SATCOM DSSS NETWORK HUB SDR SUPPLANTS HARDWARE LEGACY, FOR LOWER COST, HIGHER-PERFORMANCE & RELIABILITY, AND ADDS AJ

Cameron Pike (CP Communication Engineering Corp., Mount Vernon, Ohio, cmpike1@direcway.com)

ABSTRACT

Increasingly powerful and dense interference, combined with imminent failure of discontinued legacy hardware, threatened a safety-related SATCOM messaging service company with catastrophic failure. CP Communication Engineering Corporation (CP Comm) was contracted to design and build new CDMA hub equipment to replace the ailing legacy hardware, under extreme deployment time constraints, provide improved performance as well as implement anti-jam (AJ) to overcome the new and emerging interference environment, while maintaining uninterrupted service to the thousands of mobile terminals in the field. CP Comm's seven years of SDR experience combined effectively with timely availability of new DSPs from Texas Instruments and digital radio components from Analog Devices to provide beta units to the customer within six months, in time to avert the business catastrophe.

This paper describes the implementation of the DSSS CDMA network hub, with particular emphasis on the tremendous value of the SDR paradigm both for rapid development and deployment, as well as ease of functional enhancement through upgraded software.

1. THE NEED

The *Semper* company (a pseudonym, for privacy purposes) has been operating a mobile messaging satellite-based service, providing a safety-related service in its niche market for some 10 years. The mobile-to-earth station link (*reverse link*) employs a direct sequence spread spectrum modulation exhibiting 24 dB processing gain to facilitate both multipath fading robustness and multiple access for the several thousand unscheduled transmitters. Each message is comprised of both a spread-spectrum preamble, or acquisition sequence, followed by a spread-spectrum data payload.

Satellite bandwidth is, indeed, a scarce resource, and leasing a several-megahertz band to accommodate a net data throughput of just a few tens of kilobits per second is not sensible. Thus, an arrangement was struck between *Semper* and the satellite operator to permit dual-use of the spectrum with narrowband traffic, exploiting a certain degree of interference rejection of the spread spectrum system while recognizing the low power spectral density of the spread spectrum signal. This arrangement worked well for several years, neither system suffering harm from the other; but as the satellite operator's business developed and expanded, so did the narrowband traffic. As it did so, the narrowband interference would exceed the processing gain of the spread spectrum receiver more often. Furthermore, the operator began to offer service to higher data rate customers (occupying approximately 5% of the spread spectrum bandwidth each, some 30 dB above the noise floor), wherein the presence of even a single one within the spread spectrum band would bring the *Semper* system to a halt.

As if this changing and challenging interference environment were not enough, the legacy demodulating equipment installed in the earth station was beginning to fail. Unless something was done to address these catastrophic problems, *Semper* was faced with going out of business in short order.

2. THE LEGACY

When *Semper's* custom demodulator system was first designed a decade ago, the demanding aspects of the signal design pretty well dictated a certain architecture, and was state-of-the-art at the time, packaged in a standard 15" 6U CPCI chassis. Although implemented mostly in FPGAs, it was essentially immutable because of the complexity of this demanding application, and, because of the challenge of making the high-performance hardware meet the extra-demanding application, a number of compromises were made in the receiver.

The basic approach was to operate a matched filter and peak/threshold detector at a multiple of the chip rate seeking occurrences of the $O(10^4)$ -chip acquisition sequence, followed by a despreader/demodulator for message payloads. The signal feed from the earth station equipment is a ~30 MHz bandwidth signal centered at 70 MHz IF; a complex baseband data stream was generated from this by an analog downconverter stage and quadrature (dual) 4-bit ADC, sampling at four times the chip rate. This high-rate data stream was processed by a large bank of 20 fast FPGAs

(10 each on two 6U CPCI cards) programmed to perform preamble detection and some simple form of AGC function, with feedback to the sampler circuitry. It is noted that even the task of passing fast data streams among 20 FPGAs contributes significantly to the complexity of the system. The fast data stream was also bused to a backplane, to be available for further processing.

The preamble detection signal was provided as a daisy-chained backplane signal; when the detector asserted this line, the first available (non-busy) "demodulation processor" began operating on the synchronously-clocked backplane data stream. The despreading operation was performed in an FPGA in real-time, with reactive delay-locked-loop chip clock tracking, and producing four partially despread samples per encoded symbol. These samples were passed to a DSP that performed the final demodulation functions, namely phase tracking, symbol integration, symbol deinterleaving, Viterbi decoding, and message format / checksum validation. A maximum of six despread/demodulation units were integrated into the system, and thus only six message detections could be processed at a time. Lastly, the data messages were picked up over the PCI bus by a CPCI Pentium SBC running DOS (later upgraded to Linux), and delivered to the *Semper* infrastructure via TCP/IP. At least one complete chassis was required for each satellite beam providing coverage in the *Semper* system.

By today's standards, this implementation was nightmarish. It was slow, unwieldy, providing little in error diagnostics, and mostly impossible to modify when the RF environment changes. The analog section needed constant adjustments, and bursty noise caused "storms" of false detections, thus masking any legitimate messages because the relatively few (~6) demodulators were busy processing noise! Furthermore, due to a range of factors, the signals from the transmitters could vary by as much as 12 dB, and higher-powered signal receptions were not only distorted beyond recovery, but likewise tended to mask traffic for some period of time afterwards. All of these factors combined to produce a system that could only boast 84% probability of detection, except when being totally jammed by other legitimate satellite traffic!

3. THE SDR SOLUTION

CP COMM had been assisting *Semper* for several years in various capacities, and we were quite familiar with the application and the system. Our experience and work in software radio for a number of customers served to keep us well-equipped with the newest hardware and software available, as well as building a sizable quantity of reusable IP. When *Semper* contacted us regarding their need, they indicated they needed a real solution in a matter of a few months. Given CP COMM's readiness, we undertook the task. Our solution is comprised of a minimal analog section, an RSP, two DSPs, and an embedded microcontroller uClinux system for network connectivity. Altogether, the system fits within a 1U chassis, with an abundance of elbow room. It is far quieter, more reliable, easier to maintain / modify, uses less power, is far better performing, more robust, addresses the present and emerging jamming threats, and far less expensive than the legacy system.

We analyzed the problem, and determined that a new architecture was in order. In pursuit of our desire to minimize signal distortion, we always try to convert the analog signal into digital as early in the signal chain as possible. A new signal converter was designed employing direct-sampling the IF feed at 100 MHz with a 14-bit ADC from Analog Devices. The minimal analog signal path is as follows: upon entering the new system, the analog signal is immediately converted from the single-ended 50-ohm analog input to differential 200-ohm; an off the shelf 30 MHz SAW filter is sandwiched between two differential amplifiers and drives the ADC in differential mode. This provides superior noise reduction and better system performance while reducing EMC shields, etc., within the chassis. The ADC provides a nominal 80 dB dynamic range, thus obviating any analog AGC for the well-behaved earth station IF feed. Local oscillator (LO) signal generators and clock sources needed by the legacy system are obviated, and the new system simply locks the ADC sampling clock to the earth station's main 10 MHz reference feed. (After sampling, no other clocks are critical, with respect to noise, performance, purity, etc.,)

The 100 MHz, 14-bit undersampled data stream is processed through a Receive Signal Processor (Analog Devices AD6624A) that performs digital mixing with a complex sinusoid, and filtering / downsampling. The last stage of this device is programmed with a critically shaped pulse matched filter, whose bandwidth is equal to the chip rate, and the data stream is thus critically sampled at the chip rate at 16 bits each for I and Q.

This data stream is produced from the RSP at a "real-time" rate, i.e. clocked synchronously. Thereafter, the data is all buffered in FIFO or deep memory, which is more accommodating to interruptible, multi-tasking DSPs. Furthermore, this permits the processing algorithms to be optimized for signal processing performance since the entire data record (causal / non-causal) can be used to demodulate or derive signal metrics – long the dream of DSP engineers! The raw channel data is simultaneously sent to two destinations: the acquisition processor DSP, and to a deep memory accessible to the demodulation processor DSP.

The acquisition function is very straightforward, and is appropriately classified as a traditional DSP task: matched filter processing. This is performed by a Texas Instruments TMS320C6416, executing at 1 GHz, and using only internal

resources (contrast with the 20-FPGA legacy solution). The processing algorithm is dominated by multiply-accumulate (MAC) operations, and is fairly deterministic in its execution. Thus, it is not surprising that this application boasts a processor utilization over 90%.

Besides an amazing 20-fold reduction in package count, a significant improvement is found in the increased dynamic range of the internal data representation, making overflow errors nearly impossible (14-bit ADC, 16-bit I, Q, 32-bit accumulators in the DSP). Taking advantage of the full flexibility of the DSP, the detection algorithm is made more sophisticated to improve performance in both the analytically attractive AWGN environment, but also the real life interference environment of the satellite channel, the most important factor being overlapping transmissions of widely differing powers. Rather than a single logic signal indicating a threshold detection (as in the legacy system), the new detection reports now consist of a time-tag (actually, a numerical address in memory) indicating the start of a new possible message, and the actual weight of the decision metric. Furthermore, a more useful local (temporal) noise estimate is made to continuously adjust the detection threshold, thus implementing a constant-false-alarm rate (CFAR) detector in the highly non-stationary environment. Other information is also made available for improving performance of the demodulator, if necessary, such as correlation lag weights that can be used in an equalizer. Since the detector algorithm is implemented in software, it is easily altered in order to optimize performance in the actual deployment environment.

The demodulation processor is a TMS320C6713, operating in more of a multi-task environment. External SDRAM is sufficient to store approximately two seconds' worth of raw data. One process receives raw data from the RSP and enters it in a large circular buffer maintained in the external SDRAM. Much of the complicated memory management becomes invisible to the programmer by enabling full L1 and L2 caching, and permitting the clever DMA engines to figure out whether the data is cached or not.

A second major process receives detection reports from the 'c64, and enters the report and its metric(s) in an internally maintained fixed-length queue. Clever management of this queue is critical in boosting system performance, and "intuitive" operation, and it is acknowledged that if the detection threshold is set properly, some detections will not be processed, i.e. fall off the bottom of the queue. An explanation is in order. Recall that the transmitters are uncoordinated, each transmitting whenever it is ready, and because of varying conditions will arrive at the satellite (earth station) with varying power levels, as high as 20 dB of link margin, to as low as undetectable; all below the noise floor. Because of the number of transmitters and their respective average transmission rate, overlapping transmissions from multiple mobile units regularly occurs. The high-powered detections are very likely legitimate messages, but many of the lowest powered detections are simply false, or even intended for an adjacent satellite antenna beam. From the perspective of the users, the most powerful transmissions, i.e. the most likely to be genuine messages, should have a higher priority of being processed than the low-powered ones. This agrees with a system-resource utilization perspective, since higher-powered receptions do not require as much processing (time) to properly decode the message as do marginally-powered receptions.

Thus, the queue is ordered by precedence of the detection weight rather than arrival time, so that higher-powered detections have higher priority than (i.e. will be processed before) lower-powered detections.

If this is the only consideration, however, there is still some room for improvement. Detection reports are generated by the acquisition correlator at the end of the preamble, and beginning of the message payload; the complete message has not yet been stored in memory, nor even been received! Hence, there's a second level of dequeuing: the highest-powered detection whose complete message payload is certain to have been received is selected as the next one to process. We then set the detection threshold low enough such that the demodulation processor is always busy attempting demodulation, and succeeding on even very low-power signals, but not at the expense of missing larger ones. As detection reports are entered in the queue, and reports taken off, the oldest / weakest reports are eliminated without ever being processed. This method has been demonstrated to make most effective use of the system resources, as well as significantly boost performance. (Under current and near-future traffic projections, a single processor is more than sufficient to keep up; however, provision has been made for additional processor resources to be tasked from a single queue.)

One aspect of the architecture is worth noting here. It matters not that the detection reports may indicate messages that overlap in time, nor in what particular order the reports are processed; all the data for some period of time is available, and each new demodulation task simply begins requesting data from the SDRAM at the beginning of the new message. The caching and DMA hardware of the processor is free to make the best use of bus bandwidth, etc.

Traditional approaches to demodulating the spread-spectrum messages are used, with some economies due to system constraints. The steps are despreading, Doppler correction, deinterleaving, Viterbi decoding, and checksum validation. One economy is that since the messages are short duration, and the start time has already been determined by the acquisition processor, time-tracking is obviated. A second economization is that the highest rate processing is only at the chip rate; sub-chip time recovery is

unnecessary for higher-powered signals, and is transformed to an equalization problem for lower-powered signals.

Signal processing performance is greatly improved by having the ability to "remember" raw data samples. For example, the entire message data record is available to extract tracking metrics (e.g. signal strength, SNR, Doppler shift, correlation function), and thus can be calculated to a high degree of reliability. Armed with all of these metrics, the algorithm returns to the beginning of the data record, and applies suitable corrections to the signal from the very beginning of demodulation. This is in stark contrast with traditional radio architecture, which is compelled to estimate these metrics "on-the-fly", and perform the corrections after demodulation is already in progress. It was observed in the legacy system that quite often the first few bits of the demodulated message were garbled while the tracking circuits were still in their initial transient response; whenever this occurred, the entire message was rendered erroneous because of the few bit errors.

Finally, completely demodulated messages are collected by the uClinux embedded processor and communicated to the *Semper* infrastructure via TCP/IP.

A comment is warranted regarding latency. In this application, timeliness of message processing is critical; but only to the degree of some significant order of magnitude. Given that the messages must pass through the internet at some point, it is recognized that "timeliness" is measured in tenths to a few seconds. The possible delays introduced by the methods of processing herein described are clearly below this order of magnitude, even when messages are produced "out of order"; i.e. earlier arriving (weak) messages produced after later (strong) ones.

3.1. Anti-Jam

The main interference sources in this system are benign, well-behaved, but extremely powerful and each occupying roughly 5% of the spread spectrum bandwidth. Dynamic range is not an issue, due to the 14-bit quantizing and bit-growth room. Two approaches are attractive, with differing levels of sophistication: coherent subtraction of the interfering signal, and simpler spectral excision. The former, requiring demodulation and remodulation of the interfering signal, would be required if the entire spread spectrum band becomes occupied. The latter is adopted in this system, and performed in the transform domain, just prior to despreading correlations, at a relatively low cost in system resources. The acquisition processor and the demodulation processor each perform their own spectral excision.

4. THE BENEFITS

A beta system was delivered to CP COMM's customer on time, and performed as promised. An immediate benefit resulted from the increased sensitivity and performance, and probability of detection was over 99% (conversely, probability of missing a message for any reason – receiver "malfunction", transmission error, misaimed TX antenna, environmental problem, etc. – was less than 1%.).

The customer stress level reduced immediately, and he was able to discard his old failing equipment in favor of the new.

This new architecture has opened up new applications. The hardware can easily be repackaged for portable use in terrestrial applications, and can accommodate continuous data demands due to its higher performance.

5. IS THIS SYSTEM <u>REALLY</u> SDR?

In this rapidly growing and changing field, I would say "yes". Although the SCA played no part in the development, and the platform is not the "Pentium machine" of radio, the key qualities of SDR are present:

1. All salient receiver functions are performed in software
2. The software components are reusable (or, in fact, have been recycled into this project!)
3. The hardware is largely unrestrictive, and can serve as a platform for a multitude of different modulation / communication methods

6. CONCLUSION

Working knowledge of SDR and a having built a repository of SDR-related IP enabled CP COMM to very quickly design, build, and deploy a sophisticated network hub receiver, to the great benefit of the customer. Improved performance and reliability, as well as documentation and flexibility all provide not only a lower-cost replacement, but a path for growth in business opportunity and service offerings.