

THE CUBESAT PARADIGM: AN EVOLUTIONARY APPROACH TO SATELLITE DESIGN

Richard P. Welle

The Aerospace Corporation, welle@aero.org

ABSTRACT

Traditional satellite development programs involve extensive careful engineering analyses, modeling, testing, and mission assurance processes. Development programs can span many years; in some cases more than a decade can elapse between program initiation and first launch. In contrast, agile space programs, and particularly CubeSats, reduce the development effort to levels that allow program durations measured in months rather than years.

The key to the short development cycle of CubeSats is the CubeSat paradigm, which is distinct from the CubeSat standard. The CubeSat standard is defined by the CubeSat Interface Control Document (ICD), which defines the interface between the CubeSat and the launch vehicle, and sets tight constraints on such factors as dimensions, mass, and hazardous materials. For the launch provider the ICD ensures that the CubeSat, as an auxiliary payload, will do no harm. The point of the CubeSat ICD was to ensure access to space at a cost where a mission failure was not intolerable.

The original concept of CubeSats was as a teaching tool in a university environment; there was a high value placed on innovation and risk taking. However, many non-university programs recognized the value of the CubeSat ICD for supporting a program of rapid technology development. The opportunity to fly a high-risk mission at nominal cost encouraged entrepreneurs to establish programs that required multiple generations of spacecraft design on a very short cycle, with the understanding that there will be failures on orbit. However, these failures will all be lessons that lead to improved designs in the next generation of the satellite. The program is defined not by the capabilities of the first satellite to fly, but by the capabilities of the *n*th generation of satellite. The CubeSat paradigm is the recognition that failures are a natural part of the development cycle, and that as much, or more, can be learned by having these failures in early flight models than through extensive ground modeling, development, testing, and mission-assurance programs. The CubeSat paradigm is a natural (although possibly not foreseen) outcome of the CubeSat ICD.

INTRODUCTION

A CubeSat is any satellite that conforms to one of various CubeSat design specification documents. The original CubeSat design specification was developed at California Polytechnic State University in San Luis Obispo (Cal Poly) in 1999 and is based on the Poly Picosat Orbital Deployer (P-POD) that was developed at the same time¹. Since then there have been alternative CubeSat standards developed, each based on an alternative deployer. What all have in common, though, is that the deployer provides a standard interface with a launch vehicle, and, at the same time, is designed to carry a secondary payload while ensuring minimal risk to the launch vehicle and the primary payload.

All CubeSat deployers are designed around a standard "unit" volume that is approximately a 10-cm cube. This dimension was selected based on the concept that a volume of one liter was a reasonable working volume for an experimental satellite and provides adequate surface area for solar cells on each face¹. Most deployers are designed to deploy three one-unit (1U) satellites, each 10 cm square and a little over 10 cm long, with the three satellites

configured in a single stack that is 10 cm square and about 34 cm long. Although the original intention was to launch three 1U satellites, a 3U CubeSat deployer can also carry a single satellite that is 10 cm square and 34 cm long, two satellites each 17 cm long, or any combination of satellites that total 34 cm in length. Figure 1 shows a photograph of a 1.5U CubeSat with deployed solar panel wings.




Figure 1. The Optical Communication and Sensor Demonstration Satellite; a 1.5U CubeSat with deployed solar panels.

The original intent of the CubeSat standard was to provide a simple and repeatable (and reliable) interface with a launch vehicle so as to reduce the effort (and cost) of integrating a secondary payload. The goal was to enable inexpensive flight opportunities that could be used by universities for educational purposes. A large fraction, if not the majority, of CubeSats have been developed by educational institutions for research purposes. The utility of the CubeSat standard has been recognized beyond the university, and the form factor has been adopted by industry as a vehicle for technology demonstrations. In addition, there is a growing interest in CubeSats for commercial applications.

The CubeSat is generally defined in terms of the CubeSat Design Specification, or Interface Control Document (ICD) which defines the interface between the CubeSat and the deployer, and sets tight constraints on such factors as dimensions, mass, and potentially hazardous materials. The design specifications, intended to enable development of educational satellites, also enabled and encouraged a tolerance to risk, not for the primary payload or launch vehicle, but for the CubeSat itself. This tolerance for risk, far beyond that normally acceptable for traditional satellite programs, led to what has been called the CubeSat Paradigm, essentially a new approach to development and operation of spacecraft. The CubeSat Paradigm, and its associated risk tolerance, is applicable to certain classes of CubeSat missions, but it is not applicable to all classes of CubeSat missions. The first part of this paper will review the CubeSat ICD and the constraints it places on satellite developers. The second part will focus on the CubeSat Paradigm, its implications for possible satellite development strategies, and the types of missions where application of the CubeSat paradigm is appropriate.

THE CUBESAT ICD

CubeSats almost invariably fly as secondary payloads on launch vehicles that are carrying one or more primary payloads. This arrangement ensures that CubeSats can get rides into space without having to buy an entire launch vehicle, and allows the launch provider to gain a little income from the excess capacity often available when launching satellites. The CubeSat ICD provides for a standard interface between the CubeSat and the CubeSat deployer. The deployer, in turn, has an interface with the launch vehicle. Because it is standardized, the deployer (such as a P-POD) can be qualified to ride on a given launch vehicle without the need to fully specify the design of the satellite inside the deployer. Thus, the effort of qualifying the deployer to ride on a launch vehicle needs to be undertaken only once, and the cost of this effort does not fall to each individual CubeSat. The qualification to ride on the launch vehicle carries with it various constraints on the CubeSats that will ride inside the deployer, primarily in the form of "do no harm" requirements that protect the launch vehicle and the primary payloads. These requirements are passed on to the CubeSats through the CubeSat ICD, and are nominally the same for all CubeSats, but with some minor variations depending on the specific deployer design. There are also some variations from one launch to another based on specific requirements that may be imposed by the primary launch customer.

A basic understanding of the requirements can be obtained by considering the specifics of the ICD developed for the P-POD². The P-POD is a simple box with a door and a spring mechanism. Figure 2 shows a photograph of a 3U deployer. The door is opened on a signal sent from the launch vehicle and the spring mechanism pushes the CubeSat(s) out of the box with an ejection speed on the order of 1 m/s. The specific requirements imposed on the CubeSat fall into three categories.



Figure 2. Photograph of a P-POD CubeSat deployer with a mock-up of a three-unit CubeSat. (NASA image)

The first category of requirements is aimed at ensuring that the CubeSat will fit in, and easily exit, the box. Requirements in this category include dimensions, mass limits, center-of-mass limits, and materials considerations that ensure it will slide out of the box when the door opens. Generally the mass limit is 1.33 kg per CubeSat unit.

The second category of requirements are designed to ensure that the CubeSat will present no hazard to the launch vehicle or primary payload. This category includes limitations on hazardous materials, prohibition of

pyrotechnics, limitations on propulsion systems and stored energy, and materials restrictions to limit out-gassing. There are also requirements that the CubeSat be powered off until deployed, and that no deployable components (such as solar panels) be release from the CubeSat for 30 minutes after it leaves the P-POD. Also in the mission-safety-related category are testing requirements. These include vibration, shock, thermal, and vacuum bake-out tests. This is not an exhaustive list of mission-safety requirements, but provides a basis for understanding the CubeSat development process.

The third category of requirements includes those aimed at regulatory compliance. In the United States, this category includes licensing for radio-frequency transmissions, approval from NOAA for any Earth-observation systems (including all cameras), and compliance with orbital-debris limitations.

The requirements imposed on the CubeSat as a secondary payload have evolved somewhat since the first CubeSat launch in 2003, and vary somewhat from one launch provider to another, but the essence of the requirements remains unchanged. Having a standard set of requirements is beneficial to both the launch provider and to the CubeSat builder. For the launch provider the ICD ensures that the CubeSat, as an auxiliary payload, will do no harm. Furthermore, any launch provider can establish the capability to launch CubeSats by qualifying a CubeSat deployer (the P-POD, or its equivalent), without having to delve into the details of each CubeSat that might be launched. For the CubeSat builder, the ICD provides a set of rules that must be met. However, more significant for the CubeSat builder is that the ICD provides for a standard interface that, if met, allows the CubeSat to ride to space on a broad range of launch vehicles with minimal integration effort. This means that a CubeSat complying with the standard will have a broad selection of ride opportunities at competitive prices, and that these ride opportunities will be frequent.

The CubeSat ICD was created with the intention of encouraging flight opportunities for educational purposes. The first university CubeSat flights took place in 2003. By 2006, the CubeSat standard was recognized by industry as presenting a valuable tool for technology demonstration flights. Not long after that, the CubeSat form factor was recognized as a potential vehicle for commercial applications, and by 2010, the first commercial venture developing a business based on CubeSats was founded.

THE CUBESAT PARADIGM

The CubeSat ICD, as described above, establishes the defining characteristics of the CubeSat itself; the size, mass, etc., and the limitations that ensure it presents no risk to the host vehicle or primary payload. What is not immediately evident in reviewing the CubeSat ICD is the new approach to space systems enabled by the CubeSat ICD. A key outcome of the CubeSat ICD is that the space launch business, at least for kg-class spacecraft, has effectively been containerized. In the simplest approximation, the launch provider (in this case including both the launch vehicle provider and the CubeSat deployer provider working together), 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.

The CubeSat deployer is analogous to the standardized shipping container used for a large fraction of cargo transportation around the world. These containers, called "intermodal" because they are designed for ocean, rail, and highway use (see figure 3), have revolutionized the cargo transport industry over the past several decades by making it possible to pack any cargo into a standardized container, then ship the container to a destination, by one or more modes of transport, then unpack the cargo at the destination. The containers are moved over the road, rail, and ocean transport networks without regard to their contents, so the transport providers can focus only on efficient transport of the containers without having to develop efficient means of handing all the diverse cargos that might

be shipped in the containers. The standard container provides a benefit to both the shipper and the cargo owner by simplifying the interface between the cargo and the transport provider.



Figure 3. Uses of a standardized shipping container. (U.S. DOT images)

In a similar manner, the CubeSat container provides a standard interface between the launch provider and the satellite. 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 effect of this is that both the number of CubeSats being developed and built, and the number of launch vehicles capable of carrying CubeSats, has rapidly increased. For the CubeSat developer, the choice of launch providers means that the development effort can be undertaken without immediate concern about identifying a launch opportunity. Instead, the developer can assume that a ride will be found when necessary, and that the cost is reasonably well-known.

In contrast with shipping by intermodal containers, which travel over various shipping modes from and to the destinations selected by the cargo owner, the CubeSat developer has a limited choice about the destination of the satellite. As a secondary payload, the CubeSat will be delivered into an orbit that is determined by the requirements of the primary payload. The CubeSat developer has the option of waiting for a ride going to or near the orbit required by the mission, but this may have the effect of severely limiting the launch opportunities. An additional factor constraining launch opportunities for CubeSats is that the proliferation of CubeSat projects has exceeded the proliferation of launch opportunities. As a result, there is significant demand for launches, and available launches may be booked out more than a year in advance. At this time it is not clear whether the rate of growth in the number of CubeSat projects will continue, or whether the growth in available launches will be able to keep up with the demand. In principle, if the number of CubeSat programs increases sufficiently, it may be possible for a CubeSat launch consolidator to contract for an entire launch vehicle that will carry nothing but CubeSats.

Although not completely transparent, the path to space for a CubeSat is vastly simpler than for traditional satellite programs. It is this launch availability, combined with the original goal of the CubeSat as a teaching tool, that leads to what has been termed the CubeSat paradigm. The Oxford English Dictionary defines "paradigm" as a typical example or pattern of something - a model. In modern usage, the term has taken on a connotation of something more fundamental, perhaps a model that has broader implications for the way we do things. The CubeSat ICD is a standard, a model for a way to make a satellite. But more fundamental is the concept that there should even

be a standard model of a satellite. This concept, that there should be a standard model, is the true innovation of CubeSats. By conforming to the CubeSat specification, a satellite builder has vastly simplified the interface with the launch provider - so much so that there are now very frequent and very inexpensive opportunities to get into space.

The low cost of the CubeSat and the availability of high-frequency, low-cost rides to space encourages a culture of tolerance to risk. The original concept of CubeSats was as a teaching tool in a university environment. As originally conceived, 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. However, many non-university programs recognized the value of the CubeSat ICD for supporting a program of rapid technology development. For example, technology demonstration missions can often be flown with a high risk tolerance, particularly when the missions are part of a series of tech demo exercises. Under these circumstances, 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.

Similarly, the opportunity to fly a high-risk mission at nominal cost encouraged entrepreneurs to establish programs that required multiple generations of spacecraft design on a very short cycle, with the understanding that there will be failures on orbit. However, these failures will all be lessons that lead to improved designs in the next generation of the satellite. The success of the program is defined not by the capabilities of the first satellite to fly, but by the capabilities of the nth generation of satellite. The CubeSat paradigm, which is enabled by frequent, low-cost launch opportunities, is the recognition that failures are a natural part of the development cycle, and that as much, or more, can be learned by having these failures in early flight models rather than in extensive ground modeling, development, testing, and mission-assurance programs. Because one can now fly frequently, there is a more important benefit; the opportunity to learn by experience. The CubeSat paradigm is a natural (although possibly not foreseen) outcome of the CubeSat ICD.

This is not to say that all CubeSats can be built using the CubeSat paradigm; application of the CubeSat paradigm requires a tolerance to risk that is not always appropriate. For university satellite programs where learning is the primary goal, the CubeSat paradigm should be a given. For programs focused on technology evolution and/or maturation where the ultimate goal is an operational system or process that may take several years to develop, the CubeSat paradigm is a potential approach. However, for programs that are one-off technology demonstration missions where the typical CubeSat risk tolerance is unacceptable, aspects of the CubeSat paradigm need to be reevaluated.

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 low Earth orbit (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.

A corollary to this is that 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, most CubeSats are not designed with lifetimes in excess of one year in mind.

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.

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.

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 payload performance can be collected and transmitted to the ground. Again, this approach is only suitable for missions with a very high risk tolerance.

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.

An extension of this approach can be termed "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 operations 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 more 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 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. 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 satellite development model, the final satellite in the series, and any copies, should be as capable as the traditional satellite.

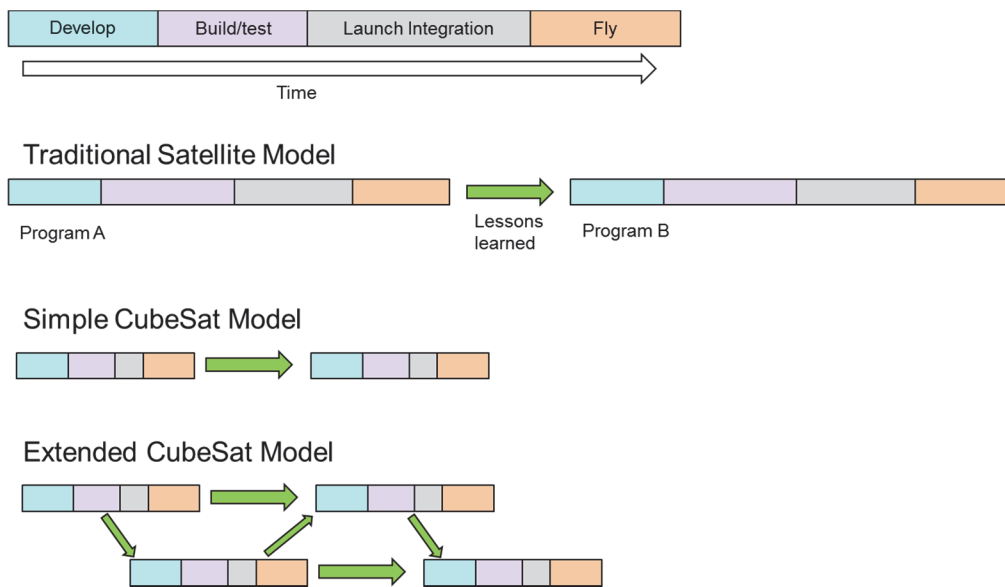


Figure 4. Sequential redundancy.

This model assumes that experience in building and flying CubeSats directly translates into improved designs. Although there is no rigorous way to quantify any increased capabilities that come about as a result of design experience, there is existing data to support the assertion that experience leads to higher reliability. Data on CubeSat reliability have been collected and correlated with the experience of the development team. The results³, shown in figure 5, clearly indicate a significant upward trend in reliability with increasing team experience. First-time builders (not uncommon in the CubeSat world) are not as successful as teams that have built and flown one or more previous satellites. One benefit of the CubeSat approach to satellite design is that the fast cycle time typical of such projects means that a team can be kept together through many projects. Thus, a small, dedicated team of engineers can build up the experience of multiple satellite projects, on a timescale short enough that there is not a lot of turnover on the team, and any experience gained by the team is retained.

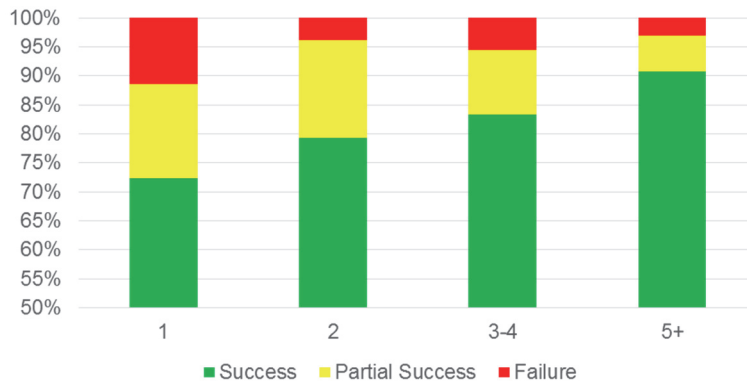


Figure 5. The effect of experience on mission success; the percentage success is shown as a function of the number of satellites previously built by the developer (from reference 3).

One further topic that should be considered under the CubeSat paradigm is standardization. In one sense, standardization enables the CubeSat paradigm; without standardization of the launch interface CubeSats would not have access to the frequent and inexpensive rides that are a prerequisite for the CubeSat paradigm. On the other hand, although the idea of standardization of various systems and components in CubeSats (power, flight computer, sensors, actuators, radios, etc.) is often suggested, this type of standardization is not really compatible with the CubeSat paradigm. CubeSat development under the paradigm is aimed at steady advances in the capabilities of the satellites. Although there are a number of excellent systems and components available, there is still plenty of room for improvements, both in increasing capabilities and in reducing the power and mass requirements. Adoption of significant degrees of standardization within CubeSat designs will stifle creativity in the long run, but it may lead to improved reliability in the short run.

Application of the CubeSat paradigm

The CubeSat paradigm is a valuable tool for advancing the state of the art in CubeSat capability and reliability, but it is not applicable in all cases nor, perhaps, even in the majority of cases. The CubeSat paradigm assume a tolerance to risk appropriate in only two circumstances; either a single satellite is being developed in an educational setting where the process of designing and building the satellite has as much or more value than actually flying the satellite, or a satellite is being developed as part of a long-term series of satellites where the end goal is a satellite design with capabilities well beyond what can easily be achieved in a single development stage, and anomalies (or outright failures) in intermediate satellite designs are taken into consideration in the overall plan. The sequential redundancy approach is appropriate if the program goal is either technology maturation for its own sake, or the

development of a satellite that is far beyond the current state of the art and cannot reasonably be achieved in a single design effort.

Although one may be tempted to take the approach of the CubeSat paradigm in the development of a technology demonstration mission, one must be very careful in this approach if the technology demonstration mission is not one of a larger series. Specifically, if a program has the goal of a flight demonstration of a specific technology, and only one flight is planned, then the tolerance for risk is likely to be far smaller than appropriate for application of the CubeSat paradigm. The expectations for the mission must be clearly understood, both by the CubeSat developer and by the customer, before assuming that the traditional CubeSat approach is appropriate.

Implications of the CubeSat ICD

There are two significant implications of the CubeSat ICD that must be considered in selecting the CubeSat form factor as a baseline design for any mission. The first, and most obvious, is the strict limitation on size and mass. The second, and not as obvious, is the constraints on mission driven by the fact that all CubeSat launched (at least to date) have been as secondary payloads.

The impact of launching as a rideshare payload varies depending on the mission. For most educational missions, and many tech-demo missions, the orbital parameters are a secondary consideration. In some cases, particular orbits are desired, but a large orbital range will still satisfy the mission requirements. In such cases, there are likely enough launch opportunities to choose from that a launch opportunity satisfying the mission requirements will be available within a reasonable timeframe. In a few cases, a tech-demo mission may require a very specific orbit, in which case the wait for launch may be long.

Most operational missions, on the other hand, are likely to have more stringent orbital requirements. In this case, the opportunities for rideshare might be insufficient. One approach to this issue is to design the mission with the intention of building an ad-hoc constellation using quasi-random orbits⁴. This can be a very effective strategy for large constellations of satellites; subsets of the constellation are deployed from multiple launch vehicles going to orbits that are selected based on their relative value to the overall mission of the constellation. An alternative, if the constellation is large enough, is to purchase a dedicated launch. However, this works only if all the satellites of the constellation are in the same orbital plane, or if the program duration is long enough to allow for altitude changes and variable precession to populate different planes.

The more obvious implication of the CubeSat ICD is the constraint on mass, volume, power, etc., that derives from the requirement to launch within a small box. This constraint is well understood and, in most cases, can quickly be used to determine whether a given mission can be accomplished using a CubeSat. In general, many missions will be constrained by simple physics; one cannot squeeze a 2-meter telescope into a CubeSat. But one should be careful in applying the physics constraint to any given mission. For example, it is straight forward to demonstrate that a 3U CubeSat is not going to be able to provide ground imaging at 50-cm resolution. So if the mission planner starts by assuming that the mission is imaging at better than 50-cm resolution, then a 3U CubeSat is precluded by definition. But is the resolution really the mission? Most often the mission involves determining something about the ground being observed; land use, vegetation, cloud cover, water quality.... Does it really require 50-cm resolution? Can the actual mission be better served by lower-resolution, more frequent observations? This is something that must be resolved in defining the mission. The requirements are about the information to be obtained, not how it is obtained. The CubeSat paradigm applied to the problem may find an unexpected way to get the data.

CubeSats in deep space

CubeSats are often seen simply as deployable satellites, and there have been proposals that they can also be used at distant destinations in the solar system. However, the point of the CubeSat ICD was to simplify the interface between the CubeSat and the host vehicle so that CubeSats can be launch agnostic. This was required, in part, because the CubeSat would be flying as a secondary payload that would, in most cases, be unrelated to the primary payload. In going to deep space, both the CubeSat and the host vehicle are designed specifically for one application. Great care is normally taken in any deep space mission to ensure that every bit of available mass is put to the best use. A CubeSat with a standard deployer comes with a significant mass penalty (the CubeSat deployer can weigh as much or more than the CubeSat). Since both the primary mission and the CubeSat belong to the same customer, and since the cost of getting the mass to Mars or Venus, or any other distant destination, is so high, it makes sense to put extra effort into decreasing the mass of both the CubeSat and the deployer. Such an effort, if taken to the extreme appropriate for deep space missions, is likely to modify both the CubeSat and deployer so much that they will no longer resemble anything envisioned in the original CubeSat ICD.

Equally importantly, the CubeSat paradigm cannot readily be applied to deep space. The cost of any deep-space mission is high, so mission reliability is critical, even for the deployables. At the same time the mission frequency is low, limiting the opportunities for reflight. The CubeSat paradigm, calling for sequential redundancy and incremental improvements with a high tolerance for failure, is not an appropriate option. The fast-turnaround aspect of CubeSats is limited to near Earth, or, at most, Earth-lunar space. Flights to other planets are rare and necessarily long, so fast-turnaround is simply not possible. Without that, the CubeSat paradigm is not applicable. One possible exception to that would be an effort aimed at incremental development of deployable subsatellites where they are launched and tested in LEO to settle on a design for use on deep-space missions. It is not clear that this would present significant value over conventional methods of designing deployable subsatellites, but it would not be outside the scope of the CubeSat paradigm.

So, do CubeSats have a place in deep space? Certainly there is value in deployable subsatellites; this approach has been used a number of times in the past. However, with the high value placed on mass savings, a subsatellite built to the CubeSat ICD probably does not make sense. And with the high value placed on reliability, applying the CubeSat paradigm directly to the deep-space mission certainly does not make sense.

SUMMARY

The creation of the CubeSat standard has led to a very rapid proliferation of experimental satellites in the one-to-ten-kilogram mass range. Primarily these satellites have been developed for educational purposes, but their potential value in technology demonstrations and in operational missions has not gone unnoticed. In general, the approach to mission assurance taken by the CubeSat developer community has been much more relaxed than with traditional satellites. For educational and, to some extent, tech-demo projects this tolerance to risk is appropriate. The potential for high-frequency flight opportunities has led to the concept of serial redundancy in CubeSats. This is the recognition that a high risk tolerance for any given flight is acceptable if the flight is part of a series of flights aimed at incremental technology advances; for any given flight in the series the risk of failure is offset by the potential gains across the series as a whole. This concept of serial redundancy, part of the CubeSat paradigm, can lead to very rapid advances in the capabilities of space systems, but is appropriate only for programs involving multiple flights over an extended time period.

In contrast, technology demonstration missions that are not part of a series, but are instead simply taking advantage of the CubeSat form factor to enable a flight opportunity, are typically not as tolerant of risk as the flights that are part of a series. It is certainly appropriate for one-off missions to take advantage of CubeSat bus capabilities that have been developed as part of a sequential program, but the planning for a one-off mission should not baseline any systems that have not previously been proven in space.

Based in large part on risk tolerance, CubeSat development efforts should fall into either of two categories. In the first category are educational development efforts, and long-term development programs involving a sequence of increasingly sophisticated satellites. In these cases the risk tolerance can be high and the CubeSat paradigm can be applied. The second category includes one-off tech demo missions or operational missions where there are no plans for follow-on satellites. In these programs, the risk tolerance is likely to be much lower, and the CubeSat paradigm is not likely to be applicable.

-
1. Bob Twiggs, "Origin of CubeSat," in "Small Satellites, Past, Present, and Future," H. Helvajian and S. W. Janson, eds., 2008.
 2. The P-POD ICD is maintained at Cal Poly San Luis Obispo. The latest version can be obtained from: <http://cubesat.org/images/LaunchProviders/mkIII/p-pod%20mk%20iii%20icd.pdf>
 3. Greg Richardson, Kara Schmitt, Mary Covert, and Christa Rogers, "Small Satellite Trends 2009-2013," SSC15-VII-3, presented at SmallSat 2015.
 4. Joseph W. Gangestad, James R. Wilson, Kristin L. Gates, and John V. Langer, "Rideshare-Initiated Constellations: Future CubeSat Architectures with the Current Launch Manifest," Presented at the 31st Space Symposium, Colorado Springs, Colorado, April 2015