LINK LEVEL ADDRESS MECHANISMS IN AMATEUR RADIO PROTOCOLS

by

H. S. Magnuski, KA6M
311 Stanford Avenue
Menlo Park, CA 94025

In October, 1982, agreement was reached on a new Link Level protocol for amateur packet radio networks. One of the unique aspects of the protocol is the set of address fields used at the beginning of each frame. This paper reviews the types of addressing used prior to the adoption of the new standard, and explains in detail how the new address mechanism works.

On October 9th, 1982, representatives of the active packet radio user groups throughout the U.S. agreed to establish a new standard for Link Level connections in the packet radio service. This agreement unified the development activities of several diverse groups, and provides for a point-to-point protocol which can be used both in terrestrial and satellite networking, This paper will review how addressing was done prior to the adoption of the new standard, and will then describe, in some detail, how the new address mechanism works, and what advantages and disadvantages are gained by adopting it.

1.0 Background of HDLC Numerical Addressing

The first byte of an HDLC frame following the opening flag is always an address byte. This byte permits a maximum of 256 stations in a network to be addressed, and sometimes less, as byte 00 hex is often reserved for special purposes, and byte FF hex is usually a broadcast (all-stations-addressed) address. Since 254 stations is a reasonable number for active stations on a single frequency metropolitan area net (metronet), the initial implementor of HDLC-oriented transmission, Doug Lockhart) VE7APU, decided that no more complicated addressing scheme was required and he programmed his Terminal 'Node Controller board to deal with a single byte address. Doug's scheme initially utilized dynamic address assignment) and works as follows:

When a station initially came on the air, it would send a special sign-on packet to a central control station, which would assign the next available numerical address to the new user. If the user logged off or timed out, the address would be put back into the free pool and handed to the next station to sign in. Once assigned an address, the station would use that address in all outgoing packets, and the central control station would know by that address who originated each packet in the metronet. All traffic flowed through the central control station, so there was no need for one remote to address another,

Several groups in Canada and the U.S. desired to use Doug's code without a central station, and so the following form of static addressing was adopted:

Each user in a specific geographical area was assigned a numeric address, and that address was hard-wired into the outgoing address field of each transmitted frame. After a connection was made, the receiving station would lock onto the transmitting station's address, and all other packets on the channel would be discarded.

The author implemented a variation of this idea that would allow a simplex repeater to 'be used as part of a metronet. Some range of the address space, say 80 to 9F hex would be dedicated as repeater input addresses, and in repeating the packet the repeater would transform the address to the range A0 to BF. Thus a station would use address 01 for point-to-point work, and would use addresses 81 and A1 for repeater uplink and downlink, respectively.

The dynamic and static addressing; schemes described above served the early packet experimenters well, but it soon became clear that there were significant problems with each of the methods.

The dynamic addressing scheme failed if the central control station failed, and many users did not want to be dependent on a single central repeater, The scheme did not allow for multiple repeaters serving a given area, and failed if there was any overlap in central or remote station RF domains,

The static assignment method was just as bad, requiring an individual or club to prescribe addresses, and limiting the number of users to 62 if a repeater was involved. Visitors to an area might conflict with already assigned addresses, and using statically assigned addresses on a global satellite channel would be impossible,, Since neither of these methods would be workable over the long term, several proposals were put forward to correct the problem.

2.0 HDLC Extended Addressing

It became clear that if one could substitute a callsign for the hardwired numeric code, some of the addressing problems mentioned above would go away. The radio callsign is a unique identifier, and could be used to tag each packet being transmitted. So Terry Fox and others at AMRAD in Washington, D.C., proposed to utilize the extended addressing feature in HDLC to incorporate the transmitting callsign.

The HDLC standard permits an arbitrary extension of the address field through utilization of the least significant bit of the address byte as an extension flag. Basically, if the least significant bit (lsb) is zero, then the next byte forms part of the address field too, and one continues to extend the address field until a byte is found where the lsb is set to one.

This idea would be fairly easy to implement, and did not require much overhead in each packet. It also did not violate any standards related to the HDLC spec. Unfortunately, there were deficiencies in the method, particularly with regard to point-to-point addressing in a network, and flexible use of multiple repeaters in one RF domain. With just one address, a station could not target the receiver of an outgoing packet. Also, no provisions were made for dealing with multiple repeaters on the same frequency, and it became clear that this approach would have to be revised.

3.0 The New Link Level Standard

When the packet radio groups got together in October there was clearly an urgent need to agree upon a uniform link level standard, and to adopt one which would be sufficiently flexible to accomodate a variety needs and environments.

Two proposals were put forward at the meeting, and it turned out that there was more in common in the proposed solutions than there were differences. The first proposal was by the AMRAD and New Jersey packeteers, and it was a result of extensive discussions on how to implement an X.25 type of service suitable for the radio community. This proposal, called AX.25, called for a destination and source callsign in an extended address field. The other proposal was from the author, and it also specified use of two or optionally three callsigns in the packet, but positioned them immediately after the control field. Each party was willing to compromise, the callsigns were kept in the extended address field, an optional repeater address field was added to the AX.25 spec, definitions for bits in the sub-station and protocol ID fields were rearranged, and thus emerged the new standard protocol.

4.0 Components of the Address Field

As stated before, the standard protocol relies on the HDLC extended addressing bit to mark the end of the address field. The address field is composed of two or three seven-byte callsign address fields constructed in the following manner:

CHAR1 CHAR2 CHAR3 CHAR4 CHAR5 CHAR6 SSTD

where each call is left justified in the field, padded with blanks, uppercase, and each CHAR is shifted left so that the lsb is zero, except for the last byte of the address fields, where the lsb is set to one, indicating the end end of HDLC extended addressing. The first callsign field is for the destination or receiving station, the second callsign field is for the source or transmitting station, and the optional third field is for a repeater.

The SSID byte is of the form:

R 11 SSSS X

where R is the "repeated" bit, set to one only in the optional repeater address field when the packet has been repeated. The '11' field is reserved and set to ones. It may be used for control purposes if required by a local area net. The SSSS field is the sub-station code, normally all zeros except for situations where a person has more than one station on the air. The X is normally zero, unless this field is the last of the callsigns, in which case it is set to one, indicating the end of the variable length address field.

5.0 Uses of the Address Fields

The simplest case is a point-to-point connection between two stations. Station A puts CALL B in the destination field, and CALL A in the source field. Similarly, Station B puts CALL A in the destination field and CALL B in the source field. When A receives a packet it first verifies that the frame check sequence (FCS) is correct. Next, it scans the address field to determine if the address field length is either 14 or 21 bytes. Assuming that the field is exactly 14 bytes long, it checks for CALL A in the destination field and for CALL B in the source field, and discards the packet if there is not a correct match.

If a repeater is being used, the optional third address field will be utilized, and the following sequence of actions will take place:

The transmitting station extends the address field to 21 bytes and places the callsign of the desired repeater into the third field. The I've-been-repeated bit (R bit) must be set to zero. The repeater is constantly monitoring the channel for packets and discards all packets which do not have a correct FCS or do not have exactly 21 bytes in the address field. If the third address field is present and if the call exactly matches the repeater's call and SSID, and if the R bit is zero, then the repeater sets the R bit to one and retransmits the packet. The receiving station determines that the address field is exactly 21 bytes and checks that the R bit is set to one. If the R bit is zero it has received an uplink packet and should discard it. The receiving station then checks for proper source and destination callsigns, and accepts the packet on a proper match.

The receiving station knows how to reply to an incoming transmission because an examination of the address field will inform it whether or not a repeater was used, and if one was used, its callsign. This information is used to construct the reverse path to the transmitting station.

6.0 Limitations

This protocol is intended only for point-to-point connections directly or through a single repeater. The situation where multiple repeats are required is not covered and is the task of higher level protocols.

The protocol also does not specify when the reverse path should be constructed. An implementor may decide to establish the reverse path only on information contained in the initial connect (SABM) packet. If so, and if the connector changes paths during a session, the two stations might be using different forward and reverse paths to talk to each other.

The call signs add overhead to each packet (20 bytes at most), and these extra bytes will obviously slow down throughput. In doing link calculations, however, one finds that radio and modem turnaround delays are probably the most significant factors affecting overall efficiency, and that the extra overhead of the callsigns are justified in terms of the elimination of the requirement for a central control station, or elimination of any kind of connection-state information in the repeater.

7.0 Summary

The address mechanisms in this new protocol are fairly simple to implement and provide for a point-to-point connection in local area metronets, in backbone terrestrial networks, and in satellite channels. The addition of a single repeater option makes the protocol useful for communications in a limited geographical area. The protocol will serve as a building block for higher level connection-oriented protocols such as AX.25, and can be easily used for connectionless, datagram-oriented protocols such as IP/TCP.