# Design of a Next-Generation Packet Network

Bdale Garbee, N3EUA

Current amateur packet radio experience centers around 1200 baud half-duplex AFSK operation for both local and long-haul use. Technologies are discussed that have the potential to effect a fundamentally different environment for the next generation of packet networking. A preliminary proposal is made for an example network configuration in the Rocky Mountain Region. Applications, ramifications, and problems facing the new network are discussed.

## Background

While the existence of 1200 baud half-duplex as a defacto modem standard for packet radio has undoubtedly proven extremely beneficial to the rapid growth of this operating mode, the fact that our standard has been 1200-baud has probably caused many would-be packeteers from the communications and software industries to go look for something more exciting to play with. This has probably delayed the development of higher-performance packet links, since the existing packet user base has for the most part appeared content with the status quo.

There is a substantial difference in performance between existing wired networks, and the technology that we are using on our RF links. So substantial, in fact, as to make the applications possible with contemporary packet radio qualitatively different from those possible on commercial and academic networks. In particular, we have for the most part limited ourselves to remote login, file transfer, and electronic mail capabilities.

As a group, we have become so accustomed to 1200-baud operation, that there now exists a substantial mental hurdle to be overcome before any real advancement can be made in packet radio's future. We need to open our minds to the possibilities that exist with technology that is available **today**. We also need to break the traditional "I'll do it my way" reputation that surrounds amateur radio, and learn how to work together if we intend to plan and build substantially faster networking resources for amateur use worldwide.

## Goals

If we are to discuss methods for dramatically increasing the throughput of packet radio networks, it is important for us to have a consistent vision of what our goals are. When we are evaluating the potential impact of any specific change or improvement in our network, we must do so with clear knowledge of what we are trying to achieve.

The most significant reason for wanting to increase the speed of our networks is the potential for experiencing qualitatively different, really exciting applications. One of the problems with our current 1200-baud technology is that interactive applications are **slow**. This leads us to speak of how inappropriate packet radio is for keyboard to keyboard QSO's, and to talk about how wonderful background tasks like electronic mail are. In reality, we are first and foremost amateur radio operators, and that tends to imply that we like to be communicators. For many hams, communication equates with "live and direct", real-time radio communications.

Many of the modes we already enjoy in amateur radio have digitally oriented counterparts that we can and should explore. Amateur TV, or fast-scan, enthusiasts might enjoy a network of live HDTV transmissions, all running on suitably fast shared digital backbones. Repeater cross links might become much more common with higher fidelity through application of digital voice transmission over the same shared digital network.

But even if all we can envision is the same set of tasks that we perform with packet radio today, wouldn't it be neat to have them run 50, 100, or even 1000 times faster?

## Technology

If we are going to make any qualitative difference in packet radio operation, we will need to migrate to higher speed modem and radio technology. A lot of work has been underway in this area in the last year that is worth reviewing.

Perhaps the most widely publicized project for higher-speed packet links has been the TAPR development project aimed at producing a 9600-baud direct FSK packet modem and radio combination. First unveiled at the Dayton Hamvention, the initial units promise black box operation, with HDLC data in and out on one side (or RS-232 with an optional built-in TNC board), and an antenna connector on the other side for 2 meter operation. Compatible with the K9NG design TAPR has offered in kit form, and which has formed the basis of at least one commercial offering, the TAPR unit is destined to set a new standard for "low-end" packet radio operation. If operated in conjunction with a full-duplex digital repeater, 9600 baud can form the basis of a metropolitan packet channel that will provide substantially improved performance with existing 1200-baud applications.

The WA4DSY 56 kilobit modem design has seen gradually increasing acceptance in the last year. Part of the reason this modem has not seen wider use (also true of the K9NG 9600 baud design, and a fundamental driving force in the TAPR 9600 baud project) is that it is not a complete, turnkey solution. In order to put a 56kb station on the air, one needed to be willing to modify TNC-2 hardware or experiment with PC plug-in cards to get a digital data stream, and then had to locate and purchase an external transverter to move the modem's output from the 28Mhz region to some frequency of local use, such as 220Mhz or 430Mhz. While none of these problems are insurmountable (witness the number of stations on the air in Georgia and elsewhere!), they have undoubtedly hindered widespread use of this modem design. The prospects for this unit will improve with the availability of other hardware that will be described shortly.

The development effort that is likely to have the most significant impact on packet radio in the near term is the N6GN project to design RF modems (or true "packet radios") for the 900Mhz and 1.2Ghz band, that provide 250 kilobits per second using direct FSK. With an estimated parts cost of under $200 per unit, these may well be a good choice for our next long-haul backbone step. And eventually, they show great promise as the technology of choice for end-user connection to the network. By the time this paper is published, prototype units should be in the demonstration and testing phase.

When we are investigating backbone links, another project worth knowing about is another that N6GN has been working on. With nudging from N3EUA, N6RCE, and others, the design of a 10Ghz microwave digital link, using cheap surplus gunnplexer modules, and providing data rates from 1 to 10 megabits per second, is now in the PC Board layout phase. Prototypes were shown at the Dayton Hamvention in April 1989, operating at 1 Mbps, and work is progressing on development of kits for proliferation of this technology. Microwave links are particularly attractive for backbones, because of the relatively low power levels required, the availability of very large amounts of system gain with moderately sized antennas (dishes), the fundamentally point-to-point nature of the technology, and the resulting possibility of frequency reuse. Whenever the geography and environment allow, use of microwave bands for amateur packet radio backbone links should be strongly encouraged.

Now that we've talked about all this wonderful RF gear that's just around the corner, and which can provide substantially faster raw data link speeds, how do we generate and accept bits that fast? It's a combination of hardware evolution, and software innovation. In the digital hardware arena, there are two projects underway that are worthy of note.

Two years ago, the PS-186 packet switch board was unveiled. For a variety of unfortunate reasons, the units have never been produced. That is about to change. AEA is expected to be shipping units in the very near future, and a port of the KA9Q software is underway by N3EUA. Bluntly, the PS-186 is the most likely packet switch candidate for the next round of packet network improvements. Based around an 80186, the unit provides 4 high speed HDLC channels with full DMA, capable of over 1 Mbps per channel, full duplex. In addition, the unit draws very little power, and includes features like a "firecode" reset circuit for "pushing the big red button" remotely, and it's single-board configuration makes it ideally suited for mountaintop packet switch operation.

Simultaneously, Mike Chepponis K3MC has completed the design of a plug-n card for the IBM PC and compatibles that will provide two or more channels of fast HDLC, capable at least of handling the N6GN 250kbps radios. Sporting an onboard V40 microprocessor, the design allows migration of part of the networking code onto the card and away from the main computer, making for faster and more efficient network operation. In addition, it appears that it will be possible for this unit to run in a standalone configuration, providing some subset of the PS-186 functionality when less capability is required for packet switch sites.

With these two pieces of digital hardware, arriving at almost exactly the same time as the N6GN faster RF solutions, we are suddenly faced with a complete hardware solution for faster packet operation at both the end user and network backbone levels! And most excitingly, the cost of a PS-186 based packet switch site with a 900Mhz/1.2Ghz or 1 0Ghz RF setup for the backbone link(s), and one or more channels of contemporary or new-technology RF for local access, has the potential to cost about the same number of dollars as a site with an equivalent number of ports of 1200-baud TNC, firmware, and commercial FM radios!

It is probably not surprising that most of the high-performance digital development is happening hand-in-hand with creation of device drivers and improved networking capabilities in the KA9Q TCP/IP software implementation. To date, the KA9Q package is the only software that has been available to test packet hardware at speeds faster than the 9600 baud range. This is in part due to the ready availability of full source code. But even more importantly, the TCP/IP protocol suite is a mature set of protocols, that have evolved in an environment of heterogeneous computer systems connected by a wide assortment of networking technologies at various speeds. Since this is a pretty accurate description of the emerging amateur packet network, we should not be surprised that the TCP/IP protocols also work well in our environment! Particularly since the KA9Q package provides all of the common AX.25 functionality, compatibility with existing NET/ROM networks, and a fairly reasonable programmer's interface for development of new applications. At least initially, the KA9Q package will be the software of choice for operating on our next generation network.

## Proposal

In order to illustrate how we might combine the emerging hardware and software technologies to the problem of building our next generation "dream network", let's take a look at a real situation. Without being overly specific and detailed, here is a proposal that will be made for the Rocky Mountain Region, centered around Colorado.

The existing packet backbone is composed almost entirely of NET/ROM compatible nodes operating either single-ported on 145.01 Mhz, or in some cases dual-ported to 446.8 Mhz as a backbone frequency. There is a dense population concentration along the Colorado front range of the Rocky Mountains, including relatively isolated RF communities in Denver, Colorado Springs, Northern Colorado around Fort Collins, and so forth. The average distance between these clusters of humanity is between 40 and 50 miles. To the west, there are scattered clusters of activity in and around the high-altitude ski country, and at Grand Junction on the western edge of the state. Paths between switch sites on the East/West backbones are frequently far enough apart to preclude easy application of microwave technology such as the N6GN 10Ghz design. To the South, a couple of long hops connect New Mexico and Texas, and to the North, very reasonable paths connect into Wyoming.

Two problems then exist. One is fitting technologies to needs, and sequencing changes to achieve maximum benefit early in the process. The second is gaining sufficient support (monetary, political, and volunteer labor) to make the plan succeed. This is a tough battle, particularly when the technical solution proposed is esoteric enough (at least compared to contemporary amateur packet technology) to seem "scary".

First, existing relay sites should have their digital hardware replaced with PS-186 boards. Equipped with a standalone version of the KA9Q NOS software, these boards will interoperate with the existing NET/ROM backbone, and provide a foundation for building better links.

Second, links along the Front Range that are of moderate distance should be upgraded to 1 Mbps links on 10Ghz, if possible in parallel with the existing 1200 baud technology until the new links are stabilized and proven. Because there is a potential for tremendous data transfer needs among the high population densities in this area, the technology providing maximum data rate over point-to-point links is the obvious choice.

Third, the sites along the East/West paths and to the South, where distances are perhaps too great for reliable 10Ghz operation, should be equipped with 1.2Ghz N6GN units at 250kbps. We choose 1.2Ghz over 900Mhz primarily because part of Colorado is included in the area closed to amateur use of the 900Mhz band. These radios, combined with gain antennas (perhaps using power splitters or duplicated TX gain blocks for multiple point to point links), have the potential of operating reliably over the path lengths involved (on the order of a hundred miles, between 14,000 foot peaks...). The 10Ghz technology is favored because of cost (less than $100 parts cost per end including a 2-foot dish!), data rate, and fundamental point to point operation... but there are undoubtedly locations where dishes or other microwave bits and pieces will not be appropriate.

Simultaneously with installation of these faster backbone links, we should investigate better local access ports. The advantages of full duplex repeaters for true CSMA/CD operation have been widely discussed. In Colorado, a mixture of 1200 baud and 9600 baud full duplex ports is probably a good initial step, with 56kb cross-band full duplex worthy of consideration.

### Making It Happen

Before we go out and start building things, we ought to be prepared to spend some time assessing the current usage patterns on our existing packet network links. As appealing as it might seem at times, removing an existing facility to make room for "bigger and better" is rarely a wise move. We need to make sure that any new facility we install, and any facility that we upgrade, provides a growth path for existing users.

If we are going to successfully upgrade our major packet facilities to a substantially higher level of performance, we're going to need to enlist all the help we can get. In particular, some energy should be expended investigating applications that are outside of the traditional uses of packet radio. Armed with ideas, we can then approach groups that are not known for being strong packet supporters, with some hope of gaining cooperation and support.

The N6GN 10Ghz microwave links have as an option an audio sub-channel, that might be used to provide phone patch capabilities to a remote repeater site that we want to equip with fast links. The expansion bus of the PS-186 might be utilized to support custom hardware providing remote control and telemetry for existing RF gear at a switch site. I'm sure you can think of other such possibilities.

It is a fact of life that many of the best RF sites for high speed packet backbone hardware are currently controlled by FM repeater or fast-scan video user groups. The key to gaining access to these sites is finding ways to add value to their operation when they allow us rack and tower space. In Colorado, the emerging relationship between the Rocky Mountain Packet Radio Association (RMPRA), and the Pike's Peak FM Association (PPFMA) is an example of how this can really work.

### Keeping It Happening

Even if we do our homework, plan our links carefully, and get as many groups as possible involved in implementing our network, there are still several problems that can crop up that deserve discussion.

An initial burst of enthusiasm followed by a loss of interest before any milestones are reached can be a real problem. The best way to counter this phenomena is to sequence the installation of switch and local access improvements such that high visibility sites are upgraded first. This way, the maximum number of folks will have a chance to experience the improvements, and get excited about the possibilities. And if you can grow the excitement by pulling in more and more new folks to help with the effort, so much the better!

Despite our tradition in the amateur hobby of working with gentleman's agreements, we need to consider putting some of the decisions made, particularly between groups that don't have a history of beneficial cooperation, into writing. This will help to make sure that the loss of a key member of the effort won't leave the two groups wondering what "the other guys" were supposed to do, or even worse leave two groups that won't speak to each other! The things that most need to be documented are who is going to pay for and implement each piece of the network, and what their expected schedule is.

And finally, as W0YNE has been heard to question on innumerable occasions: "Who's going to take care of all this stuff?" It's easy to find people to drive half a day, climb a mountain, and install a new piece of gear. It's not too hard getting those folks or folks like them out the first or second time something breaks or gets hit by lightning. But what happens a year or two down the road when the initial support groups have all burned out and gone back to chasing 40m DX, or rag-chewing on 80m? Considering carefully how to handle long-term upkeep of our networking facilities is a need that cannot be overstated. If we're going to go to the trouble of building a real network, we need to be able to maintain it!

### Conclusion

Sitting back and thinking about all of the issues that surround the design, acquisition, installation, and maintenance of a next-generation amateur packet radio network may make the problems seem to be insurmountable. There are, in fact, some really difficult questions that need to be answered. But I hope the message conveyed by this paper is that none of the problems are impossible to solve! We can, and should, strive for a major upgrade in quality and performance of our packet network.

I've talked briefly about some of the technologies that might play a part in building our "dream network". It seems particularly exciting that the PS-186, the K3MC fast I/O card, and really fast RF hardware by N6GN are all coming together at about the same time. We've asked many times "where is my high-speed packet?" The answer is that it's here, it's now, and we no longer have any excuses for not improving our networks!

It is hoped that by the time this paper is published, further details will have been worked out regarding the development of a faster packet network in Colorado. The author would be pleased to hear from individuals involved in similar efforts elsewhere. Electronic mail is preferred, bdale@col.hp.com on the Internet, 76430,3323 on Compuserve, or N3EUA @ KA0WIE via packet BBS.

References:


Brock, M., Antonio, F., and Lafluer, T., *A High Performance Packet Switch* ,Proceedings of the Sixth ARRL Computer Networking Conference.

Chepponis, M. and Karn, P., *The KISS TNC: A Simple Host-to-TNC Communications Protocol,* Proceedings of the Sixth ARRL Computer Networking Conference.

Finch, R. and Avent, *S., A Duplex Packet Radio Repeater Approach to Layer OneEfficiency*, Proceedings of the Sixth ARRL Computer Networking Conference.

Avent, S. and Finch, R., *A Duplex Packet Radio Repeater Approach to Layer OneEfficiency - Part Two,* Proceedings of the Seventh ARRL Computer Networking Conference.

Heatherington, D., *A 56 Kilobaud RF Modem,* Proceedings of the Sixth ARRL Computer Networking Conference.

Kantor, B., *Software Design Issues for the PS-186 Advanced Packet NetworkController*, Proceedings of the Sixth ARRL Computer Networking Conference.

Chepponis, M. and Mans, B., *A Totally Awesome High-Speed I/O Interface for the IBM PC XT/AT/386 and Macintosh II Computers,* Proceedings of the SeventhARRL Computer Networking Conference.

Garbee, B., *More and Faster Bits: A Look at Packet Radio's Future,* Proceedings of the Seventh ARRL Computer Networking Conference.

Garbee, B., *Digital Dreams,* 7 3 Amateur Radio, October 1989.