MORE AND FASTER BITS
A Look At Packet Radio's Future

Bdale Garbee, N3EUA

Hewlett-Packard Colorado Springs Division

ABSTRACT

The biggest problem facing amateur packet radio today is the inability of the ham community to envision the breadth of possibilities that exist once higher speed modems become available. This paper attempts to survey some of the applications popular in other networking environments, and comments on their possible use in the amateur service. In addition, a preliminary report on work in progress to develop multi-megabit per second connections on the microwave bands is presented.

1. Background

Much of the effort expended to develop the current amateur packet radio network has come from people whose primary orientation is towards radio and *'communications", and not towards formal networking, or even "computer communications". While we owe a debt of gratitude for the efforts and commendable successes of this group, it should be equally obvious that we need to take a 'networking view" of our future if we are to successfully advance the state of the art in amateur packet radio.

Even in the world of professional networking, there are many different application scenarios, each involving a different set of priorities and solutions. We must be very careful to make sure that decisions we make about protocols are reasonable given the environment in which we operate. Because our environment is unique, some techniques, benefits, and restrictions of wired networks do not apply to us. In addition, we need to be aware that sometimes a solution used in one kind of wired network applies closely to our needs, while at other times another kind of wired network more closely resembles our situation. More about this later.

It is unfortunate that the evolution of grass-roots BBS networks such as Fidonet have involved as many false starts and wrong turns as they have. This is particularly true when we consider that these efforts have in many ways paralleled the early efforts of such seminal groups as the Arpanet research community. Here, as in amateur radio, it is unfortunate that the movers and shakers are primarily *not* networking people, and therefore do not have at their disposal the set of historic knowledge and lore that would have allowed them to avoid many of the pitfalls they have encountered.

What impact does this have on us in the amateur packet world? Merely that we should not be afraid to evaluate and adopt existing hardware and protocol standards *when* and *where* they make sense. Lets *not* fall into the "Not Invented Here" trap, dooming ourselves to repeat all the mistakes of those who have come before us.

We can, and should, do it better in the amateur world. What we need is to expand our vision of the future to include applications and technologies that may not at first seem particularly relevant. In order to do this, we need to know what has already been done, and build from there. This paper attempts to provide an overview of applications, and some incentive for further study. The rest is up to you !

2. Application Overview and Bandwidth Requirements

In the world of wired networks, it is common to segregate network applications into two categories: those that require noticeable bandwidth, and those that do not. The term "noticeable" is wonderfully vague, and has come into wide use as a result of the fact that the base bandwidth available is fairly high, typically 10Mbits/sec. Using Ethernet, it is not uncommon to find an entire software development lab running Telnet and SMTP-based mail on a single shared cable with little degradation in response time.

In the amateur packet radio world, we are currently somewhat less fortunate. The defacto standard data rate is 1200bits/sec half duplex. This is over 800 times slower than coaxial lans at 10Mbit/sec. It is therefore not surprising that we should be substantially more concerned about available networking bandwidth on our RF channels than most users are on wired networks. In order to better understand why we need faster data rates, let's take a look at some existing applications, and make some qualitative assessment of their characteristics.

Electronic *mail* is wonderful from a networking standpoint, because it is typically both a low bandwidth application, and a fully non-real-time application. By this, we mean merely that the volume of electronic mail generated by an average user it typically very small with respect to the available network bandwidth, and that the user in a rational electronic mail scenario is able to locally create and queue mail, and thus is not constrained to wait in real time for a mail transfer to occur. The computer takes care of it in the background. From this we can conclude that electronic mail is unlikely to be a burden on network bandwidth. The high degree of success achieved by the current PBBS mail forwarding network, using only 300 baud HF and 1200 baud VHF link, supports this conclusion.

The use of a *virtual terminal* protocol is at almost the exact opposite extreme. Here, the essence of the application is real-time use of a remote computer by a user as if on a local terminal. In this case, the response time to individual keystrokes typed on the keyboard is the limiting factor. We have all experienced the frustration of trying to access a remote PBBS system, and the agonizing delays associated with waiting for the system to "come back". This is therefore a class of application which we should consider supplanting with different ways of operating (local mail editting and queueing as opposed to online entry of mail text into a PBBS, for example). If we really need to continue supporting virtual terminal services, and I believe we do, then this may be an ideal application of much higher speed modems. Not because the volume of data we wish to transfer is high, but because of the need for fast round trip times, and minimized channel access delays.

File transfer protocols fall somewhere in between mail and virtual terminal access. This is due in part to the possibility of using either an FTP-like real-time user interface, or a background batched mode of operation. An interactive interface to a file transfer makes the user aware of the time involved to complete the transfer, and therefore the network bandwidth. In this case, it is the need for the network to move large blocks of data quickly that causes a desire for more bandwidth, not necessarily a need for short round trip times. On the other hand, a batched, background file transfer protocol allows the user to specify transfer parameters, and then proceeds to completion without the need for user level intervention. If our use of file transfer protocols is primarily real-time, we care a great deal about the speed at which they occur. If they are primarily background tasks, we are likely to be less concerned, since we can go do other things while the transfer takes place. Thus, file transfer is an application where our need for speed can be very high, depending on our usage patterns.

There exists an entire class of protocols which have not to date been used by the amateur radio community, doubtless because their data transfer requirements are so high as to make their use on a 1200 baud network essentially impossible. Perhaps the most well known of these protocols in the wired network world is NFS, the Network File System. In this application, a file system on one computer can be made available to another computer system transparently. A user on machine A can be given full operating system level access to a storage device on machine B as if it were directly attached to his own system. This is a natural evolution of, and very desireable replacement for, simple file transfer protocols in a homogeneous computing environment.

When we talk about local file repositories in the amateur world, we tend to mean PBBS's with the ability to 'upload" and "download" files, or the idea of an FTP server with massive amounts of disk space available for storing files. I would like to propose that we instead try to consider the notion of a local NFS server, making available sets of related files as file systems, allowing a multitude of varying access mechanisms based on need. In order to make this work, we need *much* higher speed networks. This is because NFS in essence replaces a local disk and controller, and so is expected by the user to perform at rates approaching those of fixed disk controllers.

If it is difficult to imagine using NFS in an amatuer environment, substitute any of a number of other exciting high bandwidth-requirement applications (digital voice, digital video, etc), and the analysis will be very similar.

3. So, Where Do We Need Speed?

Given the preceding brief evaluation of existing high-usage protocols, what should our direction be with respect to application of higher data rates? The traditional model of "the network" assumed in the amateur packet radio community seems to be one of high speed backbones connecting lower speed local area networks. This is in direct opposition to the existing model in use in the wired network world, leading one to wonder if it is correct.

In a typical corporate or university network, fairly high speed local area networks (LANs) are built within a functional area or building, typically using Ethernet technology at 10 Mbits/sec. When it becomes desireable or necessary to link these LANs toegether, gateways are installed which operate at the full LAN speed on one port, and at some reduced speed on the linking port. Typical data rates connecting these gateways are those available from commercial common carriers. *2400* or 9600 baud via dialed phone lines, 56kbits/sec or approximately 1.5Mbits/sec via leased lines, and so. obviously, these data rates are much lower than the LANs, leading us to recognize a model of high speed LANs connected to low speed backbones. Why the discrepancy with amateur packet radio's model? To answer this, we need to look at how the previously mentioned protocols are typically used in a working environment.

Mail is equally likely to be sent to an associate on the local LAN cable, or an associate at a remote site, perhaps with a tendancy towards higher local usage in environments where electronic mail is used as a primary communications tool within a working group. Regardless, since it is a relatively low network bandwidth application, the impact on required speed is very low.

Virtual terminal access is interesting because while it tends to not produce very large quantities of data transferred, there is a strong tie to perceived response time at the user level. Thus, we may require very high available network bandwidths in order to produce short round trip times, but not be using a very large percentage of that available bandwidth. This is the classic example of an application where speed is an issue not because of the volume of data to be transferred, but because of the need to transfer small pieces of data very quickly. Virtual terminal sessions are typically weighted very heavily towards local use, since a programmer might have access to several local machines within his group, but only infrequently need to log in to a computer system at a remote site.

File transfers are heavily weighted in actual practice towards local usage. A typical software engineer working in a corporate R&D environment might use several different computer systems to accomplish his work, frequently moving files back and forth between the various local systems. That same engineer might find it useful on occasion to obtain a piece of software or a document from a remote division of the company, or from some other remote site. While it is difficult to characterize the actual split between local and remote file transfers, it is fairly obvious that more bandwidth will be needed locally than on the connecting backbone.

The NFS protocol draws both virtual terminal and file transfer concerns. Because it serves as a virtual disk drive interface, users are apt to expect it to perform with a response time similar to that of convential fixed disk drive interfaces. In addition, since it is frequently used to load executable programs and/or data files from a remote disk subsystem, the amount of data transferred can be very large over short periods of time. Thus, in order to make use of protocols like NFS, packet voice, or packet video, we need ***lots*** of bits per second locally. The remote use of NFS is almost non-existent.

It appears that those discussing high speed backbones for amateur radio are for whatever reason anticipating that almost all traffic will pass out of the local LAN. This is without precedent in the wired network world, leading one to suspect that it may not be an accurate assumption.

4. Faster Modems on the Market

As mentioned earlier, the defacto data rate in amateur packet radio is 1200 bits/sec. Some time back Kantronics introduced a 2400 bit/sec modem option on their TNC's, but this has been slow to catch on probably because a factor of two increase in data rate,is more than offset by the non-standard nature of the modem, concern over connection with existing stations.

Work done by K9NG and others has resulted in the availability of 9600 and 19200 bits/sec modems, but these have not come into wide use as a result of the need to modify radios in order to work effectively at these speeds. This situation may change if TAPR follows through with its plans to develop RF transmit and receive modules specifically for use with K9NG-style modems.

Commercial vendors such as GLB and AEA also produce 9600/19200 baud modem-s, complete with radio gear. Unfortunately, the relatively high cost of these unit, coupled with the recently justified uncertainties about the future of the 220Mhz band, have prevented their widespread use.

A year ago, WA4DSY published his design for a 56kbit modem, which has met with great enthusiasm. However, the cost of a fully configured WA4DSY modem including RF gear and antenna can easily top $700, and there is very little available in the way of digital hardware to drive the modems at 56kbits.

5. A Project in Progress

Recognizing the need to develop much higher speed data transmission technology in order to expand the base of available applications for packet radio, Glenn Elmore, N6GN, and I have begun a project to integrate Ethernet controller chip technology with microwave transceivers. The eventual goal is to produce a unit which attaches directly to an off-the-shelf Ethernet controller card via the 15-pin connector, and which terminates in perhaps a 4-foot dish at 10Ghz or 24Ghz on the RF end.

The immediate advantage of this approach is the ubiquity of the 15-pin Ethernet connector, which appears on everything from PC plug-in controller cards to high-end engineering workstations and mainframes.

Projected cost of each unit is about half of a fully configured WA4DSY 56kbit station, providing 10Mbit/sec data rates point to point. Because many Ethernet controllers are already supported by software such as the KA9Q TCP/IP package, these units will be immediately useful without the need for software development, which has been one of the problems with digital interfaces for the WA4DSY design.

While the characteristics of microwave transmission force some restrictions on the application of this concept, we feel that it will be a substantial boost to application developers to have this kind of data rate available.

6. Issues at Microwave Frequencies

6.1. Directivity

In the wired network world, there tend to be many networking hardware technologies in use, but they can typically be broken into two distinct categories. Some are busses such as ethernet which may contain large numbers of hosts sharing a single connection media, with the resulting ability to broadcast, and the need to share available bandwidth. Others are point-to-point links, which typically connect two systems only, with no broadcast capability, but with the full channel capacity available for each link.

Modems such as the K9NG and WA4DSY designs are typically operated on the VHF and UHF bands, where omnidirectional antennas are quite feasible. Because of this, they are ideally suited to broadcast-oriented network design, and will continue to be the best solutions where many stations need to be able to communicate with a central repeater site.

At the microwave frequencies required for truly high-speed data transfer, improving signal fidelity usually involves increasing the gain, and therefore directionality, of the antenna system. This is because of the relative difficulty of generating large transmitter output powers at microwave frequencies. The extreme directivity of this approach, coupled with the line-of-site nature of microwave communications, means that point-to-point network topologies are necessary. This is ideal for backbone links, since they are typically "mountaintop to mountaintop" anyway, and we get the added benefit that local users can't interfere with the backbone since the line of sight path is typically beyond their reach. It is also satisfactory for specific high-bandwidth requirement local links, such as between two users involved in packet video experimentation.

6.2. Addressing

In order to make this project economically feasible, it is necessary for us to base the digital interface design on the readily available, fairly inexpensive VLSI Ethernet interface chip sets available from National Semiconductor and others. The result of this is that we will be forced to use the 48-bit Ethernet address and packet framing format as the basic bit protocol in this project. Widespread application of this technology may dictate a change in the FCC rules to allow non-AX.25 addressing in packets transmitted on amateur radio. Gross hacks are no doubt possible that would allow some encoding of the user callsign to serve as the 48-bit address, but this seems both counter-intuitive and counter-productive, in part because it might entail modification of host networking software.

There is an existing standard for the issuance of Ethernet addresses that causes them to be tied to specific hardware. This should allow them to serve as unique identification for monitoring purposes, assuming that FCC or volunteer monitoring of a 10Mbit signal on 24Ghz is even possible.

6.3. Data Rates

While our initial experiments are targetted at data transfer rates of 500kbits/sec to 2Mbits/sec, development of a digital transceiver capable of 10Mbits/sec that will function over a 100 mile line-of-sight path seems feasible. An initial prototype at the lower data rate may in fact be operational by the time this paper is published. In any case, we hope to successfully complete the project sometime in the next year.

7. Conclusion

The ham community has embraced a set of applications on packet radio that mirror the capabilities of other familiar modes, and other similar services such as the telephone BBS network. If we want to remain on the cutting edge of technology, thereby eliminating the now-too-common "packet burnout" syndrome, we need to expand our horizons to encompass applications that will soon be possible with advances happening *right* now in the packet community.

The earlier qualitative look at bandwidth utilization on wired networks, backed up by some quantitative analysis of the HP Corporate Internet, leads to the conclusion that we need lots of bandwidth locally, and progressively less bandwidth as we get farther and farther away in a networking gateway sense. On the other hand, hams are notorious for wanting to participate in personal DX, and therefore we may want more than a mere minimum data rate on our long haul networks. However, we should concentrate on our local networks, because that's where the bulk of our utilization will be, and because it's going to be easier to rationalize spending big bucks on hardware that we use personally every day than it will be to rationalize big bucks on mountaintops.

For the majority of our local network activity, modems of the K9NG and/or WA4DSY variety are probably* the best choice. This is primarily because they operate on bands where omnidirectional antennas are possible, supporting the idea of multiple-access repeater and central facility sites without the need for lots of duplicate hardware. In order to make these units practical, we need to develop some easily duplicated, low cost RF modules. We do not, as yet, have a "complete solution."

We may soon have available RF gear that plugs into an Ethernet controller and provides lOMbit/sec data transfer on line-of-sight, point-to-point links. This may be the ideal technology for backbone links, and will probably also find application in local nets housing "power users", much as the 386 AT-clones have infiltrated corporate environ men ts largely dom inated by slower speed machines.

So, what will our future network look like? Why don't we try for 9600 baud "average ham" hardware based on K9NG modems and the forthcoming TAPR RF hardware, 56kbit local capability for power users, 10Mbit microwave for the esoteric high-bandwidth local applications and backbone links, and an applications mix that includes remote file access, digital voice, and digital fast scan TV! Anything is possible, with enough bits!

8. Acknowledgements

The material in this paper has been influenced by several individuals. First and foremost is Glenn Elmore, N6GN, of the Hewlett Packard Network Measurement Division. Glenn is developing the actual RF/Microwave hardware that we hope to run at 10Mbits/sec in the near future. Mike Chepponis, K3MC, has demonstrated considerable insight about the possible futures of amateur packet radio during our electronic mail discussions of the last year or so. Thanks also to Bob Gobrick, WA6ERB, and the many members of RMPRA who provided input at the 1988 RMPRA PacketFest about the kinds of applications that they hope to see running on packet radio in the future.

9. References

Goode, S., *Modifying the Hamtronics EM-5 for 9600 BPS Packet Operation,* Fourth Computer Networking Conference, 1985.

Heatherington, D., A *56 Kilobaud RF Modem,* Sixth Computer Networking Conference, 1987.

Ingram, D., *An Expandable Microwave Network for Multimode Communications,* First Computer Networking Conference, 198 1.

Tannenbaum, A., *Computer Networks,* Prentice-Hall, 1981.

National Semiconductor, *Series 32000 Databook,* Data Sheets on the DP839X Ethernet Controller Set, 1986.