

Part 1 - Introduction

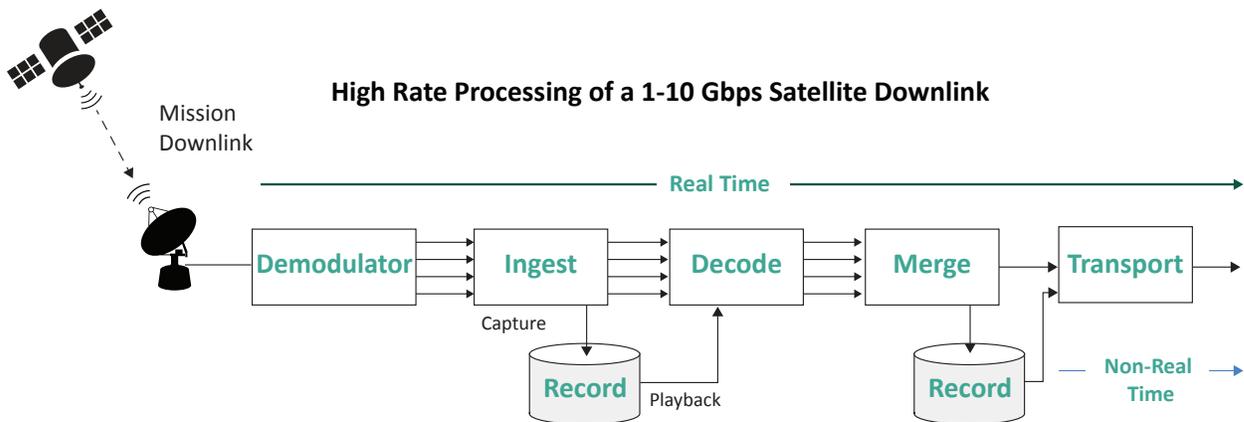
Introduction

Satellite downlinks are real time data streams. A real time data stream is one that is continuous and often has no means for retransmission. Fall behind in the receipt and processing of the data and there's no catching back up because there's never a break or gap in the data.

In the late 1970's, 1 Mbps was high speed for a real time data stream. Over the next thirty years, it was 5 Mbps, then 10 Mbps, then 100 Mbps, then 400 Mbps. You get the idea. Some of today's payload downlink data rates are excess of 1 Gbps are common, with some as high as 10 Gbps.

The processing of high rate real time data streams (>1 Gbps) can be particularly challenging. Catching all of the data, decoding it, processing it, and getting it moved from point A to point B, relies on parallelism, transaction rates, and latency.

High Rate Data Flow



Part 2 - Ingest, Recording, and Processing

Ingest

High rate downlinks may rely on phase shift keying to split the data across multiple channels in the RF transmission. The demodulator often has a separate physical output for each channel. For example, a downlink that is QPSK modulated would have an I channel output and an Q channel output. Many high rate demodulators have four separate outputs. For example, a 3.2 Gbps downlink might arrive on four 800 Mbps channels.

On some legacy demodulators, these outputs are serial data/clock. At these data rates, the electrical protocol is differential ECL, usually Negative ECL (or NECL). There is a migration to using the UDP protocol over dedicated point-to-point Ethernets where there is a dedicated Ethernet connection for each channel.

The first step in the high rate processing is to ingest these data streams without any data loss. Serial streams are parallelized in hardware (e.g. turned into 32-bit words), blocked into data frames, and then DMA'd into memory. A network stack handles the UDP protocol.

One important design element for ingest is to have larger rather than smaller packets. Smaller packets require more network transactions per second. There is both header overhead and transaction overhead for each packet. For example, transmitting a 2 Gbps data stream in UDP packets with a packet payload size of 1,024 bytes requires ~244,000 packets per second. The same 2 Gbps data stream with a maximum UDP payload size of 9,000 bytes generates ~28,000 packets per second.

Data Recording

High rate processing systems may immediately archive the raw data. If there's a drop-out downstream, the data can be processed after the fact.

Data recording also relies on parallelism. The data is stored to disk arrays using RAID controllers. A RAID controller stripes data across multiple disks. There is no single low-cost disk drive that sustains continuous writes at 10 Gbps, but an array of 24 disks only requires each disk drive to support writes at ~500 Mbps (accounting for some recording overhead).

Decoding and Data Merging

The spacecraft performs data encoding/interleaving of the high rate downlink to support decoding and error recovery on the ground. Encoding techniques include Convolutional Encoding, Interleaving, Reed-Solomon, and Low-Density Parity Check (LDPC).

Data encoding techniques have asymmetric processing requirements—the processing needed to decode the data is much higher than the processing needed to encode the data. Given the constraints on processing in space, this asymmetry is a favorable one. Data decoding and error recovery on the ground side can consume significant processing resources. In many cases, the demodulator performs some or all of the work. When that is not the case, the high rate processor gets to fill in.

Here too, parallelism makes the processing possible. High rate systems can use low-cost Graphics Processing Units (GPUs) for block-oriented processing. These GPUs have thousands of processors and the decoding/error recovery algorithms work on blocks of data. Multiple blocks are processed in parallel on the GPU. This does induce latency in the system throughput as multiple data blocks must be collected so they can be sent to the GPU in groups.

Data merging may be required across the separate data channels, putting the separate data streams back into a composite data stream. The decoded and merged data may be stored (recorded to disk) in parallel with it being sent downstream for processing.

Network Transport

The processing of the data itself, turning it into useful information, often takes place away from the antenna. The high rate processor can rate buffer the stored data across a wide area network if the network bandwidth is constrained. In addition, the high rate processor can implement a forward error correction (FEC) technique over the WAN to reduce the network latency induced by network transmission errors.

Post Processing

Playback processing pulls the raw or partially processed data from the disks and performs the required post-processing to include any image or signal processing and the corresponding data analysis. With the data captured and archived, this post-processing can take place in non-real time.

Can we help? AMERGINT's expertise is available to assist in your systems engineering and design

Randy Culver
randy@amergint.com
719-522-2802