A CubeSat-Based Optical Communication Network for Low Earth Orbit

Richard Welle, Alexander Utter, Todd Rose, Jerry Fuller, Kristin Gates, Benjamin Oakes, and Siegfried Janson
The Aerospace Corporation
PO Box 92957
Los Angeles, CA, USA, 90009
welle@aero.org

ABSTRACT

A common challenge in developing small-satellite-based Earth-observation (EO) missions is getting the data to the ground. Typically, most satellites download data via space-to-ground radio-frequency (RF) links, communicating directly with fixed ground stations as the satellites fly within range. For many next-generation EO missions, such as hyperspectral imaging or SAR missions, the volume of data generated is large enough to tax most conventional RF downlink systems. An alternative approach to the problem is to develop a network of small optical relay satellites in LEO that allow for short range optical or RF communication from EO satellites to the network. The LEO network satellites then relay the data around the Earth to a network satellite in view of an optical ground station. Implementing such a system requires the development of both optical downlinks and optical crosslinks for small, preferably CubeSat-scale satellite. The key challenge in implementing a high-rate optical communication system is the pointing and tracking of the laser beam. Most free-space laser communication systems incorporate a complex two-axis gimbal to control beam pointing. An alternative approach is to hard mount the laser transmitter to the satellite body and point the laser solely with the attitude control system of the spacecraft. For real-time communications through a constellation, a means will be necessary to allow each node to point simultaneously at both the source and destination. This paper describes two alternatives to the two-axis gimbal for this application. In one approach, beam steering is accomplished with a single-axis gimbal, combined with rotation of the satellite about the receive axis. In the other solution, requiring no gimbals at all, the node consists of two satellites flying in close proximity, with one satellite acting as the receiver and the other as a transmitter, and with a short-range omnidirectional link between them.

BACKGROUND

A common challenge in developing small-satellite-based Earth-observation (EO) missions is getting the data to the ground. For many next-generation EO missions, such as hyperspectral imaging or SAR missions, the volume of data generated is large enough to tax most conventional radio-frequency (RF) downlink systems.[1,2] For other EO missions, such as hazard monitoring or data collection for weather forecasting, data latency is a key issue.[3] It is well known that laser downlinks offer the potential of multi-gigabit-per-second download speeds,[4] which are typically two to three orders of magnitude faster than RF downlink speeds. On the other hand, optical downlink systems can easily be interrupted at any given ground station (possibly for long periods) due to cloud cover.

The European Data Relay System is being developed to address both the data-latency issue and the cloud-cover issue by placing a set of data relay satellites in geosynchronous Earth orbit (GEO). These satellites receive optical communication from EO satellites in low Earth orbit (LEO) and relay those communications to the ground using a Ka-band RF downlink.[5] Because of the range from LEO to GEO, however, this system requires that LEO satellites carry a large (~35 kg) laser terminal.

An alternative approach to the problem is to develop a network of small optical relay satellites in LEO that allow for short range optical or RF communication from EO satellites to the network. The LEO network satellites then relay the data around the Earth to a network satellite in view of an optical ground station not obscured by clouds. This approach allows for high data rates and low latency, and can be implemented, in principle, with a much smaller optical terminal on the EO satellite.

Implementing such a system requires the development of both optical downlinks and optical crosslinks for a small, preferably CubeSat-scale, satellite. The first optical downlink developed for a CubeSat is the NASA Optical Communication and Sensor Demonstration (OCSD) program. These satellites, developed by The Aerospace Corporation and currently scheduled for launch in late 2017, have been designed for optical downlink rates up to 200 Mb/s. Follow-on optical downlink demonstrators are being developed by Aerospace for rates approaching 1 Gb/s. In addition, we have developed concepts for CubeSat-based optical crosslink nodes that would support the eventual deployment of a LEO network of optical relay satellites to enable high-volume, low-latency downlink from new Earth-observation satellites. The availability of such a system would also enable new capabilities for small EO spacecraft with limited mass and power budgets for downlink, further supporting next-generation distributed EO satellite sensing systems.

Recent progress in sensor technology has allowed low-Earth-orbit (LEO) satellites to shrink significantly in size, disrupting a legacy industry where traditional satellites cost $500 million to $1 billion to build and launch. Major investments are being made to address the new opportunities this provides for data collection, and many companies are launching nanosatellites and/or microsatellites into LEO to capture this opportunity. The rapidly expanding satellite infrastructure is generating vast amounts of data, with no signs that the trend will level off. To bring all that data down from LEO requires an average communication rate of several Gb/s, continuously, and that demand will continue to grow.

Typically, most satellites download data via space-to-ground radio-frequency (RF) links, communicating directly with fixed ground stations as the satellites fly within range. The current ground station infrastructure has several key limitations that present significant challenges as the satellite industry continues to grow. Satellite-to-ground communications are "line-of-sight," meaning that ground stations are able to receive data directly only from satellites that are above the local horizon. The duration of a satellite pass over a ground station depends on the altitude of the satellite and the distance between the ground station and the ground track of the satellite. With satellites in LEO, the maximum pass duration is typically less than ten minutes.

The frequency of passes is strongly dependent on the satellite orbit parameters and the location of the ground station. For example, a satellite in equatorial orbit will pass over an equatorial ground station on each orbit. With a typical orbital period of 90 minutes, that means 16 passes per day. Similarly, a satellite in a polar orbit will pass over a ground station located at the North Pole once per orbit. On the other hand, the satellite in polar orbit will pass over the equatorial ground station between two and four times per day depending on the alignment of the ground track with the location of the ground station. However, the satellite in equatorial orbit will never pass over the polar ground station. Most LEO satellites are in orbits at some inclination between equatorial and polar, and most ground stations are located at latitudes well south of the North Pole. As such, the pass frequency for any given satellite over any given ground location will typically be three to five times per day for ground stations that are not at high latitude (above about 60 degrees) and not at latitudes higher than the orbital inclination of the satellite.

The consequence of limitations on pass duration and frequency is that a satellite will be within communication range of a given ground station for no more than 10% of a day, and typically for less than 2% of a day. These constraints on pass duration and pass frequency are driven by orbital dynamics and can be overcome only by increasing the number of ground stations or locating the ground stations at very high latitudes. Avoiding downlink constraints requires a large number of geographically diverse ground stations that are inherently underutilized.

One method of compensating for the limitations on ground contact time is to increase the data transmission rate during what contact time is available. High data rates in the RF require some combination of high transmitter power and high-gain antennas on the satellite and the ground station. High power transmitters and high-gain antennas on the space segment are constrained by power and mass limitations on the satellite. High-gain antennas on the ground are not mass limited, but tend to be very large (10 meters or more in diameter) and require significant capital investment.

Significant further increases in downlink capability can be obtained by developing an in-space relay network configured to allow continuous communications from space to ground, as illustrated in figure 1. With such a network in place, an EO satellite in LEO with data to download could simply transfer that data to a nearby node in the network. From there it would be forwarded through crosslinks to a network node within view of an available ground station. Such a network could allow high-volume, low-latency download from anywhere in LEO.

For new satellite companies leveraging advances in satellite capabilities, capital investment for an extended ground station network is particularly burdensome because the size and cost of the ground network does not scale with the size of the satellites. Ground station costs have not scaled at the same rate as satellite costs, requiring significant further investment to match growth in satellite capacity. An available LEO network could minimize, or even eliminate, the need for new satellite companies to develop their own ground network.

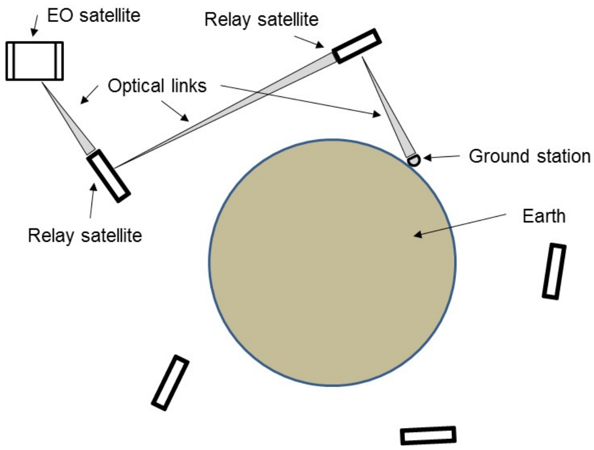

Figure 1. Schematic illustration of a LEO optical network.

OPTICAL COMMUNICATIONS

The key challenge in implementing a high-rate optical communication system is the pointing and tracking of the laser beam. Laser communication achieves high data rates with moderate powers by focusing the transmitted power into a narrow beam, which must be directed with sufficient accuracy and precision to ensure that the intended receiver is reasonably centered in the beam profile.

The importance of pointing is illustrated in figure 2, which shows the data rate achievable as a function of transmit pointing accuracy for three different ranges. Besides pointing accuracy, the other factors that affect data rate are transmit power, the size and efficiency of the collection optics and receiver, and, of course, the range to the target. Ultimately, the goal is to get as many photons as possible into the collector, and to convert those photons to data bits. For a given laser power and a given detector system, the key factors are the beam divergence (which must be wide enough to compensate for uncertainty in pointing) and the range to the target. The data rate is approximately proportional to the optical power falling on the collector. As such, the data rate will be inversely proportional to the square of the beam divergence (which is limited by the pointing accuracy), and also inversely proportional to the square of the distance to the target.

The three cases illustrated in figure 2 cover three different ranges and assume a constant laser power, constant collector area, and constant detector efficiency. The shortest range, 1,000 km, is typical of LEO-to-ground links. The intermediate range, 5,000 km, is approximately the longest range possible with a crosslink in LEO. The longest range shown, 40,000 km, is typical of a link from LEO to GEO. In this simplified illustration, the effect of atmospheric distortions on the downlink have been neglected. To compensate for this, the undistorted optical power directed at the terrestrial receiver would have to be increased, perhaps by as much as an order of magnitude. This will shift the blue trace to the left, but not beyond the green trace. Furthermore, for ground-based receivers, there is the option of using larger collection optics than may be practical on space-based receivers. There are also options for active compensation for atmospheric distortions. As such, a LEO transmitter with sufficient pointing capacity for a 5000 km crosslink should have adequate pointing capacity for a 1000 km downlink at the same data rate.

Optical Communication and Sensors Demonstration

Most free-space laser communication systems incorporate a complex two-axis gimbal to control beam pointing. Such gimbals are typically too massive to consider their use on small satellites, and particularly on CubeSats. The NASA-sponsored Optical Communication and Sensor Demonstration (OCSD)[7,10] is a CubeSat laser-communication demonstration mission that takes a different approach by hard mounting the laser transmitter to the satellite body and pointing the laser solely with the attitude control system (ACS) of the spacecraft. Obviously this simplifies the construction of the transmitter, but it does place a burden on the ACS. On the other hand, continuing improvements in CubeSat-scale attitude control systems have led to the development of commercially-available systems that could support very useful data rates in body-mounted optical communication systems.

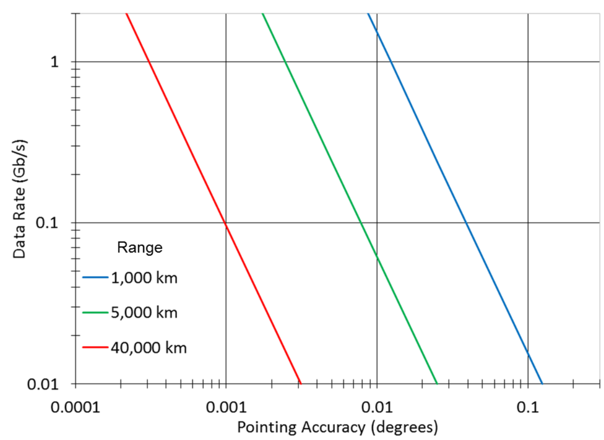

Figure 2. Data rate capacity of an optical link from a 4-W laser to a 10-cm-diameter receiver as a function of range and transmitter pointing accuracy.

The OCSD program included the development of two flight units and an engineering model. Fabrication and testing of the flight units was completed in mid 2016, and they were slated to fly on a Falcon 9 in October 2016 to be deployed by the SHERPA mission. This mission was delayed and ultimately canceled as a result of the Falcon 9 launch pad fire in September 2016, and the OCSD flight units are now scheduled for launch in late 2017. The OCSD engineering model was flown as a risk reduction effort in 2015.[6,7] This flight was unable to achieve all its flight objectives because of a software anomaly in the ACS, but it continues to provide a useful testbed for a number of flight systems that were first flight items, including a software-defined radio, and star cameras needed for the flight units.

Although the OCSD flight units have yet to fly, the experience gained with flying the engineering model, and with the build of the flight units, combined with the expected gain in downlink capacity generated by the optical communication system, has led to a preference for including laser communication where appropriate on future CubeSat missions being developed by Aerospace. We are currently in the final stages of building R-Cubed - a 3U CubeSat testbed for optical imaging systems that will include a communication laser to get the expected image data to the ground.[11]

The key development that is enabling CubeSat-based laser communication systems is the rapid advance in the capabilities of CubeSat attitude-control systems. NASA published data[8] on trends in CubeSat pointing capabilities in 2014 that indicated pointing accuracies on the order of 1 degree should be possible by 2017. Instead, however, the trend is much steeper than that, and there are already commercially-available CubeSat attitude-control systems advertising accuracies on the order of millidegrees. As can be seen in figure 2, pointing accuracies in the millidegree range are sufficient for gigabit rates for LEO crosslinks and LEO-to-ground applications.

CONSTELLATIONS

Even though optical communication systems are capable of very high data rates, they are limited by the availability of optical ground stations, and by the inability of optical communication signals to penetrate clouds, which further limits the utility of existing optical ground stations. Looking to the future, these problems can be addressed by using orbital optical relay systems to move data from the point where it is generated to a location where there is clear access to a ground station. Ultimately, it will be possible to establish an all-optical backbone in space that would allow EO satellites to download data not by transmitting it directly to the ground, but by transmitting it to a node in the optical network, as illustrated in figure 1. A single constellation would be able to provide downlink services for a number of client satellites, or even provide a means of moving data optically from point to point on the ground. With a sufficient number of nodes in the network, it will reach a point where there is always a node within range of any given LEO satellite, as well as at least one node within range of an available ground station. This will enable download of data generated in LEO in essentially real time, with little to no latency beyond that driven by the speed of light.

The number of nodes that would be required in such a constellation to ensure that one is always visible to a client satellite depends on the altitude at which the constellation is flying. At one extreme, the European Data Relay System places optical relay satellites in geosynchronous Earth orbit (GEO).[5] In GEO, only three relay satellites are required to cover the entire globe, as well as all of LEO. On the other hand, by placing the relay satellites in GEO, any client satellites in LEO have to satisfy the pointing requirements for a 40,000 km link, which can be very challenging, as illustrated in figure 2.

An alternative is to put the relay constellation in a relatively low orbit. This will ensure that the crosslink range can be kept short in comparison to the GEO distance, but will lead to a requirement for a large number of satellites in order to provide full coverage. Assuming, for simplicity, that all the relay satellites are at the same altitude, the maximum separation between relay satellites is constrained by the requirement to keep the crosslink beams above the ground, and preferably above any part of the atmosphere that might cause distortion of the beam. Figure 3 shows the maximum separation range between satellites as a function of the orbit altitude and the minimum tangent altitude of the crosslink beam as it passes above the atmosphere.

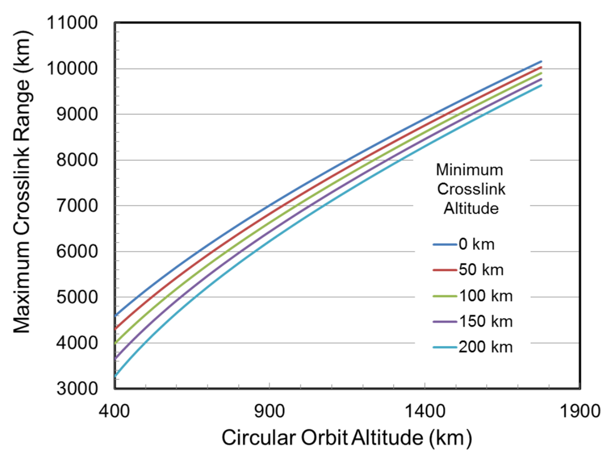

Figure 3. Maximum possible crosslink range as a function of orbit altitude and minimum tangent height.

While there are a number of factors that should be considered in selecting an orbit altitude, there are benefits of keeping it low, with the limitation that atmospheric drag should not be so high as to limit the lifetime unnecessarily. This can be satisfied with an orbit altitude above about 500 km. At the same time, keeping it below about 650 km would ensure that the satellites can meet the 25-year deorbit rule without active deorbit devices. For an orbit altitude of 600 km, and assuming a minimum crosslink tangent altitude of 100 km, the maximum possible crosslink range is about 5200 km. To fill a single orbital plane at 600 km with evenly spaced satellites without exceeding this crosslink range would require nine satellites, and they would be spaced at about 4800 km. Depending on details of constellation geometry, a complete constellation at this altitude would require something between 50 and 100 satellites. Using a higher orbit, say in the 1500 km range, would allow the satellite numbers to be reduced by a factor of two to three, but satellites at this altitude would necessarily be far more complex in that the crosslink range would be longer by about a factor of two, there would be a need for an active deorbit capability, and the radiation environment is much more demanding. Furthermore, opportunities for rideshare to this altitude are rare, so the launch costs would also likely be higher.

Nodes

In the simplest, conceptual, version, a node in a communications constellation must be able to take in data from a source and then re-transmit it to a destination (which may be another node in the constellation, or a point on the ground). If the intention is for the node to operate in real time, with the data being transmitted as it is being received, then the node must be able to point simultaneously in two directions; the receiver portion of the node must point at the data source and the transmitter portion of the node must point at the data destination. Since the source, the node, and, possibly, the destination are all in orbit, they will be moving relative to one another, and the node will have to track both source and destination with time.

It is possible to avoid the need for simultaneous pointing at both the source and the destination if the node acts in a store-and-forward mode, receiving and storing a message while tracking a source, and then adjusting the pointing to track the destination while the message is passed down the line. While this is not true real-time communication, it can be closer to real time than systems that simply keep the data on board until a ground station comes into view. Depending on the data volume, a typical message duration may be only a few seconds long, and will not likely be longer than a few minutes simply because orbital dynamics will limit contact durations. If it also takes a few minutes to redirect the node from receive pointing to transmit pointing, then the total message transit time is still no more than a few minutes. Even if the full path to an available ground station involves multiple transfers through nodes, it still should be possible to get any data down from orbit within an hour.

For true real-time communications, store-and-forward is not an option, and a means will be necessary to allow the node to point simultaneously at both the source and destination. The obvious solution is a two-axis gimbal, but these are, as noted above, not readily available for CubeSat-scale systems. Aerospace has developed two alternatives to the two-axis gimbal for this application. In one approach, beam steering is accomplished with a single-axis gimbal, combined with rotation of the satellite about the receive axis. In the other solution, requiring no gimbals at all, the node consists of two distinct satellites flying in close proximity, with one satellite acting as the receiver and the other as a transmitter, and with a short-range omnidirectional link between them.

Single-Axis Gimbal

Assuming the incoming beam does not carry information in the form of polarization angles, a satellite receiving an optical signal is free to rotate about an axis defined by the line of sight to the signal source. This freedom of rotation, combined with a single-axis gimbal for the transmitter, is sufficient, in principle, for simultaneous transmission in any direction relative to the incoming beam. In practice, it may be beneficial to concede some limits on transmission direction in order to simplify satellite design. Consider the design of the CubeSat Optical Relay Demonstration (CORD) satellite illustrated in figure 4. This is a 3U CubeSat with about half of the volume dedicated to a combined optical receiver and transmitter. The receive axis is aligned with the long axis of the CubeSat, with the incoming signal entering the satellite through one end. The signal is focused with an 8-by-10-cm off-axis-parabolic reflector to a detector as illustrated in the figure. The receive portion of the payload also includes a receive beacon camera used to align the receive axis with a beacon from the transmitter, and a receive beacon laser co-aligned with the receive axis and providing a tracking point for the transmitting satellite.

The transmit portion of the payload includes co-aligned data and beacon lasers as well as a transmit beacon camera. The transmit beacon camera is used to receive the beacon from the receiver satellite and inform the pointing of the transmit lasers. The optical path from the transmit payload components includes a flat mirror mounted on a single axis gimbal so as to rotate the pointing direction of the transmit components about an

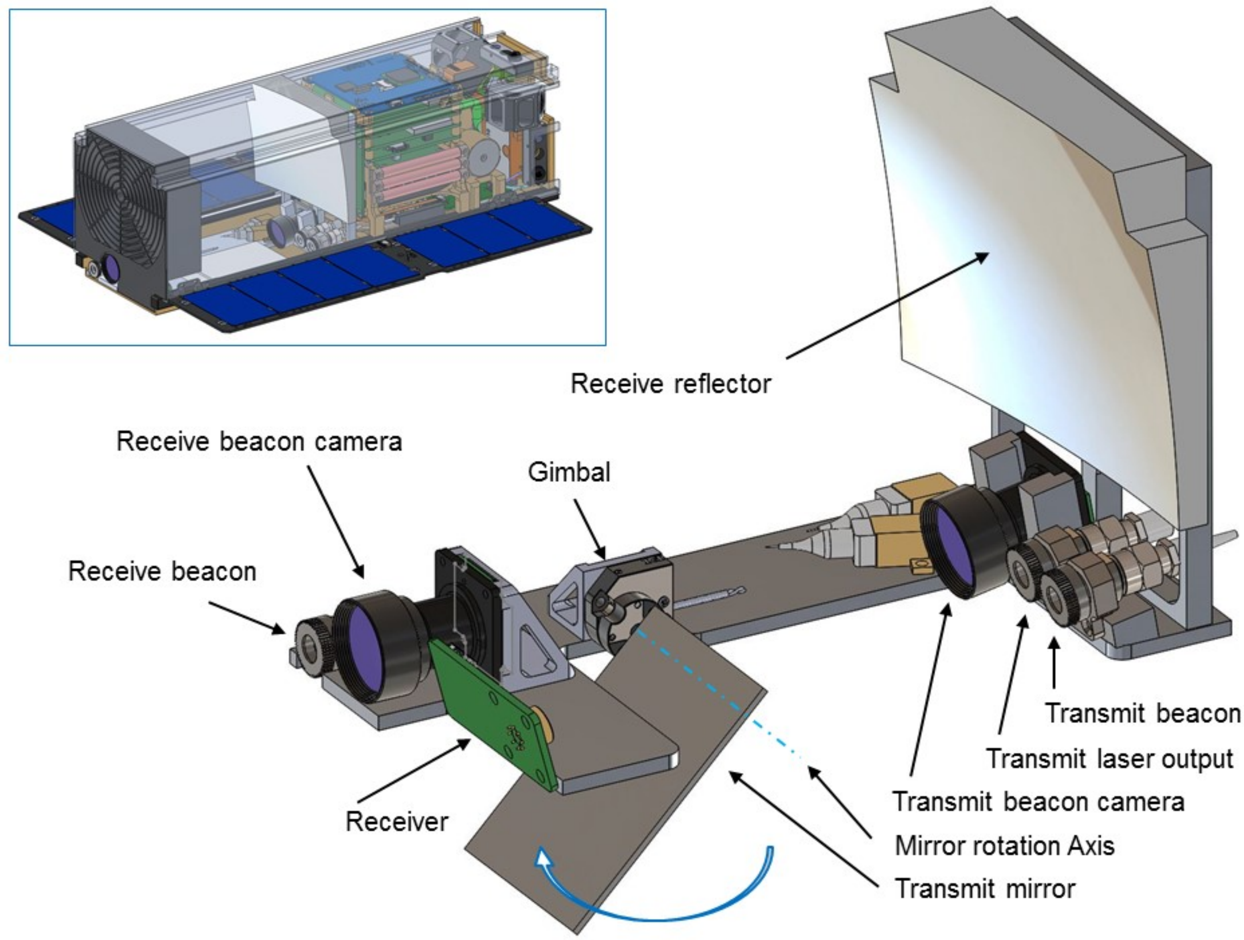

Figure 4. Conceptual design of an optical relay satellite using a single-axis gimbal in a 3U form factor.

axis perpendicular to the receive axis of the satellite. This rotation, combined with the rotation of the whole satellite about the receive axis, enables simultaneous reception and re-transmission of an optical signal to most of space. While two-axis gimbals are typically complex and massive, and not readily compatible with CubeSats, single-axis gimbals are much easier; compact, vacuum-rated, high-precision gimbals are readily available catalog items.

The pointing geometry using a single-axis gimbal is illustrated in figure 5. The coordinate system is oriented such that the incoming signal arrives from -x direction. The axis of rotation of the gimbal on the relay satellite is perpendicular to the receiver axis, and therefore perpendicular to the x axis of the coordinate system. The gimbal rotates a mirror, which reflects the beam from the transmitter on the relay satellite. The transmitter is mounted such that it transmits along a direction perpendicular to the gimbal axis. For most directions relative to the incoming beam, the required pointing of the

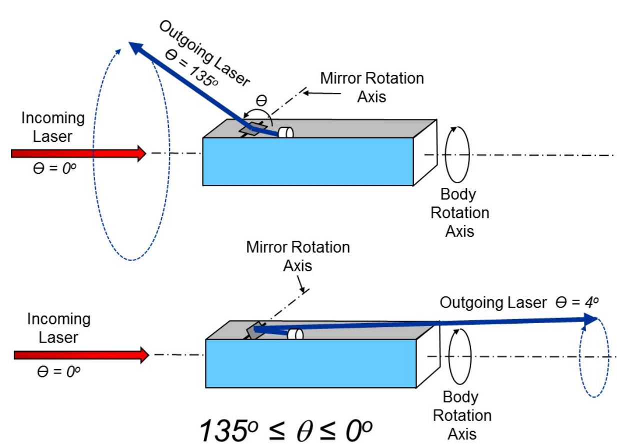

Figure 5. Pointing geometry of an optical relay satellite using a single-axis gimbal.

outgoing signal can be achieved by rotating the mirror about the single gimbal axis and rotating the whole satellite about the body axis. Provided the satellite body does not move other than to rotate about its long axis, the receiver will remain pointed at the source signal while the transmitter is pointed at the destination.

Because of limits on the size of the mirror, CORD is limited to transmission angles between 0 and 135 degrees from the +x axis. With a full 360-degree rotation of the satellite about the receive axis, this leaves a solid cone with a half-width of 45 degrees centered on the -x axis where the laser cannot transmit (see figure 6). While alternative satellite designs may be able shrink (or even eliminate) the exclusion cone, we chose to simplify the satellite design in exchange for some loss of the ability to transmit generally back toward the original source. It is anticipated that most scenarios involving a relay would have the signal being relayed to a location further from the source than is the relay itself. In that case the rotation angle of the transmit vector would never be more than 90 degrees, so this is a minor limitation.

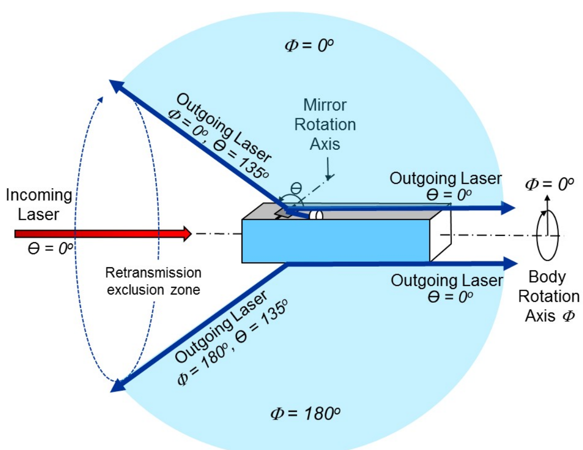

Figure 6. Available pointing directions for the CORD satellite concept.

The CORD satellite design was developed to demonstrate an optical relay capability in a 3U CubeSat. The bus is essentially identical to the bus portion of the R-Cubed spacecraft and closely similar to the bus of OCSD. In addition, the transmit laser of CORD is identical to the laser in R-Cubed, which is, in turn, only a slight upgrade to the laser in OCSD. As such, the CORD demonstration mission can be flown almost entirely with components and systems that are, or soon will be, flight proven.

Two-Satellite Node

It is also possible to avoid the use of gimbals altogether. A relay node in a space-based optical network can be provided by a two-satellite system where one of the two satellites operates in the receive mode, the second operates in the transmit mode, and the two satellites fly in close proximity to one another. Data transfer between the two satellites is provided by an omnidirectional, or nearly omnidirectional RF or optical system. In this configuration, both the receiving satellite and the transmitting satellite can point to the required degree of precision at their respective targets without interfering with one another and without requiring a mechanical gimbal between them.

The first satellite, the receiver satellite, can be similar to the CORD satellite, but without incorporating the transmit section. The attitude-control system of the receiver satellite points the satellite at the signal source to maintain signal quality. The receiver satellite also includes a short range transmitter that is either omnidirectional or has a relatively wide beam to minimize pointing requirements. This short-range transmitter can be either an RF transmitter or an optical transmitter.

The second satellite, the transmitter satellite, includes a laser transmitter for sending data to another node in the network or to the ground. The attitude-control system of the transmitter satellite is used to point the satellite at the intended signal receiver. The transmitter satellite also includes an omnidirectional short-range receiver compatible with the short-range transmitter on the receiver satellite.

In operation, the two satellites work together to transmit data continuously through the node. Data received at the receive satellite is transmitted immediately, using the short-range transmitter, to the transmitter satellite. The data received by the short-range receiver on the transmitter satellite is then transmitted to the next node in the system, using a long-range, precisely-pointed laser transmitter.

A variation on this concept would include an isolated transmitter satellite node in a stand-alone mode that would fly in close proximity to an EO satellite, essentially acting as an auxiliary to the communication system. If this transmitter satellite were equipped with an RF receiver configured to receive from the EO satellite's normal downlink transmitter, then it would provide a pathway for linking the EO satellite directly (and continuously) into the LEO network. Both of these concepts are illustrated in schematically in figure 7.

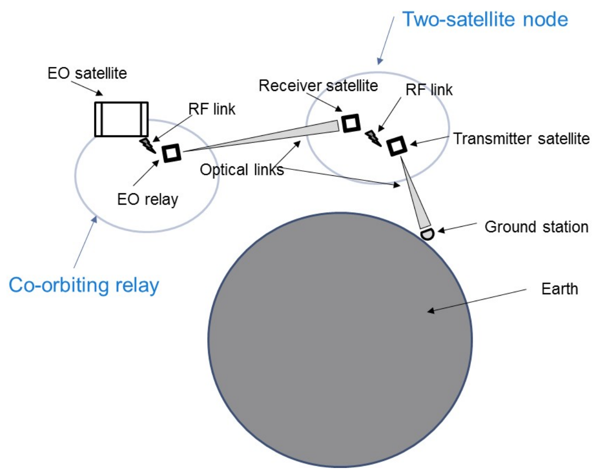

Figure 7. Schematic representation of the use of two-satellite nodes and co-orbiting relays.

Operation of this system requires that the two satellites of the node remain in close proximity to one another to enable the use of the short-range communications link between the two. To minimize the propulsion requirements, the two satellites should have similar mass and drag profiles, which would minimize their tendency to drift apart due to atmospheric drag. In addition, each satellite preferably would be sufficiently irregular in shape that they can fly in either a high-drag or low-drag mode, enabling the use of variable drag to maintain proximity between the two satellites. While variable drag has not yet been used to demonstrate high-precision station keeping, it has been used for approximate station keeping with the AeroCube-6 mission, and will be used for proximity operations in the OCSD mission. The AeroCube-6 satellites, which have limited attitude control capabilities, have been using variable drag since late 2014 to maintain station within a few tens of km.[9] Pre-flight modeling of the OCSD orbital dynamics indicate that it should be possible to use the higher-fidelity ACS on these satellites to control the separation to better than 2000 m.[10]

Experience gained with OCSD will indicate whether it will be possible to maintain separation in a two-satellite node within the range of the short-range omnidirectional crosslink. If the use of attitude-driven variable drag is insufficient to maintain relative separation requirements, either due to operation at altitudes where drag is too low, or due to excessive communication time requirements that interfere with variable-drag operations, then a propulsion system would be required to maintain separation.

Because the relay satellites can be small and simple (a 3U CubeSat would be adequate, and 1.5U may also be sufficient, depending on requirements), initial deployment of the satellite network can be relatively inexpensive. An individual two-satellite node could, for example, be launched as two CubeSats from the same deployer, or even launched as a single CubeSat (3U or 6U) that would then separate into the two node components. For a symmetric system operating with bidirectional communication, both satellites in a single node could be identical. Even for a unidirectional system, there would be only two types of satellites. A large system of nodes could thus be deployed in LEO for a very modest cost.

SUMMARY

Continuing advances in the capabilities of CubeSats, particularly in the area of precision attitude control, support the eventual deployment of an all-optical LEO communications network that would provide continuous, high-volume, low-latency download of data generated by LEO Earth observation satellites. While GEO-based optical relay systems can provide complete coverage of LEO with only three satellites, the range from LEO to GEO makes pointing a challenge for very small satellites. A LEO optical network capable of complete and continuous coverage of the whole of LEO space, would require on the order of 100 satellite nodes, but would enable short-range links from LEO Earth-observation satellites, greatly simplifying requirements for satellites designed to use the network. If built around the CubeSat standard, the cost of the nodes in the LEO network would be low enough that the entire constellation could be competitive with the cost of putting a single relay satellite in GEO.

The key challenge in free-space optical communication systems is pointing and tracking to ensure that the narrow communication beams reach their intended target. Traditional systems use complex and massive two-axis gimbal systems, but these are not compatible with many developing small-satellite EO missions. Alternative designs for relay nodes using single-axis gimbals, or no gimbals at all, are compatible with current CubeSat technologies, and can enable the deployment of an all-CubeSat optical relay network in LEO.

REFERENCES

1. L. Fang, N. He and H. Lin "CP tensor-based compression of hyperspectral images", JOSA A (34) p.252.
2. M. Villano, G. Krieger and A Moreira, "Data Volume Reduction in High-Resolution Wide-Swath SAR Systems" APSAR Conf. Sept 2015.
3. Goldberg, M., "NOAA Direct Broadcast Data Initiative to Meet NWP Latency Requirements," 41st Meeting of the Coordination Group for Me-

teorological Satellites (CGMS-41), Tsukuba Japan, 8-12 July 2013.

4. Fields, R., Kozlowski, D., Yura, H., Wong, R., Wicker, J., Lunde, C., Gregory, M., Wandernoth, B., and Heine, F., “5.625 Gbps Bidirectional Laser Communications Measurements Between the NFIRE Satellite and an Optical Ground Station,” Proc. of the 2011 Int. Conf. on Space Optical Systems and Applications, pp.44-53, Santa Monica, CA, 11-13 May, 2011.

5. F. Heine, H. Kaumpfner, R. Lange, R. Czichy, M. Lutzer, and R. Meyer, “Laser communication applied for EDRS, the European data relay system”, CEAS Space J (2011) 2:85–90.

6. T.S. Rose, D.W. Rowen, S. La Lumondiere, N.I. Werner, G.A. Maul, D.H. Hinkley, D.H. Chien, R.P. Welle, S.W. Janson, “Laser Downlink Demonstration from a 1.5U CubeSat,” Photonics West Conference, San Francisco, February 2016.

7. S.W. Janson, R.P. Welle, T.S. Rose, B.S. Hardy, R.M Dolphus, A.Y. Chin, C.M Coffman, S.D. La Lumondiere, N.I. Werner, D.A. Hinkley, “The NASA Optical Communication and Sensors Demonstration Program: Initial Flight Results,” Small Satellite Conference, SSC16-III-03, Logan UT, August 2016

8. C. Frost, E. Agasid, et al. “Small Spacecraft Technology: State of the Art,” NASA Technical Report TP-2014-216648/REV1, NASA Ames Research Center, 2014.

9. Gangestad, J., Rowen, D., Hardy, B., Coffman, C., and O'Brien, P., "Flight Results from AeroCube-6: A Radiation Dosimeter Mission in the 0.5U Form Factor." 12th Annual CubeSat Developers Workshop, San Luis Obispo, CA, 2015.

10. S.W. Janson, R.P. Welle, T.S. Rose, D.W. Rowen, D.A. Hinkley, B.S. Hardy, S.D. La Lumondiere, N.I. Werner, “The NASA Optical Communication and Sensors Demonstration Program: Preflight Update,” Small Satellite Conference, SSC15-III-1, Logan UT, August 2015.

11. Dee W. Pack, David R. Ardila, Eric Herman, Darren W. Rowen, Richard P. Welle and Sloane J. Wiktorowicz, "Two Aerospace Corporation CubeSat Remote Sensing Imagers: CUMULOS and R3," Small Satellite Conference, SSC17-III-05, Logan UT, August 2017.