33rd Space Symposium, Technical Track, Colorado Springs, Colorado, United States of America
Presented on April 3, 2017

UNDERSTANDING SPACE WEATHER EFFECTS WITH DISTRIBUTED SENSOR SYSTEMS

Richard P. Welle
The Aerospace Corporation, welle@aero.org

INTRODUCTION

Terrestrial weather is defined by variations in the atmosphere, in the form of, for example, temperature, pressure, wind speed and direction, moisture content, precipitation, cloud cover, and fog. Humans have for millennia worked to understand and predict weather phenomena. That understanding, and the ability to generate useful forecasts, has been enhanced significantly by the proliferation of a wide range of sensors that continually monitor atmospheric conditions. Such sensors may make in-situ measurements of local conditions, or may measure conditions remotely by observing emission and/or propagation of electromagnetic waves. Although more data is always desired, the high spatial and temporal frequency of the data already generated by atmospheric monitoring systems has facilitated rapid advances in the understanding of terrestrial weather dynamics.

In a similar manner, space weather represents variations in the environment in near-Earth space as defined by electric and magnetic fields, by the flux of electromagnetic radiation, and by a variety of particles including molecules, atoms, ions, and electrons, all with a range of kinetic energies. The environment can affect spacecraft health (and human health) in near-Earth space through various mechanisms including radiation-induced damage caused by energetic particles, build-up of surface and volume charges that lead to discharge-induced damage, and variations in atmospheric drag. In addition, the space environment can affect the propagation of electromagnetic signals passing through it. In extreme cases, variations in the space environment can also affect the performance of various terrestrial infrastructure facilities such as electric power distribution systems.

Terrestrial weather, of course, can be highly variable, with variations being driven primarily through uneven solar heating, combined with the rotation of the earth and variations in Earth-surface composition (primarily water) and topography. The space environment is also highly variable, with the variations being driven primarily by variations in the output of solar UV and x-ray photons, and variations in the solar wind including, in extreme cases, coronal mass ejections that reach the Earth. Variations in the space environment are also affected by the rotation of the Earth and its magnetic field.

Terrestrial weather is associated almost exclusively with variations of atmospheric conditions in the troposphere, which extends up to about 20 km above the surface in the tropics and to lower elevations at higher latitudes. Space weather, on the other hand, has no hard upper boundary, although most interest in space weather is focused on the Van Allen radiation belts, which extend up to about eight Earth radii. Thus, the volume that must be considered in developing an understanding of space weather phenomena is at least four orders of magnitude larger than the volume of the troposphere that contains terrestrial weather.

While significant effort is invested in monitoring variations in solar activity, which drive much of the variations in space weather, a thorough understanding of the underlying physical processes in the near-Earth environment can be achieved only with detailed measurements of weather phenomena throughout that environment. Specific targets of such measurements include electric and magnetic fields, as well as charged and neutral particle number densities and energy spectra. Some of these measurements can be obtained with remote sensing techniques, such as GPS radio-occultation (GPS-RO) methods, but many are best achieved with in-situ measurement systems.[1]

The challenge with in-situ measurements is, of course, the volume of space involved. To fully instrument near-Earth space with a 10-km resolution would require something like half a trillion sensors - an absurdly large number that is clearly impossible for the foreseeable future. Even at a very modest 1000-km resolution, a complete sensor set would require half a million elements - no longer completely absurd, but still impossible with current technology. And a 1000-km resolution is quite coarse compared to many structures in the radiation environment in near-Earth space.

Another challenge, of course, is that it is physically impossible to place a stationary in-situ sensor anywhere in near-Earth space except on the circle of geosynchronous orbit. In most of near-Earth space, an orbiting sensor will be moving at several kilometers per second relative to the Earth below it.

So the obvious approach is to have a set of sensors in various orbits that map out the environment over space and time. The goal of this paper is not to review options for the sensors themselves; rather it is to consider options for how those sensors are deployed.

SENSOR DEPLOYMENT OPTIONS

Flying instruments to measure the Earth's radiation environment has a history as long as the space age; the first successful launch of an artificial satellite from the United States, on January 31st 1958, was Explorer 1, whose payload included a cosmic-ray detector designed to respond to energetic protons and electrons.[2] Data from this mission hinted at the presence of the Van Allen belts, the existence of which were confirmed by the Explorer 3 satellite launched less than two months later. [ref. 2] These first two satellites to explore the near-Earth radiation environment were quite simple by modern standards, and were not particularly large, with a mass of about 14 kg each.

A more recent example of satellites making in-situ measurements of the radiation environment is the Van Allen probes - a pair of satellites launched in 2012. These are much larger satellites (approximately 650 kg each), launched for the express purpose of exploring the Van Allen belts, and carrying a suite of instruments including electric and magnetic field sensors and sensors for evaluating the number density and mass and energy spectra of charged particles.[3]

Another mission currently flying is the Magnetospheric Multiscale (MMS) mission, a set of four spacecraft (approximately 1350 kg each) launched in 2015. The four spacecraft fly in a tetrahedral formation, with the individual spacecraft separated by about 10 km. This allows the mission to explore the spatial structure of the environment by correlating time-variable data from all four spacecraft.[4]

In contrast to the Van Allen probes and the MMS mission, both of which consist of dedicated free-flying satellites with the sole purpose of exploring the near-Earth radiation environment, there are a number of instruments designed for the same purpose that, instead of flying on a dedicated satellite, are hosted on satellites that have an unrelated primary purpose.

This approach is illustrated by the GEOScan payload, designed to be included as a hosted payload on the Iridium-NEXT satellite constellation.[5,6] This 5-kg payload design incorporates several instruments, including two space-weather-related instruments that illustrate the distinction between in-situ measurements and remote sensing. The dosimeter instrument consists of a pair of Teledyne micro dosimeters, one configured to measure electrons with a 100-keV energy threshold and the other configured to measure protons with a 3-MeV energy threshold. Both micro dosimeters would measure only particles that impact the sensing elements, so would be in-situ measurements, measuring only the local environment at the spacecraft. In contrast, the Compact Total Electron Content Sensors (CTECS) consist of a GPS receiver that monitors the refraction of GPS signals along the path from GPS satellites to the CTECS receiver. The data thus obtained would provide an integrated total electron content (TEC) measurement along the observation path. Compilation of data from the full set of sensors distributed across the entire Iridium-NEXT constellation would allow a continuous global snapshot of the Earth's ionosphere in the volume below the altitude of the GPS satellites (about 20,000 km). Thus, while the dosimeters would provide local data that can be compiled over time to build up a dose map in the Iridium-NEXT orbit, the CTECS instrument would provide a global measurement.

These two approaches to flying space-weather missions are complementary: The dedicated, space-weather-only, satellite mission will fly in an orbit optimized for the space-weather mission, and will be able to carry an extensive suite of instruments with a comparably large mass, power, and downlink budget. The hosted-payload mission, on the other hand, should be substantially less expensive, but will be constrained in mass, power, downlink budget, and, perhaps more importantly, in the selection of the orbit. The Iridium-NEXT satellites were launched into circular orbits at just under 800 km altitude, which means that any in-situ measurements would be confined to that orbit. Unfortunately, the GEOScan payload illustrates another challenge with hosted payloads; that of getting a ride. In this case, the GEOScan was not selected for hosting on the Iridium constellation and will not be flying.

A third approach to flying space-weather instrumentation takes advantage of the recent proliferation of CubeSats and CubeSat-related technologies. An example of a space-weather mission in this category is the AeroCube-6 mission, launched in 2014. This mission had, as its original purpose, a flight demonstration of next-generation chip-scale radiation dosimeters, similar to the dosimeters planned for the GEOScan mission. While initially designed as a technology-demonstration, AeroCube-6 was also designed to enable an extended science demonstration mission using those chip-scale dosimeters to measure fine structure in the low-Earth-orbit (LEO) radiation environment.[7]

The mission consists of two satellites built to the CubeSat design specification, each with a size of one half unit (where a CubeSat unit is a cube that measures 10 cm on a side). Figure 1 shows a photograph of one of the satellites. The AeroCube-6 satellites each fit within a volume of one half liter, having dimensions of 10 by 10 by 5 cm, and had a mass under 1 kg. The two AeroCube-6 satellites launched together and were attached to one another as if they were a one-unit CubeSat as they were deployed from the host vehicle. This CubeSat then split into two halves a few minutes after deployment, and the two resulting 0.5U CubeSats have been flying independently since then.

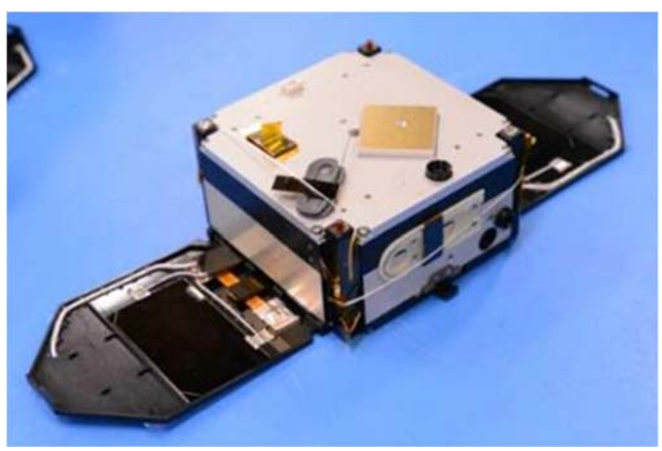

Figure 1. One of the AeroCube-6 satellites. The body of the satellite measures 10 by 10 by 5 cm.

The primary flight payload on each satellite consists of one current-generation chip-scale radiation dosimeter and two next-generation chip-scale dosimeters, as shown in figure 2. For the science mission, the dosimeters outputs are recorded continuously and downloaded on a regular basis. Part of the science mission requires distinguishing between time-varying phenomena and specially-varying structure in the radiation environment encountered by the spacecraft.[7] Since spacecraft in LEO move at a speed in excess of 7 km/s, data from a single spacecraft cannot easily be used for this purpose. Thus, the two AeroCube-6 spacecraft are flown in the same orbit with a separation of a few tens to a few hundreds of km. The data from the two satellites can then be correlated to distinguish between time-variable phenomena and spatially-varying structure.

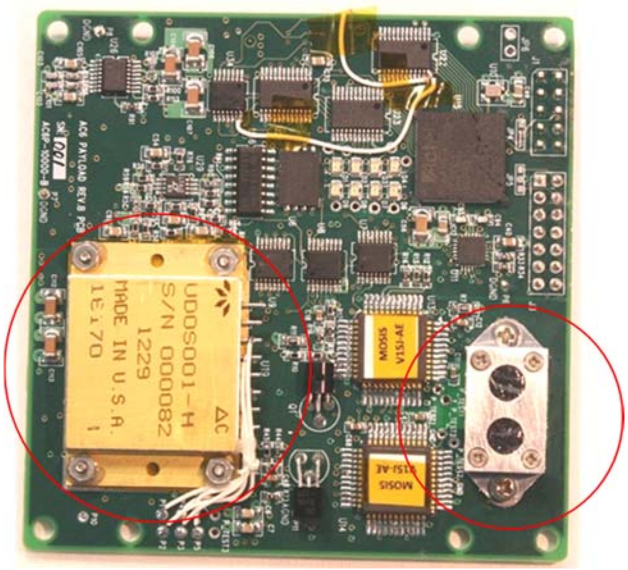

Figure 2. The radiation dosimeter payload on the AeroCube-6 satellite. The larger circle encloses the heritage sensor while the smaller circle encloses the two next-generation sensors. The overall size of the board is 8 cm square.

Since the AeroCube-6 satellites do not carry propulsion, the separation between them is maintained using variable drag, as follows. The two satellites were launched initially into an elliptical orbit of about 600 by 675 km altitude. At this altitude they encounter enough atmospheric drag to cause the altitude to fall a few km per year. The satellites will tend to separate from one another when one of the satellites falls more quickly to a lower orbit and thus orbits faster and moves ahead. The satellites, as depicted in figure 1, normally fly in a sun-pointing mode, spinning about the shortest axis. As such, the nominal drag of the two satellites is essentially identical, but small fluctuations do accumulate over time, causing the satellites to drift apart. When one of the satellites gets too far ahead, the attitude of that satellite is adjusted to reduce its drag, while the attitude of the trailing satellite is adjusted to increase its drag. The relatively-higher drag of the trailing satellite causes its orbit to decay more quickly than that of the leading satellite, allowing it to speed up and begin to approach the leading satellite. As the satellites approach the desired separation, the drag of the two satellites is adjusted to allow them to settle into orbits with equal semi-major axes, at which point they should no longer drift relative to one another.

Figure 3 shows the orbit history of the two satellites since launch, showing both the orbital separation and the difference in the semi-major axes. The initial large separation of the two spacecraft came about because they separated from one another with a larger separation speed than expected, and because variable drag operations were not implemented until after initial satellite checkout. Once the process was properly implemented, the two satellites were brought back together and typically kept within 100 km of one another. Precision formation control using this technique is more of a challenge than propulsion-based systems because of natural (and hard to predict) variations in the atmospheric density at this altitude. For the purposes of this mission, precision separation knowledge is far more important than precision control. The two satellites are equipped with GPS receivers that are used to obtain orbital ephemeris information that is precise to less than 20 m, or about 3 ms at orbital speeds.

The AeroCube-6 mission may, at first, appear to be a scaled-down version of the large, dedicated, space-weather satellites. However, it is probably more appropriate to think of it as simply a free-flying space weather sensor (although it does carry a secondary experiment to gain on-orbit experience with a set of experimental solar cells[8]). It is, of course, entirely possible to fly a CubeSat mission using methods and processes typical of traditional large spacecraft, in which case the CubeSat would effectively be a scaled version of some larger satellite. However, the approach used for CubeSat-based missions can be very different under the right circumstances. For the space-weather mission, where a distributed collection of in-situ sensors can provide data not available by other means, the CubeSat can represent a fundamentally-different approach to satellites, through the application of the CubeSat paradigm.[9] For a typical large space mission, mission assurance is critical; loss of a satellite when there is only one or a small number of satellites represents a significant loss of mission. On the other hand, for a distributed sensor system consisting of a large number of discrete sensors each flying in its own satellite, the loss of a satellite represents only the loss of data points. The mission, which is the collection of a large data set, is only marginally compromised by the loss of a small fraction of those data points. Thus, it is not necessary to take the degree of care that would be required to ensure that each and every satellite will work, provided that the satellite design in general has been proven to be suitable for the mission.

While this approach can lead to a substantial reduction in the unit cost of each satellite, more important is that it opens the opportunity to apply evolutionary processes to the development of the overall sensor network. While the space-weather mission, including both the sensors themselves and the satellites on which they fly, is an ideal target for the CubeSat paradigm, not all aspects of the CubeSat paradigm can be applied.

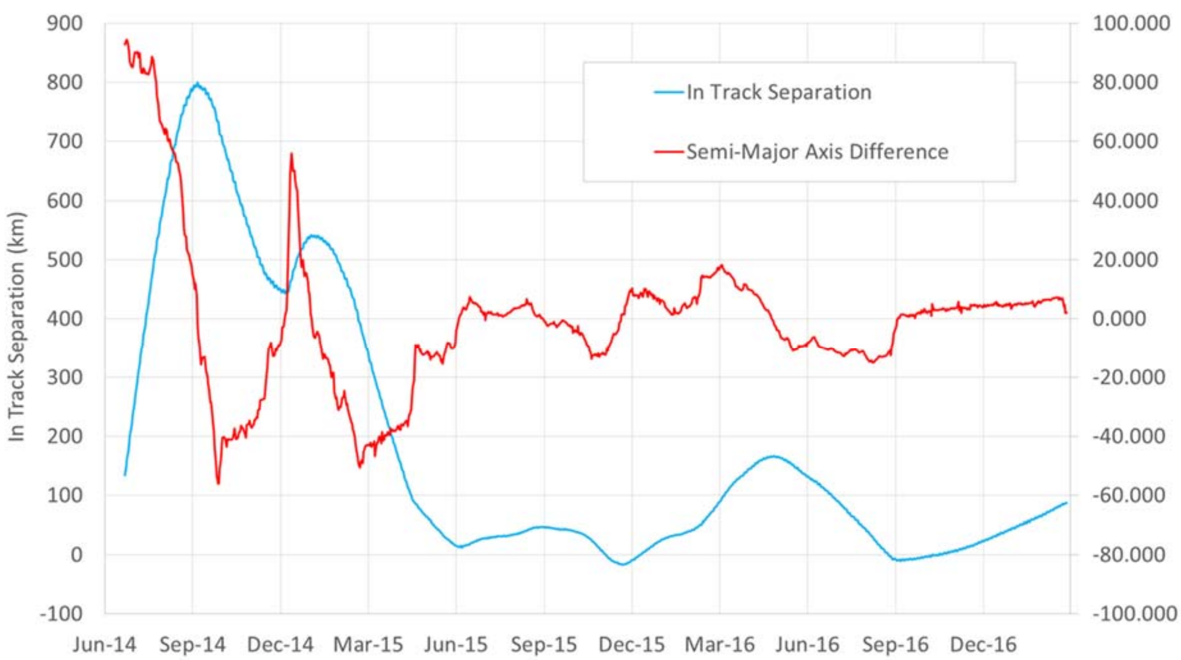

Figure 3. The in-track separation of the two AeroCube-6 satellites as a function of time since launch

THE CUBESAT PARADIGM

The CubeSat paradigm derives from the standardization inherent in the CubeSat Interface Control Document (ICD), which establishes the defining characteristics of the CubeSat itself; the size, mass, etc., and various limitations to ensure that it presents no risk to the launch vehicle or primary payload. A key outcome of the development of the CubeSat ICD was that the space launch business, at least for kg-class spacecraft, was effectively containerized such that the launch provider delivers a box to orbit (the P-POD, or equivalent) and the satellite developer need only design and build a satellite that fits in the box and conforms to standard rules to protect the launch vehicle and primary payload. In providing a standard interface between the launch provider and the CubeSat, the consequent simplification of the integration process reduces costs for the launch provider and, for the satellite developer, provides a set of standards which, if satisfied, will allow the CubeSat to ride on any of a number of launch vehicles.

The low cost of the CubeSat and the availability of high-frequency, low-cost rides to space encourages a culture of tolerance to risk, a key component of the CubeSat paradigm. The original concept of the CubeSat was developed with the intention of providing teaching tool in a university environment, where there was a high value placed on innovation and risk taking. For education purposes, this makes sense; one can learn as much (or more) from a failure as from a success. Many non-university programs also recognized the value of the CubeSat ICD for supporting a program of rapid technology development. In particular, technology demonstration missions can often be flown with a high risk tolerance when the missions are part of a series of tech demo exercises; an anomaly encountered on one flight can serve to inform the design of subsequent flights. Since the development cycle can be very short, a goal of demonstrating a particular technology in space can be applied to a series of flights rather than to a single flight. Under this approach, any single flight can have a high tolerance to risk under the expectation that a series of flights spread over a reasonable time interval will ultimately be able to satisfy all the program objectives. When applied to the space-weather mission, where the number of sensors can potentially increase to very large values, the CubeSat paradigm can be used to foster both the development of the fundamental design and the growth of the overall sensor constellation.

Elements of the CubeSat paradigm

Many aspects of the CubeSat paradigm derive from a goal to keep costs low enough that a satellite failure would not be intolerable. An example of things that CubeSat programs can do to keep costs down is the use of commercial off-the-shelf (COTS) electronics. Typical satellite programs will incorporate a large fraction of space-rated (meaning principally radiation-tolerant) electronics. Typical COTS electronics can tolerate a limited amount of radiation, however, and this limit is rarely reached in LEO where most CubeSats fly. As such, typical CubeSats will fly exclusively or almost exclusively with COTS electronics. To some extent, issues encountered due to radiation in one satellite project will be mitigated through elimination of suspect components in subsequent flights, but there is typically no systematic effort to evaluate the radiation tolerance of electronics components selected for a flight project.

For a space-weather monitoring system, the desired orbits are likely to be substantially higher than the LEO orbits typically used for most CubeSats, leading to a requirement for electronics that are more tolerant to radiation. The traditional-satellite approach to this would be to implement a requirement that all electronics on the CubeSat be space-rated, which would come with some (probably substantial) cost. The CubeSat-paradigm approach to this would be to fly the CubeSats in successively higher orbits while monitoring response to radiation. These CubeSats would have to have sufficiently-robust monitoring and communication systems to allow errors to be detected and relayed to the ground. Information obtained from those flights would be used to modify designs for subsequent flights.

CubeSats are often designed with short lifetimes in mind. For an educational project, the bulk of the education comes out of the design and build, with an additional gain during initial on-orbit checkout and operations. Beyond that, the marginal value of the satellite for educational purposes is not large. Similarly with tech-demo missions; once the technology has been demonstrated (unless on-orbit lifetime is part of the demonstration) there is little marginal value in continuing to operate the satellite. As such, many CubeSats are not designed with lifetimes in excess of one year in mind. For space-weather monitoring, on the other hand, longevity will be much more important. Application of the CubeSat paradigm would call, again, for an evolutionary approach. When first designed, it is not always clear what systems or components of a CubeSat will be life-limiting, but with an evolutionary program having good health monitoring, it will be possible to gradually extend the life of future missions by upgrading those components or designs found through experience to be of limited lifetime.

Another approach to minimizing costs is to limit the testing regimen throughout the program. If the risk tolerance of the program allows, a large portion of the environmental testing is deferred until completion of the initial satellite build. This approach is partially justified by the overall simplicity of most CubeSat programs; issues encountered late in testing can often be corrected within a short time frame because the entire satellite can be disassembled and reassembled in a few days. This approach can, however, lead to missed launches if there is insufficient margin built into the schedule to allow for modifications to correct issues discovered late in testing. Some CubeSat developers will build an engineering model that is a nominal duplicate of the flight model. Ideally the engineering model will be built before the flight model, with the experience gained through its build and test being available to inform the build and test of the flight model. The engineering model is then available on the ground for testing, software checkout, and anomaly resolution after the flight model is delivered. For a space-weather monitoring system that will eventually see the deployment of large numbers of identical or nearly-identical satellites, the testing regimen can evolve with time. The early models can be tested extensively to evaluate the design, but testing with the later models can be limited to that required to catch manufacturing errors.

Similarly, CubeSat development programs often give up extensive careful modeling of spacecraft performance, particularly in the area of mechanical integrity. This can be partially justified in that the CubeSats are so small that they become rugged simply by being more compact. Nevertheless, there are still mechanical issues that crop up during testing that could have been caught with careful modeling. However, with CubeSats it is not always clear whether it is less expensive to avoid testing anomalies through careful modeling, or simply to expect the testing to catch some issues that are then corrected through redesign after testing. A planned build of very large numbers of identical satellites may benefit from more extensive modeling, but it is not clear that this would be more useful, or less expensive, than adding more flight testing in the early models.

Another approach sometimes used is to work to an aggressive schedule with a hard delivery deadline and to deliver whatever is ready at the deadline. Using this approach, payloads may be only partially functional at the time of delivery, with the expectation that more will be learned by flying a partially-functional payload, followed by a fully-functional payload on a later flight, than will be learned by simply waiting for a later flight. Of course, the satellite has to be sufficiently operational that data obtained on satellite and payload performance can be collected and transmitted to the ground. Again, this approach is only suitable for missions with a very high risk tolerance. It is interesting to note that the first space-weather mission, the original Explorer 1 satellite in 1958, is an example of this approach; at that time, there was an urgency to fly for political reasons, and the payload was not fully ready (the data recorder was unfinished). The program flew anyway (limited to only real-time download of sensor data) and followed up two months later with another flight carrying a fully-functioning payload with a data recorder (Explorer 3).[2]

A CubeSat developer also often chooses to give up redundancy in the various satellite subsystems; CubeSats typically fly with a much larger compliment of potential single-point failures than traditional satellites. A corollary to this, however, is that the low cost of CubeSats, particularly the marginal cost of building spares, makes it possible to approach redundancy from an entirely new direction; instead of building redundancy into each subsystem on a satellite, the CubeSat developer may have the option of flying an entire duplicate satellite for redundancy. Of course, this approach will not mitigate design issues, but will serve to mitigate workmanship issues, some radiation-induced issues, and operational issues. A space-weather monitoring system can take this approach to the extreme of treating each satellite as individually expendable as long as the aggregate set of satellites has an acceptably-high success rate.

A variation on this approach is called "sequential redundancy." In this approach, a CubeSat is designed to be one of a series. The first of the series is delivered and launched, and the experience gained through the design, build, test, and operation of the first model are then applied to the development of the second model. In this approach, the success of the mission is defined by the success of the first satellite in the series that accomplishes the full mission. Of course, the mission goals may evolve as the satellite series progresses, leading to a continuing development effort reaching for ever advancing goals. The point of the CubeSat ICD was to ensure easy access to space at a cost where a mission failure was not intolerable. The redundant-satellite approach or, even more so, the sequential-redundancy approach means that a satellite failure is not necessarily a mission failure.

The sequential-redundancy model is an interesting approach to the development of complex satellites in that it involves continual on-orbit testing of satellite systems as they are developed. The CubeSat developer, instead of going through the traditional development process involving extensive careful design, pre-screened electronic components, built-in redundancy, and extensive testing at the component, subsystem, and system levels, will elect to simply build the best satellite that can be designed and delivered within a predefined time period and then move on to the next satellite in the series. The information obtained from the first satellite is used in the design of the next satellite in the series. Certainly this process is used to some extent in traditional satellite programs; lessons learned in the build of one satellite are applied to any future satellite where they are useful. However, the time cycle for this in traditional large-satellite programs is typically several years long; in some cases a satellite may be in the development phase for a decade or more, and the final design frozen several years before launch. With CubeSats, the development cycle can be measured in months rather than years.

An extension of the CubeSat sequential-redundancy model is to develop multiple satellites in parallel, but offset in phase, as indicated in figure 4. In this model, satellite A starts development, going through a design, build, and test phase before launch. Satellite B will start development while A is waiting for launch. The design of satellite B can incorporate modifications that correct issues discovered during the build and test phase of A. Satellite A will then launch, providing on-orbit operational experience that will be used to inform design modifications for satellite C. In addition, issues discovered during the build and test of B can be corrected in the design of C, which will take place while B is waiting for launch. This process can continue as long as necessary for the satellites to reach the desired level of capability (or may continue indefinitely as the requirements evolve along with the experience of the satellite builders). While any individual satellite will not be as carefully designed, or as thoroughly tested, and will not have as much built-in redundancy as a single satellite developed using the traditional

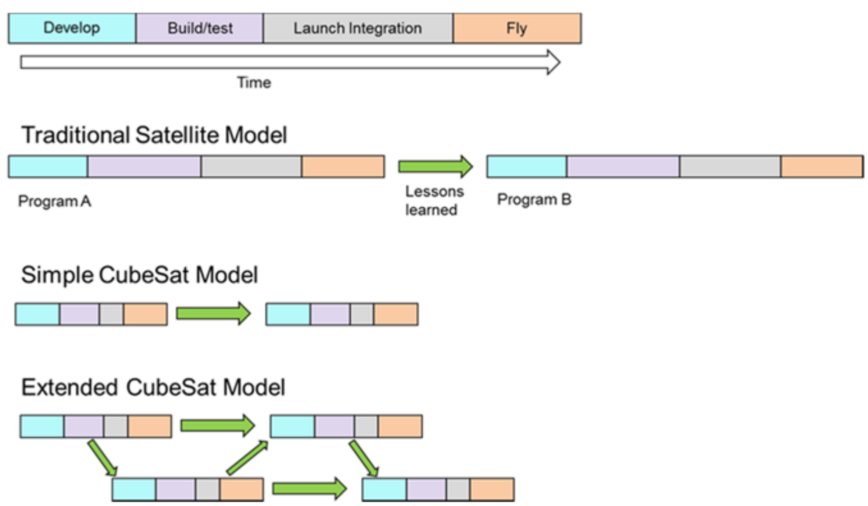

Figure 4. Graphical representation of sequential redundancy.

satellite development model, the final satellite in the series, and any copies, should be as capable as the traditional satellite.

In applying this method to space-weather monitoring, the first few satellites in the series would be experimental to some degree as the design is refined. Ultimately, though, after a few generations, the satellite design would become more fixed and the numbers would increase. However, the satellite design would never be completely fixed. In parallel with the engineers using flight results to improve the satellite design, the scientists analyzing the mission data will inevitably recognize opportunities to improve the utility of the data by making incremental changes in the instrument design, in the selection of the orbits, and in the desired data to be measured. This feature of the CubeSat paradigm allows the science mission to evolve incrementally in the same manner as the satellite design; as the scientists gain a better understanding of the space-weather environment, they will realize that there are other data that would be useful - different particle energy levels for example, different data grid spacing, or different aspects of the electromagnetic field.

There is a corollary in this evolutionary approach with terrestrial weather monitoring systems. For example, terrestrial temperature records go back hundreds of years, and during that time there has been continual progress in the technology used for those temperature measurements. New stations have come on line and old stations cease to operate. Failure of any given temperature monitoring sensor will not cause a significant loss of data; it can be replaced with data from a nearby sensor, or with a new sensor in the same location. Similarly, other terrestrial weather measurements are evolving. Wind measurements take place at higher elevations above ground. Radar is used to measure precipitation. Satellite-based measurements now constitute a large fraction of the data used in weather forecasting. All of these advances represent an evolutionary approach to weather monitoring. A similar evolutionary approach could be applied to space-weather systems.

While many aspects of the CubeSat paradigm support a distributed space-weather monitoring system, there are, of course, issues that will be encountered. Foremost among these is, perhaps, the tight limitations on payload volume, mass, and power in CubeSat-scale systems. On the other hand, though, there is a clear trend for shrinking of instrumentation. Figure 5 shows three generations of radiation dosimeters. The first generation is clearly too large for a CubeSat, while the latter two generations both flew (together) on AeroCube-6. The radiation dosimeters in the GEOScan payload designed for Iridium-NEXT are of the same size as the second generation dosimeters in this photo. In addition, the CTECS payload in GEOScan is derived from a demonstration flight of the CTECS on the second PicoSat Solar Cell Testbed (PSSCT-2) in 2011. The PSSCT-2 satellite, although not conforming to the CubeSat standard, was a small (<5 kg) containerized spacecraft deployed from the space shuttle on its last flight (STS-135).[10] The GPS-RO payload on PSSCT-2 was flown as a risk-reduction effort in support of potential future missions. Similar GPS-RO payloads are currently flying on a number of 3U CubeSats being used for both space and terrestrial weather measurements.

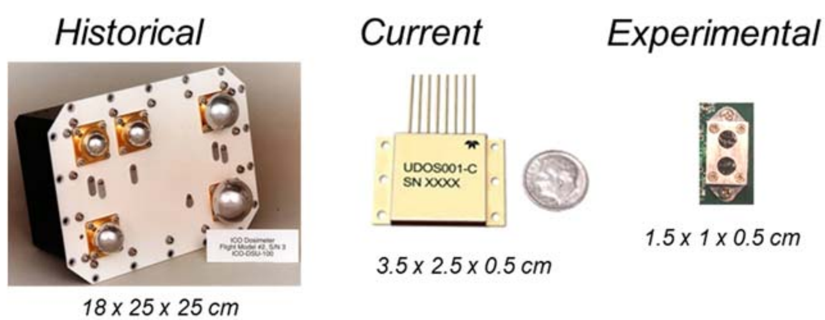

Figure 5. Three generations of radiation dosimeters.

Another key issue that may be more problematic for high numbers of free-flying space weather sensors is space debris. CubeSats, at least those launched from the United States, are required to comply with the 25-year deorbit rule; either the satellite must be able to maneuver into a disposal or reentry orbit at end of mission, or it must be in an orbit that will naturally decay such that the satellite will reenter and burn up in the atmosphere within 25 years of end of mission. For CubeSats in LEO, at altitudes less than about 650 km (depending on the mass and drag profile of the satellite), natural decay is the normal disposal mode. Satellites in higher orbits will need either a deployable drag-enhancement device, or a propulsion system. Since most of the relevant space-weather environment is well above 650 km altitude, this can be an issue.

An interesting option is suggested, again, by the Explorer 1 and Explorer 3 missions. These two satellites were launched into elliptical orbits with apogees, respectively, of 2,550 and 2,799 km, well above typical LEO CubeSats. However, the perigees were only 358 and 186 km respectively. The low perigee resulted in substantial drag, and the satellites reentered the atmosphere after 12 years and only three months, respectively. Clearly satellite lifetime is strongly influenced by the perigee, even when the apogee is high, so an interesting option that would span most of the Van Allen belts would be to drop the CubeSats into a geosynchronous transfer orbit (GTO) with a perigee low enough to ensure that the satellite meets the 25 year rule. The apogee would initially be at GEO, but would fall over time as the orbital energy is dissipated by the drag at perigee. By this means, a network of in-situ sensors could be deployed to orbit throughout the Van Allen belts and still meet the 25-year reentry rule.

Another advantage of flying in this highly-elliptical orbit is that it is possible to collect data over the entire orbit, but to confine data download to the periods when the spacecraft is relatively close to the ground, limiting the required range of the communication link. The disadvantage of this orbit, of course, is that the radiation environment is particularly severe, and it will be a challenge to design a CubeSat capable of surviving for a significant lifetime in that environment. However, that is the region of space where sensors are most useful, so it provides a motivation to apply the CubeSat paradigm to development of a free-flying sensor that can survive and operate for a useful duration in this environment.

SUMMARY

The CubeSat paradigm has evolved in response to the implementation of the CubeSat standard for containerized launch of secondary payloads and the consequent increase in the availability and decrease in cost of launch opportunities for satellites that conform to the CubeSat standard. The CubeSat paradigm cannot be applied to all CubeSat missions, but is an ideal approach for missions, such as in-situ space-weather monitoring, that will ultimately involve deployment of one or more large constellations of identical or nearly-identical satellites. Continuing advances in miniaturization of space-weather sensors, and progress in radiation-tolerant CubeSat buses are likely to make such an approach to space-weather monitoring far more feasible in the near future.

REFERENCES

1. Schrijver, Carolus J., et al. "Understanding space weather to shield society: A global road map for 2015–2025 commissioned by COSPAR and ILWS." *Advances in Space Research* 55.12 (2015): 2745-2807.

2. Van Allen, J. A., et al. "Observation of high intensity radiation by satellites 1958 Alpha and Gamma," Journal of Jet Propulsion, September 1958.

3. NASA press kit, Radiation Belt Storm Probes, http://www.nasa.gov/pdf/678135main_rbsp_pk_final_hi.pdf

4. "Magnetospheric Multiscale: Using Earth's magnetosphere as a laboratory to study the microphysics of magnetic reconnection" (PDF). NASA. March 2015.

5. Dyrud, Lars P., et al. "GEOScan: a geoscience facility from space." *SPIE Defense, Security, and Sensing.* International Society for Optics and Photonics, 2012.

6. Dyrud, Lars, et al. "GEOScan: A global, real-time geoscience facility." *Aerospace Conference, 2013 IEEE*. IEEE, 2013.

7. Blake, J. B., and T. P. O'Brien. "Observations of small-scale latitudinal structure in energetic electron precipitation." *Journal of Geophysical Research: Space Physics* 121.4 (2016): 3031-3035.

8. Lee, Justin H., et al. "On-orbit characterization of space solar cells on nano-satellites." *Photovoltaic Specialists Conference (PVSC)*, *2016 IEEE 43rd*. IEEE, 2016.

9. Welle, Richard P., "The CubeSat Paradigm: An Evolutionary Approach to Satellite Design," *32nd Space Symposium, Technical Track*, Colorado Springs, April, 2016.

10. Bishop, Rebecca, et al. "First results from the GPS compact total electron content sensor (CTECS) on the PSSCT-2 nanosat." Proceedings of the AIAA/USU Conference on Small Satellites, SSC12-XI-2, 2012.