with both the Russian Rockot and Dnepr-1 launch vehicles. The former has a lift capacity twice our needs, the later has a capacity three times our needs. Costs are comparable.

3. MISSION DESIGN

3.1 Requirements

There are three mildly competing science and one programmatic requirement which affects the mission design:

- Perform a global survey which will explore, discover, and map the position of **all** shallow water coral reefs (the viable ones all lie between latitudes of ±30 degrees);
- Establish a baseline spectral data set for all known reefs; and
- Closely (in a temporal sense) monitor the health of selected coral reefs and their environs.
- Design life is to be five years with no consumables preventing an extension to 10 years of operation.

The good news, from the flight mission design point of view, is that these are the only requirements. This is a dedicated mission, not subject to trades and costs associated with competing science goals on a single platform.

3.2 Orbit

Remote sensing almost always calls for a sun-synchronous orbit, and this mission is no exception. Maximum illumination would be good, but after 11am local time the clouds start to build, and this is no radar mission; we choose to start at 10:30am local time and trim that crossing time if necessary. It isn't obvious how the trade off between better coverage at a poorer sun angle versus the opposite case will work itself out, and the answer may change over the life of the mission. An even more interesting trade involves the repeat pattern desired. Our footprint is about 600 km wide. One could choose to revisit the same reef every day for a week, but then not see that particular patch of ocean for a month. Alternatively, one could drift 450 km per day and thus have a constant revisit time of 6 days. We have chosen the later as our baseline design.

Then there is the question of altitude. From the science instruments' point of view, lower is better. This is mostly a matter of background radiation. As described below we are planning to use Charge Coupled Device (CCD) sensors in the instrument cameras, and there are two radiation effects to consider: total dose and particle background. The total dose for a 5 year mission as a function of altitude is shown in Table 2 (these results use SPENVIS, the ESA-modified AE8 model, 2.5 mm Al sphere). The SEU rate calculation (SPENVIS again) is a proxy for the rate at which the columns in the CCDs will be corrupted by trapped proton events; the absolute numbers are not significant here. Higher would also demand more glass on the front end of the cameras.

Altitude (km)	Total Dose (Krads)	SEU Rate	Aperture
600	12	4.2	1
700	16	7.4	1.4
800	20	13	1.8
900	26	19	2.3

Table 2: Radiation Effects

From the spacecraft bus' point of view, higher is better. Below 700 km, and most certainly below 600 km, orbit maintenance will require a fair amount of propellant and ground intervention.

4. SCIENCE PAYLOAD DESIGN

4.1 Science Payload Design Approach

The Science Payload will be designed and fabricated by CSR. It is composed of 14 separate instruments which, by design, have a high degree of commonality. Each instrument consists of an optical camera, cooled CCD focal plane, and supporting electronics. Within these 14 instruments, there are three different types whose names reflect their scientific function *(see Table 3)*. Each camera operates independently from the other, which results in a high degree of failure tolerance for the mission.

Quantity	Descriptive Name
12	**Picture** – a monochromatic n x 1024 pixel image; 10 m resolution
1	**Pan** – a panchromatic n x 1024 pixel image; 5 m resolution
1	**Palette** – a 1024 element spectral dispersion of a central pixel

Table 3: Science Instrument Nomenclature

4.2 The Picture Instruments

The 12 instruments which comprise the Picture suite, are distinguished by their center wavelength, which is determined by an optical bandpass filter *(see Table 4)*. Other than the filter, they are identical in construction and function.

Nomenclature	Center Wavelength	Full Bandwidth
Picture/410	410 nm	10 nm
Picture/443	443 nm	10 nm
Picture/485	485 nm	10 nm
Picture/510	510 nm	10 nm
Picture/530	530 nm	10 nm
Picture/550	550 nm	10 nm
Picture/570	570 nm	10 nm
Picture/675	675 nm	10 nm
Picture/683	683 nm	10 nm
Picture/700	700 nm	10 nm
Picture/765	765 nm	20 nm
Picture/865	865 nm	20 nm

Table 4: Optical Filters for 12 Picture Instruments

4.3 Optical Design

The camera optics are based on a catadioptric design with a 680 mm focal length and 85 mm aperture (*i.e.*: f/8). (This entrance aperture really is the minimum possible for our resolution requirements; we would run into diffraction problems if it were smaller, even if we might not need quite as much signal as this is going to give us). A multilayer interference filter at the aperture defines the center wavelength and bandwidth of the particular camera.

At the focal plane is a CCD such as the Fairchild Imaging CCD525. This is a standard commercial device designed specifically for use in Time Delay Integration (TDI) cameras and is fabricated as an array of 2048 x 96 pixels. The CCD has 13 x 13 micron pixels which the above optics map to a ground resolution of 10 meters. The image is clocked every 1.3 milliseconds from one row to the next – as the spacecraft advances 10 meters over its ground track – until 96 integrated samples of the scene have been collected. At the end of this integration process the entire row of 2048 pixels are read out. (Only the central 1024 pixels are used in our application to image the desired 10 km wide swath.)

In this application, we have an input signal of 10^7 photons/s/nm/pixel at 100% albedo. The albedo will not, in general, be that bright, and there are optical inefficiencies, but with a 10 nm bandwidth, we will collect a charge of at least 10^6 electrons/s/pixel in the CCD. In 1.3 milliseconds that is getting to be a rather small number, but TDI lets us integrate 96 times, so that our nominal signal is about 125 thousand electrons. To achieve 0.1% optical measurements reliably, we set our noise floor requirement to a part in 4000, or 30 electrons RMS. There are two components to this noise: the signal electronics and the detector itself. At CSR we routinely fabricate CCD electronics for astrophysics missions with noise levels on the order of two electrons RMS, so that is not an issue.

The commercial CCD detector, on the other hand, has significantly higher noise; Fairchild specifies this particular device at 70 electrons RMS, but that is at room temperature and operating 20 times faster than we need. By cooling the CCD to less than –20C and operating at a more modest shift rate, we will easily achieve the required noise levels.

4.4 Electronic Design

The instrument electronics must clock the CCD (shift the individual pixels as the scene passes beneath the spacecraft) and process the resulting stream of video pixel data. The initial analog interface with the CCD is a pair of Correlated Double Sampling amplifiers that remove a variety of systematic noise sources. (Although the Fairchild CCD has four outputs, we only use two). The resulting waveforms are digitized by a 12 bit Analog-to-Digital (A/D) converter and presented to digital logic contained in a Field Programmable Gate Array (FPGA). The FPGA formats the stream of pixel information into standard Consultative Committee for Space Data Systems (CCSDS) Type 1 telemetry packets and forwards these packets to the spacecraft bus for storage and subsequent transmission to the ground. Separately, telemetry packets containing engineering data are generated at a lower rate to monitor the health and calibration of the instrument.

Since the CCD is being advanced one 1024 pixel row every 1.3 milliseconds, a single instrument generates just under 10Mbits of primary science data per second. There is nothing inherent in the instrument design which constrains the length of the image (the width is fixed at 1024 pixels x 10 meters = 10 km); the instrument simply takes data until told to stop. If a reef is known to be small, one second of data will yield a 10 km square image. If a reef is long – and lies along the path of the satellite ground track – one could in principle record a 10 km x 1000 km image. The ultimate limitations are only the size of the spacecraft bus data storage and the time available for downlinking the data.

The instrument FPGA is also responsible for handling the command interface with the spacecraft. It does this by listening on a party-line MIL-STD-1553B data bus for commands addressed to it. Because the instruments are quite simple, the command dictionary is equally simple. The current design recognizes only 12 distinct commands for setting CCD clocking voltages, trimming TDI time intervals and row counts, and gating the CCD shift and telemetry generation process.

4.5 The Pan Instrument

The Pan instrument has an optic with twice the aperture and focal length of the Picture instruments, and a filter which simply cuts out the near infra-red wavelengths to which the CCD is sensitive but would simply add noise to our image. The result of these optical parameters is a panchromatic camera with a characteristic pixel size of 5 by 5 meters. This over-sampling of the spatial dimensions of the Picture instruments provides diagnostic information for the ground data analysis process.

Again, we are only processing the central 1024 pixels from the CCD, so our swath width is only 5 km wide (centered on the 10 km swatch of the Picture instruments). This CCD must be clocked at twice the rate of the Picture CCDs, thus generating twice the data rate. If we had a narrowband optical filter in our optics, the signal level on the CCD would be the same for this instrument as it was in the Picture case – twice the aperture size for half the pixel size. However, here we are looking wide band, say 300 nm versus the previous 10 nm, giving us 30 times the signal. In fact, we have to reduce this signal somewhat, probably by integrating for a shorter interval on the CCD, to keep from saturating the detector.

4.6 The Palette Instrument

The Palette instrument has the optic of the Picture instruments, but only looks – spatially – at the central 10 meter pixel being simultaneously observed by the other instruments. The light energy in this pixel is spectrally dispersed by a grating so that the entire visible wavelength band covered in discrete steps by the Picture instruments is continuously imaged on the 1024 pixel CCD. The complete spectral signature of the center pixel yields diagnostic information for the ground data analysis of the Picture data.

We clock the CCD at the same rate as the Picture instrument, and the photon flux is the same. Nominally, each pixel has the energy from a 10 meter ground image contained in a 0.33 nm spectral band (spreading the 380 to 720nm visible band over 1024 pixels). That signal would be rather faint, and is too narrow a spectral band to be useful. One could simply send all the data to the ground and let the analysis process add 6 pixels together to get a value for each 2 nm. Unfortunately, this means adding each channel's noise, too, so that signal-to-noise only improves by the square root of 6 in this process. Fortunately, one can add the signal noiselessly on the CCD (properly implemented, the noise is all in the CCD readout stage). The strategy, therefore, is to sum 6 pixels together within the CCD and read out the spectral information for every 2 nm band over the visible range. The fact that our signal-to-noise is a factor of four worse than the standard Picture instrument isn't significant for this application and the small degradation is worth the return of over-sampled spectral data.

4.7 Science Payload Accommodations

The accommodation requirements for the science payload – all 14 instruments – can be summarized in Table 5:

Item	Requirement
Mass	75 kg
Power	75 Watts orbital average 250 Watts peak
Thermal	Provide radiator for CCDs
Data	150 Mbps acquisition rate 1 GByte/day downlink

Table 5: Resources required for Science Payload

Although the data rate from the instruments to the spacecraft is high when an observation is being made, many observations will be only 10 km in length (only 20MB total), and most will be less than 100 km in length (200MB total).

5. SPACECRAFT BUS DESIGN

5.1 Spacecraft Design Approach

The fully-redundant spacecraft bus is based on the *Flex*Bus spacecraft series successfully applied for the Champ (DLR, Germany) and Grace (JPL/NASA) missions, and provided by EADS Astrium, Friedrichshafen, Germany. The *Flex*Bus standard architecture consists of a backbone made up by a set of fixed core elements that is used for all missions, *e.g.:* S-Band RF equipment, on-board computer, power control & distribution unit, heaters and thermistors, and system software. Spacecraft structures, of course, are almost always a custom design.

5.2 Configuration & Structure

The configuration philosophy for the CRSM centers on the optimum accommodation of the instrument telescopes onto a thermally stable common optical bench, pointing nadir during in-orbit operation (*see Figure 1*). This 'Nadir Platform' is made of carbon fiber composite and guarantees low thermal deflection and good co-alignment between the different payload channels; it is run cold biased. The camera electronics are located directly below the optical bench, in the "warm" part of the bus. Most of the core bus components are mounted on the ±Y load carrying panels of the box-shaped structure. The camera optical axes are aligned with the Z axis; the spacecraft is flown with the velocity vector nominally in the X direction.

We are helped in small ways by the constrained nature of the mission. Although the solar arrays -- body mounted on the ±X sides of the spacecraft – are not at optimum collecting angles to the sun while we are collecting data (typically at low latitudes where the arrays get almost no input), we are free to rotate the spacecraft to put one of the X faces normal to the sun at higher latitudes. Thus we avoid deployable, rotating solar drives. Similarly, most ground stations are where the reefs aren't, so we are free to point the body-fixed X-Band antenna by steering the spacecraft during telemetry passes.

5.3 Pointing Requirements

Taking pictures of reefs introduces some fairly stringent requirements on the spacecraft bus attitude and control system. (Unlike most land-oriented remote sensing imagers, we have no serendipitous features to help locate the reef.) The science payload has a field of view of 10 km, so we require that the image data be taken with a pointing accuracy of ±1 km. After the fact we need to know the pointing to ±100 meters (10 pixels). A feel for the later number may be had by noting that the spacecraft will move 100 meters over the ocean surface in 13 milliseconds, and an angular error of 100 meters at an altitude of 600 km is about 30 arc-seconds. These are not the extreme requirements of an astronomical telescope, but a fair amount of care is required to achieve a trustworthy result.

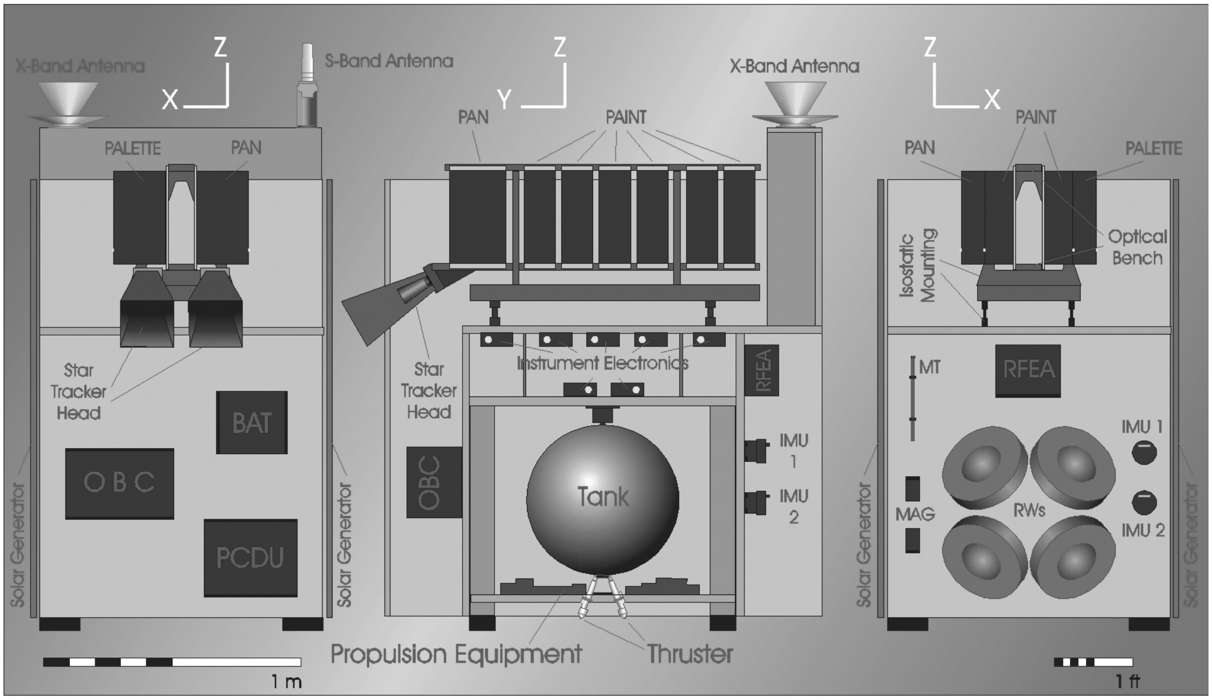

Figure 1 : Satellite Configuration & Structure Concept

5.4 Bus Margins

The total satellite mass adds up to 492 kg (Bus: 339 kg, Payload: 75 kg, Propellant: 78 kg). With a launcher capability of 1100 kg and 1610 kg for the two reference systems (Rockot and Dnepr), the mass margin yields 123% and 227%, respectively.

The required orbit average energy of 345 Wh is provided by a power system with an available energy of 475 Wh, thus providing a margin of approximately 38 % (assuming **no** reorientation of the bus at high latitudes).

The total estimated delta-V for the orbit injection, station keeping (at 550 km to give us margin) and de-orbiting calculates to 159 m/s. With an initial total spacecraft mass of 500 kg and a typical specific impulse of 210 s for the monopropellant propulsion system, the required propellant mass is about 37.2 kg. Adding 0.8 kg residuals, the final propellant required is 38 kg. The selected tank holds up to 78 kg thus providing the necessary margin for a mission extension to 10 years.

6. GROUND OPERATIONS AND DATA ANALYSIS

Mission Operations for this satellite is straightforward and needs no particular elaboration. A single 6 minute pass at 30 Mbps would suffice to return our nominal science data for a single day’s operation. Science operations, on the other hand, are rather more complex.

On the planning side one wants to maximize observation efficiency. This is important, since it has been estimated that just “blindly” scheduling observations would result in as few as 20% returning prime science. We need to monitor both cloud cover and sea state, together with calculations of the more predictable sun and instrument view angles, to piece together a daily observation load. Target planning has to operate on a daily time scale.

Quick look data processing can operate on a time scale of several days. If the repeat track for the orbit is 6 days, we simply need to know – before the next visit – if the previous data was good and/or the reefs health is stable. Making these decisions does not require the full data set, nor does it require unusual data processing resources. The full processing of the data can be accomplished on a month time scale. In all cases the data will be publicly accessible as soon as the data products are ready.

ACKNOWLEDGEMENTS

Much of the text in this paper has been excerpted from the CRSM Mission Description Document, a volume with many authors, which lays out all aspects of the Mission. The principal organizations associated with this Mission are

- College of Charleston
  Dr. Phil Dustan, Science PI
- EADS Astrium
  Spacecraft Bus
- General Dynamics, Inc
  Mission Operations
- Massachusetts Institute of Technology
  Flight System Manager
- Planetary Coral Reef Foundation
  Sponsor; Education & Field Research
- Scripps Institution of Oceanography
  Science Operations Center
- SeaSpace Corporation
  Data Analysis
- Steven Institute of Technology
  Data Analysis
- Sound Seas
  Conservation & Public Policy
- University of Southern California
  Science Applications

REFERENCES

1. D. R. Bellwood, R. P. Hughes, C. Folke, and M. Nyström, “Review”, *Nature* **429**, pp. 827-833, (2004).
2. T. Agardy, “America’s Coral Reefs: Awash with Problems”, *Issues in Science and Technology* **20,2**, pp. 35-42 (2004).