Fermi Finds More Than 100 Gamma-Ray Pulsars

P.S. Ray,[1] J.E. Grove,[1] K.S. Wood,[1] A.A. Abdo,[2]
T.J. Johnson,[3] and D. Parent[2]
[1]*Space Science Division*
[2]*George Mason University*
[3]*National Research Council Postdoctoral Research Associate*

Pulsars: Neutron stars are the supremely compressed remnants of massive stars, left over after supernova explosions. They contain a mass of about 1.5 times that of our Sun in a sphere of diameter approximately equal to the Washington, D.C. Beltway. The conditions of extreme density, gravitational field, and magnetic field cannot be replicated on Earth, so they are superb physics laboratories. They spin with periods ranging from about 10 seconds down to about 1.4 milliseconds per rotation, which is faster than a kitchen blender. As they spin, they accelerate particles to high energies and emit beams of radiation, most prominently radio waves, and appear to pulse as the beam sweeps across the Earth, much like a lighthouse. Thus the name: *pulsar*.

Gamma Rays from Pulsars: The last major gamma-ray telescope, the EGRET instrument on NASA's Compton Gamma-Ray Observatory (which operated from 1991 to 2000) detected gamma-ray pulsations from a total of six young pulsars. Evidently, for the most energetic pulsars, particles accelerated in the magnetospheres could reach high enough energies to emit gamma rays, the most energetic form of light. However, with only these early results, it was unclear exactly where the emission came from or how many pulsars were gamma-ray emitters.

The Fermi Gamma-Ray Space Telescope: In June 2008, Fermi was launched and began surveying the sky in the gamma-ray band. The prime instrument aboard, the Large Area Telescope (LAT), is a gamma-ray telescope vastly more capable than its predecessors. NRL designed and built the LAT calorimeter, which measures the energies of the gamma rays, and both the instrument-level and observatory environmental testing were done at NRL (see Fig. 1). In three years of surveying the gamma-ray sky, the LAT has discovered gamma-ray pulsations from over 116 pulsars using a variety of powerful strategies,[1] as we describe below.

Known Radio Pulsars: For pulsars already known in the radio band, the best way to find gamma-ray pulsations is to make use of precise timing observations from ground-based radio telescopes. By collaborating with radio astronomers around the world, LAT scientists were able to detect gamma-ray pulsations from 35 additional young radio pulsars, beyond the six known from EGRET.

FIGURE 1
Photo of the Fermi spacecraft undergoing testing at NRL. The Large Area Telescope is the silver rectangle at the top. The other instruments and spacecraft subsystems are in the lower part.

Blind Frequency Searches: One important discovery using the LAT is that many gamma-ray pulsars are radio quiet, almost certainly because their radio beams are much narrower than the beams of emitted gamma rays and thus miss our line of sight. These pulsars can only be found by "blindly" searching for their pulsations directly in the gamma-ray data itself. But the gamma-ray data are exceptionally sparse — the LAT might detect only a handful of photons per *day* from a typical pulsar. Doing these searches required the development of specialized search algorithms and the application of high-performance computing. But the searches were successful, yielding 35 new pulsars discovered purely via their gamma-ray pulsations (see Fig. 2). While four of these new pulsars have since been detected in the radio band, the rest appear to be truly radio quiet.[2]

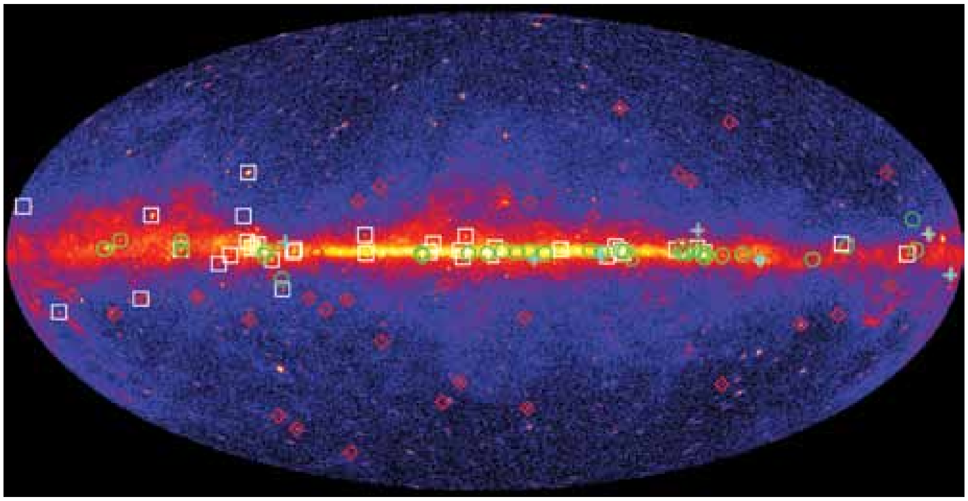

FIGURE 2
Map of the sky in gamma rays with energies above 100 MeV, made with three years of Fermi LAT sky survey data. The map is in galactic coordinates so the center of our galaxy is in the middle and the diffuse emission along the galactic plane is easily seen. The symbols mark the 116 pulsars whose gamma-ray pulsations have been detected with the LAT. The cyan crosses are the previously known gamma-ray pulsars. The green circles and white squares are young pulsars found based on radio timing, or blind searches, respectively. The red diamonds are millisecond pulsars.

Millisecond Pulsars: The fastest spinning pulsars are those that have been spun up to millisecond periods by the matter falling on them from a companion star in a binary system. These pulsars are nature's most stable clocks, with pulsation periods that change by only about 1 microsecond in a million years. These millisecond pulsars were not known to be gamma-ray emitters before Fermi. Shortly after its launch, the LAT detected pulsations from several known millisecond pulsars; that number has now grown to 40. Perhaps more exciting, a group of radio astronomers, led by NRL, has been searching for new millisecond pulsars by pointing their radio telescopes at unidentified gamma-ray sources discovered during the LAT sky survey. These searches have been spectacularly successful,[3] discovering 38 new millisecond pulsars so far. This represents an increase of more than 50% in the number of millisecond pulsars known in our galaxy, many of which are extremely interesting in their own right.

Conclusion: The Fermi LAT, built by an international collaboration with major contributions from NRL scientists and engineers, has been incredibly successful in its studies of gamma-ray emissions from pulsars. It has increased the number from a small handful of objects to a large population with which important statistical studies are possible. It has revealed new classes of gamma-ray pulsars that were only suspected previously. Its prowess at pointing out new gamma-ray sources that turn out to be radio millisecond pulsars has yielded more than a half dozen new sources now being used as part of a pulsar timing array with the goal of detecting gravitational waves. Studies of these precise clocks have potential applications ranging from understanding the physics of matter at super-nuclear densities, to tests of general relativity, and even autonomous spacecraft navigation.

[Sponsored by NASA]

References

[1] P.S. Ray and P.M. Saz Parkinson, "Pulsar Results with the *Fermi* Large Area Telescope," in *High-Energy Emission from Pulsars and their Systems*, Proceedings of the First Session of the Sant Cugat Forum on Astrophysics, eds. N. Rea and D.F. Torres (Springer-Verlag, Berlin, 2011), pp. 37–56.

[2] P.S. Ray et al., "Precise Gamma-ray Timing and Radio Observations of 17 *Fermi* Gamma-ray Pulsars," *The Astrophysical Journal Supplement Series* **194**, 17 (2011).

[3] S.M. Ransom et al., "Three Millisecond Pulsars in *Fermi* LAT Unassociated Bright Sources," *The Astrophysical Journal* **727**, L16 (2011).

Blossom Point Tracking Facility — A Unique NRL Asset

G. Fuller[1] and L.C. Brown[2]

[1] *Space/Ground Systems Solutions*

[2] *Formerly of the Space Systems Development Department*

Introduction: Founded in 1956 as a segment of the first satellite tracking system in the United States, the Blossom Point Tracking Facility (BPTF) has played a significant role in a number of national space-related successes. Over the past 50 years, the BPTF has undergone a series of evolutions that have culminated in it becoming a state-of-the-art command and control facility, capable of supporting launch through end-of-mission-life operations. The foundation for this broad range of capabilities is the NRL-developed, government-owned, nonproprietary, Neptune Common Ground Architecture (CGA) software system, which has contributed to the successful launch and mission operations of more than 25 satellite systems that comprise more than 80 different spacecraft (see Fig. 3).

FIGURE 3
Blossom Point satellite tracking facility today.

Initially developed to support an existing major program, the capabilities of CGA have matured to include spacecraft support through all phases of program life — system development, ground system and spacecraft integration and test, launch and early-orbit checkout, mission and payload operations, and ultimately, spacecraft decommissioning. This broad range of capabilities is implemented as a layered architecture in CGA, which has been developed to maximize scalability, minimize future development, and provide highly efficient ground control and spacecraft operations.

Modular Architecture Reduces Development Costs: This work began in the 1980s, with the core of the software built as generic command and control (C2) code. Such code is mostly used to manage production machinery in factories and can be found in most supervisory control and data acquisition (SCADA) systems, such as are used in the management of most utility control stations. The software core comprises three layers (see Fig. 4):

- *Infrastructure* — managing process monitoring and control, symbol processing, data distribution, and message passing.
- *Services* — managing the scripting language, telemetry, data analysis, logging, recording, commanding, and memory processing.
- *Human–Computer Interface (HCI)* — managing text and graphics, system status, hierarchy, and tracking trends.

Most notable about this core is that nothing in it is specific to satellites, ground stations, or any other specific application. It is pure, generic C2.

At the next layer, applications are added to deal with the specifics of handling ground station equipment, test equipment, and satellite equipment and data itself. This follows a modular approach, which enables the addition of new spacecraft, new ground equipment, or new integration and test equipment, without the need for extensive software development. There is enough commonality that even when a new application needs to be written, much of it can reuse existing design and software. Through this modular approach and commonality, BPTF realizes a 95% software reuse rate for new programs, as measured in software lines of code (SLOC).

Automation of Spacecraft Operations: In 2006, under the direction of the program sponsor, the Automated Ground Operations software (AGO) was added to the CGA system. AGO, also known as Lights Out, uses the core and applications within the CGA system to effectively determine the state of the spacecraft, ground system hardware and software, and the environment tasks to be performed. AGO then uses this information to schedule future operations and tasking (see Fig. 5). The operations include correcting anomalies on the spacecraft and ground systems, re-allocating and rerouting ground processing, and notifying the appropriate engineers via text message or email when necessary. This eliminates the need for operators to man consoles 24 hours a day, 7 days a week, realizing a significant cost savings to sponsor programs.

AGO uses simple rules and hierarchical databases to ascertain that equipment is operating as planned. The hierarchical databases are used to generate information about the spacecraft and ground system from the raw data collected. These databases effectively determine if everything is within predefined limits. If an out-of-spec condition is detected, AGO follows the procedures described in the database for each known or previously unseen anomaly. The remedy can be as simple as ignoring the condition to as complex as "safing" the spacecraft and notifying the engineers of a catastrophic condition.

The procedures executed in response to these conditions are identical to the procedures previously executed by the human operators when they were charged with monitoring the spacecraft and ground systems' health. The procedures executed were written by the requisite engineers and have been fully tested in the CGA environment. The AGO software does not

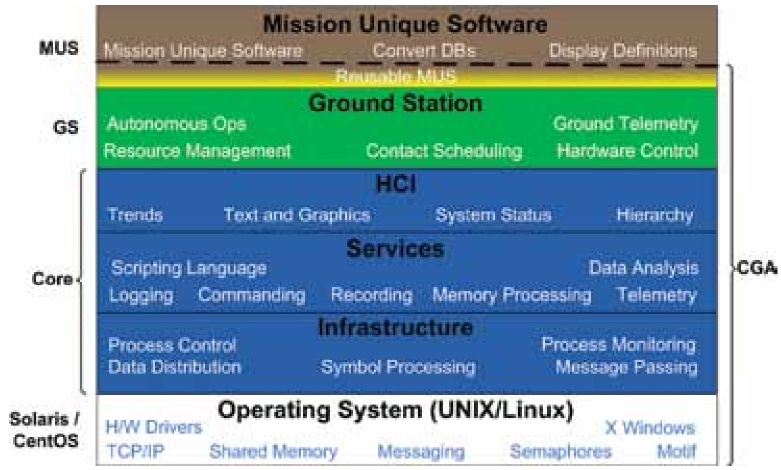

FIGURE 4
CGA's layered software architecture is a key enabler of the scalability and high software reuse rate achieved at NRL's Blossom Point Tracking Facility.

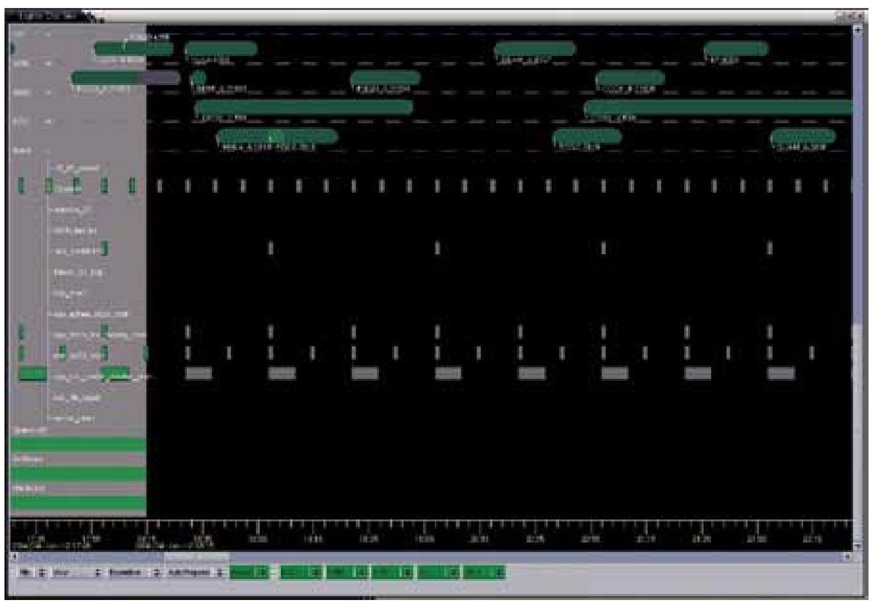

FIGURE 5
CGA display of the top level status of all spacecraft and ground systems.

make any specific determination about what to do in response to an anomaly; it just identifies the conditions and executes the proper procedure in response. This is exactly what operators were trained to do.

This approach is possible because the CGA system has access to, and controls all aspects of, the ground system from pass scheduling to dynamic resource planning, allocation, and configuration, to pass plan execution. All aspects of the hardware and software are reconfigurable by the system. The system has been engineered such that no manual patch panels, switches, or human input are required. This allows the software to dynamically reconfigure and replan tasks in response to changing conditions.

NRL's AGO system has resulted in major cost savings for Blossom Point. Staff positions have been reduced from 28 operators to 4 spacecraft engineers. Currently, no operator positions are staffed at BPTF. All contacts are performed by the AGO software. In addition, the software is better suited to perform the more mundane tasks of satellite monitoring. It never gets bored, takes a break, or comes to a different conclusion based on the same set of facts.

CGA and AGO continue to evolve. NRL is currently under contract to deploy this technology to the Air Force Research Laboratory at Kirtland AFB, Albuquerque, New Mexico, to support a new satellite program being launched in early 2013. Under another current initiative, the system is also being deployed to several universities across the United States for the operation of CubeSat Colony 2 satellites. To date, AGO has supported more than 20 spacecraft on orbit, taken over 400,000 satellite contacts, and saved more than 125 man-years of effort. Work performed under NRL contract.

[Sponsored by the NRL Base Program (CNR funded)]

■

Orbit and Mission Design for the TacSat-4 Satellite

W.J. Barnds and K. Akins
Spacecraft Engineering Department

Introduction: On September 27, 2011, at 15:49z (11:49 EDT), the Tactical Satellite-4 (TacSat-4) spacecraft successfully lifted off from the launch pad at Alaska's Kodiak Launch Complex en route to high-elliptical Earth orbit. TacSat-4 provides 10 ultra-high-frequency (UHF) channels usable for multiple permutations of communications, data exfiltration, and Blue Force tracking. The unique mission orbit augments geosynchronous communications by allowing near-global, albeit noncontinuous, coverage, especially at the high northern latitudes. Key to demonstrating improved access to space was establishing the capability to rapidly deploy the satellite in orbit, activate key onboard systems, and deploy the main payload's antenna reflector dish within a 3-day period so that the mission could continue into engineering evaluation and checkout (EE&C) in preparation for transition to user demonstrations within 7 days of launch.

The launch itself represented the culmination of an extensive astrodynamics effort undertaken by the Naval Research Laboratory to first identify the optimal mission orbit (see Fig. 6), then to implement a scheme to achieve that orbit as rapidly as practical, and lastly to maintain that orbit flexibly to maximize operational life while ensuring compliance with deorbit requirements. To satisfy these mission imperatives, NRL developed (1) a launch day maneuver plan that used a set of "blind" maneuvers to rapidly raise perigee altitude while efficiently expending available propellants, and (2) an orbit evolution plan minimizing propellant expenditures, guaranteeing successful deorbit of the space vehicle (SV) within 25 years of launch, and allowing the greatest flexibility for extended operations.

The most challenging phase of the first 3 days was the 3-hour period directly following launch (Fig. 7). Further, under worst-case assumptions, a potential existed for the most critical spacecraft activities to be further compressed into a 30-minute window. Essential activities during this period included validating proper execution of the automated activation sequence (via stored onboard commands); initial subsystem checkouts; deployment of the solar arrays prior to battery depletion; validation of thruster orientation; and execution of a perigee-raising maneuver at the first apogee passage. This last activity, one of the most crucial, had to await completion of the others.

Maneuvering Into Orbit: The earliest possible start time for the first maneuver was effectively at the same time as the optimal start for that maneuver (after the checkout burn; see Fig. 7). Moreover, the processing timeline for performing an orbit determination to fix the actual insertion orbit would not support a maneu-

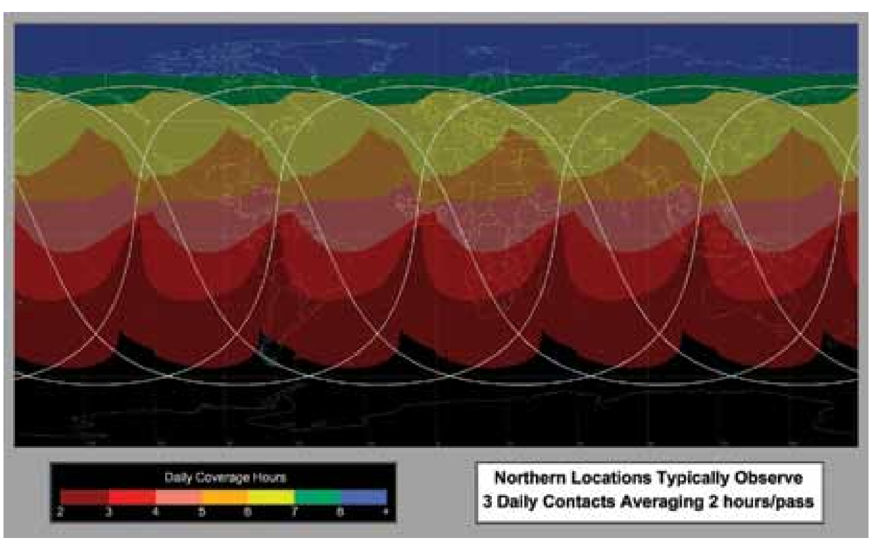

FIGURE 6
TacSat-4 daily coverage.

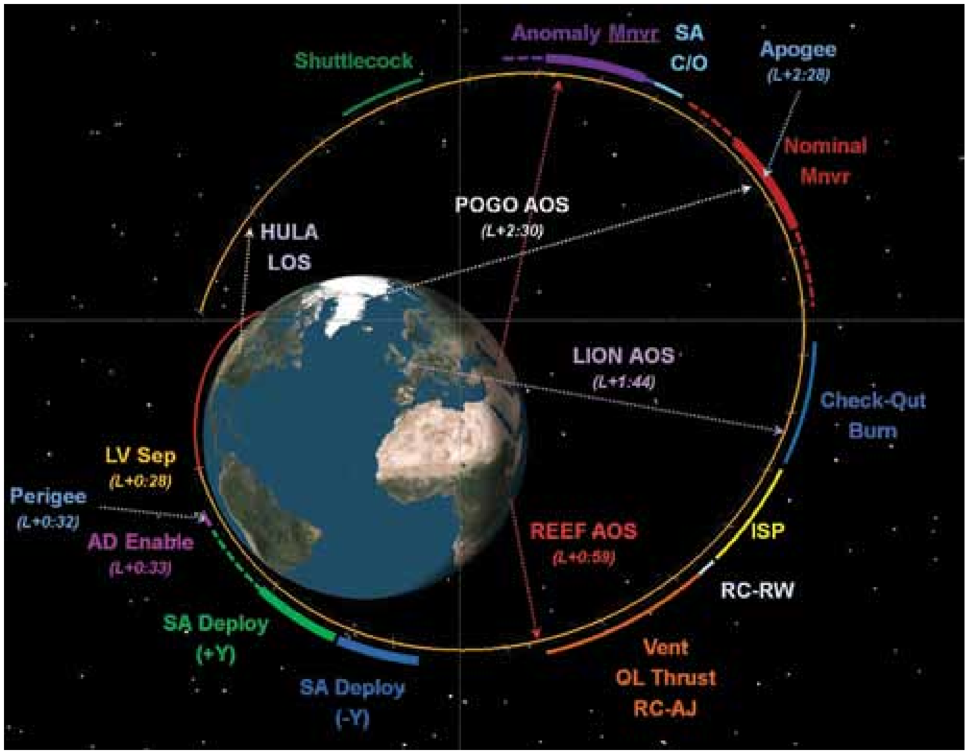

FIGURE 7
Revolution One operations plan: activities in the first 3 hours of orbit. LV Sep = launch vehicle separation, AD = attitude determination, SA = solar array, AOS = acquisition of signal, OL = open loop, RC-AJ = rate control with attitude jets, RC-RW = rate control with reaction wheels, ISP = inertial sun pointing, SA C/O = solar array checkout, LOS = loss of signal. REEF, LION, POGO, and HULA are stations in the Air Force Satellite Control Network.

ver until approximately 40 minutes after the start time for the optimal burn. To avoid this later limitation, an analysis was performed to assess the use of a priori insertion orbit accuracy (from the Minotaur-4+ booster) for maneuver planning efficiency. Booster-estimated errors are typically substantially worse than the accuracy achieved after processing on-orbit data, but in this case, the booster predict was sufficiently accurate to allow maneuver planning "in the blind." Using the expected insertion covariance of the booster, it was determined that the first two perigee-raising maneuvers (on the first and second apogee transits) could be performed efficiently without any need for orbit determination to refine knowledge of the insertion orbit.These two maneuvers would nominally be centered on apogee, to raise perigee altitude optimally from an initial 185 km to 350 km and 500 km altitudes. The maneuver inefficiencies associated with off-nominal insertion would incur less than a 5% propellant penalty, which was acceptable under mission rules. A third perigee-raising maneuver and an apogee-correction maneuver completed the initial orbit transfer plan and placed the SV in the proper orbit to proceed with deployment activities.

Orbit Maintenance: The TacSat-4 orbit maintenance strategy was comparatively straightforward (see Fig. 8). To avoid degradation of the germanium coating on the main reflector, the satellite needed to maintain a 600 km minimum operating altitude. To optimize coverage for several candidate user groups, a fixed ground track was required with longitude of the ascending node (LAN) at 12 deg east longitude. Although the desired ground track was not at a stable point, applying a small altitude bias created a drifting ground track that closely approximated the desired track with LAN fixed in a range from 9.3 to 15.8 deg east. The altitude was initially biased low to cause the ground track to drift eastward until natural orbit pumping caused a westward drift back to the initial longitude. At that point, the orbit was adjusted higher to restart the eastward drift. These annual maintenance maneuvers can be seen in Fig. 8 as the sudden changes in perigee altitude 1 and 2 years into the mission; there are larger decreases in apogee altitude, but the trend is more difficult to discern given the scale.

The nature of TacSat-4's deorbit plan can also be seen in Fig. 8. The orbit inclination is initially at 63.6 deg (above the ideal critical inclination of 63.4 deg shown as the green dashed line) and decreases due to

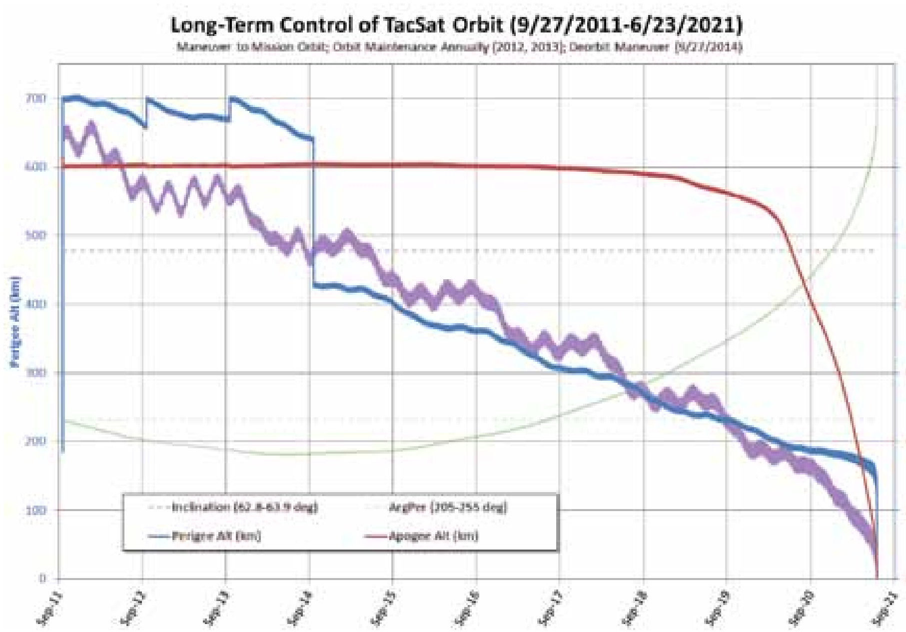

FIGURE 8
TacSat-4 orbit evolution and maintenance, 2011–2021.

higher-order gravitational effects. The initial inclination offset drives a decrease in the argument of perigee, reducing it from its initial value of 210 deg to 205 deg. Once the inclination moves below 63.4 deg, the argument of perigee begins to rotate in the opposite direction, increasing back to 210 deg approximately 5.5 years after launch. The main benefit of starting with the biased inclination is that the observed lunar-driven decrease in perigee altitude for critically inclined orbits drives perigee altitude decay for a much longer period of time (compared with starting at the nominal critical inclination). The extended period of perigee-decay reduces the amount of propellant required to achieve the requirement to deorbit within twenty-five years of the end of operations.

Deorbit Replanning: In March 2009, analysis of the predicted post-mission decay of TacSat-4's orbit indicated that TacSat-4 would naturally deorbit in 11 to 23 years. Subsequent to that analysis (May 2009), the NASA Marshall Space Flight Center issued a substantial downward revision to their solar activity forecasts for Solar Cycles 24 and 25, which increased predicted orbit lifetime to 20 to 200 years. During January 2011, the predicted intensity for Solar Cycle 24 continued to drop, exhibiting 50% lower sunspot activity than predicted in May 2009.

While NASA-Marshall noted in 2011[1] that "no generally accepted physical solar model is available to accurately predict future solar activity" and did not issue explicit guidance for Solar Cycle 26 (then antici-

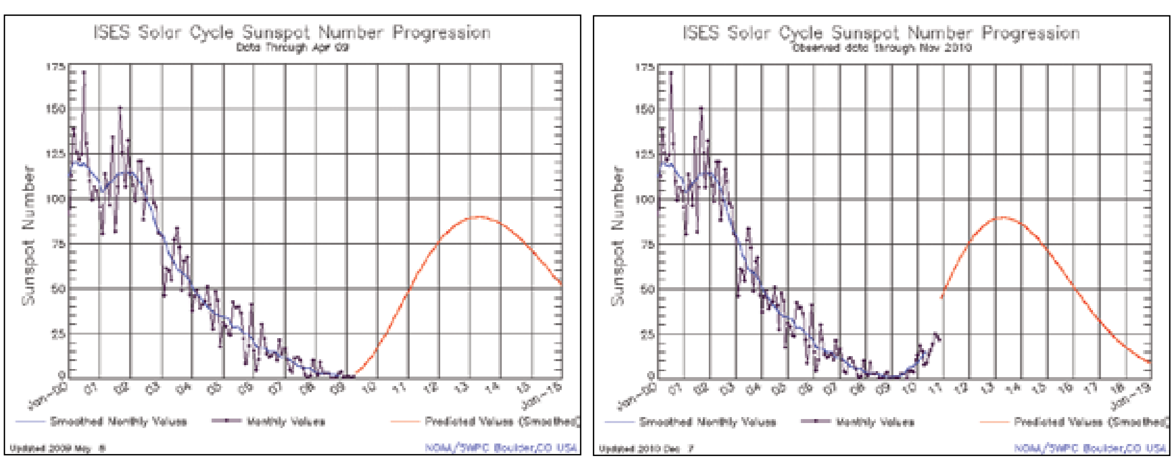

FIGURE 9
Changes in NOAA solar activity predictions.

pated to peak around 2022), their published predictions extended through 2030 (as opposed to the data available in March 2009, which only extended through 2023). NASA-Marshall supported this extension by noting that the most recent solar cycle is generally a good indicator of the level of activity to be expected for the next solar cycle (Fig. 9). Although we note that this guidance failed for the transition from Solar Cycle 24 to 25, this guidance represented the most reasonable course of action for Solar Cycle 26 predictions.

To achieve deorbit within the required 25-year time frame, the TacSat-4 team changed its approach to achieving and maintaining the mission orbit and achieving reentry. Based on consideration of the long-term dynamics of TacSat-4's post-mission orbit, the delayed launch date, and solar activity uncertainties post-2022, the TacSat-4 mission operations manager incorporated an end-of-mission perigee-lowering maneuver into the baseline mission plan. Under the current concept of operations, after the third year of operations, the SV will execute a maneuver to lower perigee below 435 km altitude and command the SV to rotate its solar arrays to a posture designed to maximize the average area-to-mass ratio. Reentry is expected to occur within 7 years of this maneuver.

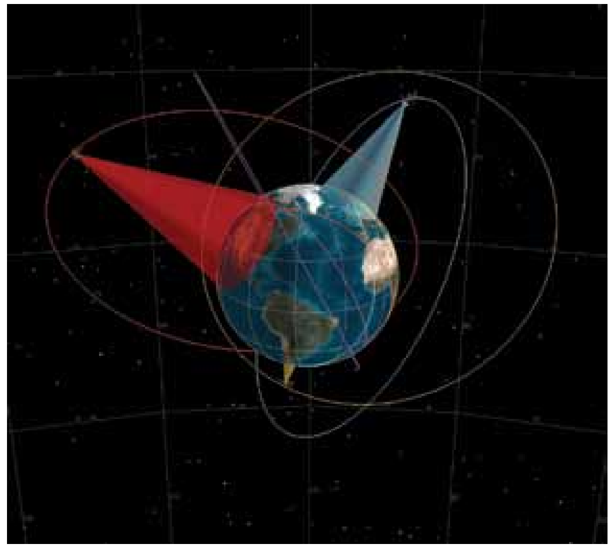

FIGURE 10
Four-satellite constellation.

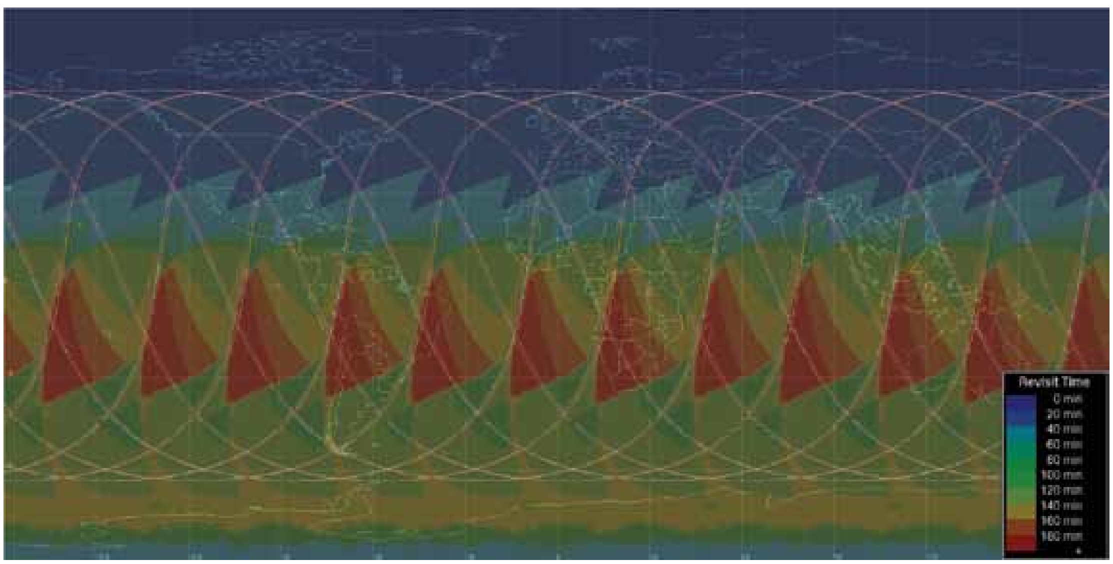

FIGURE 11
Four-satellite constellation coverage.

Multiple Satellite Constellations: NRL's astrodynamics effort has extended to exploring options for multiple-satellite constellations, one of which is shown in Fig. 10. Using a combination of surrogate, genetic, and gradient search algorithms, NRL optimized a four-satellite constellation design to minimize average revisit time for northern latitudes. The constellation is capable of delivering near-continuous coverage for users north of 62 deg latitude (Fig. 11).

[Sponsored by ONR]

Reference

[1] R.J. Suggs, H.C. Euler, Jr., and S.W. Smith, "Future Solar Activity Estimates for Use in Prediction of Space Environmental Effects on Spacecraft Oribital Lifetime and Performance," National Aeronautics and Space Administration, Marshall Space Flight Center, Huntsville, AL, Jan. 2011. Monthly report available at http://sail.msfc.nasa.gov/.

Autonomous Release of a Snagged Solar Array

G. Creamer,[1] G. Henshaw,[1] R. Hartley,[2] S. Roderick,[3] J. Hope,[3] and J. Obermark[4]
[1]*Spacecraft Engineering Department*
[2]*Information Technology Division*
[3]*PTR Group*
[4]*DCS Corporation*

Introduction: Many of satellite failures have been documented over the years,[1] some classes of which could potentially have been serviced if an existing servicing infrastructure were in place. An example of a possible servicing mission is the release of a snagged solar array or antenna. Though rare in occurrence, arrays do get snagged, due likely to jammed spring release mechanisms or blockage from shifted thermal blankets, which may occur as a consequence of the high loads imparted during launch. Just recently, a C-band antenna failed to deploy on the Intelsat New Dawn satellite and a solar array failed to deploy on the Telstar 14R satellite.

A concept for autonomous release of a snagged solar array has been developed in the Naval Center for Space Technology using a two-armed robotic servicing spacecraft and a variety of navigation and imaging sensors. A description of the technologies, hardware suite, and laboratory validation results is presented here.

Terminal Concept of Operations: The terminal concept of operations (CONOPS) begins with the assumption that the servicer spacecraft has achieved close-proximity rendezvous with the target spacecraft to within tens of meters. The CONOPS is divided into three phases: terminal approach, servicer preparation, and array release. During the terminal approach phase, the servicer follows a commanded approach path towards the target based on relative position and orientation measurements from its navigation sensor. Transition to the servicer preparation phase occurs once the target spacecraft range approaches the workspace of the servicer robot arms. During this phase, the servicer is maneuvered into a station-keeping capture box near the mounting location of the snagged solar array, and the two arms are positioned in preparation for array release. Lastly, after the arms are appropriately posed, the array release phase commences using localized machine vision to perform the final servicing task.

Localized Machine Vision: For the vast majority of satellites with large arrays stowed during launch, there is a small gap of a few centimeters between the edge of the folded array and the main bus structure. The goal of the machine vision logic is to identify the gap using a camera mounted to one of the arms (the imaging arm), and subsequently guide a fiducially marked tool mounted to the other arm (the servicing arm) inside the gap, wherein a separation force is imparted to release the array. To accomplish this

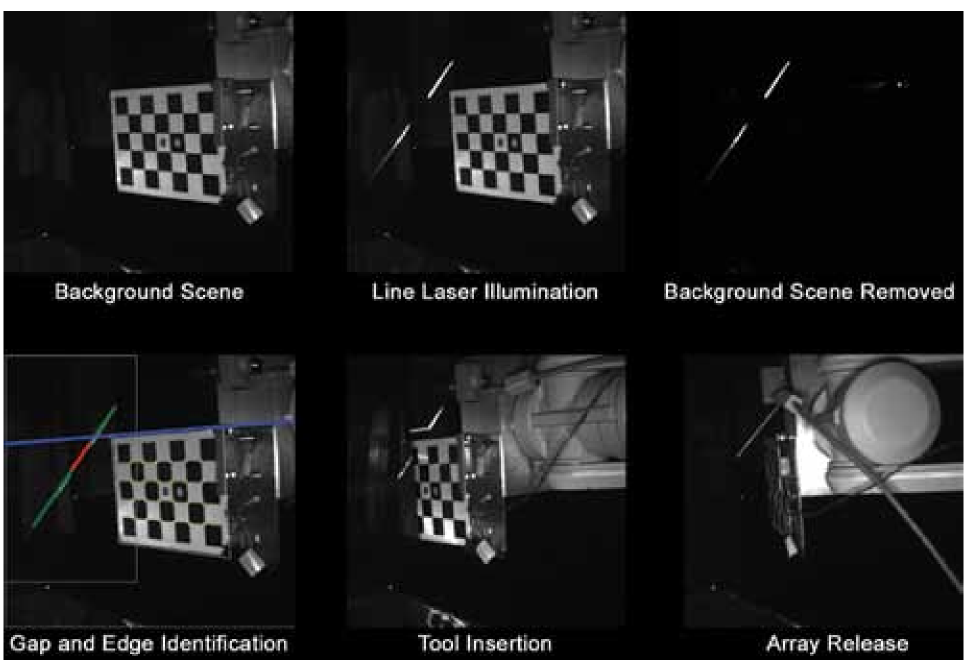

FIGURE 12
Typical local scene images obtained during array release.

task, a background image of the gap and the fiducial tool is first acquired from the camera. The servicing arm then illuminates a 45 deg laser line across the gap and a second image is acquired. The two images are differenced to highlight the laser line, and its broken edges are identified using an edge detection algorithm. Additionally, the top edge of the fiducial tool is detected from the raw background image, and a line defining that edge is extended towards the target to identify the location of the tool relative to the gap. The imaging arm then guides the servicing arm tool inside the gap to perform the array release. Typical camera and processed images corresponding to each of these steps are shown in Fig. 12.

Laboratory Validation Tests: NRL's Proximity Operations Testbed was used to validate the CONOPS and array release concept described above. This facility consists of two industrial-sized 6-degree-of-freedom robotic platforms operating in a workspace of approximately 25 × 10 × 3 meters, and is capable of computerized emulation of realistic close-proximity spacecraft motion including relative orbital dynamics, attitude dynamics, and control. The target platform represents an approximate full-scale mock-up of the back end of a typical geostationary satellite, including a launch vehicle upper stage marman ring and an apogee kick motor nozzle. Attached to the edge of the target is a simulated snagged array, with Velcro strips holding the array against the target edge and a pneumatic cylinder providing spring release. The servicer platform consists of an imaging arm, a servicing arm, a scanning laser navigation sensor, and end effector hardware including a camera, line laser, release tool, and LED illuminators. Figure 13 shows each platform in its final integrated configuration.

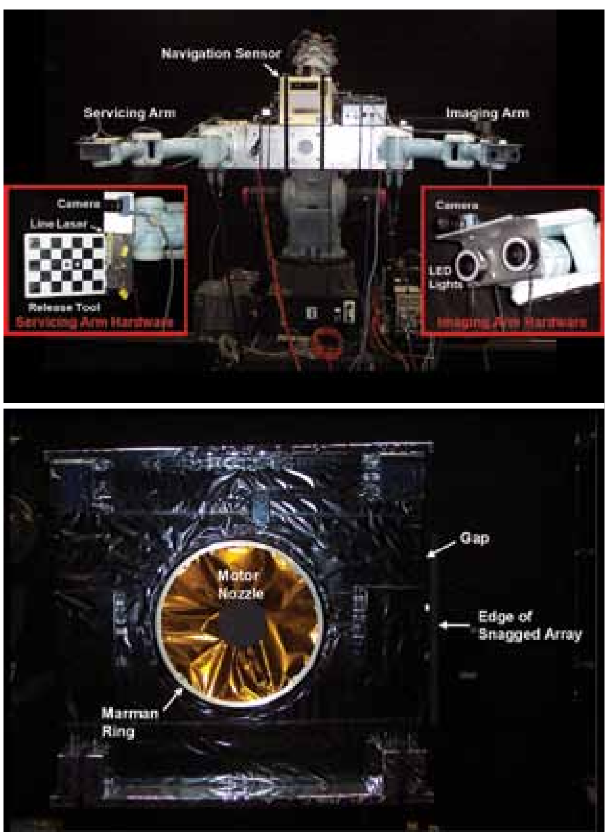

FIGURE 13
Integrated servicer platform (top) and target platform (bottom).

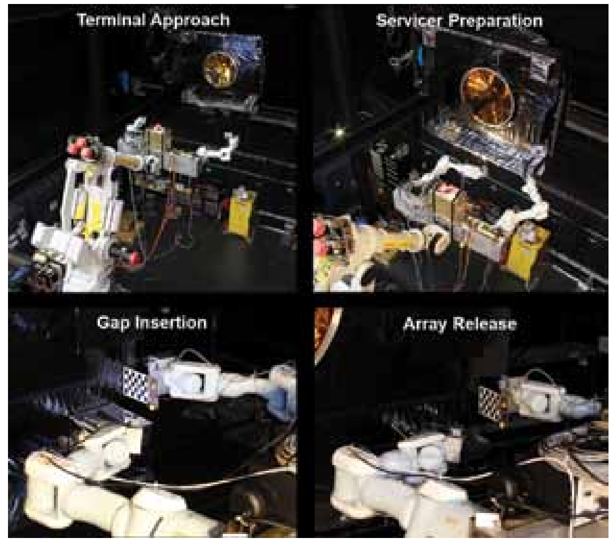

FIGURE 14
Progression of the concept of operations phases in NRL's Proximity Operations Testbed.

A series of tests was conducted to demonstrate the performance of the array release concept. Each test consisted of autonomous implementation of the three CONOPS phases, beginning with terminal approach and ending with array release. These phases are depicted in the series of photographs in Fig. 14, obtained during a typical test.

Conclusions: A geosynchronous servicer with a general set of tools and sensors could offer servicing capabilities for every high-value asset, with the potential to save years of critical operational life. This novel work has demonstrated the challenges and promise of a class of on-orbit servicing, the autonomous release of a snagged array.

[Sponsored by the NRL Base Program (CNR funded)]

Reference

[1] D.M. Harland and R.D. Lorenz, *Space Systems Failures: Disasters and Rescues of Satellites, Rockets, and Space Probes* (Praxis Publishing Ltd., Chichester, UK, 2005).