The "Cloverleaf" Performance-Oriented HF Data Communication System

By Raymond C. Petit, W7GHM
PO Box 51
Oak Harbor, WA 98277

News Item: In March of 1990, AKØX and W7GHM conducted a series of on-the-air tests of a new HF data communication system design over the 1500-mile path between Boulder, Colorado, and Oak Harbor, Washington, On five different days, and on the 80, 40, 30, and 15-meter bands, Ed sent text from Isaiah 55 to Ray, W7GHM, at higher speeds than any other data mode was capable in the same band conditions, and Ed's "Cloverleaf" signal was so compact that twenty of them could have been packed, without mutual interference, into the same 2-kHz space now used by one packet channel. On 15 meters, Ray **was** printing data at 75 bit/s free of errors. On 30 and 40 meters, during a time when the packet link between AKØX and W7GHM was nothing but retries, Ray got 50 bit/s from the Clover link. (The design limit of AMTOR is 33.) On 80 meters, when it was necessary to repeat single letters several times on CW to be understood, the Clover link delivered 15 bit/s, free of errors.

During a second series of tests by Willard, N7JTQ, and Ray, W7GHM, a one-way Clover link delivered throughput-per-bandwidth performance from about 50 to over 1000 times better than HF packet in the severe multipath environment of a night-time 50-mile path on 80 meters.

The first test began at 10 PM on April 29. The Clover link transferred data at 37 bit/s on a sustained basis without losing or garbling any data. Willard and Ray were unable to establish a packet link because the TNCs retried out making connect requests.

The second test began at 9 PM on April 30, 1990. The Clover link delivered 75 bit/s with one in ten of the data blocks lost. (The two-way Clover protocol will provide for retransmission of lost blocks without requiring retransmission of blocks already correctly received.) Shortly after 10 PM, a file transfer on packet was aborted after a few minutes by a link failure. The average data rate was 0.7 bit/s. A second attempt produced a nearly identical result. The Clover one-way link then delivered 117 blocks of data at 50 bit/s, losing only two blocks.

On May 2 at 7 PM the third test **was** conducted. The conditions were much more stable, and a 2 kbyte file transfer on packet averaged about 31 bit/s. Afterwards, the Clover link delivered 75 bit/s. At about 8 PM, another packet file transfer averaged about 17 bit/s and it was followed by a Clover data transfer at 50 bit/s.

In the final series of tests on May 31, Logan, KL7EKI, sent the entire test message at 95 bit/s from Wasilia, Alaska, in ideal 20-meter conditions. His 9-watt signal was received 2000 miles away without a single error requiring correction. In late evening on 30 meters, when Logan and Ray's packet link was immobilized, Logan sent Clover data at 15 bits/s and 15 lines of text were received before the first uncorrectable error occurred.

The performance "figure of merit" was obtained by dividing the number of correctly received bits per second by the spacing (in hertz) required to guarantee that two signals of the same type do not cause mutual interference. even if one is 60 dB stronger than the other. For a Clover signal, this spacing is 100 Hz. Packet and AMTOR require at least 20 times this space.

Introduction

"Cloverleaf" is the name I have given to a modulation and coding scheme which offers very worthwhile improvements over HF packet, AMTOR, and even CW. The specification for this system will be a gift to the Amateur community. Here are its main performance features:

1. It is exquisitely compact in bandwidth. For practical purposes a Clover signal is entirely contained in a channel only 100 Hz wide. Clover signals are designed for channel spacing of 100 Hz; they need no guard bands. The actual channel width of the Clover receiver is this same 100 Hz. A Clover signal in the neighbor channel will not cause interference even if it is 60 dB stronger than the signal to which the receiver is tuned. A network of 10 Clover links can be maintained without mutual interference in a 1-kHz band. Each channel is a "clear channel": no collisions, no "splatter," and no key clicks come from the other users, even if they are much stronger.

2. A Clover link communicates data at the highest speed the RF propagation path permits. As the band conditions change the system adapts to them automatically. Under the best conditions, it operates at above 100 correct data bits per second delivered to the user. Under the worst conditions, conditions in which even CW data rates approach zero, a Clover link can communicate a few bits per second.

3. The Clover design uses Reed-Solomon error-control coding to correct errors in transmission, rather than rejecting a block of data on account of the errors. The coding is set such that it will recover blocks which have as many as 10% of their data symbols lost. If too many errors occur, the receiving station obtains a retransmission from the sender. It is never necessary to retransmit a block of data which has been received successfully on account of errors in previous blocks.

4. The data link is completely transparent to the data user. No restriction exists on the alphabets or data sequences. There are no illegal data codes. Input is accepted as a sequence of bytes and delivered to the output of the link in the same format. Programs using a Clover link are

totally free to define their data in the ways best suited to their needs.

5. The Clover system requires and takes advantage of the extreme frequency precision and stability of its transceivers. It also requires accurate knowledge of time. The transceivers obtain both the frequency and time information automatically from observations of WVW (for Western Hemisphere applications). A Clover channel can be centered 50 Hz from a band edge with confidence.

Overview of the Design Specification

The Cloverleaf system uses channels 100 Hz wide with their edges at multiples ot 100 Hz. The Cloverleaf modulator generates 25 pulses per second in Mocks which vary in length from 3 seconds to about 30 seconds. The pubes are very carefully shaped in amplitude such that they produce no sidelobes in frequency. The demodulator has dynamic range high enough to guarantee that adjacent Cloverleaf signals do not interfere with each other even if one is 60 dB stronger than the other.

For the higher speed modes, the data is transferred via the difference in the phase and amplitude of the successive pulses. Depending on the stability of the radio channel, each pulse conveys from I to 6 bits of data using from two to sixteen distinct phase levels and up to four amplitude levels. When the channel conditions are too poor to support the "phase" modes, the data is sent in interleaved binary pulse-position modulation to make possible, by signal averaging, the recovery of data at low speed when the signal is below the noise level.

The data IS encoded into Reed-Solomon error-correcting codes which permit the receiver to recover all the data in a block-most of the time-despite the inevitable sudden losses of signal due to multipath effects and noise. The number of phase levels and the Reed-Solomon code block size used are adjusted automatically according to the channel conditions such that loss of blocks at the receiver is kept to about 10% of the blocks.

The Original Cloverleaf System

In May, 1988, I ran a simulation which satisfied me that the system described here was feasible. On June 19, 1988, I conducted an on-the-air test of a breadboard Cloverleaf modem with W7HHU on 80 meters, and it confirmed my expectations. The remainder of 1988 was spent in developing the operating system software for a 6809-based demodulator and primitive error-control coder. On January 7, 1989, AKØX successfully sent me test patterns on 15 meters from Boulder, Colorado. I spent the first half of 1989 developing an automatic synchronizing protocol and a very flexible Reed-Solomon encoder and decoder. In May, 80-meter tests with W7ZEG at a time of intense solar activity showed that there are times when nothing gets through! I spent the second half of 1989 working on the hardware for the second version.

The original breadboard version has separate transmitter and receiver modules, Both are implemented with 6809 processor boards and hand-wired expansion boards, and the software provides one-way data transfer only, for the present, The transmitter generates the pulses in software as a succession of 16-bit data words and presents them to a 16-bit D/A converter and filter. The output of the converter/filter is centered at 500 Hz and sounds like a rapid, smooth and steady hooting.

The receiver downconverts its audio input to a pair of baseband channels in quadrature, passes both signals through a 10-pole linear-phase active low-pass filter, then through digitally gain-controlled dc amplifiers, and finally to 8-bit ND converters. With this arrangement, AGC range is about 90 dB, but distortion in the audio output stage of the radio makes only about 60 dB of it usable. The audio signal used for the downconversion is synthesized in 0.1-hertz steps and the receiver software can continually adjust the frequency over a small range to maintain zero beat.

The receiver software includes an analyzing routine which can distinguish between normal background noise and this noise with a weak carrier present. The synchronizing protocol first presents a steady carrier. This "wakes up" the demodulator which the tunes itself to zero beat. Then a stream of phase-reversing pulses establishes "framing" of pulses. This is followed by blocks of 9 pulses with inter-block gaps of one pulse length. With these blocks the receiver obtains "byte synchronization." Then follows information about the modulation and coding formats in the upcoming data blocks, and a "countdown" anticipating the first block of communication data. An enormous amount of redundancy is used in this process to ensure that the synchronizing data is correctly received even under marginal conditions.

The data, text from Isaiah 55, is sent in Reed-Solomon coded blocks in the selected format. At the end of each block is a gap of one pulse length. The receiver measures the power level in this gap and the signal power level to estimate received signal-to-noise ratio. It also measures the phase jitter in the received signal and the number of errors it was required to correct in the Reed-Solomon decoder. All this information is used to keep a running estimate of the channel capacity. In a two-way protocol yet to be implemented this information would be used to request updates in the modulation and coding format employed by the remote station. The time and frequency discipline maintained by the protocol will shorten the process of establishing a connection to a few seconds.

The Second-Generation Transceivers

A pair of Clover transceivers is in development. These units will provide means for extended evaluation of this design as a practical alternative to packet under conditions requiring the highest levels of performance. Waveform generation, channel filtering and demodulation will be handled by a Motorola DSP56001 digital signal processing chip with 16-bit A/D and D/A conveners. Data coding and the two-way protocol will be managed by a 6809 chip. All the system clocks are phase locked to a single master reference oscillator which is checked against WWV automatically to obtain both frequency and time synchronization. The RF portion of the transceivers feature an ultra-low noise frequency-synthesized local oscillator system which will provide reliable (and required) frequency accuracy to within a fraction of a hertz. The IF crystal filters are I kHz wide, permitting observation and perhaps operation on IO Clover channels simultaneously. The linear RF power amplifiers will deliver 20 watts. An RS-232 serial port will provide the data link to the host application and computer.

INTERVIEW: Questions and Answers

Question: Why have you chosen to use PSK instead of the universally accepted narrow-band FSK?

Answer: The modulation scheme here can't really be called PSK as that term is used in amateur applications. PSK, like FSK, is supposed to be a constant-amplitude modulation method. To rapidly change the phase of a carrier is to generate enormous out-of-band radiation, and that is precisely what I wanted to avoid. FSK is better in this respect, and that is one reason it makes sense to use it on HF. But it is even better not to change the phase at all when the carrier level is nonzero. This modulator never shifts the phase of a carrier or a pulse, but generates separate smoothly amplitude contoured pulses, each of them appearing at a new phase angle as required by the data being sent. This concept is not at all new, as any communication-systems engineer can attest. But it has not been used on HF before to my knowledge, at least in the ham bands. In fact, it would have been illegal until just recently!

Question: So a lot of the motivation for this design was to make a data signal which was just as compact as possible, right?

Answer: Yes. A lot of people during the eighties worked hard on this problem and their work has now been canonized in the engineering textbooks. They were usually thinking of how to make high-speed microwave data links more efficient, but the same principles apply for low-speed systems, too.

Question: Why do you regard this extreme spectrum compactness to be so important?

Answer: What would you give for a tenfold increase in the size of the CW and data portions of every ham band? Imagine the tiny 30-meter band being expanded to the size of the entire 80-meter band! What a thought! The Clover design will accomplish the same effect, not by expanding the space, but by allowing us to pack ourselves ten times closer than we do now without mutual interference. And the specification for this system puts real meaning to the idea of "no interference." A Clover signal 60 dB stronger than the one you are listening to and only 100 Hz away won't degrade your copy. How far away must such a signal be now in any of our data modes to get the same protection? What this all amounts to is simply making better use of existing resources.

Question: Are you sure that phase-encoded data methods will work on HF? After all, a lot of people have done a lot of work, and they always have ended up using FSK.

Answer: Yes, very true. All the classic work which gave us what we have now, as far as modulation methods are concerned, was done before microprocessors existed. And before chips which were designed specifically for digital signal processing. And before ultra-stable HF radios were practical and inexpensive enough for the home builder. It's time for a reevaluation of the conventional reasoning on this subject.

Question: But what about the fundamental problem of multipath, selective fading, and Doppler shifts? Phase data is certainly going to be lost, no matter how sophisticated the hardware.

Answer: Yes. A great deal of my work with the original Clover machine was devoted to on-the-air observation. One of the neat things the Clover demodulator permits me to do is to watch the phase of any HF signal I could tune to with my ICOM 735. I began by extended observations of the 10-MHz carrier of WWV in Boulder, Colorado, about 1500 miles from here. I wrote a program which could track the carrier on a long term basis while I observed the short-term phase behavior by watching a dot move around in a circle on the scope. I also wrote a program which could display the phase spectrum of the signal on my computer screen. An another program could count the number of bit were times when WWV's phase really jumped around, but much of the time it was very stable.

Then I listened, with the same setup, to many different foreign broadcast stations. I soon discovered that some had very unstable carriers! Some were using frequency synthesizers which had residual low-frequency phase modulation in their outputs. (That was detectable because their carrier amplitude was quite constant but their phase went jiggling around wildly-phase variations caused by propagation are always accompanied by amplitude variations.) I observed many amateur CW signals and I could distinguish the transmitters using keyed oscillators from the ones using unkeyed oscillators. This all gave me a feel for what to expect.

The textbooks broadly agreed that the multipath dispersion on moderate-distance HF channels is almost always under 4 milliseconds. The degrading effect of this dispersion is reduced if the pulse duration is much greater than that figure. At bandwidihs I was considering, selective fading doesn't exist; when fading occurs, virtually the entire signal fades in unison. I decided to use a 40 millisecond pulse duration, 25 pulses per second, and chose four phase levels. This would put out 50 bit/s. After a lot of analysis and experimenting with computerized Fourier analysis of signal waveforms, I found a shape that was compact with no sidelobes and was practical to generate. I built the receiving filter to match it and verified that I could recover the clock and data information from the signal after it emerged from that filter. Now to avoid errors at 25 bit/s and four phase levels in the modulation, the propagation-induced phase error during one bit intervals cannot exceed one-eighth of a cycle in 40 milliseconds, which amounts to an instantaneous frequency deviation of plus or minus just above 3 hertz. At times of intense solar activity or over some very long paths I saw that figure exceeded: I saw phase scattering which I knew would make Clover signals unreadable. But I also saw signals which were so well behaved that I decided it was worthwhile to make the modulator go to as high as 16 phase levels and to make the number of leveis adjustable to get as high a throughput as the channel permitted. So I wrote a program which could appraise the phase spectrum and recommend what format would get me the best data speed at acceptable error rates.

Concurrently with all this was the question of error checking. One of the things that really got me going on this was a statement I read in one of my textbooks to the effect that while in VHF or microwave channels the improvement due to error-control coding was at most a few dB, over HF channels the improvement could be spectacular. This really whetted my appetite! So I decided to implement a Reed-Solomon coding system, a task which took me many, many months, because I knew nothing about the math of these kind of things. My first coder operated on data blocks of 15 4-bit symbols, a block length of 60 bits. I soon discovered that a whole family of RS code block lengths would be needed in order to cope with the widely varying conditions I observed. In due course, I found out how to match the coder to the conditions.

The on-the-air tests verified my expectations. It was really beautiful to see data coming through "solid copy" even though the signal regularly dropped out or was obliterated for a moment by noise! When conditions were poor, we just hunkered down and got the data through at a slower rate. When conditions were good, we went faster! Under nearly perfect conditions it delivered 95 bit/s. If the data is coded in S-level Baudot, we're talking about 18 characters per second. No one I know can type that fast! By comparison, AMTOR gives you about 7 Baudot characters per second maximum by design. A packet HF BBS under ideal conditions on a 200-mile path on 80-meter daytime ground-wave conditions averaged about 10 ASCII characters per second on one occasion while sending me long messages. That was exceptional.

Question: What about the weak-signal performance?
Answer: A 10-dB signal-to-noise ratio at the receiver is perfectly adequate for data modes, but most of the time on HF we operate way. way above that level! When Ed, AKØX, and I began the 15-meter test, he was coming in something like S-8. It was way too strong. I asked him to reduce his power to as low as he could get it. I was still copying him fine when his power was so low he could not see any indication on his RF output meter! Clover, like AMTOR, works really good when the signal is weak. In the biphase mode, it can read signals which I can barely hear in the noise.

Question: How about two or more Clover signals on the same channel?
Answer: A carrier or other Clover signal 20 dB below the one you want to hear will begin to slow things down if, without it, the system is running at top speed. One of the rules the transceivers will follow is not to transmit any channel on which another Clover signal is readable. There will never be any need to, at least for a long, long time!

Question: It appears that you are turning your back on the TDMA scheme used by packet in favor of FDM. .
Answer: Yes. Everyone knows the AX.25 protocol is really best for higher-speed work al VHF and above. Bandwidths are wide and amateur data sources, so far, are very slow in comparison. It makes good sense to "time share" a channel since any given QSO needs only a fraction of the available channel time to pass data. On HF, we don't have the bandwidth (except on 10 meters) for things like 1200 bauds on voice band FM. And wideband high-speed data modes just can't take the beating from HF propagation. So I say, instead of forcing 20 stations to coexist on one voice channel, give each one a clear channel in that same space. No collisions, no QRM! And if it really comes to push and shove. within each of these channels it is possible to establish a discipline in time, like the way CW or voice nets operate, even if the stations aren't talking to each other.

The Version 2 Transceiver

Question: Have you considered multitone modulation formats?
Answer: Yes, quite a bit, and on several occasions. My IF filters on the Version 2 system are 1 kHz wide, centered at a point 2.5 kHz below the BFO. The DSP chip can easily generate more than one tone. But unless the linear amplifier intermod products are kept down. those products will interfere with signals in the nearby channels. Several times I pondered that some good power FETs are specified at - 50 dB third-order products in class A operation at about 50 watts. That's not bad. This is a promising area.

Question: Will the Version 2 transceivers be capable of operating on the other data modes?
Answer: The heart of the transceiver is a Motorola DSP56001 digital signal processing chip. A 6809 handles everything else. The transceivers should be able to do anything. For the present, I'll stick to Clover formats, but the hardware of the transceivers is very "general purpose," and I'd invite any enterprising programmers to make this matching jump through whatever loops they please.

Question: What about the RF hardware for these transceivers? The frequency synthesizers look a bit complicated for a home builder.
Answer: These transceivers are designea for just one purpose: high-performance HF data communication. They have no extras, no bells and whistles. The RF paths are incredibly simple in comparison to your typical commercial rig It is single conversion, and the IF frequency, 18 MHz minus 2.5 kHz, works with the 20-30 MHz synthesizer to provide coverage from about 2 to 12 MHz. They were designed with frequency precision and stability as a primary goal. The IF amplifier has no AGC, but instead is designed for high dynamic range. The DSP engine gets its data from a 16-bit A/D converter-that provides 90 dB range-and all the channel filtering is done in software

More than half of the circuitry is in the frequency synthesizer. It is a multiloop design which will provide one-tenth hertz steps over its entire 10-MHz range. Four of the five circuit boards of the synthesizer are identical. Most of the parts for the transceiver were obtained from a well-known mail-order house in the Midwest (USA). The crystal filter is made from inexpensive microprocessor clock crystals. The linear amplifiers are straight out of a Motorola application note. I think any experienced home builder could do it fine.

Question: It still looks pretty expensive, out of amateur range. . .
Answer: Yes it is, but in ten years one of these transceivers will be possible to build at very attractive prices. My versions are synthesized for flexibility, but in the applications I see for these things, a single-frequency crystal-controlled version at half the cost makes good sense. And just watch the prices for DSP chips fall as they are more widely used!

Question: Where do you see this system fitting into the whole of amateur data communication?
Answer: My experience with Coherent CW says that I shouldn't expect something like this to gain wide acceptance quickly. AMTOR showed what could be done when someone was designing specifically for the conditions we get on HF. Packet is gaining a wide following in spite of its performance on HF. I see that some people are thinking about making incremental improvements to HF packet: It's needed and all to the good. But at some point, a whole new approach has to be taken. It's like the conversion from am voice to single sideband back in the 50s.

I expect that the Cloverleaf system wili find its best use in the lower-frequency HF bands where it is well suited to the formation of regional networks covering areas the size of Western Europe. Link-level protocols are being worked out, but networking software is a project awaiting a taker I doubt if Clover will replace packet, but it will provide a high-performance alternative. Getting twice the throughput in one-twentieth the bandwidth at some point just has to make sense!

Question: Your introduction indicated that you intended to give this whole system to the amateur community Haven't you considered patenting this thing and making a commercial enterprise of it?
Answer: Yes, I've thought about it a lot Back in 1972. I made a commercial enterprise out of a fun project and I survived for about **3** years. I'm a very small budget operation, your proverbial garage outfit. Getting a patent costs more than my entire annual budget. it takes several years, no protection exists until you have it, and then in order to protect it you have to spend hundreds of thousands and all your emotional energy in the courts I don't want to live that way. It's more important to me that the machine gets used. that practical Clover networks get established, and that we all get to enjoy its advantages I expect to have diagrams. circuit boards, and parts kits for home builders pretty soon. I expect that eventually some companies will start making these things. I plan to be one of them.