NCPA Downlink

The Official Journal of the Northern California Packet Association
Serving Amateur Radio Digital Communications in Northern California

Summer, 2002 **Issue number 32** **$2.50**

President's Message

Gary Mitchell, WB6YRU

The NCPA annual meeting will again be at Pacificon (see notice elsewhere in this issue). The main item on the agenda is about the NCPA organization itself.

As some of you know, we've had trouble in recent years just getting a quorum at meetings. And that's not all, the number of people carrying the load can be counted on one hand–with fingers to spare.

A few of the "higher-ups" have been suggesting that the Association be reorganized in a way that would make it easier to function with minimal participation. The leading idea is to go to a committee format. Among other things, this may include altering the requirements for a quorum, scaling back this newsletter, making greater use of the internet (remailer, web page), and having two classes of members.

None of this is carved in stone, the point of the meeting will be to discuss options and ideas. Eventually, this will probably mean significant changes to the bylaws, but that will come later. The important thing is that any of the changes made must preserve the NCPA's representative nature.

Here are some specific ideas already mentioned: Have two classes of membership, one votes, the other merely advises—this would ease the quorum requirements. Change the number of directors on the board from a minimum of seven to the number of packet special interest groups. In other words, instead of seven directors, there would be one for BBS, one for keyboard, one for APRS, etc. Allowance would be made for an increased number, but the result is the minimum would drop to about five. Reduce the newsletter to bi-annual or annual, perhaps not even that much (i.e. publish only when needed). The Downlink could be only published on the NCPA web page. (TAPR, a much larger packet organization, already does this, so there is a precedent.) Allow for the possibility of holding general meetings on the internet (i.e. via remailer or whatever). This wouldn't be as good as face-to-face meetings of course, but it would make it easier for members to participate without traveling a long distance.

This meeting will be your chance to help shape the NCPA. This one will be more like a round-table discussion, so come and bring your ideas and suggestions. (A copy of the bylaws is on our web site: www.n0ary.org/ncpa.)

We'll also be electing directors to the board. If you have any nominations or would like to be on the board yourself, here's your chance to get more involved.

NCPA

In This Issue

Annual Meeting at Pacificon 2002 in Concord

The NCPA will hold its annual meeting at 11 AM in the Pilot's Cove room on Saturday, October 19, at the Pacificon convention. Pacificon will be at the Airport Sheraton Hotel in Concord (www.pacificon.org). Among other things, this is when the directors are elected by the membership. Admission to Pacificon is not needed to attend the NCPA meeting.

Directors

Mel Gregonis (high speed), W6BNG
e-mail: gregonis1@attbi.com
Howard Krawetz (Keyboard), N6HM
408-739-7936
home BBS: N0ARY
e-mail: n6hm@arrl.net
Bob Fahnestock (TCP/IP), WH6IO
home BBS: WH6IO
e-mail: wh6io@benicia.ampr.org
Bob Vallio (DXPSN), W6RGG
510-537-6704
home BBS: N0ARY
e-mail: w6rgg@sbcglobal.net
Dave "Zonker" Harris (APRS), N6UOW
home BBS: N6QMY
e-mail: n6uow@baker-to-vegas.org
Gary Mitchell (BBS), WB6YRU
408-265-2336
home BBS: N0ARY
e-mail: wb6yru@aenet.net
Mike Fahmie (BBS/Keyboard), WA6ZTY
home BBS: N6EEG
e-mail: fahmie@earthlink.net

Officers & Staff

President:	Gary Mitchell, WB6YRU
Vice-President:	
Secretary:	Dave Harris, N6UOW
Treasurer:	Howard Krawetz, N6HM
Editor:	Gary Mitchell, WB6YRU
Freq. Coord.:	Gary Mitchell, WB6YRU
Mail:	Howard Krawetz, N6HM

The NCPA Downlink is published quarterly by the Northern California Packet Association, P0 Box K, Sunnyvale CA 94087, for the entertainment and education of amateur Radio operators using digital modes, and those with an interest in them. A one-year membership in the NCPA, including a subscription to the NCPA Downlink, is $10.00 in the U.S. and its possessions.

The digital band plan and other information about the NCPA is available on the Web at: **http://www.n0ary.org/ncpa**

The NCPA Board of Directors meets electronically in order to transact association business and meet with members and interested amateurs. The address for the board remailer is: ncpa@kkn.net. Anyone can subscribe by sending e-mail to ncpa-request@kkn.net with the command "subscribe" (without the quotes) in the body of the message.

News from the ARRL

From *The ARRL Letter*, August 16, 2002

UNITED PARCEL SERVICE NOW NEUTRAL ON SAVI PROPOSALS FOR 70 CM

United Parcel Service (UPS) now says it's neutral on SAVI Technology's petition to deploy RF identification tag devices at 433 MHz at much greater duty cycles than current Part 15 rules permit for such devices. UPS clarified its position this week in an ex parte filing to the FCC.

"UPS takes no position on the rule changes proposed in the SAVI Petition because they will have virtually no impact on UPS's shipping operations and are inconsistent with efforts to promulgate international standards for RFID tags," the shipping company said. The change in position is doubly significant because UPS has an equity interest in SAVI through its UPS Strategic Enterprise Fund.

RFID tags are used for tracking shipments and packages, among other applications. The ARRL has said that adopting SAVI's proposals would result in severe and harmful interference.

ARRL Chief Executive Officer David Sumner, K1ZZ, said the League was pleased to learn that UPS had "done the right thing." Sumner had pointed out UPS's support of the SAVI petition in his "It Seems to Us . . ." editorial in the December 2001 issue of QST.

"The ARRL is very gratified that, upon careful consideration, UPS has changed its position and now recognizes that the SAVI proposal for 425-435 MHz offers no benefit," Sumner said. "We are confident that if the FCC devotes the same attention to considering the issue, it will come to the same conclusion."

UPS said it wanted to clarify its position in light of the many comments filed in response to the Notice of Proposed Rule making (NPRM) in ET Docket 01-278 that cited the shipping company's initial support of the SAVI petition. UPS has not directly commented on the NPRM previously.

UPS now says that, after further consideration, it sees no particular advantage to the changes SAVI has proposed. "UPS now does not envision any of its applications requiring a transmission duty cycle in excess of what is currently permitted under Section 15.231," UPS said.

UPS also cited concerns that the proposed operating frequencies "are not fully compatible with frequency allocations" in many of the more than 200 countries and territories in which it does business. "Thus, it is of limited benefit to global companies such as UPS should the FCC adopt the proposed 10 MHz-wide RFID band from 425 to 435 MHz."

More than 130 amateurs filed comments in opposition to SAVI Technology's RFID tags proposal, and most supported the League's position that the proposed rules are flawed and should not be adopted.

A copy of the UPS ex parte filing in ET Docket 01-278 is available on the FCC Web site <http://gullfoss2.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id_document=6513287285>.

From *The ARRL Letter,* Sept. 13, 2002

ARRL RESPONDS TO IMPLIED 222-225 MHZ THREAT

The ARRL has taken issue with a suggestion made in a non-Amateur Radio-related FCC proceeding to turn the 222-225 MHz amateur allocation over to commercial interests. In reply comments filed this month, the League urged the FCC to "do nothing" with the proposal of Data Comlink (DCL), a consortium of 20 electrical coops and allied companies.

"ARRL presumes that the proposal by DCL for reallocation of the 222-225 MHz band will not be seriously evaluated by the Commission, as it is well outside the scope of this proceeding," the League said in its September 5 filing with the FCC. Until DCL raised the 222-225 MHz suggestion last month in its own comments in WT Docket 02-224, the ARRL had remained silent in the proceeding.

DCL claimed in its comments that the amateur allocation at 222-225 MHz "is being underutilized" and that the band "would be much better utilized for commercial use."

ARRL asserted that the band, far from being underused, "remains a critical VHF allocation" for amateurs. The League noted that the ARRL 2002 Repeater Directory--albeit not a comprehensive listing--lists 1690 repeaters throughout the US, indicating an even larger number of individual users. "Indeed the number of individual amateurs using this band has increased steadily since 1989, when the amateur allocation at 220-225 was reduced by 40 percent," the ARRL said, "and now much commercially manufactured equipment is available to amateurs."

DCL had claimed that "only handfuls [sic] of individuals in the Amateur Radio Service even use this spectrum, while hundreds of thousands of potential commercial users wait with no alternatives." The League characterized as "invalid" DCL's arguments in favor of reallocating 222-225 MHz from the Amateur Radio Service and noted that the FCC earlier this year had set aside an additional 8 MHz of spectrum for Land Mobile Service operations.

The League's reply comments in the DCL proceeding are on the ARRL Web site http://www.arrl.org/announce/regulatory/wt02-224/arrl-comments.html.

The ARRL has not commented in an unrelated Petition for Reconsideration filed by Warren C. Havens on behalf of himself and Telesaurus Holdings GB LLC, in which he holds a majority interest. Filing last month under PR Docket 92-257 and RM-9664, Havens is seeking to have the FCC reconsider its decision to auction certain AMTS spectrum and instead adopt his "Advanced Technology Land Infrastructure and Safety Service" (ATLIS) proposal. Under that plan, Havens wants to see 222 to 225 MHz reallocated from amateur to public safety use. His ATLIS plan proposes to share 902-928 MHz on which amateurs are secondary.

From *The ARRL Letter*, Sept. 20, 2002

DIGITAL AFICIONADOS TURN OUT FOR 2002 ARRL/TAPR CONFERENCE

More than 100 of the most active Amateur Radio digital enthusiasts from around the world turned out in Denver, Colorado, September 13-15 for the 2002 ARRL/TAPR Digital Communications Conference. This year's event marked the 21st conference. Agenda topics ranged from APRS (Automatic Position Reporting System) to high-speed digital networking and software-defined radio (SDR), among others.

Friday's forums were dominated by discussions of APRS. Topics included a discussion of single-wire APRS weather stations, high-altitude balloon tracking and recovery—presented by representatives from Edge of Space Sciences <http://www.eoss.org/> APRS in the Sydney Olympics and the versatile Findu.com <http://www.findu.com/> on-line APRS database.

Saturday's sessions included forums on the prospect of using consumer wireless devices (popularly known as 802.11b or "Wi-Fi" devices) to create high-speed Amateur Radio digital networks. A forum on HF digital voice also drew considerable interest.

One of Saturday's highlights was a demonstration of the new ICOM D-Star <http://www.tapr.org/tapr/dv/DStar brochure.pdf> digital radio system. At the heart of D-Star is the ID-1 transceiver, which ICOM had on display at the Dayton Hamvention last spring. The ID-1 operates on 1.2 GHz and can communicate using FM analog voice, digital voice and data. The transceiver can be programmed with a desktop or laptop computer, or it can be operated in a more conventional manner via a remote front panel. ICOM's Ray Novak, KC7JPA, said D-Star will be available in the US in November. (Click here for a sample of D-Star audio recorded at the conference.)

Bruce Perens, K6BP, <http://perens.com/> was the featured speaker at the Saturday evening banquet. His entertaining presentation stressed the notion that individuals, not just corporations, still can innovate and invent. Perens called for grassroots development of Amateur Radio software and hardware according to the Open Source model. He also encouraged the audience to become educators, because, he explained, "the future strength of Amateur Radio is in our value as technology teachers."

SDR was another hot topic at the conference, and the Sunday seminar was devoted exclusively to that subject. Projects such as GNU Radio <http://www.gnu.org/software/gnuradio/gnuradio.html> promise a day when amateur transceivers will achieve extraordinary levels of flexibility. Under the SDR paradigm, software, rather than the hardware, literally will "define" the way in which a radio operates.

Proceedings of the 21st ARRL and TAPR Digital Communications Conference now are available for $20 (plus shipping and handling) via the ARRL Web catalog http://www.arrl.org/catalog/?item=8756. Order item No 8756.

NCPA

Board of Directors Electronic Meeting

Excerpts of the NCPA board remailer traffic, May 1, 2002 through August 1, 2002. Compiled by Gary Mitchell WB6YRU (Quoted material is in italic. Full text of traffic is available).

Dale Jr, William:
June 26, 2002
Can anyone comment on the status of BBS systems in the South Bay? We tried to use N0ARY at Field Day for NTS traffic but it seemed down.

The lack of information for new users is a serious reason why there is so little interest in these system.

Any interest in another forwarding BBS in the Milpitas area? What equipment would be needed? I might be inclined to set one up.

I'd like to see an article on BBS systems overview for new folk. I'd like to invite someone to come talk to our group in Milpitas some month.

WB6YRU:
N0ARY BBS has been down for many months. Bob (N0ARY) has the computer, but hasn't had time (or interest) to work on it. The plan is to bring N0ARY back up, no telling when.

N6QMY BBS (Fremont) went belly-up a while back. N6LDL BBS (Los Gatos) is still active.

The lack of information for new users is a serious reason why there is so little interest in these system.

I beg to differ. The NCPA nearly spent its treasury dry trying to do exactly that. It proved to be a futile effort. It's fairly clear most users abandoned the BBS network for the internet.

APRS, DX spotting, and keyboard are still hanging in there, as far as I know.

Any interest in another forwarding BBS in the Milpitas area? What equipment would be needed? I might be inclined to set one up.

It's not like setting up a node, BBS's need regular care and feeding. But if you might be interested, let's talk off the remailer.

Dale Jr, William:
The NCPA nearly spent its treasury dry trying to do exactly that.

Sorry to hear your efforts were fruitless. Did you investigate why? Where did the money go? The NCPA web site is nice but could use more information on the status of the BBS / DX Clusters and any nodes/gateways.

N6UOW:
This one is my fault...I took on the task

DX Spotting Nodes

September 2002

Location	Call	Alias	Frequency	Coverage
California City	K6ZZ		144.490	Antelope Valley area
	EARN8		144.490	Oak Peak
Castro Valley	W6RGG	DXCV	145.770	East, West, South SF Bay area
Chico	K6EL	DXC	145.670	Chico
	K6EL	DXW	145.670	Oroville, Red Bluff
	K6EL	DX	144.950	South Fork Mtn - Redding area
Hanford	K6UR	DXFRES	144.950	Bear Mtn, Fresno area
	K6UR	DX7	145.770	Mt. Adelaide, Bakersfield
	K6UR	DX16	145.770	Oakhurst
Livermore	NF6S	DXL	145.770	Tri-Valley area
Los Gatos	N6ST	DXLG	146.580	Santa Cruz Mtns, Monterey Bay
	N6ST	DXF	146.580	Santa Cruz/Los Gatos
Oakdale	K6OQ		146.580	Modesto area
Penngrove	K6ANP	DXANP	145.670	Sonoma County
Reno, Nevada	N7TR	RENODX	144.950,146.58,441.500 (2400 baud), 51.7	
	N7TR	PCDX1	146.580	Low Level in Reno
	N7TR	PCDX	144.950	Virginia City, NV
	N7TR	DX2400	441.500 (24OO baud)	
Rio Linda	K6NP	DXRL	144.950	Sacramento, Woodland, Davis

Bob Vallio - W6RGG   wsixrgg@crl.com

of updating the pages, but I have not made the time to do it (and others have been waiting patiently for me, rather then trying to take it over :-)

WB6YRU:
Sorry to hear your efforts were fruitless. Did you investigate why? Where did the money go?

The vast majority of it was spent on information tables at Pacificon (this was before they had free tables for clubs) and on "Intro to Packet" booklets. At first we sold the books. A few years ago the board decided to give our remaining supply (few 100 copies) away to amateur clubs in the region who said they wanted some. And some money was spent on copying costs on fliers and extra copies of our newsletter the Downlink. As to why... That's the $64,000 question.

New packet users are SO Lost ! When they open the box on the new TH-D7a(g)/ TM-D700 and start looking for DX clusters and BBS's they have almost no information.

Therein lies the rub. The NCPA can't afford any big publicity campaign that would be necessary to maintain visibility. New packet users don't pay any attention until the bug bites or they get the equipment and want to try it. And then there's the question of effort and man-power, which the NCPA has in very short supply.

Dale Jr, William:
County Emergency BBS ideas are seriously lame. Much more could be done and post 9/11 there are BUCK$$$ available.

"Dave Willey" KD6KWM:
Why (in your opinion) are "County Emergency BBS ideas are seriously lame."? Please give SPECIFIC & DOCUMENTED examples to back up your claim.

Dale Jr, William:
Well, first the only system I'm familiar with is the Santa Clara BBS system, KE6AGJ-1. It does no mail forwarding and is set up only to handle RIMS traffic - good plan as far as it goes. I am not at all clear on how, if at all, it is connected from the county to the state? Each city here just logs into it directly and drops off/picks up mail - no routing.

Little information exists as to nodes, digis and other cities also running a mailbox or BBS.

Digital Emergency comm. needs to look at how to get data/reports/whatever out of the hole where the commercial systems are down, i.e. phones/internet, and back into the phones/internet outside the Emergency area. All of this is missing or too much is done my hand if it is done at all. NTS is another sad story.

WB6YRU:
July 2, 2002
You need a node and BBS map.

BBS maps won't help users any. Message routing is automatic. Keyboard nodes are a different story. There's not much organization involved with those. Fortunately, any node will list the other nodes that it can connect to.

NCPA

NEWQPSK

Software TNC on a digital signal processor using half duplex, multi-tone QPSK, synchronous, linked, and error correction

NEWQPSK is suitable for HF or VHF, and on HF it is capable of a throughput far in excess of conventional packet modes. NEWQPSK operates on the Motorola DSP56002 EVM, which works just like a TNC. The unit will operate with any KISS interface packet software.

Developed by Pawel SP9VRC, NEWQPSK is a KISS/AX25 packet protocol with awesome performance. 15 individual tones are used and each has QPSK modulation at 83.333 baud. The raw throughput is an impressive 2500 bits/sec. Each packet has a two phase preamble for fast synchronisation and frequency error correction. Data is spread in time and frequency, using a Walsh function, which provides redundancy, forward error correction and therefore significant robustness with respect to impulse noise and interference. There are three FEC modes: none, light and strong, and the use of FEC significantly reduces the requirement for ARQ message repeat requests. In light FEC mode, the throughput is 833 bits/sec, while in strong FEC mode the throughput is an impressive 1833 bits/sec, with the ability to correct up to three bit errors per character without requiring a repeat.

For more information, check out the web page http://www.qsl.net/zl1bpu/ FUZZY/digital.html.

NCPA

TCP/IP on FlexNet, Just Another Layer

Gunter Jost, DK7WJ/K7WJ
Lichtenbergstrasse 77, D-64289
Darmstadt, Germany
Translated by: Don Rotolo, N2IRZ
c/o RATS, PO Box 93 Park Ridge NJ 07656

Abstract:

The goals and outcome of a project to optimize TCP/IP transport over the FlexNet AX.25 network is described. A number of optimizations, and their implementations, are described and discussed. These include header compression, resend minimization, packet age tracking and ACK consolidation, as well as platform considerations and potential uses.

Not long ago, TCP/IP was an exotic mode of operation in Amateur Radio, reserved for specialists. Since then, it has become one of the most popular protocols used in the digital modes, due to the wealth of interesting applications employing it. The project described here attempts to make TCP/IP usage in the FlexNet packet network as simple as possible.

This doesn't mean that we can dispose of the de-facto AX.25 standard. No, TCP/IP is merely a useful expansion, or just another layer, that can be handled by FlexNet.

This also includes the reconnection of disconnected links without the intervention of the operator, including re-routing to faster links. TCP/IP can be seen as an end-to-end layer (Layer 4 in ISO jargon), in contrast to other packet systems that are hop-to-hop layers; namely, to the real ends of a logical circuit (user/user or user/server). TCP/IP strengthens the robustness of the AX.25 connection to the entry node, allowing one to change user ports during a running connection without which the connection must be newly established. Instead of creating our own Layer 4 specification (plans have existed for some time), we find that TCP/IP offers exactly that and is available for most modern operating systems. TCP/IP is the standard protocol of the Internet and there is a considerable amount of software available.

TCP/IP and FlexNet make up a homogeneous unit in that FlexNet nodes need no further expansions or changes, because they are already fully transparent to all frame types offered by the protocol, which are passed on a virtual AX.25 L2 connection. The protocols are passed along cleanly, as long as AX.25 (on the RF path) and TCP/IP (end-to-end) run at the layers they are expected to be at.

In our first look, it was quickly seen that TCP/IP, as presently used in Amateur Radio, was given the (not undeserved) reputation of being slow and inefficient. This isn't the fault of the protocol, but of the implementations. It was clear where the crank needed to be turned. Quite simply, it shouldn't be like that, so we began this project. Today, a few thousand lines of source code that were implemented and tested on the first platform (Windows 95) are ready for use.

The idea of placing TCP/IP services directly on the network as AX.25 shouldn't be ignored. This is an interesting concept, but it should be pointed out that a careful TCP/IP implementation will return more than it costs.

Minimizing Protocol Overhead

TCP/IP packets contain a header of approx. 40 bytes. For a 256 byte packet, this is an overhead of some 16 percent, and when the required ACKs are considered, more than 30 percent overhead increase as compared to an AX.25 connection. This also assumes an optimal protocol implementation, as well as no unnecessary retries. If you lengthen the packets to 1500 bytes (a typical value on Ethernet and similar implementations), the overhead sinks to a more reasonable 5 percent. A lengthening of AX.25 frames to 1500 bytes, as is often suggested, isn't practical. The frame failure rate would become unreasonably large: If an RF link of 256 byte frames is 90 percent failure free (unfortunately, a realistic number), this value would decrease to only 40 percent error-free frames for 1500 byte packets. Sporadic interference such as radar impulses further worsen this value. For this reason, AX.25 segmentation was defined some time ago, so that a large IP packet could be broken into many 256-byte packets. In this way, each frame contains only its own additional overhead bytes, contributing to efficiency. A requirement for this to function is that each packet arrives in the correct order, which can only be safely realized through a VC (Virtual Circuit) connection. The FlexNet AX.25 header compression feature [implemented some time ago - IRZ] functions in any case only with such a static connection and is a source of further reduced overhead.

These improvements are most efficient when there is a relatively large amount of data to be sent (enough to permit segmentation). If this isn't the case, such as in an interactive Telnet session, the value becomes somewhat less. The worst case is when each byte requires its own packet, with its 40 byte overhead. In these cases, one needs some other mechanism. Luckily, there is a standardized header compression scheme, as seen in RFC 1144 [1], and implementation was relatively simple. Using the Van Jacobson scheme [2], the 40-byte header of a static TCP connection can be reduced to 3-8 bytes. the only assumption for this to work is, like the FlexNet header compression scheme, the connection path between both ends is static, i.e. a VC. This protocol was defined for relatively slow telephone connections. If you replace the concept of "Serial Line" with "AX.25 virtual circuit", you can see just how well the protocol would fit with AX.25 encapsulation. The compression and its control occurs on the AX.25 level for up to 250 TCP connections, as well as traffic forwarded through a router. With a Link Reset, the status tables on both sides of the path are re-initialized. This requires that both ends of the path know of these Link Resets, but not all AX.25 implementations do this, which can lead to problems in this area.

Thinking instead of Pumping

A general problem with end-to-end protocols is that the transport shell has only limited possibilities for influence. This is less important for a static TCP connection where the timers regulate the connection fairly well on the basis of the transport capacity, but nonetheless, an orgy of unnecessary retries can occasionally be observed.

The goal of this project was to develop an IP router not only for usage under Windows 95, but for usage under DOS or on the RMNC platform. This gave us a reason to invest a little more work early on to save work in the later cross-platform portings. Finally, we still have to co-exist with users and servers operating with network implementations that are sub-optimal.

With AX.25 these problems, once separated from channel control, are mostly resolved. The network nodes decouple both sides of a connection completely and traditional digipeating is no longer practiced. With this philosophy, an end-to-end approach such as IP becomes an attractive proposition.

IP is a connectionless protocol and the TCP placed upon it is laid out such that packets can take any possible path from one end to the other, arriving in any order. For this reason, it is possible to route individual IP packets in a connectionless manner. An IP router might never see all of the packets of a particular connection. It is also true that when you encapsulate something within AX.25, it cannot be assumed that traffic from one side of a connection travels over the same path as traffic from the other side. Only at the endpoints can you be sure of seeing all packets of a TCP connection.

An IP router also has possibilities for optimization. If a router knows of congestion, it can analyze packets in transit stored in memory and eliminate all unnecessary resends, by simply erasing any doubled packets; for example, the outgoing path in use is slower than the incoming path (not an unusual condition in a Packet network). A time stamp can also be put onto each packet. Packets that remain in the router for some time, say a minute (because of congestion or a broken link) can be erased, hopefully resulting in a TCP retry from the sending station that might be routed along a better path. Implementation of this requires an analysis of the complete IP and TCP headers, so that the order of packets can be determined. This can also be used to consolidate ACKs for a specific TCP connection, passing only the latest to the endpoint.

Through such actions, congestion created by unnecessary resends can be dealt with instead of becoming worse. With traditional implementations, traffic is brought to a complete halt in such situations. Through controlled 'packet loss', as well as the deletion of doubled or aged packets, the net is much better able to deal with changes, slowing the data to match the path's capacity.

This behavior functions very effectively and completes the usual improvements, with the router sending an ICMP-source-quench as well as requesting the sender to throttle back. The retention of the originator is during this not clearly specified, thus such flow control actions can lead to very different results.

Naturally, both partners should adjust their TCP timers properly, but luckily this functions quite well, even with the Microsoft stacks.

The proper adjustment of TCP/IP parameters is critical, above all when one wants to reach destinations via different networks. The Microsoft stack is a "black box," and hardly any parameters can be adjusted. So, we need to make all optimizations in the module which knows the existing transport layer the best: The AX.25 coupler.

Instead of simply pumping TCP/IP packets into an AX.25 connection or, even worse, into AX.25 UI frames, as is done with present implementations (such as xNOS and LinuX), the packets remain in the IP coupler's (e.g., router's) memory until they can actually be sent. This is comparable to the behavior of the FlexNet coupling to Layer 1, and contrasts with solutions such as KISS, which only generates frames that can be sent immediately. The known side effects of a KISS connection, i.e. multiple SABMs or RRs in a single transmit cycle, simply don't occur with FlexNet.

In FlexNet, ACKs are sent only for the latest packet, any ACKs belonging to earlier packets need not (and are not) sent. A carryover of this concept to the IP-AX.25 coupling brings some improvements. When the AX.25 connection is being established, during which no data can be passed, running TCP/IP packets are buffered, and so can cause unnecessary retries. The same happens with a user station is blocked for extended periods of time in a half-duplex connection. During this time it is possible for many TCP packets to pass to the user station, each of which is ACK'd individually by the user's local TCP stack. Instead of sending each of these ACKs, the FlexNet coupler can discard all of them except for the last one. Implementing this requires a close coupling of the layers, but FlexNet already has the call in it's API, so that little accommodation is needed here. Naturally, the coupler must analyze and compare the packets at the TCP level to be able to discard the proper packets. In the ideal case, a user station sends a maximum of one AX.25 I-frame during its time slot per active TCP connection, instead of the many L2 RRs and TCP ACKs as with the traditional solutions. The improvements thus realized are somewhat greater than those obtained from header reduction alone. While header reduction remains important, one

Packet Sysops of Northern California Packet Bulletin Board Systems

March 2002

Call	Location		User Ports
WH6IO	Benica		144.99, 145.71&+, 145.75&, 433.43&+
WA6ZTY	Berkeley		144.97
KE6I	Berkeley		145.01&, 433.43&
N2THD-1	Citrus Heights		145.07, 441.50
N6CKV	Gilroy		144.99
N6LDL	Los Gatos		144.97, 145.71&, 441.50
WA6NWE-1	North Highlands		144.93, 145.09, 145.75, 441.50
KD6DG	Redding		145.09
W6CUS-1	Richmond		145.63
N0ARY-1*	San Jose	*	144.93, 433.37&
K6YV	Sonora		144.97
WA6EWV-1	South Lake Tahoe		144.97
W6YX-9*	Stanford Univ	*	145.75+
W6SF	Stockton		144.99
K6MFV*	Walnut Creek	*	144.31, 145.71&+

Keys:
& = 9600 Baud Port
+ = TCPIP Port
* = Currently Inactive

cannot ignore the gains in efficiency from eliminating doubled or aged TCP packets.

In a typical HTTP session, many TCP connects are started, transferring only a few kilobytes of data before being closed again. TCP timers have no chance to adjust and many unnecessary retries occur. This is especially true in a slow half-duplex environment which is typical for a user station. The situation is less dramatic when users have few, long-lasting connections, e.g. a single FTP transfer.

For AX.25, FlexNet was able to make such optimizations directly in the L2 code, because of the tight coupling to L1. This avoids the need for the L1 process to have to analyze each L2 frame and decide what should be passed ahead or not. Unfortunately, it isn't so easy at the TCP level. Here, you must watch the status of possibly many running TCP connections, which requires a lot of memory and efficient algorithms. However, some coding efficiencies result because TCP compression needs quite similar data structures, which can be partially be used for these optimizations.

All this naturally increases demands upon memory space and CPU capacity. As long as the network runs at data rates below 100 kB/s and PCs continue to increase in capacity, no further effort is needed. This is especially true on the user side, where nowadays there is more than enough computing capacity. A fast 486-class machine with 16 MB of RAM running Windows 95 is sufficient to support a throughput of 76,800 baud TCP/IP through FlexNet. For a router, it would help to install a faster CPU, but reasonable performance with a 486 was measured.

Anachronisms

All the improvements described above assume IP transport over AX.25 Virtual Circuits, which should provide a reliable path between two points of an IP network. Using Datagram mode (AX.25 UI frames), each packet loss on the wireless portion of a TCP connection causes a TCP retry, which has to travel the whole way, end to end. (Remember the problems we had with L2 digipeating? The same applies here). Since this has not been generally learned and understood, an IP router must also be able to deal with the Datagram mode, where IP packets are sent as AX.25 UI frames. With this, one becomes stuck on the problem that only IP packets with a gross length of 256 bytes or less can be carried. AX.25 segmentation is not usable, as it depends upon the packets traversing the network in their original order.

To resolve this problem, we have available so-called IP fragmentation, which divides an IP packet into multiple smaller packets and gives (in contrast to AX.25 segmentation) each fragment it's own complete IP header. Each packet is thus autonomous and are reassembled in order only at the destination. This allows each packet to traverse the network through any available path, forbidding reassembly in routers or gateways. One problem is that when one fragment is lost, the entire packet (a series of fragments) must be re-sent. This scheme is clearly at best a temporary patch, especially when one must deal with the realities of the RF medium and the relatively high failure rate of the links. For compatibility reasons, when Datagram mode is the only way to deliver a packet, IP fragmentation must be implemented in the router.

IP Routing

IP routing is, in one regard, in competition with established AX.25 routing, especially in the FlexNet environment. In that IP addresses are mated to a call sign (dynamic address allocation for users is a subject for separate discussion), FlexNet routing is sufficient to reach a specified destination. One useful addition could be to define a layered hierarchical routing, such as AX.25 routing in the local area and IP routing for the wide area, analogous to protocol layering. An obstacle to this is that our network is not yet fully developed to offer fast connections with transit times of less than a minute over its entire range.

One argument in favor of IP routers at network facilities is that the user is freed from some work—he simply sends packets to the router and it handles the rest. At the moment this merely means that the sysop's work, in which he might not have much interest, is delegated. The user shouldn't become dependent upon this system; however, he must retain the possibility of establishing a connection with a destination directly. This means that PID transparency, reversible call sign fields and a functional AX.25 router remain important. The network can be considered a black box for the use, who doesn't need to know its internal operation to reach a destination.

With more and more users operating TCP/IP (perhaps in part due to this project), and TCP/IP is understood to be the end-to-end layer on a well-constructed AX.25 network, we can then reconsider the optimal bundling of traffic via dedicated routers. The infrastructure is, in the form of many LinuX systems, more and more available. All that was missing were the protocols and their implementation.

If the TCP traffic is optimized as discussed above, then routers can reduce network loading considerably. Routing, however, remains in the background. In any case, this is an exciting direction for development and there is considerable room for new ideas.

Until then, and possibly also after that, the user is better served with making a direct connection himself, as far as the network allows.

Which platform is the right one?

While we lean towards additions to some specific operating system, the rest of the world wants missionaries. Indisputably Linux, for example, is an outstanding server/operating system. The average user, however, wants to simply buy something from Microsoft. You just can't beat Windows 95 or NT when you want to develop a user application.

For Windows 3.x, ETHEREMU from Thomas Sailer, HB9JNX, could be installed. It emulates an ethernet card and allows it to communicate via Trumpet Winsockets. The AX.25 routes are input using text mode commands.

TCP/IP is already integrated with

Windows 95. Of course data transfer via Ethernet or "DFU-Network" is provided using SLIP or PPP protocol. Until now, what was missing was a coupler that could encapsulate TCP/IP packets within AX.25. Some solutions already exist, but are generally difficult to configure, for example NOS in a DOS-box. Others try to use protocols such as SLIP, which has the disadvantage of removing the possibility of being able to have normal AX.25 connections. Besides, one needs special and somewhat expensive hardware.

In the meantime, PC/FlexNet runs under Windows 95, all that's missing is the coupler that hands the IP packets over to FlexNet. Microsoft makes this possible only via the installation of a "Network Card," however the concept is, at least in the German translation, somewhat misleading. Of course, all of the optimizations described have been implemented using Windows 95 as a test platform. Although the timing of the TCP stacks is really set up for a fast wire connection, it keeps stations on the RF channel fairly civilized and sends practically no unnecessary packets. A list box is used for AX.25 route inputs and it also serves as a status and statistics display. No special TNC hardware is required and running multiple IP sessions in parallel with AX.25 sessions presents no problems. However, this is a solution mainly for users, i.e. network clients.

As a server system, Windows 95 is somewhat strained. While there are some simple FTP and HTTP servers that function well, the entire system in not stable enough for use as a general-purpose server. One negative is that there is (still) no possibility of forwarding IP between multiple ports, as well as AX.25 to ethernet ports. This is the domain of Linux, but NT is close behind. Linux is already available with AX.25, so the next development point lies with NT. Having a large installed user base for 95 and NT doesn't hurt either.

Uses

As already mentioned, applications and services are sufficiently available. Much that is developed for the Internet is also usable in Amateur Radio. Regardless, that shouldn't prevent someone from developing an amateur-specific program. Services such as databases or special applications like DX-Clusters are easier to access and service using TCP/IP instead of AX.25. Thanks to TCP port addressing, it's no problem to offer and try various applications and services on the same server with the same AX.25 call sign.

A further field for new applications is for image and speech transfer. While these demands today seem somewhat utopian, voice mailboxes are already using the network to transfer their messages via store & forward. It isn't a much larger step to make it possible for the user to have a direct digital connection into the system. The software already exists, though this is mostly designed for the relatively high speed of the Internet. If you don't need real-time and full duplex and can deal with PTT (amateurs know this already), one can get very good results. Modern codecs allow speech to be compressed to less than 300 bytes/sec with little loss in quality. This bandwidth is already available in most of the network. In the future, a Windows application will be introduced. With specific servers it should also be possible to have roundtable discussions like on the local FM repeater.

Image transmission also remains in the range of possibility, especially with the quality codecs available for moving picture transmission. Naturally the existing available bandwidth won't work for decent quality video, but you don't really need it to admire someone's photograph. Convenient video conference systems are realizable in the foreseeable future.

All this carries with it the ability to increase the play value of the network and through that, justify the spectrum being used. Naturally, our packet-oriented data transfer system isn't set up for real-time uses. The protocols and applications to be used require careful consideration of our unique requirements and conditions, but there is a lot of room for experimentation here.

Unquestionably, a careful optimization of all parts of the transfer chain is needed. What is lost in one component, whether in hardware or in software, cannot be recovered. For example, while increasing the baud rate is a good idea, it is not the only possibility for improving the network. Clearly, TCP/IP can bring to Amateur Radio much more than just a wireless Internet.

References:

[1] Jacobson, V., Network Working Group, Request for Comments 1144, February 1990.

[2] See also "TCP header compression according to Van Jacobson via AX.25" (Jost) elsewhere in these proceedings.

Gunter Jost's article on TCPIP over FlexNet mentions the RMNC platform. RMNC FlexNet is stand-alone FlexNet node hardware that does not require a computer as PC-FlexNet does. It has slots for plug-in modems for whatever baud rates may be required. It can handle up to 16 ports. Typical RMNC nodes in Europe have a few user ports, with the rest being dedicated to high-speed backbone links.

As Don Rotolo pointed out several years ago, there are RMNC FlexDigis in Europe that process over 10 MB of data every day.

I generally do not bring up the RMNC platform in discussions involving U.S. Hams, because we have yet to generate the kind of traffic that would require that kind of capability. The PC version of FlexNet is, I believe, adequate for our present and near-future needs, and is less expensive to set up, so it is much more likely to see use here.

For those who want to try out the RMNC platform, it is available... I just do not see it as necessary at this point for U.S. Hams. For high-speed modems and radios for PC-FlexNet visit these two web-sites:
http://www.sanco.se/baycom
http://www.baycom.org

The folks at both sites are able to handle orders from U.S. Hams.

Charles, N5PVL

NCPA

Digital Channel Allocations in Northern California

N C P A March 2002

50 MHz

50.60-50.80 (20 kHz channels, non-specific at this time)
51.12 SCA backbone
51.14 BBS
51.16 Keyboard to Keyboard
51.18 Experimental
51.62 TCP/IP, 9600 baud
51.64-51.68 (20 kHz channels, non-specific at this time)

NOTE: On this band adjacent channel interference is harder to overcome for repeaters. NARCC requests that any new six meter permanent packet installations (such as nodes) please check with their six meter coordinator. You don't need a formal coordination, but they would like to be aware of your station and have an opportunity to check for possible conflicts first.

144 MHz

144.31 BBS
144.33 Balloon & experimental
144.35 Keyboard to Keyboard
144.37 BBS LAN forwarding
144.39 APRS (U.S. and Canada)
144.41 duplex, lower half (145.61 upper half, 1.2 MHz split)
144.43 TCP/IP (OK to run duplex with 145.65)
144.91 Keyboard to Keyboard
144.93 BBS
144.95 DX Spotting
144.97 BBS
144.99 BBS
145.01 User access
145.03 Keyboard to Keyboard
145.05 Keyboard to Keyboard
145.07 BBS
145.09 BBS
145.61 duplex, upper half (144.41 lower half)
145.63 BBS
145.65 TCP/IP 9600 bps (OK to run duplex with 144.43)
145.67 DX Spotting
145.69 BBS
145.71 9600 bps
145.73 BBS
145.75 TCP/IP
145.77 DX Spotting
146.58 DX Spotting

NOTE:
Allocations from 144.31 through 144.43 are relatively close to the weak-signal sub-band–please watch your FM deviation.

220 MHz

219.05-219.95 100 kHz channels, Backbone
223.54 LAN
223.56 LAN
223.58 LAN, Gilory (GARLIC)
223.60 LAN, Sacramento Valley (SACVAL)
223.62 LAN, South Bay (SBAY)
223.64 TCP/IP
223.66 Keyboard to Keyboard
223.68 DX Spotting Backbone
223.70 LAN, Monterey Bay & North Coast (MRYBAY)
223.72 LAN, North Bay (NBAY)
223.74 Backbone, DX Spotting

NOTES:
- 219 channels are by coordination only. There are currently political problems with using 219-220, making them unavailable in most of northern CA.
- On 223.58, TCP/IP interlink (Sacramento) is secondary, not to interfere with node uplink.

440 MHz

431.45 / 434.85 Duplex (100 kHz)
431.55 / 434.95 Duplex (100 kHz)
431.65 / 438.40 Duplex (100 kHz)
431.85 / 438.60 Duplex (100 kHz)
431.95 / 438.70 Duplex (100 kHz)
433.05 TCP/IP backbone (100 kHz)
433.15 BBS backbone (100 kHz)
433.25 DX Spotting backbone (100 kHz)
433.33 Experimental (60 kHz)
433.37 BBS, 9600 baud
433.39 DX Spotting
433.41 BBS LAN
433.43 9600 baud TCP/IP
433.45 BBS LAN
433.47 Keyboard Interlink
433.49 TCP/IP
433.51 Keyboard
433.53 Keyboard
433.55 BBS LAN
441.50 Any digital

900 MHz

903.500 1 MHz wide, TCP/IP
904.500 1 MHz wide, TCP/IP
915.500 1 MHz wide, experimental
916.100 200 kHz wide, experimental
916.300 200 kHz wide, experimental
916.500 200 kHz wide, experimental
916.650 100 kHz wide, experimental
916.750 100 kHz wide, experimental
916.810 20 kHz wide, experimental
916.830 20 kHz wide, experimental
916.850 20 kHz wide, experimental
916.870 20 kHz wide, experimental
916.890 20 kHz wide, experimental
916.910 20 kHz wide, experimental
916.930 20 kHz wide, experimental
916.950 20 kHz wide, experimental
916.970 20 kHz wide, experimental
916.990 20 kHz wide, LAN links (Contra Costa County only)

NOTE:
900 MHz activity is on a non-interference basis to vehicle locator service. This sub-band is not considered suitable for omnidirectional systems. Use for point-to-point links only.

1296 MHz

1248.500 1 MHz wide, experimental *
1249.000-1249.450 Unchannelized, experimental
1249.500 100 kHz wide, experimental
1249.600 100 kHz wide, experimental
1249.700 100 kHz wide, experimental *
1249.800 100 kHz wide, experimental*
1249.870 20 kHz wide, experimental
1249.890 20 kHz wide, DX Packet Spotting
1249.910 20 kHz wide, experimental*
1249.930 20 kHz wide, experimental*
1249.950 20 kHz wide, experimental*
1249.970 20 kHz wide, experimental*
1249.990 20 kHz wide, experimental*
1250.500 1 MHz wide, experimental
1251.500 1 MHz wide, experimental
1297.000-1298.000 Unchannelized, experimental
1298.500 1 MHz wide, experimental*
1299.000-1299.450 Unchannelized, experimental
1299.500 100 kHz wide, experimental
1299.600 100 kHz wide, experimental
1299.700 100 kHz wide, experimental*
1299.800 100 kHz wide, experimental*
1299.870 20 kHz wide, BBS LAN
1299.890 20 kHz wide, DX Packet Spotting
1299.910 20 kHz wide, BBS LAN
1299.930 20 kHz wide, experimental*
1299.950 20 kHz wide, experimental*
1299.970 20 kHz wide, experimental*
1299.990 20 kHz wide, experimental*

* Full duplex channel pairs at 50 MHz separation, example:
1249.910 ↔ 1299.910

Definitions

9600 BPS Stations using 9600 baud with direct FSK (G3RUH, TAPR, etc.) modems.

Backbone No uncoordinated stations. These channels are for specific purposes as defined by the NCPA and/or affiliated groups. These are frequencies where the various BBS, nodes, and networks forward traffic and are very high volume channels. Please use the normal user entry points of the network you want to access rather than these channels.

BBS These frequencies are for user access to a full-service BBS. Keyboard-to-keyboard is tolerated. Please don't put high level nodes or digipeaters on these channels since they are local. A low-level direct link or node that links into a backbone on another frequency is the proper implementation.

Duplex Simultaneous transmit and receive by a single station, including digital repeaters. Duplex channels are intended for high-volume applications. 9600 baud or higher is encouraged, but not required at this time.

DX Spotting Northern California DX packet spotting network. No other activity should be on these channels.

Experimental Anything goes except full service BBS or any 24 Hr/Day services (nodes, gateways, etc). This is where you can test new gear, programs, etc. These channels may be reassigned in the near future, so no permanent activities please.

Forwarding same as *backbone*

Keyboard to Keyboard Primarily chat channels. These are also the primary emergency channels. No high-volume activity such as full service BBS, DX Spotting, TCP/IP, etc.

Interlink same as *backbone*

LAN Local Area Network. BBS's are grouped into LAN's for more efficient forwarding. A LAN frequency is the forwarding channel within a LAN and to the backbone. Please do not attempt to access the BBS network on these channels.

Personal mailbox/maildrop A BBS-like system, often running entirely within a TNC, with a small number of users that handles information of a personal, local or special-purpose nature. A mailbox is allowed on keyboard-to-keyboard channels ONLY if it does not forward with other BBSs. Mailboxes may forward with full-service BBSs on LAN channels at the discretion of the BBS SYSOP.

TCP/IP Stations using TCP/IP protocol on top of AX.25. Some AX.25 tolerated to communicate to TCP/IP stations if a compatible p-persistence access method used.

User Access User access to a network. This is for the next generation of packet which is expected to operate like the internet. Users would access such a network on these frequencies. The load on these channels may be rather high, like BBS channels. The activity may be any combination of BBS, keyboard, TCP/IP, or
other modes.

Procedure for changes

Send requests for changes to either the frequency coordinator or the NCPA board. The frequency coordinator will then present the request to the board along with suggested assignments. The NCPA board, elected by you, the packet user, makes all assignments.

Misc. Info.

Packet tends to splatter if the deviation is set too high. Please keep your deviation to less than 5 kHz.

Except for the 219-220 MHz segment, the NCPA currently does not coordinate individual stations, nodes, etc. leaving that to the special interest groups. BBS station coordination is done by the PSNC in Northern CA. DX spotting is coordinated by DXPSN. Some digital has been coordinated on auxiliary channels by NARCC.

The NCPA board conducts most of its meeting activity electronically by internet e-mail remailer, ncpa@kkn.net. As with face-to-face board meetings, interested persons are welcome. For more information about the remailer send email to ncpa-request@kkn.net with just the command HELP in the message body, nothing in the subject, and an email message will be sent to you. Subscribe by using the command SUBSCRIBE in the message body. Subscribing to the remailer is like attending a continuous NCPA board meeting. One must subscribe before posting messages.

Northern California Packet Association

The NCPA fosters digital communications modes of amateur radio through education, band planning, and acts as an umbrella organization for various packet special interest groups. Your annual dues helps pay for this newsletter and other educational materials activities. If you might be interested in getting more involved, please let us know.

Call: _______________ Home BBS: ________ e-mail: ________________________

Name: _____________________ Address: __________________________________

City: ______________________ State: ___ Zip + 4: _______ Phone: __________

☐ New Membership ☐ Renewal ☐ Change of Address ☐ I'm an ARRL Member
☐ One year: $10 ☐ Two Years: $20 ☐ Three years: $30
(make checks payable to NCPA)

Please indicate your area(s) of interest:
☐ BBS SysOp ☐ BBS User ☐ APRS ☐ NET/ROM ☐ TCP/IP ☐ High-speed packet
☐ DX Packet Spotting Network ☐ Keyboard to Keyboard ☐ FCC/legal issues ☐ Other:

NCPA *Downlink*
Northern California Packet Association
PO BOX K
Sunnyvale CA 94087

First Class