Improving the Packet Mail Transfer System

Immediately Realizable Goals and Posturing for the Future

Brian B. Riley, KA2BQE
Radio Amateur's Telecommunications Society

The current amateur packet radio message transfer system runs extraordinarily well in view of a myriad of conflicts, mistakes, equipment failures, et al. that plague it every hour of every day. The volume of traffic is rising almost exponentially. It could reach critical mass and 'break' the system before technological solutions can be put into place. There are, however a number of things we can do within the constraints of current software and facilities to relieve much of the pressure and streamline the system for the future. Beyond that, there are several things that the codesmiths of the packet BBS world could implement reasonably quickly to further alleviate the problems.

Forwarding Headers

The current accepted NK6K header and its WA7MBL/KA2BQE variant should be replaced by a simpler format containing just the bare essentials. This format, shown below, has already been adopted by significant number of systems across the USA and several of the prominent international mail forwarding stations.

R:910710/1234z 12345@KA2BQE.#NWVT.VT.USA

No QTH, no Postal Locator, No advertisement for your BBS or software writer, nothing but time/date stamp, message number on that system and the full hierarchical address of the system. (The missing continent indicator is not a typo, but is a matter to be addressed in this paper.)

The justification for the removals/modifications is:

a. The major consideration for the header is the 'accountability' for the handling of the message demanded of us by the FCC and its analogous bodies in other countries; to wit, that the message clearly show each station that handled it and be identifiable in each station's record keeping. With the hierarchical addressing universally being used and included plus the local message number these requirements as well as our own routing analysis is fully served.

b. It would be most efficient to simply state the. time and assume 'Zulu" time and annotate if it is other than that. However, there is no strong standard on a worldwide basis so for the moment we should all try to set "Zulu" time and annotate with the "z" to so indicate it.

c. The indicator for '@:' becomes '@' and the '#:' is no longer needed.

d. The WORLI Packet Bulletin Board System introduced to us the concept of White Pages (WP), which, in short, snoops through every message passing through a system. It takes note of the originating individual and the station from which the message was entered. It also takes note of every change in a local user's status in terms of QTH, zip-code, home BBS, name. This WP operation then formats messages once per day making note of changes and sends them, generally, to a regional WP server, who then summarizes all changes he received from the region and send them onpv to the national server as well

as back to the local stations. This information not only categorizes and records individual users but quite obviously the BBSes themselves.

For purposes of this paper I saved about 3800 messages that passed through my system over a few months. I ran a utility program that I wrote to convert headers to the format shown above against the message base. The least shrinkage to any bulletin was 36 bytes (a locally originated bulletin quite obviously); the greatest shrinkage was 1898 bytes (56 headers, originating in Europe and making several 'laps' around the continent before coming to Canada and to NWVT.) The average shrinkage was 801 bytes. More important was the average shrinkage as a percent of the total message size which was 29%; that is the messages shrunk to an average of 71% of their former size.

While it would seem that by removing this information from the headers we are removing WPs source of information we must remember that this information also flows in from each system based on USER Data files entries as well.

Querying the WP system will get any information that is needed on any BBS. The updates that get sent out based upon changes to the user data bases will be sufficient to keep the local/regional/national WP Server stations apprised of QTH and ZIP of the various BBSes.

Note for the future: we might wish to add some info fields to the WP database for BBSes to permit some additional information like, adjacent network node:, bbs code type and version. The WP system is ripe for enhancement to allow considerable information to be stored at the regional level for its local users. The national level WP database might store general pointer information to the regional WP database where more detailed information would be recorded/retrieved.

Several well intentioned kluges have been enacted to send the BID and message originator in the "R:" header. This is extraordinarily redundant. If this information were stored in a RFC-822 like header once in the beginning of the message it would save dozens of bytes per message. The use of RFC-822 headers is discussed further down in the paper.

Hierarchical Addressing and Forwarding

The current hierarchical addressing and forwarding was implemented on the basis of a paper written and presented at this conference several years ago. It has provided us with a framework that has allowed the forwarding system to grow in leaps and bounds without the need for every system to know every other. Indeed, today most sysops don't know anything about systems located in 80% of the states located more than one states removed from them; Further, they don't need to know it either.

There are however, a number of problems inherent with the original system proposed and implemented which were pointed out even then by a few, and now are recognized by most. The original format being used is:

```
host.local-org.locale.country.continent
```

The problem occurs in that parsing for purposes of routing occurs from left to right, or most specific to least specific. Most of the current systems go through lists seeing if they know anything about a given system, then only failing that does it move right to look at

local organization, state, then country and continent. The problems that occur are manyfold:

- There is a region in France that had the identifier “MA”. Messages from the US that were supposed to head there were routinely being routed to MAssachussetts because the "MA" was evaluated before the "FRA"

- We had to get the kluge '#' prefixing numbered districts like Japanese prefectures and ' ' prefixing a British zones in order to keep American zip-codes routing from being applied. Once again because the numbers were evaluated before the country code.

- A ham like WINPR winters in Florida, summers on Cape Cod, and runs a BBS in both places. He goes through several weeks of misrouted mail each changeover. This is because all the systems along his paths “know” that he is one place or another and mail gets halfway through its proper route before it hits a system that didn’t get the word and the message gets turned around. Again, this is because the "W1NPR" gets evaluated before the “MA” or “FL.”

The solution is obvious; the BBS code smiths get to work and get it all converted over to parse from least specific to most specific. To my knowledge PRMBS is the only system coded that way. But, even should all the other systems be recoded by Thanksgiving, the lead time on getting full world wide distribution to all systems will be horrendous.

There is however an interim solution which not only almost solves the problem but it makes the setting up of forwarding tables far easier for the sysops and most probably will speed up his forward process.

a. Most important, remove all individual user and BBS CALLS from the various places in your forwarding route files for any regions beyond your own. Don’t give the BBS a chance to send WINPR the wrong way, make it look beyond to the state first.

b. Remove all the continent designators. The original seven proposed continent designators are not satisfactory. There are numerous problems in Latin Ameirca, Oceania and Asia posed by these. There was a proposal by Tom Clark, W3IWI, to either dump the continent designator all together or go to a four character designator and enhance the continent set. It is unfortunate that his proposal has met with only modest success and that success has unfortunately caused more confusion than it has helped anything. In the original proposal the country list was essentially an incomplete ISO-3166 ALPHA-3 (3 letter alphabetic designator set). It contained everything we need now. There are but slightly over 200 countries on the list which is far less entries than most PBBSes currently keep with all of the individual calls that are recorded.

c. When the codewriters finally correct all of the code to parse properly from least significant to most significant element of the address then we can switch to the ISO-3166 ALPHA-2 list and drop yet another character and also have addressing consistent with The Internet.

Message Headers

The time has come to recognize that we need to carry a certain amount of information in the text body of the message. In some of the above you can see several places where we carry the same information over and over in the "R:" headers that could be stated

once in the main body and be done with it. Since 1986 the PRMBS BBS code (thanks to the farsightedness of my colleague Dave Truli, NN2Z) has utilized a major sub-set of RFC-987 (RFC-987 is a common subset between RFC-822 and X.400) headers at the head of the text body of the message. While some of it was more show than useful, later a further subset was defined, its utility has time and again proven out because it provides for a place to pass information for which there was no provision in the current RLI defined forwarding protocol. Since RFC-987 parsing rules were fully implemented, additional fields could be added by any system and passed transparently through being acted upon only by those system that knew what to do with them. This highlights the chief flaw in the RLI protocol, that all systems must change to pass information on the SEND line and not lose it.

The minimum mandatory sub-set should be:

Date:
To:
From:
Message-ID:

In addition these should be handled, when present:

Expires:
Reply-To:
Priority:

Parsing headers should be done within the Internet fashion; that is to say, any line prior to an empty blank line which has a text string terminated by a colon ':', and a "space" is considered a 'header' and that the string will be examined, its objects parsed and acted upon if recognized, and passed through, untouched, if not.

There currently is a very real problem with worldwide distribution of messages of time value. They often wind up being delivered 10 days after their usefulness has expired. If we alter the manner of message creation for the user such that when he enters the message he is prompted for title, then 'Expiration?' and give him the option of entering a number (14 - fourteen days) or a date (1 1/14/91) or just hit enter (no expiration) In this fashion at least the users who regularly originate dated material will have it available and when their bulletins hit out towards the fringes and their time is up, they will die without requiring human intervention. The time it saves the sysops is time that can be used to take the trouble to examine manually other bulletins to see if they are worth keeping around, instead of resorting to blanket '4 days and gone' type bulletin management.

More on Addressing and Forwarding

There is a need to modify our ability to address mail. A user on a system in NJ should be able to send a personal message to KA2BQE@VT.USA. The systems should be able to route that message to the WP server for VT (most systems will now do this). What is needed is code at the server to recognize that a particular piece of mail needs re-addressing. In past, if there was anything in the @ field we left it alone. The codesmiths need to now discern between a general geographic location and a specific address (ka2bqe@vt.usa versus ka2bqe@w1 koo.vt.usa) so the bbs system knows whether or not it is permissible to re-address the message.

If the above suggestions about paring down forwarding files are followed and if all systems come on-line with a reliable WP server, it can work right now. As part of a test, KB7UV, in the New York City area, addressed a message to "KA2BQE@VT.USA", it arrived in Northwestern Vermont in a timely fashion. The "VT.USA" was passed into WA2SPL in Alburg, VT, the regional mail server. and, most important, the regional WP server and it found its way correctly to W1KOO and then to KA2BQE.

Where the codesmiths can help out would be to modify bulletin distribution so that, say, a user in England planning to visit the state of New Hamphsire in the USA could send a bulletin SB INFO @ NH.USA and request local repeater information from hams in New Hampshire. In other words the message gets treated like a single message to be routed only until it enters a BBS system that identifies itself with NH.USA and then it gets treated as a flood bulletin.

Logically this solution needs to be extended to multiple addressees for bulletins. That same user in England may be planning to visit Maine, New Hamphsire, Vermont, and Northern New York. It is absurd to send four bulletins with four different BIDs. Likewise if the same BID is assigned to each bulletin the odds are just about 100 percent that only one bulletin would get through. Again, the RFC-822 style headers could be used:

To: info@nh.usa, info@vt.usa, info@me.usa, info@nny.ny.usa

Yet another need that would be nice to be able to handle would be the ability to include an individual user in the "To:" list and have a copy 'spun off' to him when the bulletin arrives in the area adjacent to that users address.

Improved SID Exchange

The Smart System ID (variously know as the SID, 'brackets message', or [...] exchange could be altered to produce a more efficient exchange. Currently, upon receipt of the CONNECT acknowledgement we await a BBS prompt. If during that wait we receive a SID, we continue to wait for the prompt and then return send our SID and await a BBS prompt, again. This double exchange can take up quite a bit of time. While it does not use up network bandwidth, it does chew up connect time to the BBS, at least 20 seconds and as much as several minutes. As an alternative, each system could keep a record of the other system's capabilities, as conveyed by the SID, and simply start sending as soon as it gets the "CONNECT" acknowledgement. SID would only be exchanged when a given station's records indicates that it has no knowledge of the connecting other station.

The principles of the new exchange are as follows:

a. upon receipt of a SID waiting for BBS prompt or command any station will return their own SID to the other side.

b. every station starts out with clean record of what its connecting BBSes can do. As such, when each is connected to or receives a connect from the other, it is treated as a conventional system and waits for/issues a SID. Once the exchange takes place the record for that system will reflect current capabilities and dictate future actions.

c. When the system connecting out has had a change in its capabilities, such that it must issue a new SID, it simply will do so as its first command and receive a SID in return.

This process automates registration of capabilities with corresponding systems eliminating constant re-negotiation every session.

d. When a system connecting out has lost its records or for some reason is starting over, it will connect to a system that may 'remember' it as being capable and will not send a SID. In this case when the station connecting station reaches timeout, it sends its SID and waits one more time for response. Again a waste of time but only happens in this startup/restart circumstance.

e. When the connecting to system has lost its record, the SID will show up in place of the response to the first send command or to an F> reverse poll. The station should disconnect and update that system's record. Again a waste of time but only happens in this startup/restart circumstance.

Improved Message Transmission/Bulletin Negotiation

There are several things that need to be done here;

a. We currently send an Sx, where x is the message type. The types accepted are P for private (not that there is privacy but it more or less indicates person to person, or person to serve messages not really of interest to users other than the addressee), B for general interest bulletins, and T for traffic (NTS, MARS, RACES,ARES). Some systems also support an untyped message; that is a message from one user to another but is visible and readable by others.

Given ham radio's stated dedication to public service, one thing we have been remiss in providing is a method to move messages based upon some priority. We should extend the Send command to three characters Sxp where 'p' is the message priority (Normal, Immediate, Urgent).

For those systems with untyped messaging the underscore would replace the blank so an urgent untyped message would be S U. No typing would be synonymous with N for normal. Thus SPN and SP would be identical.

The codewriters will need to provide the software systems with a series of options that will have the BBS behave in some sysop setable fashion when it detects the presence of non-routine traffic.

This will be open to abuse. It should be handled by the sysop as a hold function of some sort if he finds a user regularly abusing the function. There is no other way to manage this as anyone can download and run PBBS code and do a simple "MYCALL" change for connect purposes fooling any user record recorded privilege setting.

To handle non-compliant systems we again go the RFC-822 headers; using "Priority? buried in the message and passed through other systems recoverable by an RIFC-822 saavy system.

b. The TO field should be expanded to at least 32 characters for transmission so that bulletins can be addressed with more intelligence a' la usenet. For disk record efficiency, simply defining a total address space of 60-80 bytes and store the USER@ADDRESS or TO-SUBJ@DISTRIB with the "@" embedded in the space. When a bulletin is passed it would have a longer TO field and a shorter DISTRIB field and when a message is passed it has a shorter TO-USER field and a longer ADDRESS field. This

can be implemented transparently **as current** TO **field parsing for all** systems will truncate and 'correct' the excess without **blowing up. RFC-822 saavy systems would** recover the full TO field

c. The bulletin negotiation currently consists of;

SB TO-SUBJ @ DISTRIB < FROM-CALL **$BID**

This is then responded to with an OK/NO. I propose that we **extend the negotiation to** include the title as well

SB TO-SUBJ @ DISTRIB < FROM-CALL $BID
(title - subject matter - description}

This will permit the receiving system to have more latitude in the determination of what he wants to do with the bulletin. Of course there is no substitute for reading the whole bulletin, but a combination of the title and the remaining information makes an automated disposition decision more refined.

d. Along with the above modification we should universally adopt the 'R' descriptor in the SID from **AA4RE** for extended response which permits us to say RJ to reject a message for as opposed to simply NO which implies its a dupe. We may treat it like a NO, but it will allow us to do more should we care to code it.

In summation, by simple modification of the **"R:"** headers, removal of the continent designators, adoption of the full **ISO-3166** ALPHA-3 list and modification of the forwarding tables removing non-local individual calls, sysops can right now, without software modifications, significantly decrease message sizes and improve message throughput. Further, with some comparatively simple enhancements by the software codewriters, additional significant throughput yields can be realized.

The adoption of the RFC-822 style headers at the head of the main body of the message will provide a standard vehicle for passing message control information across any level system even older non-compliant systems. Because it is an existing international standard it provides compatibility across many messaging systems.

acknowledgements

This paper and its suggestions have grown out of numerous conversations and correspondence with many members of the amateur packet world. I wish to thank J. Gordon **Beattie, N2DSY,** Hank **Oredson,** WORLI, Tom Clark, **W3IWI,** Roy Enghausen, **AA4RE,** Tom Moulton, **W2VY,** Frank Warren, **KB4CYC,** for their contributions to the ideas presented here. I would also thank Ralph Stetson, KDIR, my sounding board and chief critic while much of the ideas were fleshed out.

brian b. riley, **ka2bqe**
po box 188, **harvey** road
underhill center, **vermont 05490-o 188**

packet:ka2bqe @ w1 koo.vt.usa
cis: 71420,3543
phone: (802) 899-9922 (msg) / (802) 899-4527 (voice)