STRATEGIES FOR IMPROVING WIDE-AREA NETWORKS

James Wagner, PhD - KA7EHK
3 1677 N. Lake Creek Drive
Tangent, OR 97389
(541) 928-7869
wagnerj@proaxis.com http://www.proaxis.com/~wagnerj

ABSTRACT

Wide-area single-frequency networks still cover large areas of this country. A number of strategies have developed for improving such networks, but these strategies are very slow to be adopted. This paper discusses some of the reasons for the continued existence of these networks, some of the strategies and their likelihood of success. KEY WORDS: WIDE AREA NETWORK, BACKBONE

INTRODUCTION

Packet networking began as a single-frequency network, often on 145.01MHz. NET/ROM software for TNCs permitted easy linking among nodes and it made it possible for users many miles apart to converse via packet. These networks are also the lowest cost, since only one transceiver, one antenna, and one TNC is required per node.

While this kind of network provided great advantages in comparison to the previous state of affairs, we now know that they have some major limitations. Yet, they persist. If one observes where they exist, at least in the west, it is possible to guess why. For example, a 145.01 network exists (with no overlaying backbone) over (almost) all of Wyoming, eastern Montana, Eastern Idaho, almost all of Nevada, most of Alberta, some of Utah, and other adjoining areas to the west, east, and south. Single frequency networks also exist in large areas of the mid-west and other parts of this country.

There are a number of reasons why these networks continue. At least some of the reasons are:

1. Population density is very low. Thus, hams are scarce and resources are very limited.
2. 2Meters may be the only VHF/UHF band node-ops are comfortable with.
3. There may be propagation problems on higher bands.
4. Land relief seems to be either very flat with sparse to no high points or very mountainous.
5. Weather conditions (especially winter) can be severe
6. Land ownership is often in federal hands. In many places, this makes access to high points difficult and/or expensive.
7. Power sources may be very limited, though this is changing with the installation of cell sites.
8. The initial step of introducing a backbone is often seen as disruptive.
9. There is often a cultural resistance to providing resources for an activity such as BBS forwarding which is seen as either no benefit or antagonistic to local users.
10. Such networks tend to inhibit the growth of packet & this can be seen as desirable by existing users.

11. Clubs are often not used to cooperating about anything more than a local repeater. Packet networking sometimes seem to be more than they can cope with.

Given this situation, what has been done and what can be done to improve these networks?

BACKBONE MYTHS

There are, of course, some myths which have grown up about backbones and these myths often contribute to the reluctance to adopt this technology.

Mvth 1: Backbones need to be on UHF. Of course this is false, but still widely believed. As a simple example, there is an excellent 6M, 2400 baud backbone which extends along the southern half of California-Nevada border to Las Vegas, then east to Phoenix. While this backbone is subject to the wide-area problems as single-frequency networks, it gets users off the linking frequency. 2400 baud also makes it less simple for users to connect directly to it and increases throughput.

Myth 2: Backbones can't be on 2M. Again, false. An interesting scheme is used in Alberta with a 2M backbone on 145.01 and user ports at the high end of the 2M band. This provides better than 2MHz of separation and standard repeater cavities can take care of this very well.

Mvth 3: You can't convert to a backbone without disrupting the network. False. Some creative strategies for setting route quality are about all that is needed. Of course, there will be disruption when nodes change from their old frequency to a new one, but experience has shown that the result is seldom catastrophic.

Myth 4: Backbone nodes are difficult to make. This may be somewhat true for a node with more than 2 frequencies, but for 2, only, its a "snap".

Myth 5: You can't teach an old dog new tricks. Sorry, but packeteers have disproven this long ago. A backbone can be perplexing to users who have never used one, but the principles are pretty basic and any club worth its name can solve this without much difficulty. Besides, packeteers are often there because they WANT to be at the cutting edge, or near it, at least.

TWO METER BACKBONES

As Alberta hams have discovered, sparse population can be an advantage. There are not many repeaters and not much simplex use. The result is that there are often repeater pairs at the high end of the band which are not used. It also means that the simplex slot from 147.42MHz to 147.57MHz is probably lightly used. The biggest problem may be that even a very lightly used simplex frequency or one which does not seem to be used at all may be "claimed" by a group of hams. Real care needs to be exercised when using the simplex option, but it can often work.

Using this scheme, a second 2M port is added to the networking node. This port is at the high end of the band. Sometimes, a duplexer allows using one antenna for both ports. Then, after a "get acquainted" time, the node parameters are set to disallow user connects on the networking frequency.

This greatly reduces the hidden transmitter problem for both users and the network. Users do not have to compete with network activity and visa versa. Users can also see, quite quickly, the benefits of a backbone and it can provide some real encouragement for technical or financial support.

This strategy can easily be used as a stepping-stone to backbones on other bands or at higher baud rates. Once users are off the backbone, the actual backbone frequency becomes less important.

INSERTING A BACKBONE IN AN EXISTING SINGLE-FREQUENCY NETWORK

While the methods which have sometimes been used to add a backbone to an existing network HAS resulted in disruption, it does not have to be very much. It only requires some thought, some planning, and some careful execution. But its NOT “rocket science”.

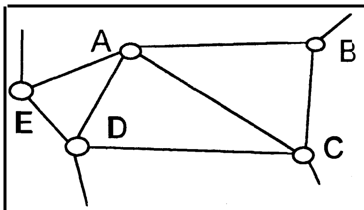

FIGURE 1 - Simple Network

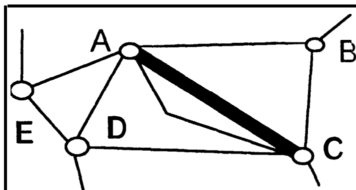

FIGURE 2 - Simple Network with Backbone link

Consider the simple hypothetical network shown in Figure 1. Suppose that node-ops want to convert the link between nodes A and C to a backbone. Technically, its not difficult to do that. A second TNC, a second radio, and (often) a second antenna at each end is about all it takes. Without changing the user frequencies at either end, the result is shown in Figure 2. The heavy line represents the backbone link.

At first, this arrangement would seem to be quite useless. But it is a very important first step. It is critical, however, to choose some carefully thought-out route qualities to get any advantage from it. Lets suppose that 192 is commonly used as the quality for good paths to neighbors. It is only necessary to select a slightly higher path quality for the backbone link, but this must be done at both ends. A value of 205 would work, so lets see what happens. The quality for any two-link path in the old network, say from E to A to B, is 144. Likewise, the route from E to D to C is 144. But E to A to C is 153. The preferred routing from E to C is now via the new backbone link even though a parallel link on the user frequency is still present. This arrangement also provides an alternate route if the backbone does not work as well at first as expected.

What are the consequences of this choice? Clearly, D will carry less of the traffic from E to C but A will carry more. The users at D will have less network competition but those at A will see more. Depending on existing traffic flow, this may be significant or it may not. But, in spite of this, there will be a net reduction in user frequency activity at C because much is now going via backbone. And, while A gets more traffic to and from E, it is transported by backbone to C, and the former traffic directly between A and C is also now hidden. Users at A will probably see little net change, and it may improve. Clearly, how the network traffic flow redistributes itself and what the consequence is depends on how much traffic and which nodes carry it. Generally, some uses will see a clear benefit and others will see anywhere from a slight benefit to a slight degradation.

But, in spite of this addition, the node-ops are left with a dilemma. A and C need to retain their old frequency in order to keep connectivity to the rest of the old network. What to do? The next thing might be to add B or D (or both) to the backbone. Adding both is not really desirable, given the links shown in Figure 1, because there is no link between B and D, and that probably means that one is hidden from the other. Lets suppose that B is the node which is added to the backbone. Figure 3 shows the result.

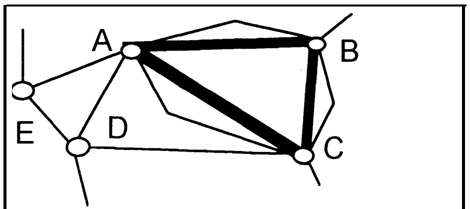

FIGURE 3 - Simple Network with Backbone cluster

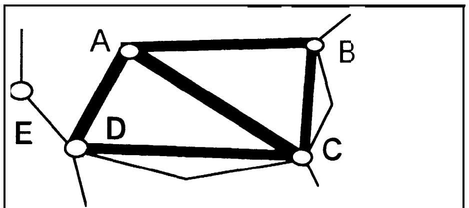

FIGURE 4 - Simple Network with Backbone cluster

This arrangement is properly called a "cluster" but the node-ops of A, B, and C still have the same dilemma. They have to stay on the old network frequency in order to retain connectivity. Users at both B and C should still see a clear improvement, but it might not be dramatic.

Adding D to the cluster, while not desirable from a hidden-transmitter standpoint, will still result in significant improvement, especially at A. Now, all network-level traffic passes through on a non-user frequency at A, eliminating all competition with users. Because there are no remaining links (without alternatives) to the old frequency for A, it can change frequency. The result of adding D to the cluster and changing the user frequency of A is shown in Figure 4. From this simple move, users at D, C, and B will see a major improvement because their nodes no longer have to compete with A. Drawbacks are: traffic between A and E must now travel via D, and D becomes a single-point-of-failure in the new network.

While no network will behave exactly like this example, it should be apparent that careful planning and execution can result in the addition of a backbone without major disruption. In all of this, the only real disruption was the frequency change at A, but I'd bet that the users would really like the results! And, all of the gain resulted from a link which, at first, appeared to offer no real benefit.

IMPROVING EXISTING WIDE-AREA BACKBONES

There are many areas covered by backboned networks which cover far too large an area. A good example is the 1.25M network in eastern Washington, eastern Oregon, and western Idaho. This network contains at least 15 nodes from Mt. Hood in Oregon to Boise, Idaho to Spokane, Washington. The east-west extent is about 280 miles (450km) and the north-south extent is about 300 miles (480 km). Clearly, hidden transmitters are a problem and a simple check of the network map shows this to be true, even in smaller sections of the network. Figure 5 shows this network and some of the connecting links to other networks.

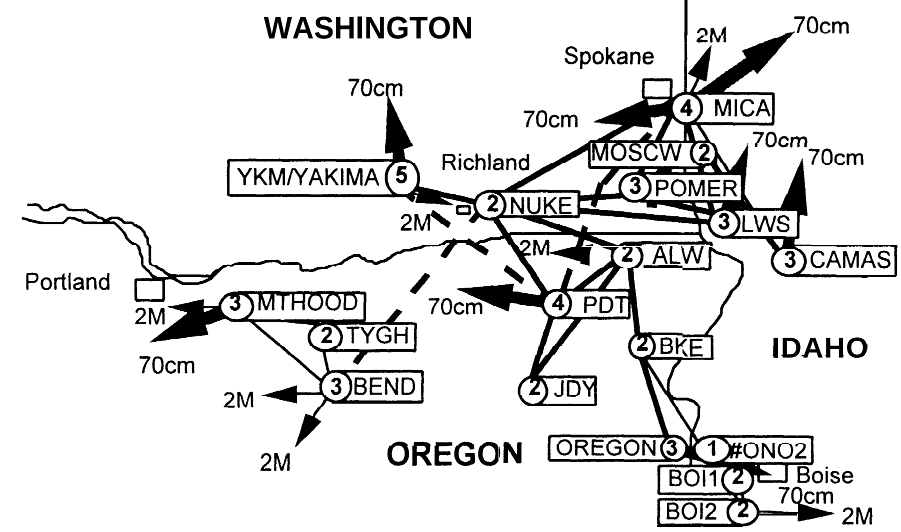

Figure 5 A Pacific NW 1.25M Backbone Network

The network shown in Figure 5 has some clear problems. Perhaps the most obvious is the group of 3 nodes at the west side which is connected only tenuously to the rest of the network. The one link works only part of the year, and then, only for part of the time. Several of the western nodes are heard occasionally by other nodes in the eastern part of the network. For all concerned, operation would be improved in both parts of the network if one section, say the western part changed frequency. The east-west extent would then be reduced by over 100 miles (160 km).

A fairly clear improvement would be to change a portion of the eastern part of the network to another frequency. This requires one or more nodes to become 3-port locations. But, several of the nodes already have 3 or more ports and those should not have to carry added burden. One good candidate would be the node called ALW; the linear path to Boise would then be away from the rest of the network. Much of the remaining at least DCDs each other (most of the time) so makes sense as a connected cluster. Here, we have one node adding one port and 5 others (BKE, #ONO2, OREGON, BOI1, and BO12) changing frequency. The cost is one radio and TNC and the result would be significant improvement for all.

This example should show that one of the real aids to network improvement is to construct a map of the network, and try to determine which paths work well, some of the time, or only rarely. Such maps often help to show how the network functions and then, how it can be improved.

CULTURAL IMPEDIMENTS

While it is not frequently discussed, culture may play a role equal to, or greater than, technical considerations in the development of networks.

For example, in the Northwest, and I suspect also in other parts of the country, there have been literal "wars" over BBS forwarding and message content. These conflicts have often concerned issues of whether a node should be "clogged" by BBS messages which some local users find worthless (ie, for-sale messages from another part of the country). In some cases, nodes have been purposely left unimproved just so that the flow of BBS traffic would not take place.

Another cultural issue has to do with keeping a packet node "for me and my buddies". If the node is not improved, newcomers become discouraged and put their TNC back on the shelf. Then, the node is not excessively congested and there is no obvious reason to "improve" it.

The biggest problem is often that this kind of thinking is seldom discussed in the open. People rarely admit that this is really why the local node has stayed single port for 10 years.

When reasons like these surface, there are only a few "solutions". One is to just leave things as they are and walk away. Another is to attempt to convince the responsible people that they, other packeteers, and ham radio in general would all benefit from an improvement. Another is to make an end run around the existing network, find a few who are interested in really improving things, and go at it yourself.

Another problem which is not well recognized is the need for cooperation. Ham groups are often not used to working with neighboring clubs, especially over the wide areas required for a functioning packet network. Often, there is suspicion or outright animosity, sometimes generated by prior conflict over repeater frequencies. Sometimes, this can be overcome by some respected group inviting others to a regional meeting about packet networking.

CONCLUSIONS

Several methods have been discussed for improving single frequency, wide area networks. Some of these ideas may work in some situations, and some won't. But above all, do not be afraid to think, question, work with other hams, and find even better solutions than the ones suggested here.

BIOGRAPHY

James Wagner is an electrical engineer & programmer employed by Kalatel Engineering in Corvallis, OR. His work is in the area of embedded controllers and design of system components for the closed-circuit video security industry. He has a BA in Physics from Oregon State Univ, an MS in Electrical Engineering, also from O.S.U, and a PhD in Electrical Engineering from Colorado State University. He has been employed by Tektronix and by the College of Oceanography of Oregon State University. His interest in ham radio began in the 1950's but did not actually get a license until 1979. He has been the advisor to the Oregon State University ARC and the node-op of their packet node. He is the author of "The Amateur Packet Radio Handbook".