ACP Downlink PHY Layer Protocol Proposal

Introduction

The ACP will have highly asymmetric links, and so completely different protocols will be necessary for the ground to space and space to ground communication systems. This document will address the Physical layer of the protocol stack for the link from the satellite down to the ground stations. The overriding concern in the design of this link is to allow for reliably transmitting fixed-length frames of data to many users with varying receive capabilities (i.e. antenna gains).

This document will address data encapsulation, modulation, and coding. Other concerns, such as the content or interpretation of those bits, or addressing are not covered here. Actual data rates are scalable will be determined by the final link budgets, also covered elsewhere.

ACP Downlink Properties

One of the key aspects which defines the ACP is that there is a single downlink stream which will be required to service users with widely varying capabilities, the most important of which for this discussion is antenna gain. Initial discussions have suggested a spread of as much as 20dB between well-equipped fixed stations with a moderately sized dish, and smaller mobile ones with something more like a patch or helix. It would be very easy, but also very wasteful, to simply send all bits with enough energy such that the weakest stations could always receive them. Instead, we need to find a way to impart as much as 20dB per bit of extra energy to those bits which are destined for the weaker stations, and this represents most of the challenge for the downlink PHY design.

Methods of varying E_b

The easiest way to vary the energy per bit ($E_{b)}$ would be to increase the power amplifier output in the satellite when necessary. This, however, is not as easy as it sounds, and is not very efficient. It is much better to run the power amplifier at the same power all of the time.

The other option is to reduce the data rate. Given a constant power, fewer bits per second results in more energy per bit. The most direct way to do this is to actually slow down the modulation, resulting in a narrower band signal. At first this might seem like an ideal choice, but certain practicalities get in the way.

It can be difficult for receivers to maintain a link when symbol rates vary widely (~100:1 in this case) and rapidly (perhaps as often as every millisecond) because symbol timing and carrier tracking which are close enough for a 25kbits stream may not be close enough for the 2.5Mbps stream which is may be there 10uS later. This problem would be worse for the weaker stations, which might not be able to track the higher-rate signal at all, and so would have no estimate of where to look for the signal when the lower rate signal is finally sent.

Additionally, if there is a desire to have the ability to vary the rates among many possible values, it becomes more difficult for stations to know which rate they should be listening for.

Instead, we propose that bits intended for weaker stations be sent at the same high symbol rate which the strongest stations will use, and that they are simply repeated (with extra coding, see below) when sent to weaker stations. As long as the weaker stations are able to track the higher rate stream, they will be able to coherently add the energy from the many copies of the channel symbols. This is similar in principle to direct sequence spread spectrum, except instead of repeating symbols consecutively (and xor-ed with a PN sequence), we repeat them blockwise.

This allows the link can be designed for the highest capability users, and throttled back for the lower rate users. By allowing for a variable number of repetitions of each block of data, we can smoothly vary the rate in accordance with the actual capabilities of the weaker stations. We could even take into account the current utilization level of the satellite to smoothly trade off robustness vs. capacity. When the system is lightly loaded, everything might be repeated. When utilization is heavy, we could reduce the repetition factor to accommodate more users.

The only real cost to this system is that synchronization and the information about what is a repetition of previous data need to be robust enough for the weaker stations to be able to track.

Modulation

The choice of modulation typically comes down to bandwidth efficiency vs. power efficiency. In the case of the ACP, we have 10 MHz of bandwidth available (3.4 to 3.41 GHz), and at most 400W of power. Given these numbers and the realities of physics (as expressed through link budgets), we could not productively make use of all that bandwidth. We have a surplus of bandwidth and a relative lack of power even in the most optimistic scenario power budget scenario. Thus, we should be designing for the most power efficient modulation we can.

Factors that play into how power efficient a modulation are its E_b/N0 performance as well as how efficiently it can be amplified. A modulation with a constant envelope can be amplified by a nonlinear and much more efficient amplifier. It is for this reason that unfiltered PSK is being proposed. BPSK is the most likely choice, but some have proposed using QPSK, OQPSK, and even Pi/2-BPSK. All should have the same link performance, but may work better through nonlinear amplifiers. Another variable which has been suggested is using some residual carrier.

The one issue with unfiltered PSK is the relatively high sidelobes (~15dB down). You can make the argument that since the signal itself is weak, and spread over a wide bandwidth, then the even weaker sidelobes are unlikely to pose any interference problems. Additionally, they would fall within amateur bands, and directional antennas a typically used in this band anyway. We need confirmation of this.

There is another option in case unfiltered PSK is deemed unsuitable. At least one source claims that with sufficient demodulation effort, GMSK can achieve the same BER as BPSK.[1] If this is true, GMSK might be an option.

Coding

Decoding turbo codes is computationally intensive, but encoding them is relatively easy, so there is little reason not to use them on the ACP downlink.[2] We can use the CCSDS standard turbo code[3], which is defined for rates ½, 1/3, ¼, and 1/6. The standard high-rate stream can be encoded using the rate ½. One nice property of this code is that the constituent codes of the ½ rate code are also the first 2 of the rate ¼ and 1/6, so when sending to the weaker users, some of repeated blocks can be the other constituent codes. This gives roughly an additional 1.2dB of coding gain at no additional cost.

All codes of this type perform better with increasing block size, but this effect must be traded off against the increased latency of waiting for more data to fill those blocks. Longer blocks will have less

---

[1] Simon, Marvin K., Bandwidth-Efficient Digital Modulation with Application to Deep-Space Communications (JPL Publication 00-17) June 2001. http://descanso.jpl.nasa.gov/Monograph/series3/complete1.pdf

[2] Patent issues notwithstanding

[3] TM Synchronization and Channel Coding – Summary and Concept of Rationale, CCSDS130.1-G-1, Section 7

overhead for the repeat flag, CRCs, etc., but are more likely to have a frame error at the same SNR. According to the CCSDS references, 1800 bits per block seems to be about as low as you would want to go.

The proposed code would be 1800-bit blocks of the rate-1/2 code (3600 channel symbols), but we could easily envision making them as long as 3600 bits (7200 channel symbols). Repeated blocks for the weaker stations should rotate between the first, second, and third pairs of constituent codes. Thus, if a block is repeated numerous times, the 1st, 4th, 7th, etc. are the output of the first two constituent codes, the 2nd, 5th, 8th... are from the second two, and the 3rd, 6th, 9th, ... are from the third pair. If the repeat count is not a multiple of 3 then the unequal energies per symbol are accounted for by the soft decision metrics.

Frame Structure

The proposed frame structure is shown in Table 1. It consists of 3 fields. The first is a sync code, most likely of 32 to 64 bits, depending on link budgets to weak stations. It only needs to be long enough for the weak stations to be able to track it over many code blocks, not just a single one. It allows receivers to track phase, frequency, bit time, and frame boundaries. Rather than use residual carrier to aid in frequency tracking, this field can be lengthened if necessary.

The second field is the Repeat Flag. This field conveys only a single bit of information, but it must be able to be reliably received by the weakest stations, or large numbers of frames will be lost. Thus, depending on link budgets, it might need to be as long as 100-200 symbols. Since it only conveys a single bit of information, we cannot do better than simple repetition. It will be modulated with a PN code just to aid symbol tracking, as in direct-sequence spread spectrum.

The remainder of the frame is the code block. Once decoded, code blocks contain 1800 bits of information. The first 1768 are a Link Layer Frame, and the last 32 bits are a CRC used to detect uncorrected errors. The link layer protocol is covered in a separate document.

Channel Symbols	**Bits of information**	**Field Contents**
32-64	0	Sync code
?	1	First/Repeat Flag
3600	1800	Code Block

Table 1: Frame Structure

Conclusion

A physical layer protocol for the downlink of the ACP has been described. Its main unique characteristic is the ease with which data destined for users with different receive capabilities can be accommodated in one constant rate symbol stream.