The Implications of High-Speed RF Networking

Mike Chepponis, K3MC
Glenn Elmore, N6GN
Bdale Garbee, N3EUA
Phil Karn, KA9Q
Kevin Rowett, N6RCE

ABSTRACT

High-Speed networking is qualitatively different from the 1200 baud omnidirectional multicasting now in widespread use. A true network of high-speed nodes require planning, and (at least from the Amateur perspective) unusual networking architectures to take advantage of the much higher speeds. Fundamental limitations on network parameters must be understood, and parameters must be coordinated for consistent operation within the network. The potential applications possible with truly high-speed networks are also qualitatively different from what is provided by slower networks, offering a tremendously exciting future for amateur radio as a whole.

1. What is High Speed?

Because of the predominance of 1200 baud half-duplex AFSK multicasting as the defacto standard for packet radio networking, it is easy to look at moderate increases in transport speed (to 9600 baud, or even 56 kilobaud) and be excited about the relatively ‘‘high speed’’ operation that such modems allow. However, from an applications standpoint, such improvements allow us to do the same things we’ve always done, albeit a little faster, but do not open the door to many applications that require real ‘‘high speed’’ networking.

High speed in this context means 250 kilobits/sec to 10 megabits/sec and above. We are intentionally excluding the WA4DSY 56 kilobits/sec modem because with that modem it is still possible to use ‘‘standard’’ networking, that is, uncoordinated stations, omni antennas, and CSMA. Because of this, the WA4DSY modems are expected to play an important part in networking systems of the future, especially as feeders to high speed networks (local access channels). But for the purposes of this discussion, we’ll consider 56kb as the upper limit of the ‘‘low speed’’ packet technology.

2. Network Architecture and Software Implications

The low capacity of the existing 1200 baud AX.25 network has restricted users to the most undemanding of applications: manual keyboard-to-keyboard chatting, the transfer of small text files, store-and-forward electronic mail, and DX spotting. But as fast digital links become available, much more powerful forms of computer networking will become practical. This requires higher level protocols above AX.25.

The ARPA Internet protocols, commonly known as ‘‘TCP/IP’’, are clearly the leading contender for a high speed amateur network. Over the past decade, the Internet protocol suite has matured in a diverse academic, military, research, and commercial setting. For several years, TCP/IP implementations have been in everyday use on hundreds of thousands of computers ranging from laptop PCs to desktop workstations to Cray supercomputers. Ethernets, twisted pairs, token rings, fiber optic channels, satellite links — and amateur packet radio channels, among other things, are regularly used to carry IP datagrams. TCP/IP has become the method of choice for providing standard network services across a wide variety of computer systems and local and wide-area networking technologies. In amateur packet radio, TCP/IP has been growing steadily in popularity since the introduction of the KA9Q Internet Protocol software package. We believe that the ARPA Internet is an excellent role model for the kind of network we can build as amateurs, and that we can learn a great deal From the Internet experience.

Most recently the original ARPANET packet switched network, operating since the early 1970s with 56 kb/s links, has given way to a new IP-based national backbone network, the NSFNet, plus thirteen regional networks. Virtually every link in this new Internet operates at the North American T-1 rate, 1.536 Mb/s. This brings the Internet into the ‘‘high speed’’ category as we’ve defined it in this paper.

The rapid growth in size and capacity of the Internet has caused some growing pains. The original routing algorithms proved inadequate; new ones have been developed and are now coming into use. A lot of work is going into the study of congestion in large datagram networks like the Internet. This work is now paying off with a set of recommended congestion avoidance algorithms for both the transport protocol (e.g., TCP) in the end system and in the interior packet switches (IP routers). The TCP algorithms in particular have also proven themselves in the slow speed amateur packet radio environment.

‘‘Network management’’ is a hot topic. Operators need to know quickly when links or routers fail, and they need to monitor traffic patterns in order to plan the most effective allocation of additional resources. Of course, the network should take care of itself as much as possible.

One might think that with a three or four order-of-magnitude increase in link speed (e.g., from 1200 b/s to 1-10 Mb/s) that our performance problems would disappear forever. This is not necessarily true! Although excess capacity is certainly an effective short term antidote to congestion, the Internet’s experience has been that network demand is gaseous — it eventually expands to fill the available capacity! Although some of this growth will come from new users and from increased use of traditional services like email and file transfer, the real kick will come from qualitatively new applications like digital voice and image transmission that high speed networking will make practical for the first time.

But even for the more traditional computer networking applications there will be challenges. For example, when the NSFNet replaced the ARPANET, the ‘‘delay-bandwidth’’ product of the Internet backbone went up by a factor of 27. No matter how large the link bandwidth gets, the speed of light remains the same. Also, the more packet switches there are in a path, the greater the total ‘‘store and forward’’ delay. For example, the distance between New York and San Francisco is 4139 km. The one-way speed-of-light delay is therefore about 13.8 ms. If a terrestrial microwave route were set up in a straight line between these two cities with 50 km repeater spacing, 84 hops would be required. If the repeaters are in fact packet switches, and if the links all operate at 1 Mb/s, then the accumulated store-and-forward delay for a 200 byte packet sent across this path would be 132.8 ms, almost ten times the speed-of-light delay. (The actual value would be larger, since we’ve ignored the processing delay of the packet switch software.)

One technique for reducing store-and-forward delay is ‘‘cut-though’’. A packet switch could make a routing decision and begin retransmission of a packet immediately after receiving its destination address — it need not wait for the packet to be completely received. One problem with cut-through is that a packet might be relayed before it is discovered to contain an error. One partial fix is to compute a separate ‘‘header checksum’’, placing it in the header itself so that the packet switch need only wait for the end of the header before making a reliable cut-through routing decision. Any errors in the data field would still be undetected, so they would have to be dealt with by end-to-end checks in the transport protocol. IP is a network protocol with a header checksum, and TCP is a transport protocol with end-to-end data checksums, so they are amenable to this technique.

It is quite likely, however, that in the near term the most practical way to handle delay-sensitive services like digital voice will be conventional circuit switching. This is standard telephone company practice, and the technology is very well understood. Digital circuit switches operating at the data rates considered here are now within the technical and financial capabilities of amateurs; DSP chips in particular are ideally suited for the task. A packet switching network intended mainly for computer networking could then be built on top of dedicated circuits, with additional circuits added or removed as needed. The NSFNet is built in just this way; MCI, their long-haul carrier, is providing them with the ability to reconfigure their links as traffic patterns warrant.

Returning to packet-switched computer networking, propagation delay is an important protocol design factor. Most protocols deal with it by keeping multiple packets in flight whenever possible. The usual method, ‘‘sliding windows’’, is used by nearly every transport (e.g., TCP and ISO TP-4) and pseudo-transport (e.g., X.25) protocol. But sliding window protocols cannot transfer more than one window’s worth of data per round trip propagation time, no matter how large the path bandwidth may be. You can improve performance by enlarging the window, but only by using additional buffer memory and by increasing the performance ‘‘hit’’ incurred by lost packets. In our NY/SF example, the window would have to be at least 84 packets in order to utilize the path's full bandwidth, even if we ignore the round trip speed of light delay and the store-and-forward delays incurred by the returning acknowledgement packets.

Another problem that needs a solution is efficient broadcasting. This is something that comes almost free with omnidirectional packet radio, even though we don't currently take advantage of it. But the use of directional point to point links will require us to come up with other methods.

One possibility is the ''flood'' algorithm used in USENET. Each multicast packet (broadcasting is just a special case of multicasting) is sent over each link exactly once. Each node keeps a list of the packet IDs it has seen recently so it can ignore duplicates. Different multicast addresses could correspond to certain regions of the network, so that data of local interest wouldn't have to cover the whole network.

As you can see by now, the availability of high speed links requires parallel developments in digital switching hardware, protocol design and software if we are to use them to full advantage.

3. Hardware Implications

In order to build high-speed networks and to develop the software meant to run on them, we need the appropriate hardware. What we need are high-speed modems and digital interface cards capable of handling these speeds. Fortunately, such hardware is now appearing.

Last year, one of us (Mike) described the ''Totally Awesome'' I/O card, for use as a plug-in for IBM PC/AT/386 machines. This card will shortly be commercially available. In addition, a version for the Macintosh Plus, Macintosh SE, Macintosh Portable, and the Macintosh II/IIx/IIcx/IIci will be available, using AppleTalk at 230.4 kilobits/second.

Briefly, this ''Awesome I/O card'' for the IBM PC/AT/386 machines features a V40 processor (code compatible with the 80x86 microprocessor family), at least two Zilog 85C30 Serial Communications Controller chips, yielding four full-duplex channels. The board can be expanded up to a total of six 85C30s, for a maximum of twelve full-duplex channels[1]. It comes standard with 256k bytes of DRAM, expandable up to 512k bytes. Also, two JEDEC-standard EPROM/ROM sockets are included, allowing up to 512k bytes of EPROM when using two 256k x 8 EPROMs. A watchdog timer, a time of day clock, a programmable interval timer, and 50 bytes of battery-backed static RAM are provided by a Dallas Semiconductor DS-1286 chip. Communication with the XT/AT/386 is accomplished by a memory window into the XT, which can be either 8K, 16k, 32k, or 64k bytes long[2]. A 20-character by 2 line LCD display provides operator interface.

One of the 8530s is totally DMA-driven. The V40 provides four DMA channels, and these four channels provide DMA on transmit and receive for both channels of one of the 8530s, allowing two full-duplex channels to operate with minimal processor intervention. The other 85C30s are interrupt driven. The DMA-driven channels are fully capable of somewhat greater than 1.5 megabits/second. The other interrupt-driven channels can operate full-duplex at 9600 bauds or slower when the two high-speed channels are used.

We expect the Awesome I/O card to dovetail nicely with the PS-186, which is now becoming available. One of us (Bdale) is porting the KA9Q TCP/IP package to the PS-186, with conditional compilation directives so that the same code can be run on the PS-186 or the Awesome I/O card. Another of us (Kevin) is preparing a standard driver software interface for the KA9Q TCP/IP package.

Concurrently with this development, another of us (Glenn) is producing the high-speed modems we will need. Development work and prototyping is presently going on to produce a very inexpensive 10W 900 MHz digital radio capable of 250 to 500 kilobit/second operation. This radio requires approximately a 2 MHz channel width and thus several clusters using such hardware can coexist in a metropolitan area without the individual cluster size becoming too large or congested. The radios are being built use Japanese 903 MHz personal radio band transmitter power train hardware and an Motorola MC13055 single chip wide band FSK receiver IC. A similar 20W radio for the 1200 MHz band is also being developed since the 900 MHz band is not available to amateurs worldwide.

4. Networking Architecture Implications For Higher Speeds

One of the biggest problems inhibiting throughput even in our current slow-speed networks is collisions. Phil Karn discussed a possible way to avoid this problem[1].

---

[1] On the Macintosh version, one of the channels is dedicated to AppleTalk.

[2] That is, some of the Awesome I/O's memory is mapped into the XT/AT/386 address space

In a high-speed network, however, additional constraints suggest to us that a Token Bus style of topology would be most appropriate. See Figure 1.

Using this technique, and using half-duplex modems, we can guarantee that no collisions occur. Within each of the cells, there is one token, which is being passed continuously by the individual stations in the cell. This happens even if a station has no traffic to offer the network; the station is responsible for repeating the token whenever it arrives at that station. When a station has some traffic to offer the network, that station listens until it sees a ''free'' token, and changes it into a packet containing the information that needs to be transmitted. At the nexus stations, traffic can flow between local cells.

Present day packet radio time multiplexing methods for the physical media result in lower and lower throughput rates on busier channels. This results From collisions, retries from collisions, and back-off delays. The result of collisions is more traffic, producing more collisions, further driving down throughput. A 1200 BPS channel can quickly drop to 100 BPS actual throughput, even < 1 BPS in some cases.

Soon much of the benefit of higher speeds is lost in collisions, retries and back-off delays.

For higher speeds that are not point to point, a different time multiplexing protocol should be used...token passing. A frame format is designed to be THE token frame. It is sent from station to station, around a logical ring. The rule is simple: only the station holding the token frame can transmit or elicit another station to transmit a frame. When the token frame holder station is finished, it transmits the token frame to the next station in the logical ring.

Using token passing, a station knows no collision can happen. Even more benefit can be gained by putting the link control in Virtual Circuit mode and getting an immediate response from the destination station. Again, no collision can happen to the acknowledgement. The source station can retry immediately if no response is heard, without fear of collisions.

Present standard token passing protocols (IEEE 802.4, 802.5) assume no hidden transmitters. This clearly will never be the case in amateur radio. Work on a change to existing protocols is proceeding. These protocols are needed to take care of the case where the token is lost or dropped, without assuming any station is a master.

Link level framing exists between stations that can all be reached directly by RF. Traffic at this level would include all applications which have a response time need, very close to the speed of the physical media. This would include applications such as real time. It would NOT include file transfers, keyboard conversations, BBS and mail reading, and remote repeater controls. It might extend to digitized voice, but not necessarily.

Connection between sites which cannot be reached directly via RF will require repeating or routing. Routing involves computerized equipment that implements a protocol allowing the equipment to select the most efficient path to send frames on to their destination. Hence, routing connects two or more physical media that can't be directly connected.

Considerable research is available on Token Bus, and particularly Packet Token Bus[2]. We intend to bring this research to Amateur Radio.

5. Application Implications

In the amateur community, one of the biggest hurdles to be overcome is the mindset of lower speed communications and current applications. While current AX.25 operation is well known, it does not offer a good beginning for envisioning and understanding a high speed network. In fact, the concept of high speed end-to-end data and the associated applications possibilities probably has no counterpart in our current amateur experience. Consequently, developing a widespread vision for such a network and coordinating its development will require a great deal of cooperation and wide-area coordination.

A high speed network no longer need be thought of as a ''computer thing''. Since virtually any signal which can be transmitted by analog means may also be transmitted digitally all of the diverse amateur modes presently in use may be supported on a network if sufficient speed is available. And because a network is not confined to communications over a single span which a single transceiver pair might operate but rather over the entire nation or the world many new social and technological applications become possible. These applications do not need to replace current exciting aspects to the hobby but rather can enhance them. DXers who enjoy the thrill of working another new one may find that a nationwide realtime spotting network, propagation and QSL server greatly enhances DXing. Similarly, chess players might find a chess ''news group and worldwide competition'' exciting. Mobilers and ragchewers might keep in touch with others in their group whether they are in one locality or spread across the nation. The possibilities are endless!

But let us offer several concrete examples, things that we intend to do once high-speed hardware is in place.

5.1. Digital Voice

This is perhaps the most obvious of the applications. Consider that if we sample speech at 7 kHz, allowing for frequencies up to 3.5 kHz to be present in our speech bandpass, and say that we require 8 bits to represent a companded sample, then we will require a 56 kilobit/second channel. With Time Division Multiplexing, we can slice up our 250 kilobit/second Token Bus network into about four voice channels. Admittedly, at the low end of the ''high-speed'' range, we cannot provide huge numbers of high-quality voice channels[3]. But, it will provide a viable testing ground for digital voice techniques, and help us build the experience to exploit the even higher speeds that will soon be available.

We should point out that there is significant research happening in packet voice technologies[3-5]. We will leverage off of that research. We are particularly interested in two technologies: Full-duplex voice communications, and packet switched conferencing.

We have a concept for a Digital HT. See Figure 2.

5.2. Video

Perhaps the most exciting new application made reasonable by high-speed networks is digital video[6-8]. Consider a garden-variety bit-mapped screen attached to an IBM PC/AT/386, consisting of a 720 x 384 bit image (Hercules Graphics resolution). This is about 250 kilobits of information, uncompressed. Even so, this proposed high-speed network will enable such an image to be transmitted in about one second(!). Of course, higher resolution and color images will take longer to transmit[4].

At these speeds, even limited animation becomes possible. But perhaps its biggest use would be in emergency applications, where we can survey the situation, scan in an image using one of the inexpensive hand-held scanners, and then transmit that image to a central command station.

[3] Although some experiments with sophisticated coding schemes achieve high-quality voice with as little as 9600 bits/second, which would yield about 25 voice channels on our 250 kilobit/second links

5.3. FAX

FAX is really an extension of Video. The differences are that paper images are usually the source, and that several compression schemes can be used[9]. Digital FAX transmission over packet radio could be a substantial asset for public service and emergency response activities.

5.4. Remote File Systems

When data can be transmitted between stations at speeds approaching the speed of data transfer to and from a local hard disk, the idea of remote file systems becomes practical. Here, it would be possible for one station to ''mount'' a file system from a remote station onto his local machine. From that point, he could then access the remote station's files as if they were local. In the commercial world, Sun workstations often come diskless, and use Ethernet and a remote file server for data storage. Not only does this make it very easy to share files, it also allows one to use storage resources of other stations transparently. For example, if KA9Q has enough disk space for the complete amateur callbook information, then others can access his database in a transparent manner.

6. Even higher speeds

As we start requiring higher speed data throughput we find ourselves needing to make much better use of our resources. At the physical layer we must use hardware which provides such communication efficiently and at low cost. To do this, we must use techniques which get as much of the original transmitter power to the receiver so that maximum path lengths and Carrier/Noise ratios may be obtained. Solutions which waste transmit power by sending it elsewhere are wasteful. In this light, omni-

[4] But not prohibitively so: a 640 x 480 x 8 bit image, such as can be displayed on the color screen of a Macintosh II is about 2.5 megabits; it will take only about 10 seconds to transmit — uncompressed. For 640 x 480 x 24 bits x 30 frames/second (full motion, full color), we need almost 222 megabits/second uncompressed; however, recent compression techniques can achieve a 10 times reduction in required transmission bandwidth in real time. Offline compression has yielded a 50 times compression, to require about 4.5 megabits/second. This is within the range of speeds of our current 10 GHz hardware. We are eager to see the day when our microwave links will handle 100 megabits/second, the standard FDDI (Fiber Optic) speed, and we will be able to handle several full motion, full color transmissions at the same time. We also note that commercial products exist today to transfer 1 gigabit/second over fiber optic links; to achieve these speeds economically via RF links will no doubt be an intersting problem to solve.

directional broadcasting on network backbones must be limited and eventually avoided per se. This is why we have suggested the Token Bus architecture, as directional antennas can and should be employed.

High speed data also will require higher bandwidths. To a limited degree, data rate can be traded for signal/noise by appropriate selection of modulation techniques but this can provide perhaps only a few orders of magnitude of date rate increase and any further increase requires increased signal bandwidth.

Fortunately the above two requirements for high speed data communications, directed transmitter power and wider bandwidths are compatible with the same amateur resource: the amateur microwave and millimeter bands. The same scaling of wavelength which allows physically realizable antennas to have high directivity and gain also occurs at portions of our allotted spectrum which have wide bandwidth.

In spite of widely held misconceptions that the microwave bands are expensive and difficult to use and only good for short range communications, these bands represent the most cost effective region in which way to exchange high speed data, in spite of the increased cost of generating moderate transmitter powers at these higher frequencies.

The ability to use low transmitter powers by effectively focusing all the energy in the direction of the receiver also offers the opportunity of frequency reuse within a local area. The current congestion on broadcast VHF packet is the antithesis of this. Rather than requiring high power and experiencing decreased C/N due to collisions with other omnidirectional transmitters, very small and inexpensive transceivers can reliably provide high C/N over moderate distances. At the same time interference-free operation of other links reusing the same channel which cross the physical path are made possible due to the highly directed beams.

As we go to increasingly higher speeds, hardware to switch packets over a network of these point-to-point links must be incorporated. To avoid collisions (QRM) and unjust sharing of the spectrum resource users will need to go to point-to-point rather than omnidirectional hardware. This will be an economic issue too since high antenna gain will be necessary to provide sufficient C/N for these wider bandwidth channels. In the interim, small local groups or ‘‘cells’’ of amateurs operating CSMA with omnidirectional antennas, and provided with a common path to a network switch, will be useful. For physical reasons, as well as the current state of technology, it is likely that these moderate power broadcast ‘‘user radios’’ will operate in the higher VHF and UHF bands. Below this range there is insufficient spectrum, and above the higher powers needed for use with omnidirectional antennas is presently too costly.

Through perhaps 1200 MHz, omnidirectional antennas with reasonable physical capture area are available. In the higher microwave regions, antennas with thin ‘‘pancake’’ patterns are more difficult/expensive to build and the resulting narrow vertical beam is at risk of being deflected by tropospheric circumstances, with a consequential reduction in system margins. Physically high level transmitters, omnidirectionally broadcasting high power, must be avoided since they are extremely wasteful of both spectrum and hardware resources.

A functioning high speed network requires good availability to end users. Because of this, hardware and sites must be selected to provide higher up-time than that which is currently tolerated for much amateur communication. If system uptime of the same order as commercial services is required, say 99.9% or better, links will generally require line-of-sight (LOS) paths and in severe situations, such as those experiencing marine air masses, extra system margin will be necessary. As an example, obstructive fading produced by nighttime movement of wet air may require shorter paths, even on VHF, to guarantee service. Generally, locations with well mixed air, like the Rockies are less likely to experience propagation abnormalities. Coastal locations, and large flat areas with water sources, are more likely to require special attention. Interference fading on a link will generally require increased C/N margin to guarantee performance. The Rayleigh distribution model probably serves as a useful guide to these margins. For typical fading at higher microwave 99% uptime may require about 18 dB excess C/N over LOS values, 99.5% about 22 dB and 99.9% about 28 dB excess. On many shorter (perhaps 30 mile or less) LOS paths, these numbers may be unduly conservative. On such short links it may be possible to run with C/N only 3 to 6 dB above those necessary for the desired BER on LOS free space paths.

It may even be possible in some situations to have extra long paths which rely on tropospheric ducting to provide communication (it’s not a bug it’s a feature!). Whether such links which are typically only available a few hours or less a day can prove useful for periodic but not continuous service (like electronic mail transfer) is a question which remains to be answered.

At 24 GHz and above, water vapor absorption is an item to consider, particularly on longer paths.

At 10 Ghz and below, it can probably be ignored when considering digital data links. Path links of sufficient length to make it noticeable probably suffer from much larger variations due to variations in the troposphere.

An approximate summary of amateur band utility for high-speed network might be:

144	best suited for low speed broadcast
220	
433	moderate speed broadcast user/cell radios
900	
1250	upper end of broadcast
2300	beginning of point-to-point
3400	
5700	
10000	High speed
24000	High speed and shorter paths. Water vapor absorption significant on long paths.

As mentioned above, two classes of networking radios would be presently useful; radios for high speed point-to-point connection between packet switches and end-user radios to give access to the network.

In the category of inter-switch radios a design using inexpensive microwave transceiver modules such as found in burglar alarms, automatic door openers and speed measuring equipment has already been presented at a national convention, and will be appearing shortly in a national magazine. These radios easily support data at 2 Mbaud and greater rates and can be built for much less than the cost of a commercial 2M FM transceiver. This and similar hardware should be able to support paths between high level switch sites, perhaps a pair of mountaintops, separated by 30-40 miles. Longer distances are possible with bigger antennas, or more sophisticated radio hardware, but propagational variations over much longer paths require prohibitively high system margins to guarantee high up-time and multiple short path links are preferred.

User access radios which can be used to broadcast within a local cell, as well as to contact a high level switch, are already available. An entry level for new users at 1200 baud or other conventional speed should probably be provided in every area to encourage network growth. Newer and more interesting applications will require greater data throughput and should provide the incentive for users to upgrade their stations to allow this.

No doubt some of the 250 kilobit/second user radios will be pressed into service initially for interswitch communications by using directional antennas. This may particularly be true in areas where interswitch distances are necessarily greater than the distances that are comfortable for microwave. But in the longer term, higher ‘‘trunk’’ data rates will be required and these radios should be replaced with dedicated microwave hardware whenever possible.

7. The Future

The opportunities made possible by a functioning high speed amateur network can fundamentally change the face of amateur radio by augmenting those aspects of the hobby which are already enjoyed by so many and by providing new applications which may entice a large force of new blood into our great hobby. In addition to these enhancements the public service aspect of amateur radio may be greatly enhanced by providing information age services to the public in time of emergency. These possibilities can greatly contribute to the vitality of amateur radio and help ensure the continuance of the hobby.

8. Wrapup

To make a high speed network a reality the biggest hurdles are organizational and not technical. Radio hardware, interfaces, and computers to handle the data are well understood and are becoming available to amateurs now. However, to make a functioning entity from all the pieces will require a great deal of common vision and participation by all parts of the amateur community. Radio amateurs have traditionally earned a reputation for an ‘‘I’ll do it my way’’ mentality within the hobby. This has, in fact, been the source of many of our greatest technical achievements! Fortunately, the development of a truly high-speed network leaves room for this, since local cells may find many ways to express their individuality, as long as means are found to ‘‘mesh’’ with the rest of the network at the backbone interface level.

However the physical area and user base for a nationwide high-speed network is very large, and we must therefore find ways to avoid large scale repetitions of ‘‘repeater wars’’ as we implement the network. As we work to design and implement the next generation of packet networking, we must retain an awareness of the need for cooperation, support, and goodwill between all of the nation’s packet groups and users. This will help to ensure that everyone can enjoy the exciting new applications made possible.

Perhaps more than at any time since the days of spark, amateur radio needs a well supported organization to further the relay (or networking) of our communications. In the midst of the many recent discouraging developments relative to our spectrum and membership, a real amateur network offers great hope and excitement for our future!

9. Acknowledgements

We would like to thank Bob Hoffman, N3CVL, for, once again, tolerating our last-minute submission and expertly typesetting this paper.

10. Contacting the Authors

The authors can be reached as:

Mike Chepponis, K3MC	
BBS	K3MC @ K3MC
Internet	k3mc@apple.com
UUCP	...!apple!k3mc
CompuServe	72117,2732

Glenn Elmore, N6GN	
BBS	N6GN @ K3MC
Internet	glenne@hpnmd.hp.com
UUCP	...!hplabs!hpnmd!glenne

Bdale Garbee, N3EUA	
BBS	N3EUA @ KA0WIE
Internet	bdale@col.hp.com
UUCP	...!winfree!bdale
CompuServe	76430,3323

Phil Karn, KA9Q	
BBS	KA9Q @ NN2Z
Internet	karn@ka9q.bellcore.com
UUCP	...!thumper!karn
CompuServe	73210,1526

Kevin Rowett, N6RCE	
BBS	N6RCE @ K3MC
Internet	kevinr@tandem.com
UUCP	...!pacbell!tandem!kevinr

References

[1] Karn, P. ‘‘A High Performance, Collision-Free Packet Radio Network,’’ *6th ARRL Computer Networking Conference Proceedings*, p.86 [1987]

[2] Wong, P.-C., and Yum, T.-S.P., ‘‘An Integrated Services Token-Controlled Ring Network,’’ *IEEE Journal on Selected Areas in Communications*, June 1989, Volume 7, Number 5, p 670 [1989]

[3] Chen, T.M., Walrand, J., and Messerschmitt, D. G., ‘‘Dynamic Priority Protocols for Packet Voice,’’ *IEEE Journal on Selected Areas in Communications*, June 1989, Volume 7, Number 5, p 632 [1989]

[4] Petr, D.W., DaSilva, L.A., and Frost, V.S., ‘‘Priority Discarding of Speech in Integrated Packet Networks,’’ *IEEE Journal on Selected Areas in Communications*, June 1989, Volume 7, Number 5, p 644 [1989]

[5] Ziegler, C., Weiss, G., and Friedman, E., ‘‘Implementation Mechanisms for Packet Switched Voice Conferencing,’’ *IEEE Journal on Selected Areas in Communications*, June 1989, Volume 7, Number 5, p 698 [1989]

[6] Karlsson, G., and Vetterli, M., ‘‘Packet Video and Its Integration into the Network Architecture,’’ *IEEE Journal on Selected Areas in Communications*, June 1989, Volume 7, Number 5, p 739 [1989]

[7] Shimizu, H., Mera, M., and Tani, H., ‘‘Packet Communication Protocol for Image Services on a High-Speed Multimedia LAN,’’ *IEEE Journal on Selected Areas in Communications*, June 1989, Volume 7, Number 5, p 782 [1989]

[8] Joseph, K., Raychaudhuri, D., and Zdepski, J., ‘‘Shared Access Packet Transmission Systems for Compressed Digital Video,’’ *IEEE Journal on Selected Areas in Communications*, June 1989, Volume 7, Number 5, p 815 [1989]

[9] Chamzas, C., and Duttweiler, D.L., ‘‘Encoding Facsimile Images for Packet-Switched Networks,’’ *IEEE Journal on Selected Areas in Communications*, June 1989, Volume 7, Number 5, p 857 [1989]