A DUPLEX PACKET RADIO REPEATER APPROACH TO LAYER ONE EFFICIENCY

Robert Finch, N6CXB   Scott Avent, N6BGW
c/o South Coast Radio Relay
4446 Rosebank Drive
La Canada, California 91011-1928

PURPOSE

In January, 1984, South Coast Radio Relay, a Southern California Amateur Radio group specifically chartered to experiment in state of the art message handling techniques, modified its two meter duplex voice repeater in Glendale, California, for full time digital communications. This operation has remained on the air since that time, under both the original callsign; N6TD/R and the since adopted call; N6GPP/R. This operation has undergone some specific changes in the past several years, but the basic user interface has essentially remained the same. This paper will examine the basic operational characteristics employed at the repeater site, specific to the repeater itself, and will aid the reader in duplicating the methods used to achieve the promise inherent in local area duplex packet operation.

THE PROMISE

Packet has been sold to the amateur community as an error free way of 'message handling'. That in and of itself, has been realized. However the initial means for a packet station to be used is through the simple LAN (local-area-network), made available to the user to actually accomplish any message transmission to the intended destination. The experimenters in our ranks are continuing to push for higher efficiency levels of inter-networking links to allow for more distant destinated traffic and higher volumes of in transit traffic. As a group, we quite often leave the 'end-user' ham to the low volume, low efficiency, simplex LAN. In some areas this LAN is on the same channel (frequency) as the high volume wide area networks. These wide area networks are double trouble for the local user, Not only is the traffic volume high, but it can also be automated and characteristically highly persistant in occuping the available time slots. To allow for continued interest and growth we have to change this, where possible. The authors hope to shed some light in this often overlooked area by presenting the method we have been using to attempt to accomplish some of these goals.

BACKGROUND

The current layer two protocol in use is dependent on the use of a CSMA-CA/CD (carrier sensed multiple access - collision avoidance/ collision detection) usage contention scheme to maximize channel efficiency in real-time. In order for this to work correctly it is assumed that every station in the network can correctly hear every other station within the network's basic transmission area. A group of stations in an ideal network topology can acheive this time sharing on a single frequency channel. In the real world this cannot be counted on. In practical terms it has shown itself not readily capable of this time and time again. As we shall touch upon later, this has proven to be a real thorn in the side of local amateur packet radio operations.

A paper has been previously presented at the fourth ARRL networking conference examining the relative merits of various layer one approaches, and makes excellent background reading. In this paper; PACKET RADIO TIMING CONSIDERATIONS, the author (David Engle, KE6ZE,) presented AX.25 specific timing information, and examined the relative channel through-put efficiency of each of three 1200 baud AFSK approaches. Of course the most efficient was the direct connect, i.e., without any form of repeating path involved. The second most efficient was the duplex repeater (in their experiments a regular voice repeater with a 500ms transmitter key-up delay). The final option examined, and the least efficient was the single frequency repeater or 'digipeater'. At the end of this earlier paper the author made several suggestions on improving the duplex repeater performance based on the experimentation he made, attempting to approach the measured efficiency of a 'direct-connection' simplex layer-one connection. There are differences in the specific approaches taken between our implementation and the previous author's recommended duplex repeater configuration. We do not always agree on specfic design 'practices', but our current repeater has exceeded our original design goals and in operation can essentially equal the efficiency of 'direct-connection simplex packet. Additionally we significantly increase the coverage area of the average user's station within his local area network.

OUR SOLUTION

The basic block diagram of N6GPP/R repeater, in 'black-box' form, would appear much like a regular voice repeater with the addition of a RS-232 port attached to allow for 'inter-network' additions. In fact this inter-netwroking was accomplished some years ago with the addition of the N6GPP-1 145.01 mhz. network access computer and a Motorola Mocom mobile radio. This computer is basically a Xerox 820 modifed to KE3Z specifications and running a modified version of his ARRL control firmware. The repeater itself is a late model commercial Motorola Micor duplex repeater. The principle advantage to this model repeater is its fully solid state, crystal controlled operation, without the use of any 'switching' relays, in the the control of its basic operation(s). This allowed us to cut the packet recognition and repeater transmitter turn on time to less than 30 ms.

The basic transmitter multiplier and modulation circuits are on continuously and only the solid state P.A. is enabled after valid packet recognition has taken place. In operation the R.F. transmitter 'turn-on time' is well under 100 micro secs., the balance of the 30 ms. being the reception circuitry. Additionally, within the transmit audio path is a custom 'zero-voltage' fet switch controlled configuration which allows us to maintain a tone free squelch tail. We also eliminated the stock limiting and clipping stages. This reduction of audio processing circuitry allows us to realize a perfectly level controlled sine wave with essentially no distortion due to any 'audio-processing' and therefore no splatter filtering is needed. The audio can be fully modulated to desired levels without worry of any undesirable or otherwise uncontrollable products in the R.F. output.

Working backwards from the transmitter audio circuits is the modulating 'modem'. In our case there is only one single modem in use for the entire demodulation, modulation and the outside world RS-232 connection. However the basic design is much more easily thought of, initially, as two modems connected back to back, so to speak. The modulating modem must be free of any distortion as previously mentioned, and some scheme for assurance that no start-up artifacts occur during the 'settling' of the audio tone during turn-on is required. To accomplish this a modem assuring zero-crossing tone switching is employed. In our case we settled on using a digitally based commercial 202 modem, which could be confiqured as separate receive and transmit sections and still allow for some intergration of hardware control signalling. The modem is set-up for 20ms total time for acquisition and 8.5ms RTS (request to send) to CTS (clear to send) turn around time. This turn around time is utilized for two basic functions. First the afforementioned settling of the modem generated audio is assured. Secondly the incoming RTS can be controlled on a first in-first out basis with the outside network and the duplex receive channel contending for control.

This high-speed 'turn-around' of request to send to clearance for transmission is the key to keeping the efficiency of channel through-put high enough to realize our goal of essentially transparent operation, while still allowing for the addition of various incoming sources (outside networks). By using hardware control via the RS-232 standard; full contention resolution can be accomplished with the addition of this simple first in first out circuitry. This still allows efficiency to remain at a high level.

By using a standard six-wire RS-232 implementation (ground, TXData, RXData, request to send, clear to send, and data carrier detect) multiple source and destination connections can be implemented via simple fan-out of detection signals to various contention control devices. Specifically the digital demodulation of received information generates a DCD (data carrier detect) signal which is fanned out to the first-in first out contention circuit and the input of the Xerox 820 inter-network routing computer. The output of the same computer is also routed to the first-in first-out circuit, which in turn generates the RTS signal to the audio modulation (portion of the) modem. Similarily the first-in first-out circuit handles the routing of RXD (RXData) and TXD (TXData) and the CTS (clear to send) signal which will allow inhibition of the Xerox's sending circuitry if the computer is not actually rquesting transmission first. After this circuit desides who has control of the repeater transmitter it then goes on to the finally key-up the transmitter and gate the audio thru to the modulator. Standardization on this RS-232 hardware control allows us to mix the non-intelligent nature of the duplex in to out signalling with the computer based routing of the xerox computer. In fact the new NETROM (tm) and other intelligent controllers are offering just this specific RS-232 standard hardware control, which allows us freedom for future changes and enhancements of wide-area-network integration within the current system (N6GPP/R along with N6GPP-1) to enhance the inter-networking of our local area network.

The receiver portion of the repeater is also very stripped down. Packet recognition is based around two distinct circuits. These circuits combine to acheive fast recognition without needless 'falsing' out the repeater output frequency due to 'spurious' and otherwise unneccessary signals at the receiver's R.F. input. The first is the use of the demodulation section of the above mentioned 202 modem for not only conversion of audio to digital signaling but also to allow the use of the inherent tone recognition circuitry within the standard commercial 202 modem. This timing is set to approximately 20 ms as a good compromise without being too long (slow). We found while this was sufficiently fast, when coupled with the transmitter timing, it was not totally sufficient in eliminating the falsing. To augment this required some additional circuitry. This circuitry is not neccessary in all implementations and some explanation is in order.

While we feel that the turnaround and recognition circuitry within a standard commercial 202 is ideal to meeting our goals as mentioned above, restraints can raise their ugly head(s) and did in our case. In a crowded R.F. spectrum as is found in a large metropolitan area, such as we are priviledged to have here in greater Los Angeles, the compromise afforded had to be enhanced by the use of a specific 'packet-recognition' circuit. This simple circuit essentially discards a packet when the repeater is initially started in order to allow verification of a large number of transitions within a short time interval. The assumption is that the demodulator with-in this time interval will not transition its RXData out stream a large number of times within a small timing window unless the audio signal is a valid 1200 baud data stream. The window is formed by the same DCD (data carrier detect) output signal that forms the basis for the input to the above mentioned first-in first-out contention circuitry and is then used to request the data transmission repeat. Making this recognition circuit re-triggerable and enabling the DCD connection to the contention circuitry as the result of valid packets, verification is a definite compromise. However with the interval for restart of this circuit set at 2 mins. 12 secs. we have found as a practical matter that it is used only to initally enable the repeater as intended even with the generally slow packeteer-typist. Making it a retriggable circuit allows arthritic typists access too.

Of minor significance is the repeater's T.O.T. (time-out-timer). All digital repeaters should employ these, as is standard practice with the garden variety voice repeater. As you can see by now, our repeater doesn't need to run with the more commonly found repeater control shelf. In fact a digital repeater's control circuitry can be easier to construct, For completeness we would like to point out we have a time out timer.

Of considerably more significance is the phase-coherent R.F. discriminator (detector) and the squelch circuitry found in our Micor receiver. The modem is directly connected to the discriminator. This direct connection mirrors the intent behind the essentially direct connection of the 202 modem's modulation output thru the fet switching to the R.F. modulator. The cleanest possible signal allows for the least amount of error rate increases due to distortions. The standard audio squelch is **not** needed in the conventional usage (but could be beneficial in reducing some remaining 'noise key-ups'), but since it is part of distortion producing audio processing cicuitry it isn't used. As mentioned above, the design goal was to keep the entire audio chain as clean as possible. We are currently investigating possible inclusion of a suitable squelch circuit not requiring any in line processing circuitry. While the use of product detector is not ruled out in the adopted layer **one** standard for two meters (1200 baud AFSK), it is a neccessary requirement to use a phase-coherent detector for any duplex digital repeater based on a wide bandwidth, true FSK, layer-one. In our case, the phase coherent detector in use is supplied with the standard Micor package we used.

In passing we'd like to mention that the control system does not maintain any on-site ID'ers. We handle this function as a beacon signal generated by one of the channel hosts on a regular basis, timed to assure that we are in compliance with the applicable regulations on repeater identification,

The use of a hang-time (squelch tail) has received much debate as used in our operation, and the various advantages and disadvantages are still discussed with some regularity after three-plus years of operation. Initially we found that 2 meter voice users couldn't tell that there was a repeater on our frequency pair without it. This was especially troublesome for both digital and voice users back then. While perhaps not a problem these days, by making the 'tail' longer than **one** retry interval, we incured essentially no problems with its inclusion. This helped back in the days when TAPR was shipping serial number 350 of the original tnc-1, and packets sounded like 'intermod' to the rest of the amateur **community.** By creating a tail after every series of packets sent, while the average ham might not understand what the signal is for, he at least knew there was something resembling a repeater on the frequency. With the addition of the newer version two of the layer two protocol to the scene, there has been some mention of changing this 'hang time' to longer than one 'response' interval. However we have found that in most all cases, setting user's response time to zero, and leaving the hang time as is, an even better solution. This allows us to accomodate all versions of user's operating firmware. The channel has such a good balance between receiver and transmitter coverage areas that any one packet user is virtually assured that his packet will be heard by all users, and we might mention that the older version one (layer two,) protocol runs especially well on duplex packet networks based on our experience.

The above mentioned balance to coverage areas of the repeater's receiver and transmitter, in practical terms, is more important than in a standard voice repeater. In fact the transmitter's coverage should be somewhat larger than the receiver%. To assure that the 'hidden-terminal' problems of simplex packet digipeaters aren't unneccessarily duplicated this balance becomes critical. When your packets are heard by all other users the number of users can be maintained at elevated levels compared to the average simplex repeater channel, which by its naturally occuring coverage topology cannot assure against easily resolvable collision avoidance and detection. (This AX.25 characteristic has been enumerated, discussed, and generally beated about in so many places and publications in the past several years, along with proposed modifications to the protocol etc. to help minimize its effects. We do not want to go into a longish discussion here. Suffice it to say that the layer two protocol, especially the original version one, is in light of this; a much better duplex protocol than simplex. In fact the apparent real-time loading fiqures, i.e. the level of absolute time in use for packet transmissions until the channel clogs up and is unusable for a time, is significantly higher in duplex compared to simplex operation. And the 'loading' curve that preceeds this clogging, in terms of the maximum number of pending packets serviced, is apparently much less steep with duplex packet.)

BACK TO THE USER

This discussion would not be complete without bringing up the enhanced operation the users enjoy. Beyond the higher efficiency local area network operation we can enjoy the novelty of true roundtable operation. By utilizing the essentially true to form CSMA-CA/CD nature of the layer two protocol that duplex affords, we connect up in pairs and hold discussions while in full channel monitoring mode with large numbers of 'conversational-mode' packet users. This is an excellent way to operationally verify the correct repeater transmission to reception balance. (As well as any one user's similar balance.) When a single user can request the retry of a 'lost' packet to the sender for the entire loacl area network, then excellent free-flow of discussion can be realized with a minimum of lost time in distributing 'information* within the network.

Another and perhaps the most important advantage to duplex's efficient use of the uniquely radio specific portion's of the layer two protocol in use is the variable 'DWAIT' and 'FRACK' timing intervals. By encouraging users to remain at minimum neccessary values, the non-real time 'channel-hosts' (in our case the N6BGW-9 autoforwarding mailbox and the N6CXB-1 database and file-server,) can use elevated values. This assumes that everyday users are at a low-volume usage per each station and therefore are assured transmission priority over the the higher volume per station 24 hour hosts. The hosts will wait for a longer free-time interval before transmission of pending packet traffic. By also keeping the host's packet sizes high and custom tailoring the FRACK intervals among the hosts, the hosts are not known as the channel hogs that some of the more popular simplex packet hosts are.

IN CLOSING

The authors are hopeful that this paper and our 'local area network' implementation will encourage other groups to help out their so often forgotten 'user communities', by taking some of the ideas mentioned here and using them to enhance the neighborhoods they live in.