OF VIRTUAL CIRCUITS, DATAGRAMS, AND THE CIRCULAR FILE

Terry Fox, WB4JFI
President, AMRAD
1819 Anderson Rd.
Falls Church, VA 22043

Introduction

This paper presents a slightly biased view of the main two types of networking concepts being discussed for amateur radio.

Overview

Amateur packet radio made a major breakthrough last year. After a couple years of development, a standard has been adopted for point-to-point packet communications, often referred to as the Link Layer, or Level 2 of the IS0 reference model.

Even as work was being completed on the link layer, amateurs were beginning to take on the challange of designing a true amateur packet network system. Two "camps" have taken shape in this stage of development work, the 'Virtual Circuit' camp and the Datagram' or 'TCP/IP' camp. Both groups are working on software, and I believe both will be used for a period of time to see which is best suited for amateur packet radio.

One thing both groups generally agree on is that what must be provided by the amateur network is a method of getting data from a source to a destination fairly reliably. Both groups agree that this should be assured by a transportation device at each end of a communication path, and that this communications path be absolutely reliable if necessary. This means both parties are actually designing systems that function at both levels 3 and 4 (network and transport layers). The result of this work should create "virtual-connections" between two interconnected devices within the amateur network. This virtual-connection exists between the involved devices at the interface between the Transport Layer of the IS0 reference model and whatever layer resides above it (such as a Session Layer). Since some may object to the term "virtual connection", I will instead use the term "logical network connection".

Unfortunately, the word "network" has come to mean many different things. It can mean the general concept of a large group of nodes interconnected so that data can flow back and forth between any nodes within the group. This type of network can be geographically small (as in Local-Area-Network, or LAN) or large (such as the Telenet Network). This size grouping can add to the confusion when discussing networks.

The term "network" can also mean the specific "Network Layer", or Level 3 of the IS0 reference model. The network layer is sometimes considered two sub-layers, which can also be confusing.

Throughout this paper, I will use the term "amateur network" when discussing the overall network concept. I will use the term "transport entity" when describing the interface between the upper IS0 layers and the amateur network access point. When discussing a single cluster of potentially interconnected stations (such as a group of VHF packet stations within communications range), I will use either the term "intranet" (thanks Paul!), or subnetwork, as the IS0 calls it. The term "internet" (note lower case) will be used to describe the potential interconnection of individual intranetworks to form an amateur network. This is different from Internet, which is a specific internetworking protocol.

Services Renderred By The Amateur Network

In the most basic terms, the amateur network should provide a means of transferring data from one amateur to another amateur. Ideally, both data integrity and transfer speed are important to all amateurs, but integrity and/or speed may be compromised in individual situations. The amateur network should be flexible enough to handle such special requests as reduced integrity to increase throughput (speed) for applications such as packetized voice. The other end of the pendulum is equally important. If an amateur wants to send a machine language program across the amateur network, speed may be sacrificed in order to insure absolute data integrity.

Since we amateurs live in the real world, and amateur radio is our hobby (it doesn't feel like it sometimes though), it is important to realize that whatever we do is on a small budget, and will likely suffer some disaster eventually. The amateur network should be designed with this in mind, and should be resiliant enough to cope with parts of it going down from time to time. Whenever possible, the amateur network should recover from difficulties without the users of the amateur network knowing something happened.

If a user of the amateur network knows what path through the amateur network is used to establish a network interconnection to the amateur he/she wishes to communicate, the amateur network should attempt the network interconnection in that manner. If, on the other hand, tha amateur doesn't know the path to the other amateur (or even the destination transport entity where the other station exists), the amateur network ideally should provide some type of directory to aid in establishing the network interconnection. Obviously, this directory is a frill that won't be around for a while, but some method of using it should be provided.

Sometimes it may be advantageous to provide some method of allowing the amateur network (or the other amateur's station) to directly read the status of, or control some parameters of an amateur's packet system. This may allow the amateur network to optimize level 2, 3, and 4 timers, control viewing of passwords, etc. This is sometimes referred to as an alternate control path to the amateur's packet system.

The amateur network should also allow some method of network management by requesting the status of the amateur ntework, along with controlling certain functions of the amateur network. This should be done in various levels of control, along with having geographical boundries. Traditionally, amateurs prefer to operate in a non-autocratic enviroment, so a single amateur network control group is probably beyond possibility. A hierarchical system of control would be called for, allowing some amateurs to manage their local intranet, while others would manage a larger part of the amateur network.

Cutting Up The Amateur Network Pie

This amateur network is not going to blossom overnite. It will probably take much longer to develop than the level 2 standard did. Part of this is due to the added complexity of having multiply-interconnected devices that are so interdependant on each other. In order to speed up amateur network development, along with conforming to the IS0 reference model, the amateur network should be broken up into several parts, each of which is responsible for a portion of amateur network operation.

Transport Layer Services and Responsibilities

The Transport Layer (OSI Level 4) provides a method of transferring data transparently through the amateur network between Session Layer entities such that the session-entities don't need to be concerned about assuring reliability or speed of data transfer through the amateur network.

The Transport Layer does this by using an end-to-end protocol between the Transport devices at each end of a network interconnection. This protocol is responsible for establishing a network interconnection between two amateurs; maintaining data integrity, proper data sequencing, end-to-end flow control, and end-to-end error recovery during data transfer between the amateurs; and the release of the network interconnection when it is no longer needed. It should be noted that some of these functions may be altered/removed if requested.

The Transport Layer is relieved of routing, relaying, and non-end-to-end flow control decisions by the network layer operating underneath it.

The complexity of the Transport Layer is very dependant on the type of network operating underneath it. Some network protocols require a large Transport protocol to correct for potential problems, while other network protocols require almost no transport protocol.

Network Layer Services and Responsibilities

The ISO defines two portions of the Network Layer. Subnetworks are of one or more intermediate systems which provide relaying of data through which end-systems may establish network-connections. A Network is considered the interconnection of these subnetworks to provide a communications path between Network end-points.

The Network Layer (Level 3) is responsible for establishing a data path between two Transport Layer entities wishing to communicate through the amateur network. The Network Layer should provide this service to the transport layer in such a way as to make invisible how the network routed the data. This includes how many hops or relays it took, how many subnetworks it went through, and how many data links were used. As such, the service provided at each end of a network-connection should be the same, even if dissimiliar subnetworks are used somewhere between the two end-points.

The quality of service provided is negotiated between the transport-entities and the network-entities at the time of network-connection establishment. If a quality of service is agreed to, that quality of service shall remain in effect throughout the lifetime of a network-connection.

The Network Layer provides the following functions:

a) routing and relaying;
b) network-connections;
c) network-connection multiplexing
d) segmenting and blocking;
e) error detection and recovery;
f) sequencing;
g) local flow control;
h) expedited data transfer;
i) service selection; and
j) Network Layer management.

The Network Layer data is transferred between individual network-entities through the use of Level 2 connections. In the amateur network, this usually means AX.25 HDLC connections between network nodes or entities. Level 2 AX.25 is responsible for providing reliable node-to-node data paths between the network nodes.

An important point is that the quality of service provided by the overall amateur network is only as good as the weakest portion of the path through the network.

Proposed Datagram Network Standard

The datagram network crowd is proposing the use of the DARPA TCP/IP or UDP/IP standards in building the amateur network. The Internet (IP) protocol would be used at the network layer, and either the User Datagram Protocol (UDP) for use in unsophisticated transport enviroments, or the Transmission Control Protocol (TCP) for more reliable transport service.

Proposed Virtual-Circuit Standard

Most of the work being done in the virtual-circuit area is being based on CCITT standards. One recommendation being proposed is as follows:

Use CCITT X.25 Level 3 protocols for the connections between amateur network users and the amateur network entry point.

Use the CCITT X.75 Level 3 protocol for the connections between devices within the amateur network.

Use the CCITT X.224 Level4 protocol for the Transport connections (if necessary) between the two end-points of the amateur network.

Head-To-Head Comarisons Of Virtual Circuits And Datagram Type Network Operation

As will soon become apparent, both the virtual-circuit and datagram network concepts have good points and bad points. It will be up to the amateur community and network designers to decide how these will be used in differing operating enviroments.

Both of the amateur network concepts will create a logical network connection between the two end-points of the amateur network wishing to communicate. Both will have the capability of providing either reliable data transfer, or reduced reliability in favor of increased speed of data transfer.

Design Philosophy

Even though both network designs provide the end users the same service (potentially error-free data transmission from source to destination), the way the two systems accomplish this goal is quite different.

The datagram type network design works much like the way mail is delivered by the post office. Each letter (packet) has all the information necessary for that letter to be delivered independently of any other letter before or after it. Each datagram packet has both the source and destination addresses in a header prepended to the user data, along with some control information. This packet is then shot out into the air independently of how other packets for the same source were sent. It is up to the Transport Layer to make sure the packets do get from source to the destination in the proper sequence and without corruption. This means that in a datagram network, the Transport Layer is relied on heavily to correct for Network Layer problems.

The virtual-circuit type network operates more like the telephone system. When a telephone user wishes to talk to another telephone user, the first user establishes what looks like a direct wire circuit between the two users by dialing the destination users number. Once the call is established, every word (packet) flows from the source to the destination over the same circuit. Since the same circuits are used throughout the connection, it is not necessary to have an overseeing device make sure the wires don't move or change during the connection (yes, I realize there is multiplexing and line switching going on these days, but lets not confuse the issue). When the users are done, one hangs up, and that triggers the tearing down of the circuit, making the wire connections available to others.

It is now time to discuss some of the trade-offs between the two types of systems.

Packet Header Overhead

There is a large discrepency in the amount of header type overhead that the two network designs require. This may or may mot be important, but should be considered.

In the datagram network, a minimum of 20 bytes of overhead is required by the Internet Protocol, with an additional amount required if options are to be selected. The Transport Layer protocol (TCP) requires an additional 20 bytes minimum, again more is required if options are selected. Keep im mind that this 40 bytes minimum is required in EVERY SINGLE data packet sent.

The virtual circuit network proposed relies on the fact that all the addressing information is loaded up in the connection establishment process. This can be up to 256 bytes of data in the first connection request packet (assuming the Transport Layer connection request is in the fast-select portion of the network connection request). Once the connection is made, and as long as no major errors occur, the overhead drops drastically to three bytes for the Network Layer header and three to nine bytes of Transport Layer header overhead per data packet.

It looks like the virtual circuit network design wins this one hands down.

Packet Resequencing

In datagram type networks, it is possible for packets sent after others to arrive at the destination before the earlier sent ones. This is similar to when two people correspond every day through the mail, sometimes a letter sent after another arrives before the earlier sent one. There MUST be some method of making sure that the out-of-sequence packets are re-sequenced before they can be delivered. While I have heard and read that some people consider this a "trivial" task, it does take up buffer space and processor time at the destination end-point.

Since virtual-connection networks always use the same path for every packet (unless there has been a malfunction), the chances of this problem occuring are virtually eliminated, reducing processor and memory requirements.

Once again, the virtual connection protocol seems to have the advantage.

Routing Selection

If the route through the amateur network is static (not altering for any reason other than network device failure), it can be argued that both types of network designs work equally well. The selection of routes for packets is in itself another argument for another time. It can also be argued that in a fully static network, the virtual connection may have a slight advantage, since the address overhead is not required if no decisions are to be made based on these addresses.

If dynamic routing is allowed (where changes in the route of packets from source to destination can occur for a variety of reasons), the datagram type network has a distinct advantage. Since each datagram contains both address, routing decisions can easily be made, in worst case on a packet-by-packet basis. Since the virtual connection reduces its overhead by sending the addresses only during the connection establishment process and uses "logical channel numbers" from then on, it cannot easily alter the path of packets. Keep in mind that dynamic routing may add more problems than it corrects. Network oscillation, delays due to routing decision time, and sequence destruction are but a few of the problems associated with dynamic routing.

Congestion Bypassing

Avoiding routes that have become congested is only viable when some form of dynamic packet routing is employed. Since virtual connections do not lend themselves to dynamic routing of any kind, the capability of bypassing areas of congestion is a definate advantage of the datagram form of network. The only method of reducing congestion problems in virtual connection networks is to provide some sort of look-ahead routine to make sure that congestion is cut-off before it becomes a problem. Admitedly, this is a poor form of dealing with this situation. The datagram becomes the big winner here.

Tolerance to Switch Failure

There are two issues to be concerned with in talking about packet switch failures. The first is what happens to the rest of the network when a switch fails, and the other issue is how does the switch itself recover from a failure (even a temporary one such as a power glitch). It appears that the datagram network is more resiliant in both these issues. If a packet switch fails in a virtual connection network, all connections through that switch must be torn down and re-established using another path (if available). The datagram network may have to do a similar process if it is totally static routed, but if some form of dynamic routing is used, recovery is made much easier by just re-routing the data around the failed switch.

The other issue is that of switch recovery. When a packet switch has recovered from a failure in datagram network, it just has to rebuild its routing table and inform the network it is back in operation. The virtual connection switch must do this plus re-initialize all the connections passing through it. An additional problem is that some virtual connections may not realize that the switch has failed, causing additional hardship for the switch.

It appears that the datagram network is ahead on this one also. Measures such as battery backup and uninterruptable supplies can help to reduce this, but again this is a kludge.

Reliability/Speed Tradeoffs

Much has been made of this by the datagram group. It appears that even though both networks can be made to allow for reduced reliability in order to improve speed when the reduced reliability isn't a concern (such as packetized voice), the datagram network won't try to force the reliability issue like the virtual connection network would. It is up to the reader to decide if this is a real or imaginary advantage. It appears to be much easier to make a solid pipe (virtual connection) leaky by poking holes into it than to try to plug up the holes of a leaky pipe (datagram). At this point in time, I think this is almost a non-issue.

Roving Station Situation

It isn't much of a problem at the moment, but some thought should be given to the concept of a mobile packet station, either in an auto or an airplane for example. First thoughts seem to indicate that datagrams have an advantage in this situation. This is NOT so. Since both network designs rely on providing a logical connection through the amateur network from a source end-point to a destination end-point, if one of these end-points was to change, both types of networks would have to re-establish the connection to the new end-point. It may be argued that datagrams may be easier to do this, since a whole connection doesn't have to be torn down and a new one errected. Since the Transport Layer devices must be changed anyway, the form of network re-establishment is not a major issue. Both forms of networks could employ similar methods of causing this reconnection to happen.

Alternate Data Path

Sometimes it is advantageous for either the network or the remote end user might want to control some parameter(s) of the user's terminal or computer. The CCITT has provided for this by allowing a method of establishing an alternate path (kind of an in-band method of out-of-band signalling). This mechanism involves the use of the Qualifier, or Q-bit. The Q-bit is frequently used to provide the capability to a host to control a user's PAD parameters (such as to turn off echo when entering passwords). As far as I know, there is no easy form to do this in the datagram network, unless options are defined to do this.

Local Subnetwork Use

One of the clear advantages of the virtual connection networks is that it does follow the ISO reference model as far as subnetworks vs networks. The datagram network is good for what it is intended, an INTERNET protocol. Even the name implies that it hooks up networks and subnetworks to each other. IT IS NOT MEANT TO BE A SUBNETWORK PROTOCOL. What are we supposed to use within local subnetworks in the datagram network design??? TCP/IP works to interconnect subentworks, not act as the subnetwork protocol itself. Are we supposed to use just link layer protocols when communicating locally. THIS IS TOTTALY WRONG! I cannot emphasize this enough. TCP/IP on a subnetwork level makes absolutely no sense. It takes up too much overhead, processing speed, channel overhead, and memory requirements. Much grumbling was heard at first about the overhead of the address field of level 2 AX.25. Imagine if every packet must have an additional 40+ bytes of overhead to accomplish the same task. Some form of subnetwork protocol should be implemented, but TCP/IP is not it. Link connections such as what we use today also are a mistake.

A layered approach such as the virtual connection network design makes more sense. For the local subnetwork connections X.25 seems to fit real nice. It is a small robust protocol whose major defects don't affect performance at a local level. Since it is connection oriented similar to the presently implemented level 2 AX.25 protocol, plenty of the work has already been done.

The internetwork protocol of a virtual connection network would most likely be based on X.75, which is a modified version of X.25. Some additional work would be needed to make a complete network spec, but this would be fairly simple to accomplish. Since X.75 is also virtual connection, and it is a version of X.25, the two can be mapped together quite nicely.

The Transport Layer (if even required) is based on the CCITTX.224 standard (see another paper in these proceedings). X.224 is a multi-class protocol, and even the most basic class (class 0) handles the major hole in X.25/X.75 network operation (that of re-establishing a connection after a switch failure). A more advanced class also provides for a checksum to eliminate the possibility of a switch with a memory malfunction corrupting an otherwise accurately transferred packet.

Each of these protocols loads the major overhead burden into the connection establishment process, and then operates on a very small header budget. One more point, either the X.25 or the X.75 protocols would be used not both. This is to say that if a packet is originated in an X.25 subnetwork and then transferred across the amateur network using X.75, both headers are not required, just the one being used at that particular network connection.

Flow Controls

Flow control throughout the network is handled differently by the two network designs. The datagram network normally does not provide any flow control at the Network Layer. Instead it relies on the Transport Layer for end-to-end flow control, and the Link Layer for everything else. Unfortunalely, if the Link Layer is relied on, when the Link is flow controlled, not just the one network connection flow is stopped, but ALL LEVEL 2 data for ALL level 2 connections are stopped. Sometimes this is alright, but at other times this can be a big problem. There is no way around this problem.

In a virtual connection network, each individual network connection can be flow controlled independently of any other connection, independent of Level 2, independent of the Transport Layer. Some argue that this multiplicity of possible controlling devices adds unnecessary processing overhead and can lead to buffer problems stacking up and rippling through the network. I would point out that this most likely wouldn't happen, since there a finite number of packets allowable in a network (in either network design), due to Transport Layer sequence numbering constraints, in addition to Level 2 sequence numbering constraints.

Circular File Philosophy

One of the comments I hear from time to time is that a datagram network is easier to implement, because of the capability of just tossing out a packet if it cannot be handled for any reason, and wait for a better time, or wait to see if the packet shows up again. I don't feel that the circular file is the place for my packets (some may disagree). I would prefer the situation that if a packet shows up, the network tries its best to get that packet through, and only if there is no other recourse (such as buffer limitations suddenly showing up) should the packet be thrown out or ignored. The datagram approach seems to rely on this "tossing the offending packet" instead of trying to correct the situation that caused the offending packet in the first place. I repeat, my packets belong in a better place than the trash heap.

Hardware/Software Considerations

An important consideration is what kind of hardware and software will be needed to run the two protocols. The biggest single requirement in both types of networks is going to be the requirement for lots and lots of RAM for buffers. The datagram type networks may need more buffers to be available at the end-points, while the need for more buffers in the virtual connection network may in the packet switches. It really depends on how the software is written as to how much buffering is required.

Another hardware/software consideration is that of processing requirements. This can be broken down into the individual devices that make up the network. The majority of the devices in the network will most likely be the packet switch. The datagram people claim that a datagram switch is easy to implement. Depending on the type of routing used, this may or may not be the case. If some form of dynamic routing is implemented, the packet switch suddenly becomes a much larger device requiring a lot more processor power to figure out the route the packet should take to reach its destination. Dynamic routing of some sort will probably be implemented in the datagram type network, since most of the advantages of the datagram network can only be taken advantage of in a dynamic routing enviroment.

A similar form of trade-off can be made in the packet switches of a virtual connection network, in a slightly different form. The first form is similar to the datagram approach. Full virtual connections are not maintained between every packet switch, but rather cross-connection tables are maintained at each switch (similar to the patch panel of an old phone exchange). This would allow very simple software to be implemented at the switches at first. The trade-off is that flow control can only be implemented at the Transport Layer or Link Layer (like the datagram network). If each packet switch implements a full X.75 network connection to each neighbor switch, processing overhead is increased, but the overall network becomes inherently more reliable.

The other device that must be considered is that of the network end-points. Here there is no question. Because of the need for a sophisticated Transport Layer protocol over a datagram Network Layer, the datagram network will require a substantially larger device with much more processing overhead. Distributed processing (one micro for each layer) may be an absolute requirement for datagrams, while an option for virtual connections.

An Ounce of Prevention...

Most amateur network users will always require that the network transfer data RELIABLY. The two forms of network designs place this responsibility in different places within the network. The datagram loads ALL this responsibility at the end-points of the network in the Transport Layer. The datagram Network Layer takes no responsibility whatsoever for maintaining data integrity.

The virtual connection network places this responsibility in small portions throughout the entire network, with the last margain of safety at the end-points in the Transport Layer. This distributed-responsibility scheme adds overhead throughout the network, but allows problems to be corrected along the way, rather than having everything look fine until it reaches the end-point, and only then finding out an error occured early in the network.

"What The Big Boys Use"

An issue that is sometimes raised is that of who is using what form of network. The research community seems to have fully adopted the TCP/IP datagram network concept, as provided by ARPANET. This is fully understandable, since they can quite often easily obtain the processing power necessary to implement TCP/IP. Also, since most of the research centers these days interract with the defense department who owns the ARPA network, there is some political pressure to go that route.

In the real world, the bottom line is the buck. The networks that are there not for research, but rather to provide the service of a data network (such as GTE Telenet) must look at how to provide a data network in the most cost-effective form, otherwise the competition will take their customers. It is interesting to note that the commercial networks use virtual connection protocols for their operation. In fact, Telenet was originally a datagram type network, but spent several million dollars to convert to a virtual connection network because they found out that the datagram network just wasn't cost effective. Some datagram people comment that the commercial data networks use virtual connection protocols because this shifts political network boundries out of the hands of the user and into the hands of the network. This seems to be based on articles in some of the computer journals around 1976. A lot has happened since then, including Telenet switching from datagrams to virtual connections. It is interesting to note just how many assumptions were made back then that are totally wrong today. Once again, the commercial networks use one yardstick for measuring their network, the biggest bang for the buck. No politics, because there is no room for politics. If they relied on political considerations, one of their competitors might not, and there goes the customers. It seems that the only people that can use political games are those that don't necessarily look at the bottom line, but can instead justify some additional costs for the sake of research. Does someone come to mind?

Conclusion

The major question I have for those implementing TCP/IP is what they are going to implement for the subnetwork (or intranet, or local network, or metropolitan network)? What are we supposed to use when packeting on a local basis to other hams in our area? Since a lot of our communications will always be within a metropolitan area, this issue MUST be addressed. Are we all supposed to support TCP/IP or UDP/IP? That won't work. You just can't shoe-horn all that on a TAPR board. Are we supposed to just continue to use Link Layer procedures when packeting locally? That isn't the right answer either. I believe that an AX.25 Level 3 machine could be shoe-horned into a TAPR board if one really tried.

As it appears from the above, I am going to continue the development of virtual connection network protocols. I do believe there will be a use for both network designs, and the best way to chose the correct one for the majority of the amateur network is to have both operate in a head-to-head competition. I do feel strongly that there is going to be a local subnetwork (intranetwork) protocol developed for local metropolitan users. This protocol does not have to be the same as the internetworking protocol used. In fact, I think there will most likely be some gateway operation to interconnect virtual connection networks with datagram networks. One point about this, I have heard some amateurs argue that if a part of a network is datagram then ALL of the network MUST be datagram (or vice versa>. This is not true!! All that must be done is that the gateway between the two types of networks must perform protocol conversions at both levels three and four. Since the two levels are so intertwined (especially with datagrams) this task must be accomplished. If it is done correctly, it should appear as if nothing out of the ordinary is happening.

My last comment is that given a piece of information that can be transferred using either method, which would you prefer and trust, the post office or the telephone system?