Results from the Planet Labs Flock Constellation

Christopher R. Boshuizen, James Mason, Pete Klupar, Shannon Spanhake
Planet Labs Inc.
490 2nd St, San Francisco, CA, 94105; 650-450-3456
chris@planet.com

ABSTRACT

In 2014 Planet Labs has – so far – launched two constellations of small satellites: Flock 1 comprising 28 satellites and 11 in Flock 1c. Additional launches are planned in the year, with Flock 1b scheduled for launch at the time of writing. With launch contracts already signed for more than 200 satellites and 43 already launched, when completed the Flock will be the largest constellation of satellites ever launched. A system of this scale would be extravagantly expensive if comprised of traditional satellite elements, yet we have attempted this with venture capital funding at a level that would be considered irrationally small by most aerospace standards. Thus, to accomplish our mission significant innovation is required in all segments of the operation. In this paper, we discuss the core engineering philosophies of Planet Labs, and the way in which our Flock is different than anything that precedes it.

INTRODUCTION

Planet Labs Inc., started in 2011 and based in San Francisco, was formed to execute a specific mission: to provide medium-to-high resolution imaging of the entire planet, on a daily, recurring basis. This will be undertaken by launching a constellation, which we call a “flock”, of some hundreds of satellites in a variety of orbits. Planet Labs is fully vertically integrated, operating all aspects of its business except launch. The company’s primary product offering is Earth imagery and imagery-derived data products.

The key differentiators of Planet’s dataset are one, complete coverage of the entire Earth’s land area, at a resolution not currently available (3-5m in our case, versus 1km for MODIS, and 15 m for Landsat), and two, daily revisit of this entire sampled land area. Figure 1 shows how Planet Labs dataset differs from those of incumbent operators, by plotting resolution against (our estimate) of the time it takes these systems to map the world once. The low-resolution, frequent update corner of this figure is represented by MODIS[1] and Landsat[2], while DigitalGlobe’s Worldview[3] system and Rapideye (now operated by Blackbridge[4]) are representative of high-resolutions systems whose correspondingly narrow field of view limits their ability to capture the entire Earth in a timely fashion.

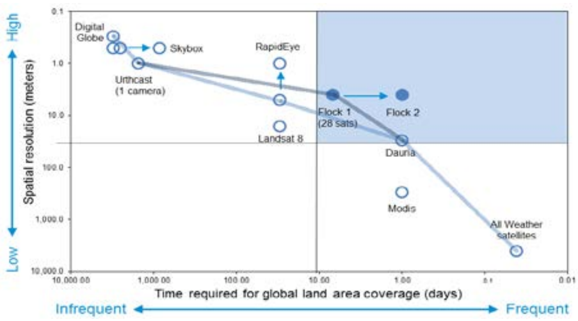

Figure 1: Planet Labs’ Flock revisit rate and resolution are compared against incumbent operators. The blue line shows the Pareto frontier prior to the launch of Flock 1a, and the grey line shows how Planet Labs’ system now defines the Pareto frontier in the high-update, high-resolution sector. The future planned capacity of some operators is shown, with Flock 2 continuing to dominate.

The Planet Flock stands out from the Pareto frontier[5] of this chart by providing a unique combination of resolution and full-Earth coverage repeat rate. Our data is different to what is currently available on the market, enabling use-cases that currently cannot be met by existing systems. Examples include the ability to monitor disaster situations and coordinate emergency response efforts, to empowering international development agencies to allocate resources more effectively, to detecting events in places that news and information cycles overlook or can't reach, to managing a global portfolio of construction assets and their development over time. Since we are monitoring the entire Earth, every day, we are able to detect world-wide events, and also provide insights on how they may be related to each other. Imaging the Earth and it's changes has not been done to this scale before, and we cannot predict all the value to be extracted from this high-resolution, daily time series dataset, but we are confident that to complete this mission, we have to do things differently.

FINDING A NEW WAY TO DO BUSINESS

There is a normal trade-off between spacecraft size and functionality, but advances in both miniaturization and integration technologies have alleviated the tension of that trade-off. In recent years small spacecraft have become more attractive due to lower development costs and shorter lead times. The established space industry is trapped in a tradition of requirements-driven approaches to design and testing that have resulted in spacecraft that have become increasingly unaffordable, in terms of cost, schedule and mass.

While Cubesats and other small spacecraft have been used in Earth orbit over the last 60 years, they have not been applied to commercial Earth imaging missions. Several recent advances have become available to enable Cubesats to be used for Earth observation missions. Advances in supply chain management along with technological innovations and commercial manufacturing practices have allowed Planet Labs to produce a remarkable capability in a Nanosat form factor. Furthermore, applying lessons from the software industry and a distributed risk approach have allowed us to build an overall constellation architecture that can scale. We talk about each of these in turn below.

A non-Aerospace supply chain

It is unfortunate, but the Dove satellites developed by Planet Labs contains no component directly sourced from the space industry. In our initial design trade studies, we considered a wide range of parts from a number of aerospace component and system suppliers, but all were rejected either on the basis of cost or performance, or both. We have succeeded instead by procuring all our components from purely Commercial, Off-the-Shelf (COTS) suppliers, and integrating and building own circuit boards. In that way, we look a lot more like a cell phone manufacturer than we do an aerospace company. (Many components were also rejected for the additional reason of being too large, which we discuss in the following section.)

The current consumer automotive and electronics industries has developed a deep and wide catalogue of COTS tools, test faculties and hardware that can be leveraged to solve the problems of Earth imaging. The same tools that perform heat transfer analysis on the latest Ford Diesel engine can also be applied to the problem of Earth observing Nanosats. The same electronic design tools used for the Playstation 4 are directly transferable to spaceflight hardware. In the Northern California area there are several electronic testing houses that are available to provide environmental tests for the consumer electronic industry that happen to be very similar to those required by NASA. Normally these houses are testing the latest video game console or other electronic gadget but they are available to test space flight hardware at a reasonable cost and schedule. The relentless drive by the consumer electronic industry to reduce price and schedule has created an available component catalogue that is extraordinary. Today there is no need to develop a specialized ASIC, as nearly all electronic functionality is available in catalogue parts. The logical functionality that is not available in catalogue parts is easily reproducible in FPGAs.

Today the lack of power is the single largest reason that Cubesats continue to be viewed as mere toys or university student training devices. Spacecraft power systems use solar cell designs that are heavy, expensive and bulky. Traditional Cubesats have extremely limited capabilities with regard to payload (sensors) and spacecraft systems. Planet has developed innovations in deployable solar arrays without cover glass that enable these vehicles to develop significantly more power that previous vehicles at a lower cost and mass than ever before.

Many technologies have benefited from advances in commercial industries. Today, C&DH systems have greater processing capability with lower mass, power and volume requirements. This general trend is enabling small spacecraft to tackle a broader range of missions. Power and reliability, traditionally the primary limiting factors, have seen significant advances due to the infusion of commercial technology and higher risk tolerance of small spacecraft.

Current satellite communication transmission strategies use VHF, UHF, and microwave frequencies. Selecting a frequency spectrum depends on a number of factors including expected data throughput, available power and mass, and licensing issues. Due to these reasons, technology development is still underway on all of these frequency spectra, and there is no "one-size fits all" solution for space-borne radio needs. Current technology shows a trend of increase carrier signal frequency and the power and mass requirements of the transmitter. Using transmitter technology appropriate for small satellites in LEO, UHF/VHF transmitters have a maximum data transfer rate of around 38 kbps, S-band transmitters have a maximum data transfer rate around 10Mbps and X-band transmitters around 100Mbps. It was conventional wisdom until very recently that X-band radios would not fit in a nanosatellite. However great advances have been made in the cell phone industry with regards to baseband components, mixers and amplifiers and software defined radios, such that we discovered that all the components for an X-Band radio are available via catalogue parts. This amazing improvement has lowered the cost of a high-performance space radio by many orders of magnitude. This is certainly an enabling technology.

Building spacecraft with lessons from the consumer electronics industry

Planet has invested significantly in the packaging and miniaturization of satellite capability, drawing many lessons from other industries including tablet PCs and smart phones. Our goal was to compress most of the capability of a traditional small satellite into the volume of a cubesat. Small satellites have, since the early 1980's, lowered the cost of access to space and allowed increased capability and responsiveness, but a typical smallsat is inefficient in many ways. For example, they are volumetrically inefficient, with the inner volume being mostly empty space and not useful capability. For optical satellites, especially those with large, traditional apertures, the satellite volume scales with the cube of the aperture diameter, which is a severe cost penalty for a scaleable, many-satellite system. The number of satellites that can fit in a launch vehicle – and so the cost of launch – thereby goes with the cube-root of the aperture. Without significant innovation in alternate optical apertures and configurations, an ultra-large scale system of sub-meter resolution satellites would be staggeringly expensive, which is why we have not seen these yet. Larger satellites are also less stiff than their small counterparts, and so they are inefficient from a mass perspective as well, being comprised of heavier structural elements and significant cabling harnesses. We have taken on this scaling challenge by trying to create an ultra-high density microsatellite.

The telescope in the Dove satellite is a Maksutov-Cassegrain design with a 91mm aperture, a tube length of 20cm and it is approximately 2.5 U (2.5 litres), which, in a cubesat, leaves little room for the rest of the satellite bus. Furthermore, the focal length (1.14m) is such that the focal array is placed approximately 32cm back from the telescope aperture, meaning the detector, an industrial CCD in our case, is at the very rear of the satellite. The electronics for the camera do not fit inside the cubesat volume proper, and so instead is housed in a "tuna can" addition to the rear of the satellite, an innovation introduced by NASA Ames in the Genesat[6] and O/OREOS satellites[7]. We needed to keep a path open for light to reach the detector. This means that the rear third of the satellite must have an narrow optical tube in the center, and so the bus has ended up being a wrap-around design of total volume of about one-quarter of a Unit.

This is not very much space! Our mission goals dictated the need for a full attitude control system, large amounts of computing capacity and data storage, and high-bandwidth communications, none of which could really be removed without sacrificing the mission itself. Thus, we needed a way to miniaturize and package all of these things in the available space. To do this, we drew lessons from the laptop PC, smartphone and tablet industries. Apart from some necessarily large components (reaction wheels) the satellite looks very much like one of these modern consumer devices: the printed circuit boards (PCBs) have a large number of layers, there is an extremely high density of surface mounted (SMT) electronics components, and almost no internal cabling.

Some small spacecraft are assembled and integrated with the same rigor as their larger counterparts, while others are integrated within a university laboratory. Effectively integrating, or sharing resources, between individual components can substantially increase the system's functionality and density, thereby reducing unnecessary mass and volume. There is a trend towards further miniaturization and higher levels of integration. Hyper-integration of hardware and software also promises to bring benefit.

The automotive and electronics industries are moving towards levels of hyper-integration or shared resources that are uncommon in the space industry. In this model many of the common resources are shared, from common power supplies to processors to components. This has created a situation where the latest computer the Surface Pro 3 has no serviceable components. The package is literally glued together. Planet Labs has adopted this approach where there are no boxes. Everything from the instrument to the radios shares the same resources from power supplies to FPGAs. This has created a high-density package that enables the small spacecraft approach.

Another revolution is the use of miniaturization technology in microelectromechanical systems (MEMS), i.e. apparatuses with microscale (µm) features. MEMS-based devices provide higher accuracy and lower power consumption compared to conventional spacecraft systems. In addition to dramatically (100-1000X) lower size and mass, MEMs are 100 X less expensive than traditional spacecraft subsystems. A key result of the combination of these factors is that spacecraft can be designed that combine multiple copies or duplicates of the same components so that individual units can degrade or fail while the spacecraft continues to function with limited system level impact.

These approaches have allowed us to create a very dense, high-capability spacecraft bus. For attitude determination, we have a star camera, GPS, a photodiode array for coarse-sun sensing (augmented by solar panel current sensors), and MEMS IMU, and for control we have four reaction wheels in a tetrahedral arrangement as well as three air-core magnetorquers for initial spin damping, coarse control, safemode-attitude recovery, and general momentum dumping. For radios, we have a 2-way UHF transceiver, and high-speed X-band downlink transmitter, S-band uplink receiver. The main flight computer is a modern, low-power x86 processor with 0.5 terabytes of solid state storage and runs the Ubuntu server operating system. None of this would have fit in a 0.5 U volume using currently available aerospace industry parts.

This dense packaging has an additional benefit: the small scale-size of the SMT components and resultant PCBs means that the fundamental modes of the spacecraft are very high, outside the regime of concern for any launch vehicle, and the overall stiffness is high. This makes sense, as cell phones, for instance, are designed to survive being dropped on marble floors, which is a similar shock impulse to the stage separation event on the Dnepr launch vehicle. Small parts make tough systems!

Building Hardware as if it were Software

In 1993, Bill Clinton's Secretary of Defense, William Perry held a meeting of aerospace execs, in which he gave the option to consolidate or wither because the government would no longer support the industry.[8] As a result the industry was distilled into a few primary contractors with high capital costs and became a shining example of a monopsonistic market, one with a single buyer and few sellers. The consolidation was further shaped by the convergence of three other events that would greatly determine the industry's capacity to innovate: 1) the cutting of costs resulted in hiring freezes and reduction of junior staff, 2) the Silicon Valley dot com boom was pulling talent away from aerospace, and 3) the founding generation of engineers began retiring. This demographic peculiarity has resulted in a monolithic industry, stifled by capital barriers, cultural homogeneity and an aversion to risk.

At Planet Labs, we work to break this cycle by embracing a more risk tolerant approach that takes lessons from the software industry: releasing early and often, rapidly iterating, and innovating to stay ahead of a fast-growing market. This has largely been enabled by Moore's Law trends in consumer electronics, which as mentioned above, has made complex and powerful components smaller and more cheaply available. The result has been more FLOPS-per-Watt-per-dollar, meaning that more power can be packed into reduced space for lower costs. The attraction is that smaller busses enable smaller payloads, which reduce the complexity, team sizes, materials, costs and development timelines.

One of the more popular software team management adages is "release early, release often", which speaks to the benefit of small iterative steps, and incorporating early customer feedback. In the context of space, this can only mean one thing, which is frequent deployment into space – putting a literal interpretation on the idea of "launching early. A single iteration of a Dove satellite takes 8-12 weeks from design to manufacture, and we have produced a total of 10 full design iterations with the production of the corresponding flight hardware in the past three years. Contrasting this with years or decades as is the case in other parts of the aerospace industry, we are able to ward off obsolescence, and it enables us to continuously integrate the latest technology improvements into our design. Flying a large number of satellite variations – including things such as alternate solar panel encapsulants, battery protection systems, and imaging sensor configurations – also allowed us to explore another software tool, "A/B testing." A/B testing is primarily used in the design of user interfaces for software products, by putting either an A or B variant in front of two different user groups, and monitoring how each is used. This "in the field" testing allows useful designs to be tested and validated quickly and cheaply, and most importantly, helps evolve toward the best possible design with a minimum of time-consuming engineering analysis.

This approach is highly informed by agile development methods, which is most recognized in Silicon Valley as an approach to software development. We adopt the following values outlined in the Agile Manifesto[9] and work to extend their application beyond software to the design and development of our hardware:

1. Individuals and interactions over processes and tools
2. Working software over comprehensive documentation
3. Customer collaboration over contract negotiation
4. Responding to change over following a plan

We will discuss each of these below.

Individuals and interactions over processes and tools

In agile development, self-organization and motivation are important. We have launched a total of 71 satellites, spanning 8 builds (builds 9 and 10 are yet to launch) and an engineering team of under 30 for the majority of this time. This has been possible because we have a highly motivated team whose work is not organized around a hierarchical distribution of tasks pulled from a plan, but rather, the work is self-organized around learnings from feedback loops and adaptation cycles. Learning is the primary goal and success metric of our missions. This is similar to the Silicon Valley approach of releasing a Minimum Viable Product (MVP)[10], which is considered the leanest possible product that can be built and released to test assumptions about the products future feature set and market needs.

For the Flock 1 mission, our team organized, selected/built tools and created processes around a narrow set of goals. Flock 1 was our first demo constellation after launching 4 demo sats. The Flock 1 mission aimed to: 1) learn to operate a large constellation of satellites, 2) test the range of performance of our latest hardware upgrades, and 3) optimize data downlink under realistic conditions. We do not make long term commitments to the tools or processes designed for this specific mission, but re-evaluate based on the goals and learnings of each mission. The experience gained is invaluable however, and many hundreds of bug fixes and software features have been rolled out between Flock 1 and 1c. This allows us to flexibly scale within a resource-constrained environment, much like any other software start up.

Working software over comprehensive documentation

Most start-ups have limited resources and are vulnerable to market shifts, this forces us to prioritize building over documenting. Our development sprints have produced 10 builds in 36 months. In both the software and the hardware, we have focused on building and testing the hardware in space rather than proving everything works on the ground and exhaustively tracing requirements and compliance. This is much like releasing a software MVP for beta-testing, then iterating on the feedback. Further, the continuous integration of new technologies and learnings is not easily documented in the traditional aerospace way at the pace we are developing. This is not to say there is no documentation in agile development, rather it is out of necessity or what is often referred to as Just Barely Good Enough (JBGE) by Scott Ambler. Modern software development environments allow the automatic creation of documentation for well-written software, and we make extensive use of this in the company: well-written code is its own documentation.

Customer collaboration over contract negotiation

Early customer engagement and collaboration reduces the time required to learn key lessons, and makes time to fix problems in subsequent flock launches, which is not a luxury most other satellite companies have. This decreases risk in uncertain or emerging markets and the company can build what the technology allows or what the market is demanding. In our case, a gap in the market called for a service to complement Earth imagery of *anywhere* on Earth, with persistent Earth imagery of *everywhere* on Earth, an approach that trades spatial resolution for temporal resolution. User-centered design is an approach in agile development that transforms end-users from simple consumers of a product into producers, by integrating their needs and knowledge of the market into the design process from the beginning. We integrate customers into our development process and this is more cost-effective than investing in contract negotiation.

Responding to change over following a plan

At the core of this value is the ability for risk to be distributed throughout the system and across stakeholders – thus modifying the traditional approach to systems engineering slightly. Instead of spending the effort to develop an exquisite system a more moderate approach has been adopted. This approach is based upon a combination of demand from the market and capabilities available in hardware and software versus the art of the possible. This approach tends to develop a system that is just barely capable of achieving the systems goals with no margin. Then during systems tests and into on orbit test as deficiencies become evident only the actual short falls that affect the system are corrected. It is our experience this approach saves considerable effort in not fixing those items that are not in the critical path of delivering the product. The core of this approach is feature requests and bug fixes, not Gantt charts and delinquency reports. This does require that all the customers have bought into this approach and are willing to work through the process. It is our experience that this approach of launching with known bugs provides the most agile and cost effective approach.

Distributed risk

Being part of a burgeoning competitive market, forces us to be adaptable and nimble. This is best illustrated by how we deal with redundancy. Rather than building redundant systems into a single satellite, which requires a larger team and more materials, thus higher costs, we build redundancy into the overall constellation. This puts the cost burden on production rather than R&D, which is proportionately less expensive. We plan for satellites in our constellation to occasionally face episodic power brownouts, latch-ups, or single-event upsets, and we deal with that on a system level by launching large numbers of them. This provides huge redundancy and reduced development time, and overall reduction in risk. It also has the additional benefit of allowing us to A/B test different satellite variants across the constellation and quickly respond to market shifts.

Additionally, being adaptable to change has enabled our satellites amortize over shorter lifetimes, which makes deploying them into lower orbits economically justifiable. Earth observation payloads typically require higher orbits and larger apertures, which makes larger satellites, bigger teams, higher costs (the vicious cycle). But smaller satellites, even if shorter lived, have a business case and can achieve similar resolutions with smaller optics (the virtuous cycle).

Ultimately, agile development has enabled a risk posture that embraces learning from failure, experimentation and rapid iteration. Doing business this way will help us to scale to the needs of a volatile market, ultimately diversifying the aerospace industry, while simultaneously attracting the talent that can innovate and disrupt the industry.

IN-ORBIT PERFORMANCE

Flock 1(a)

Flock 1 was launched to the ISS 9 January 2014 as an internal technology demonstration of the Flock constellation architecture. The 28 satellites were deployed over 3 weeks in February 2014 and it has provided an invaluable platform for technology iteration and for developing the infrastructure required to operate a large constellation of satellites. Due to our aggressive development lifecycle and cost, we are in a unique position to launch real demos instead of performing extensive ground-based analysis. This allows us to rapidly gather real flight data to qualify design variants and to test ground based systems under operational conditions. Below are some of the highlights of the mission and key lessons learned.

The majority of Flock 1 satellites were RGB imaging systems, but we included 5 variants of our imaging system containing different optical spectral bands and camera firmware, which informed filter selection for future design revisions. Images taken from these off-nominal satellites were given to early customers and partners, in a great example of an MVP. We flight qualified a number of systems that were not tested by Doves 1 and 2, including GPS, star camera, Si based solar cells.

The Flock was inserted into orbit from the ISS via the Nanoracks deployment system. Being the first constellation deployment from the ISS posed a great number of technical, operational, and regulatory challenges. By committing ourselves to addressing these early we have now qualified many of our systems and streamlined this process for ourselves (and others). The low space-station like orbits provided challenging conditions under which to acquire and commission satellites. This has resulted in us developing a number of novel tracking and ranging techniques that allow us range all of our satellites in a single pass. These techniques were used to successfully contact all 11 of the recently launched Flock 1c satellites on their very first ground station pass.

We avoided the typically arduous process of developing a comprehensive thermal model in design of the Flock 1 satellites, and instead relied on engineering best practices and basic energy budget calculations. Flock 1 provided excellent data on thermal gradients and their effect on the satellite subsystems, and this has resulted in a number of passive thermal management improvements on future designs. This *in situ* information is clearly better than any on-the-ground simulation, and is a great use of “build a little, test a little”.

We performed a number of differential drag experiments to quantify the phasing control authority and the ADCS performance while in a high drag 3-axis attitude mode. Differential drag of this kind was first proposed by the Aerospace Corporation in 1989[11] and tested for cubesats in 2012[12]. We were able to successfully replicate these results across all 28 satellites, dispersing them along-track by 360 degrees from front to back in 35 days. The differential drag algorithms developed using this data will be used for phasing and station keeping in future flocks. We also found that the canonical coefficient of drag (Cd=2.2) for CubeSats underestimates reality by up to 50% at very low orbits, and we have now invested into optimizing operations to extend orbit lifetime.

Another problem with very low orbital altitudes that the orbits provided by JSpOC-generated TLEs are highly variable and unreliable for clusters of satellites at sub-ISS altitudes. Having a ranging (or some other orbit determination) capability is critical for quick acquisition and identification of satellites. Orbits inclined like the ISS orbit precess at a rapid rate, and are traditionally not used for Earth observing missions for power and scene lighting conditions. This drove us to develop a robust concept of operations to work within these conditions and allowed us to evaluate such orbits for their relative value compared to sun-synchronous.

Perhaps most importantly, we confronted operating a constellation of 28 satellites in very challenging conditions. This led to us building custom mission control software for pass prediction/deconfliction, satellite task queuing, and telemetry monitoring. Our small operations team gained invaluable experience in anomaly detection, debugging, automating operations and in performing real time operations trades. This experience cannot be taught or learned through ground testing, and is something we can leverage as we continue to launch record numbers of satellites.

SUMMARY

Planet just recently launched Flock 1c, bringing the total number of launched satellites to 43 in just 15 months. We have contacted 42 out of 43 satellites launched (the 43rd was never deployed). Our unusually rapid design lifecycle is enabled by a new agile aerospace philosophy, which leverages innovations in the software and consumer electronics industries to get a higher return on investment. We are participating in a fast-emerging market, which demands we embrace risk as an inherent part of our engineering approach. This has allowed us to learn more lessons faster than would be possible with analysis and ground testing. This approach has been repeatedly validated by numerous technology demonstrations that have provided on-orbit data and operational experience that is unique for such a young team. Perhaps the most important Planet innovation is breaking the schedule and cost-cycles that the space industry is infamous for. Without innovation in this area, an ultra-large constellation would be impossible. With more than 150 additional satellites contracted for launch we are eager to continue pushing the envelope of agile aerospace and to provide "everyone, everywhere, everyday" products.

REFERENCES

1. http://modis.gsfc.nasa.gov/
2. http://en.wikipedia.org/wiki/Landsat_program
3. http://www.digitalglobe.com/about-us/content-collection#worldview-2
4. Blackbridge/Rapideye http://www.blackbridge.com/rapideye/
5. Pareto analysis, http://en.wikipedia.org/wiki/Pareto_analysis
6. http://www.nasa.gov/centers/ames/missions/2007/genesat1.html
7. http://www.nasa.gov/mission_pages/smallsats/ooreos/main/
8. N. Augustine, Augustine's Laws, 6th edition, AIAA, 1997, ISBN 1563472406
9. http://www.dau.mil/pubscats/pubscats/AR%20Journal/arq2001/Deutch.pdf
10. http://agilemanifesto.org/
11. http://en.wikipedia.org/wiki/Minimum_viable_product
12. Orbital formationkeeping with differential drag, Leonard et al., Journal of Guidance, Control, and Dynamics, Vol. 12, No. 1 (1989), pp. 108-113. doi: 10.2514/3.20374
13. Operations, Orbit Determination, and formation Control of the AeroCube-4 Cubesats, Gengestad et al, SSC13-X-4, Small Spacecraft Conference, 2013