Applenet Hardware:
Background and Current Status

CONFIDENTIAL

Authors:
Bob Paratore/Jan Sheehan

Preface

This document presents the current status of and thoughts about the Applenet hardware. Chapter 1 explains local area computer networks, and Chapter 2 describes the current hardware design and implementation of Apple's version of a local network, Applenet. Chapter 3 ends the discussion of current thoughts and designs with an explanation of the personality module, a low-cost way of emulating certain devices. The last chapter of this document, Chapter 4, discusses what developments and user needs future versions of Applenet must consider.

1. Network Concepts

Computer networks connect computers and related resources together so that they can communicate with one another. Each separate connection to the communications medium, called a node, attaches one or more computers and related resources to the network. Long-haul computer networks, such as ARPANET, connect computers that are more than a few miles apart; local area computer networks, such as Applenet, connect computers that are only a few thousand feet apart.

Networks, whether local or long-haul, need to control how and when the connected computers communicate with one another. Networks control communication through choice of structure (topology) and through rules (protocols) implemented in the system's hardware and software. The structure of a network defines how nodes are physically connected to one another. The system protocols are divided into sets of rules for each level of the communication process, and define the structure of packets, how and when a node can transmit or receive a message, and how the network and user software forms and interprets packets.

Communication between nodes in a network is either circuit-switched or packet-switched. Circuit-switched networks dedicate a communication line for the duration of a communication. However, packet-switched networks do not dedicate a line. Instead, computers connected to packet-switched networks transmit packets, small envelopes of information. The advantage of packet switching is that many nodes can share the same communication line and transmit virtually simultaneously due to the short length and transmission time of the packet. Usually, local networks are packet-switched rather than circuit-switched.

1.1 Local Network Topologies

Some local networks have centralized control. In these networks, one node receives and routes all messages. Common topologies for such networks are the loop and the star (see Figure 1-1). In a star network, all computers connect directly to the controller and send/receive messages only through the controller. In a loop network, the connections of the nodes form a circle. A message travels around the loop to the controller and the controller routes the message around the loop to its intended destination. These two topologies have one major disadvantage: their dependency on the controller node. Because the ability of a star or loop network to function depends upon the status of the controller, expensive redundancy schemes must be included in the network design to ensure reliable network communication.

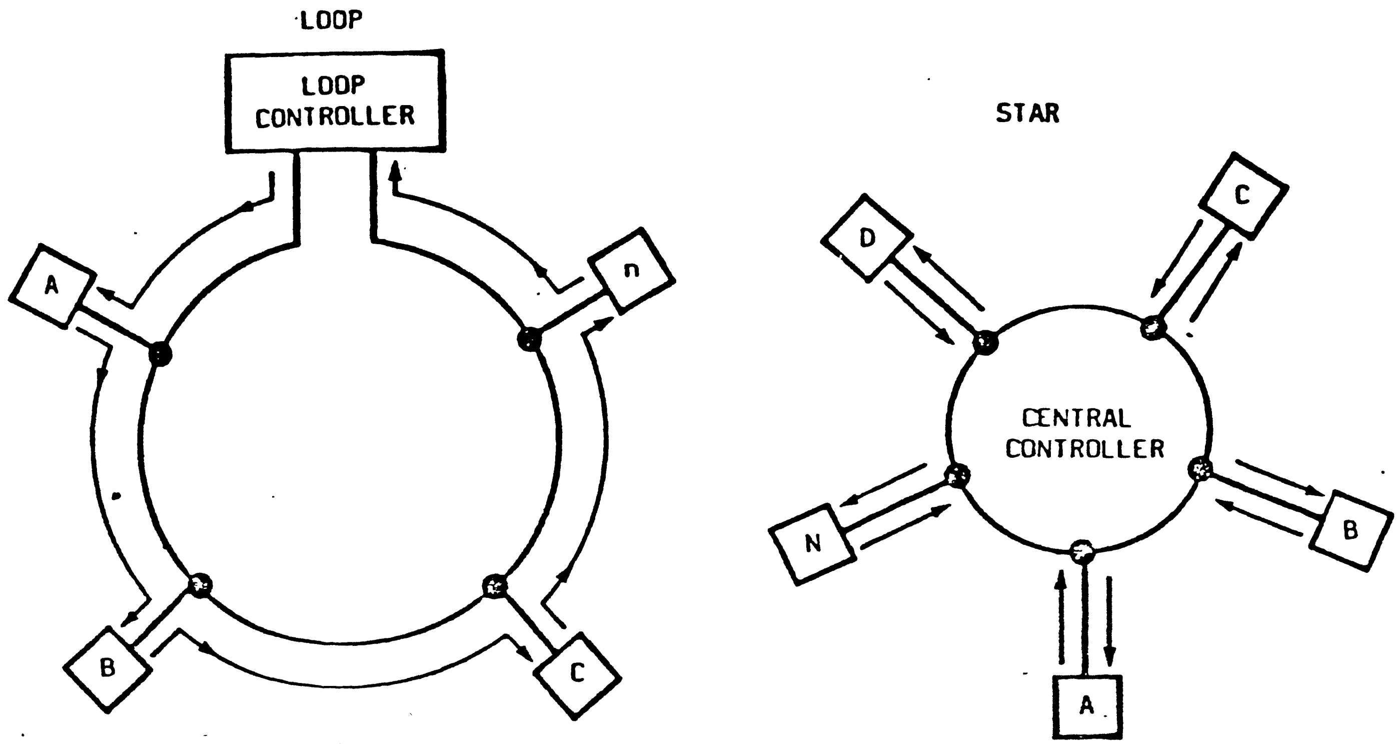

Figure 1-1. Star and Loop Topologies.

The ring and the bus are the most common architectures for networks that distribute routing control among the nodes (see Figur 1-2). As with the loop network, the nodes in a ring network are connected in a circle. However, each node in the ring is an active element of the network and passes a message on to the next node or accepts it if the destination address matches its own. Although the ring network distributes network reliability among all the nodes, the failure of one node still disrupts the whole network. Therefore, ring networks also require expensive redundancy schemes to ensure reliable communication. A common bus network consists of a passive communication bus and attached nodes that connect to the bus with a tap that pierces the shielding of the communications medium. Each node contends with all other nodes for use of the bus for transmission, and monitors the bus for messages with its address. Messages travel along the bus without interference or intermediate routing by other nodes. Therefore, the failure of any node connected to the bus does not affect network traffic because no node other than the sending and receiving nodes are involved.

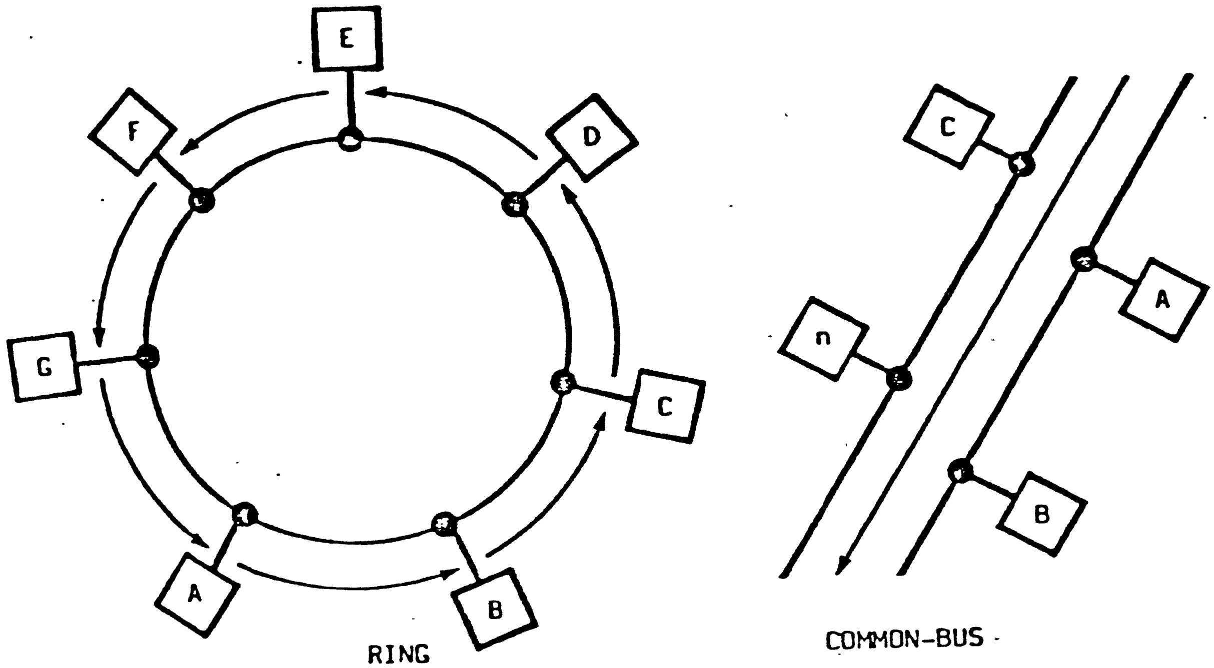

Figure 1-2. Ring and Bus Topologies.

The ring and the bus networks also differ in their low level protocols, their electrical level rules of transmission. A ring network controls access to the line with a token that passes from node to node. If the line is free when a node receives the token, it can attach a message to the token for transmission or pass the token to the next node. In a bus network, each node connected to the network monitors the line. If the node does not detect network activity, it can attempt transmission.

Due to the short time intervals involved, it is possible for two or more nodes to detect a free line and transmit simultaneously. If this occurs, the messages collide. In certain bus networks, the nodes can sense this situation immediately and stop transmitting. After detecting the collision, each node wishing to transmit generates a random wait time and repeats the line-test transmission process. Bus networks that can detect message collisions are called Carrier Sense Multiple Access networks with Collision Detection (CSMA/CD).

1.2 Connecting Networks

Although there are limits to the number of nodes connected to any one local network, local networks can be connected to other local networks of the same type by a device called a bridge. The bridge can only connect networks of the same type. For example, a bridge can connect an Applenet network to an Applenet network, but cannot connect an Applenet network to an Ethernet network. The bridge filters messages intended for a station connected to a different local network and passes them along the network connecting the bridges; it does not interfere with a message intended for a node on the same local network. Messages pass from bridge to bridge within the extended local network until they reach the bridge associated with the destination node.

Figure 1-3 below shows an extended network that consists of three local networks, Net A, Net B, and Net C, connected to one another by bridges L1, L2, and L3. The portion of the figure labeled 'traffic' illustrates the steps necessary for a message to travel between different pairs of nodes connected to the extended network.

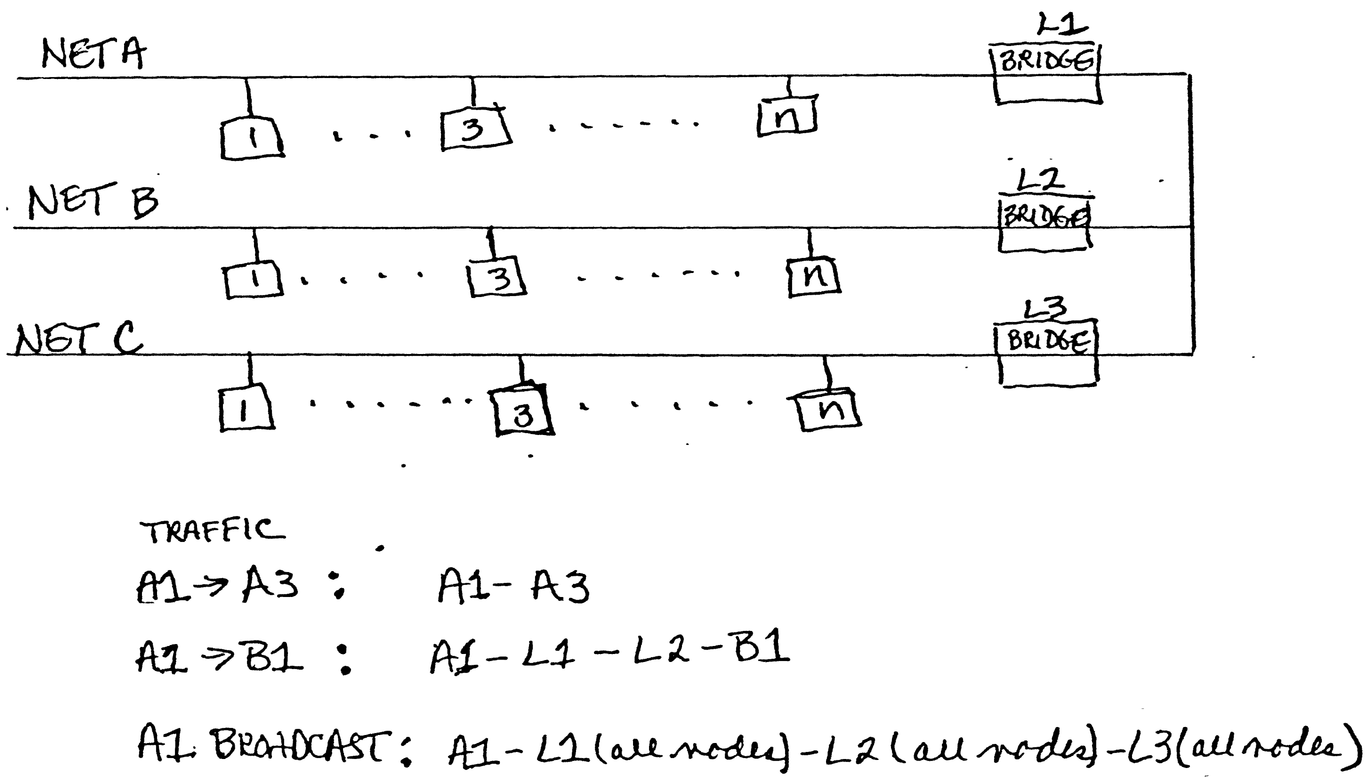

Figure 1-3. Bridge Connections.

Currently, due to a lack of network standards, many different types of local and long-haul networks exist. This creates problems when a computer connected to one network needs to communicate with a computer or use a resource connected to another type network. For example, an Applenet in one city may want to communicate with an Applenet in another city via the telephone network. For a node to communicate with a node on a different type of network, the host network must have a device called a gateway that can translate the packet's header and routing information, and the physical level protocols of the one network into those of the other and back again. In the example cited above, both Applenets would require a gateway. Figure 1-4 below shows two physically separated Applenets connected via modems that act as gateways to the network that connects the two locations.

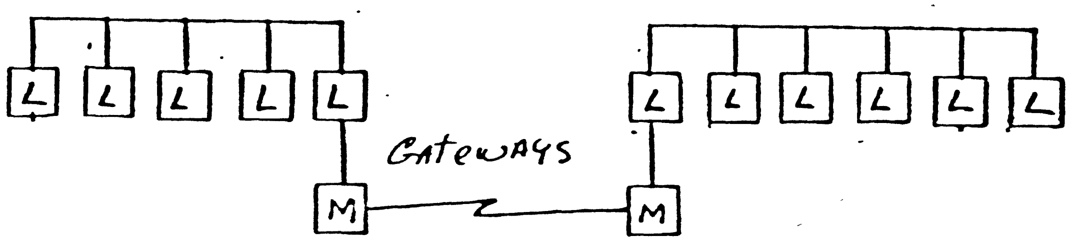

Figure 1-4. Gateway Connection.

1.3 Uses of Local Networks

Local networks, useful in many different environments, usually offer some combination of file transfer, resource and load sharing, control diagnostics, program download, and routing/transmission priorities. In the office environment, the target market for LISA and Applenet, these features allow optimal use of expensive equipment such as high-speed printers and mass storage media, and increase the ease of communication by supporting applications such as electronic mail.

Figure 1-5 shows two different Applenet configurations. The first presents a small installation that consists of six LISAs connected to the network. One of the LISAs is a printer server for the others and has an attached printer. The second configuration illustrates how bridges connect the nodes and resources of three Applenets that form a larger, extended network. The configurations are flexible so that a customer starting out with a small installation can easily expand the network later to suit future needs.

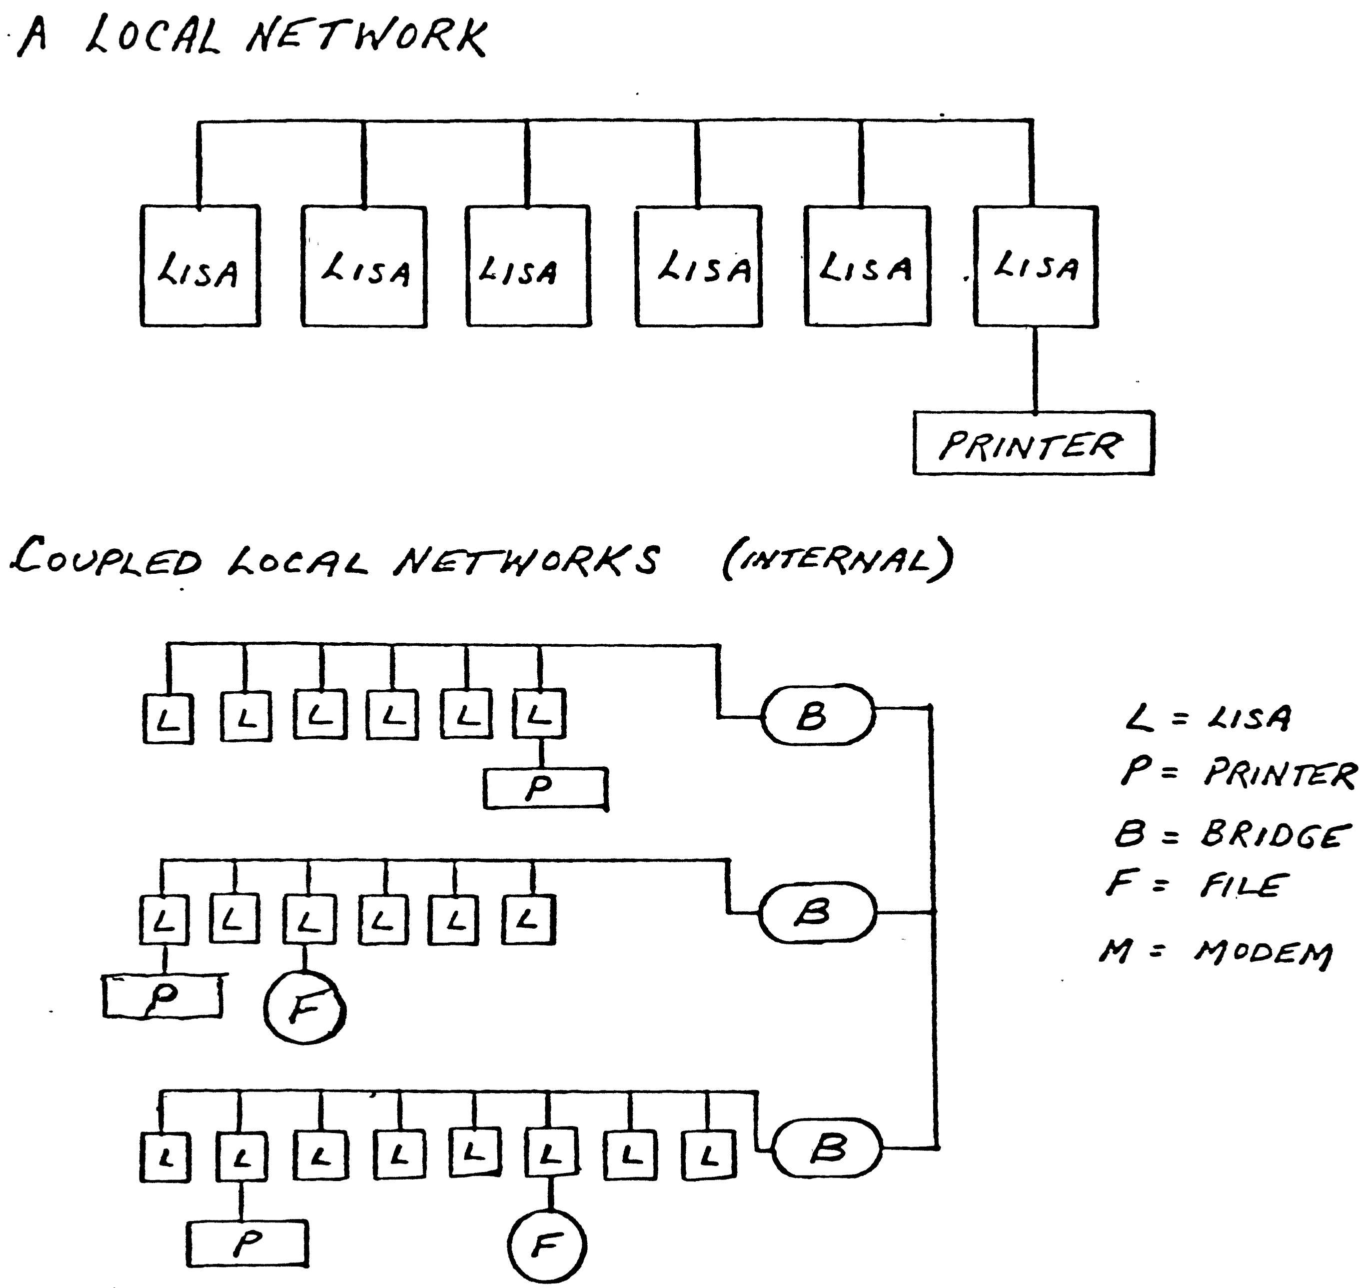

Figure 1-5. Network Configurations.

2. Applenet Hardware

Applenet is a baseband Carrier Sense Multiple Access network with Collision Detection. In other words, each node can detect a free bus and can recover from a collision if more than one node detects the free line and transmits at the same time. The network has no central controller; network access is contention-based and is governed by arbitration algorithms and collision avoidance schemes explained below that are unique to Applenet.

Network nodes attach to the network via drop cables that connect to the main bus through devices called cluster boxes that each contain approximately four drop lines. The main cable supports up to 128 nodes.

The Applenet hardware design is divided into the physical level and the link level and conforms to the following constraints:

* The cable must support a transfer of 1Mbps over a distance of one kilometer on the main line and 14 meters on a drop cable.
* The hardware must meet or exceed all FCC, EMI, and RFI requirements for Class A equipment.
* Installation must be simple and must not require skilled labor.
* The network requires an AC coupled or like line so that the failure of one node does not affect network performance.
* Adding or deleting a node on a drop line must not affect network performance.

The following two sections explain each of the two levels of the hardware design and their function within the system. The final section of this chapter describes the current implementation of the design.

2.1 The Physical Level

As mentioned above, message transmission on the Applenet conforms to bus contention arbitration algorithms and collision avoidance schemes unique to this network. When a node attempts message transmission, the associated hardware executes the following sequence of actions (see Figure 2-1):

1. The hardware listens for an interval of about 45us which represents line idle. The length of this interval, called the inter-record gap, depends on software overhead.
2. After detecting an idle line condition, the hardware sends a pulse 1us long to request the line.
3. The hardware then listens again for approximately 12us. The hardware then injects a 125ns delay to delay ensure that the requesting device hears a conflicting request if another node simultaneously asks for the line.
4. If the node detects a conflicting request, the link to physical controller randomly assigns the node to slot 1 or slot 2. Nodes assigned to slot 1 repeat steps 3 and 4 until no conflict occurs. Nodes assigned to slot 2 cannot request the line if any slot 1 node requests line. Nodes that cannot retransmit return to listening (receive) mode.

5. Once a node has control of the line, it sends an address and the packet header.
6. The transmitting node listens for a negative acknowledgement pulse, a node busy signal (NAK), from the destination node.
7. If the transmitting node does not receive a node busy signal, it sends the body of the packet. Otherwise, the transmitting node terminates the transmission.

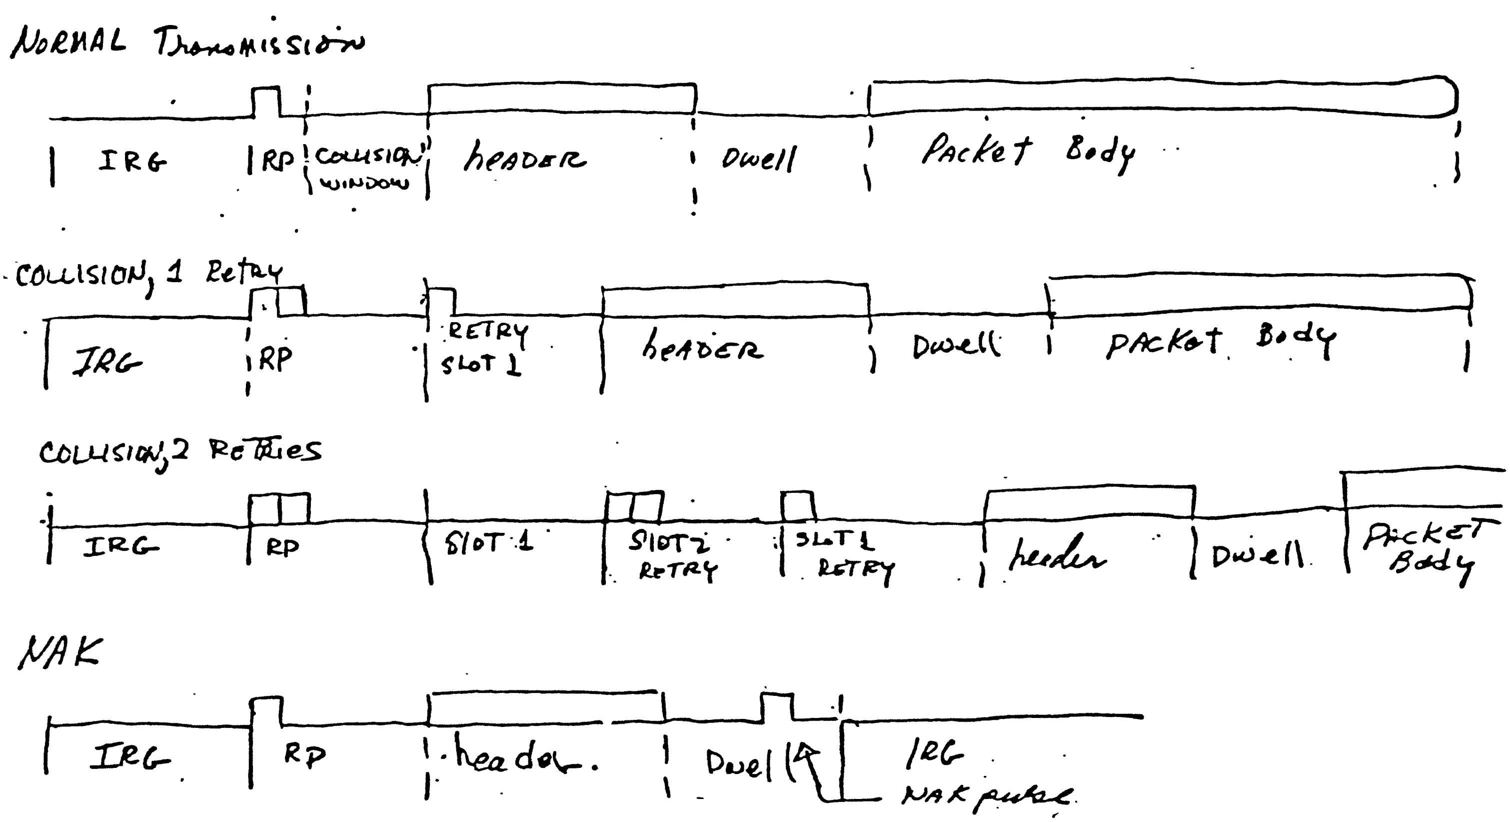

Figure 2-1. Transmission.

Applenet transmits and receives a Modified Harvard Frequency Modulation code which has no baseband DC shift and no error propagation. Figure 2-2 illustrates the differences between the Harvard Frequency Modulation code and the Modified Harvard Frequency Modulation code and shows the state definitions both use for encode and decode.

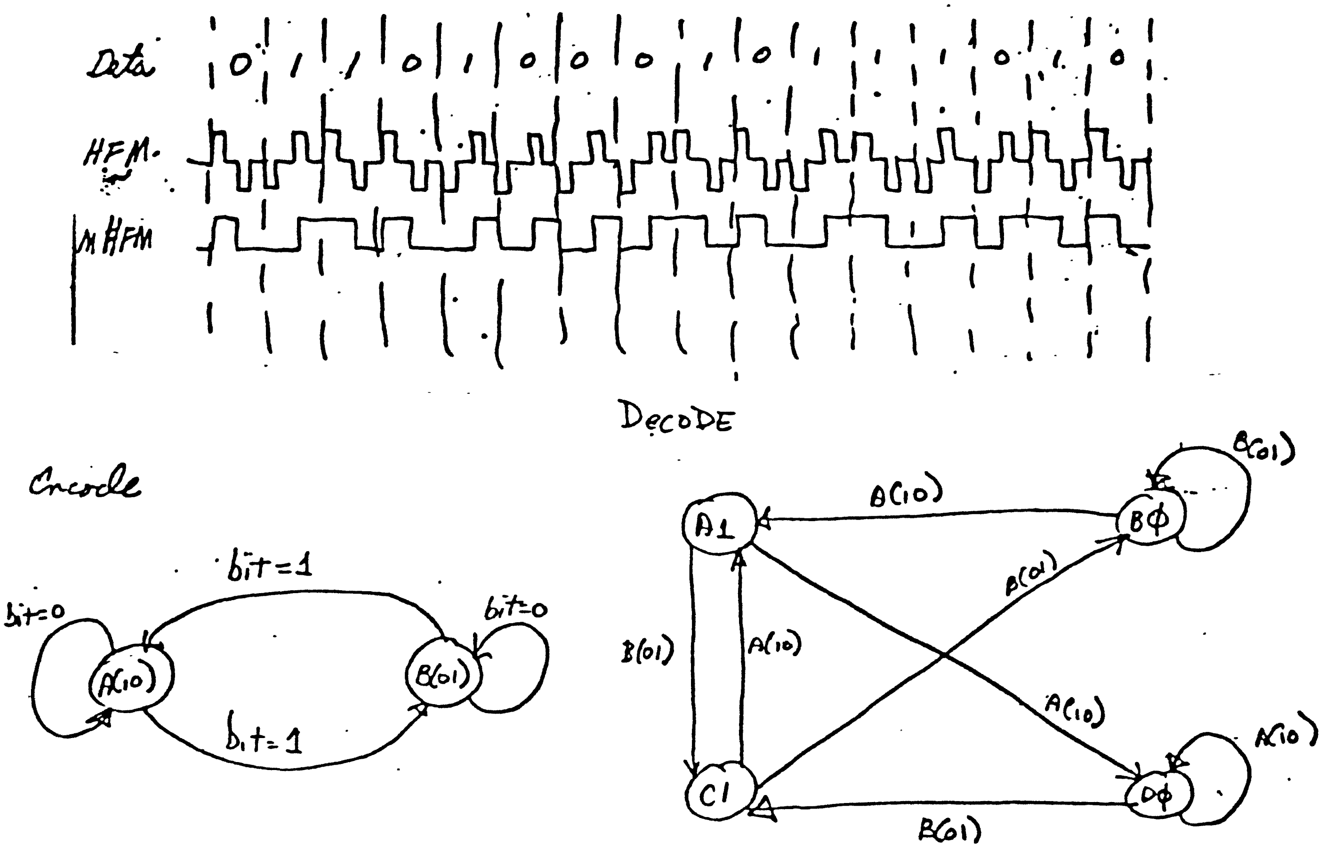

Figure 2-2. MHFM and HFM Transmission Codes.

The physical level hardware is also responsible for packet reception. The associated hardware executes the following sequence of actions whenever another node attempts packet transmission.

1. Listens for a reserve pulse.
2. Checks header that follows pulse, if any, for own address or broadcast.
3. Checks to see if a message receive buffer is currently available if the message is a broadcast message or is addressed to the node. If a buffer is available, the hardware copies the message. If no buffer is available and the message is not a broadcast message, the receiving node sends a negative acknowledgement pulse (NAK) over the network.

2.2 The Link Level

The link level is responsible for extraneous error detection and packet level flow control. It executes all link commands and interprets the high-level protocols coming from the host network-level software that controls network communications.

2.3 Logic Implementation

Five logic groupings make up the Applenet hardware logic: host interface, bus arbitration, link to physical controller, encode/decode, and analog (see Figure 2-3). This section explains each of these groupings.

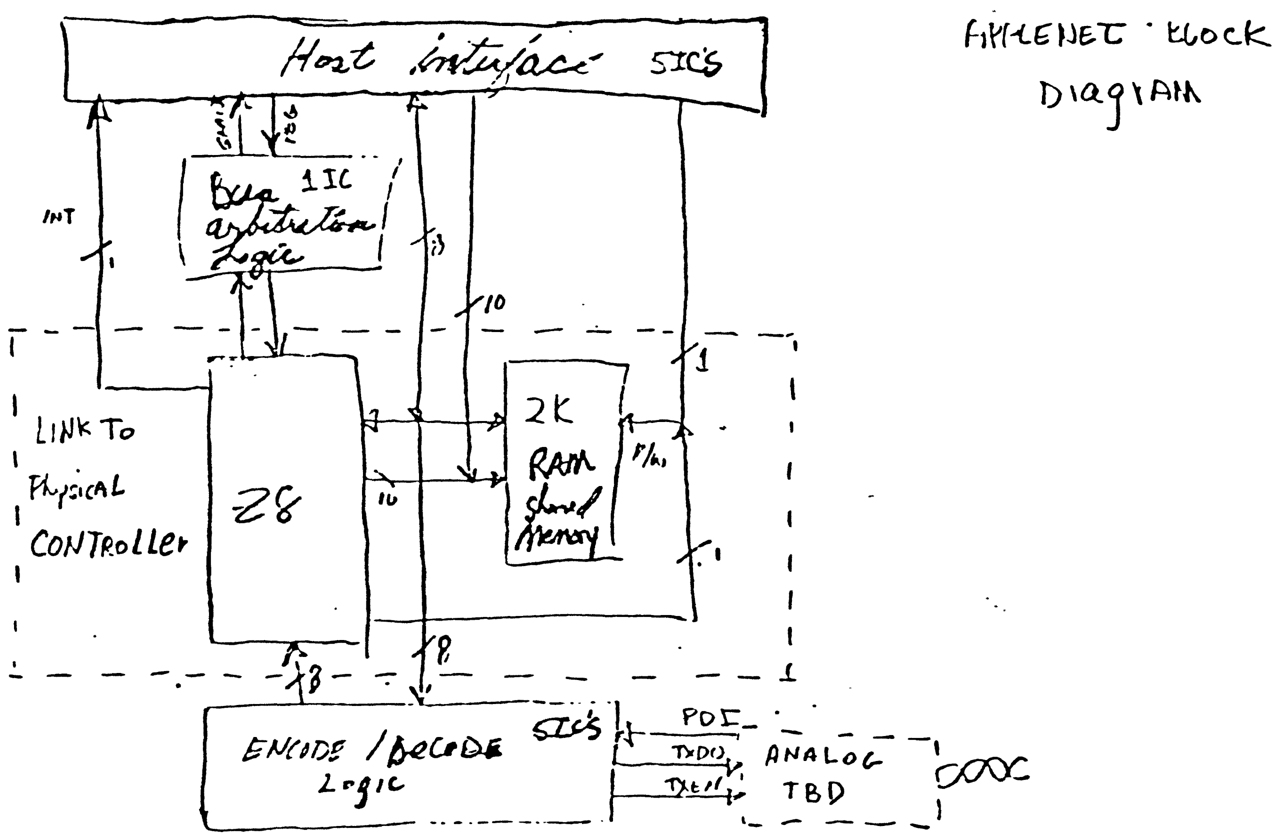

Figure 2-3. Applenet Logic.

The host interface is the only host dependant logic on the Applenet card. The host interface logic provides a 10-bit address bus, a bidirectional 8-bit data bus, and a Read/Write line. These lines, controlled by the bus arbitration logic, must support tristate logic.

As a result of a hardware/software tradeoff, the bus arbitration logic is extremely simple; it requires only one IC. The Applenet arbitration scheme is very different than others currently in use because the network is asynchronous and the link to physical controller logic must have access to the bus on demand. A hardware/software protocol implements the bus arbitration algorithm which consists of the following actions.

1. The host clears a latch in the arbitration logic and then reads the access bit from the data bus.
2. If the access bit is low, the Applenet memory is available.
3. If the memory is available, the host transmits a burst of 8 to 32 bytes of data into the net memory. The length of the burst depends on the speed of the host.
4. After transmitting the burst, the host checks the access bit again to see if the Applenet memory is still available.
5. If the memory is available, the host transmits another burst of data. If the memory is not available, the host begins the process again beginning with step 1.

Internal studies show that this bus arbitration algorithm does not produce significant bus bandwidth degradation. In network tests that used an Apple II as the host, the time required for a point to point transfer of a 272-byte packet sent in 8-byte bursts was 15ms, the Disk II transfer speed. Therefore, as expected, the host, not the arbitration algorithm, is the limiting factor in speed of data transfer.

The link to physical controller logic consists of a Z8 single-chip computer and a shared memory. The Z8 interprets single-task high level commands that the host sends to the network logic. It also queues incoming messages in buffers in the shared memory and interrupts the host when a complete packet is in memory.

The host responds to the interrupt, reads the buffers into the memory allocated to the destination process, and notifies the Z8 via semiphores that the buffers that contained the packet data are now empty. The Z8 is responsible for buffer management in the shared memory; it dynamically allocates and deallocates the buffers and informs the host via semiphores of the location and state of each buffer.

The encode/decode logic consists of five ICs that encode 8-bit binary data into a Modified Harvard Frequency Modulation code and decode the MHFM data back into 8-bit binary data. The design consists of a PAL and a latch that perform the state transactions, and a FIFO that provides temporary storage, serialization, and deserialization of data. The output of this logic drives yet-to-be-defined analog hardware.

Current thinking anticipates that the analog logic will consist of one, inter-design custom chip and some discrete components. Figure 2-4 below shows a generalized block diagram of an Applenet local loop. Although this seems to be a workable structure, there are many unanswered questions about the implementation of each component. The following paragraphs explain current thoughts about the implementation of the cable, the terminations, and the cluster box.

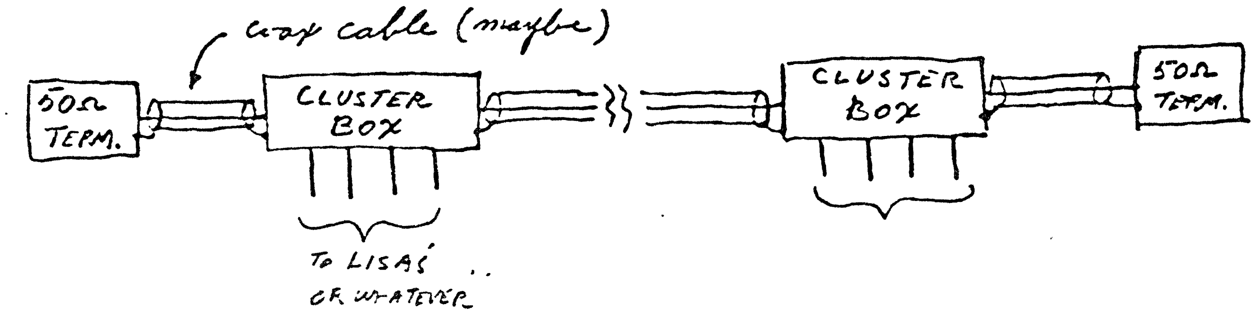

Figure 2-4. Applenet Local Loop.

Due to the practical considerations of cost and ease of installation, coaxial cable is favored over twinaxial cable. Although it may prove adequate, the coaxial cable does not have the RF radiation immunity of a shielded twisted pair cable. However, replacing the coaxial cable with twinaxial cable at a later date affects the design and schedule of both the terminations and the cluster boxes.

Whether the cable is coaxial or twinaxial, the shield of the cable should connect to ground at only one point. The best location for this connection is mid-span of the cable. Figure 2-5 illustrates the termination designs for both cable types.

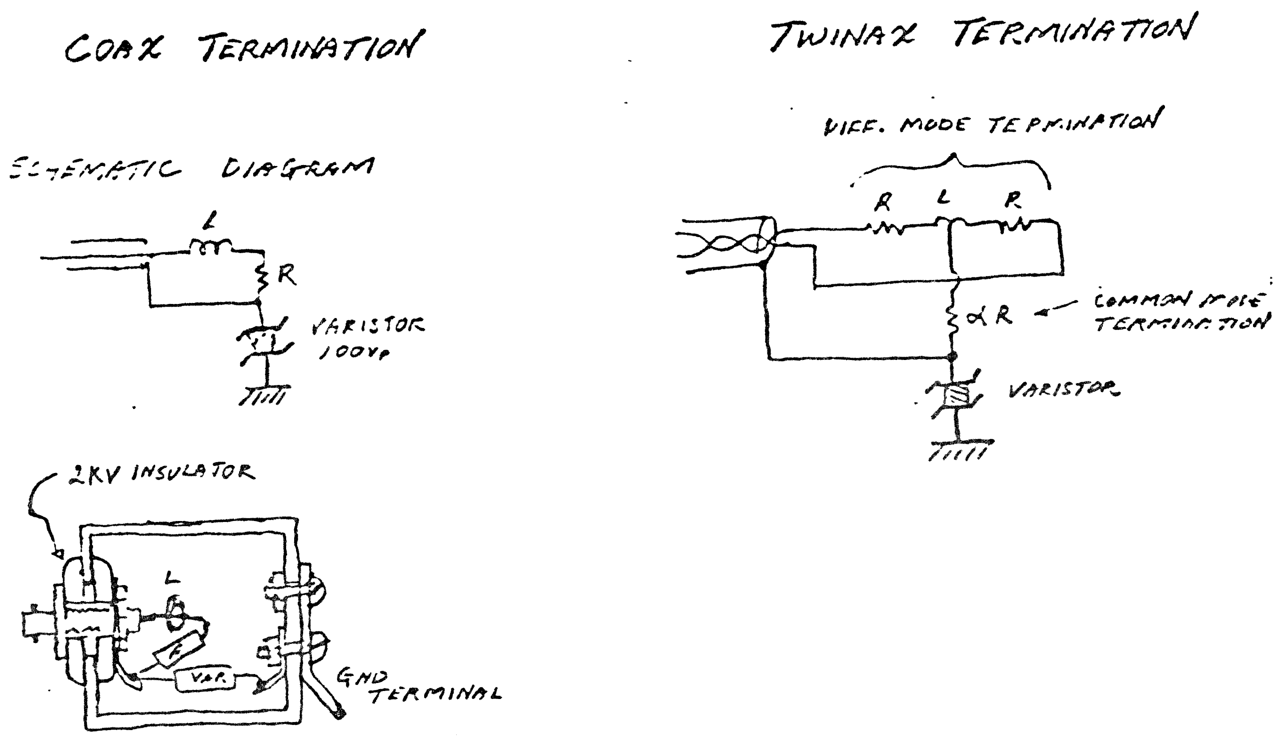

Figure 2-5. Applenet Terminations.

The cluster boxes connect the nodes via the drop cables to the main cable. The internal wiring of the boxes differs depending on the choice of loop cable. Figure 2-6 below shows a diagram of the drop cables within the box, the cluster box internals, and the details of the toroid.

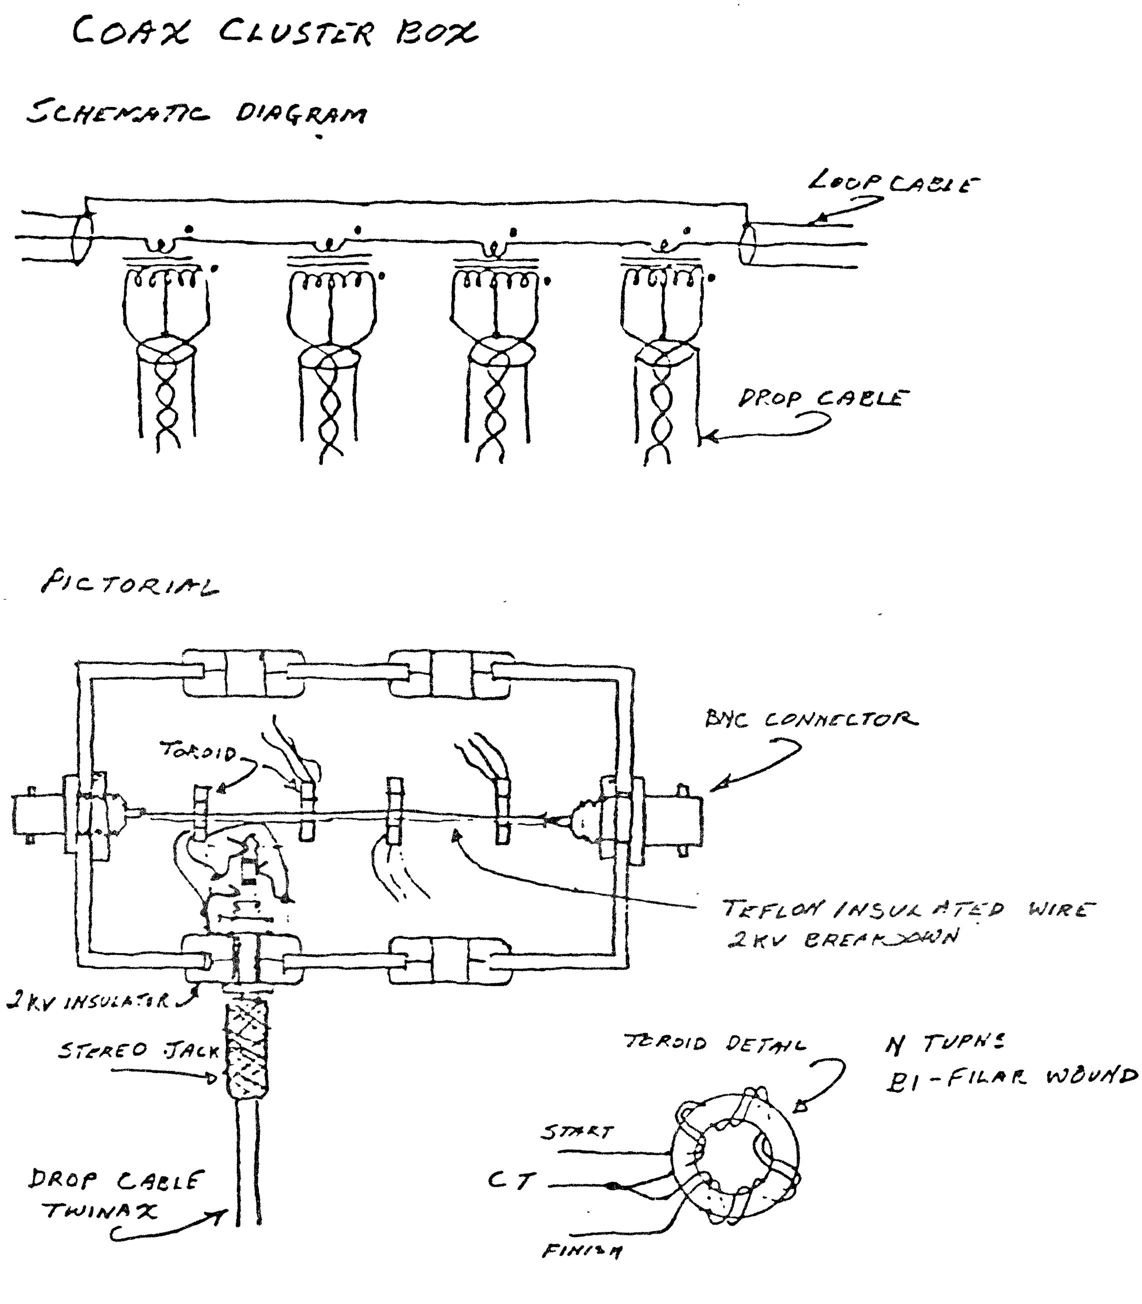

Figure 2-6. Applenet Node Connections.

3. The Personality Module

The necessity for bridges and gateways in the Applenet network inspired the concept of the personality module. A personality module is an inexpensive way to provide certain necessary network features. For example, a LISA equipped with two network boards can function as a bridge. However, this solution means the customer must buy an extra LISA and two extra network boards for every local network that connects to an extended network. A personality module, consisting of a power supply and the necessary logic cards housed in a black box, is a much cheaper solution.

The first module designed will probably emulate an RS232 interface which can link all Apple products to the network. An RS232 emulator can also interface a number of other devices to Applenet (see Figure 3-1).

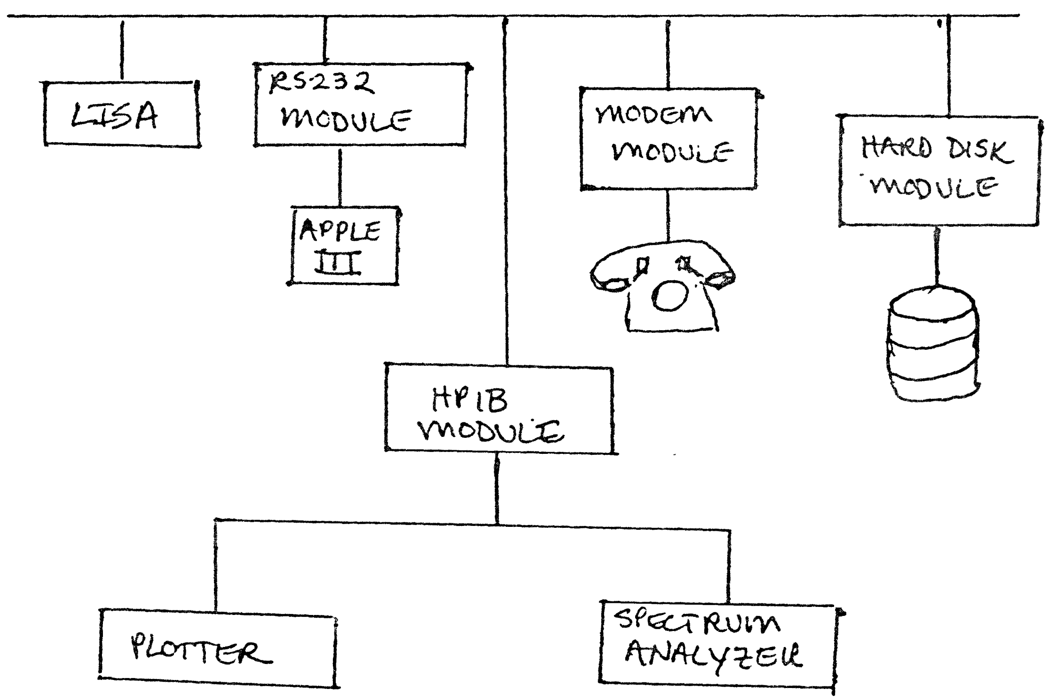

Figure 3-1. Uses of the Personality Module.

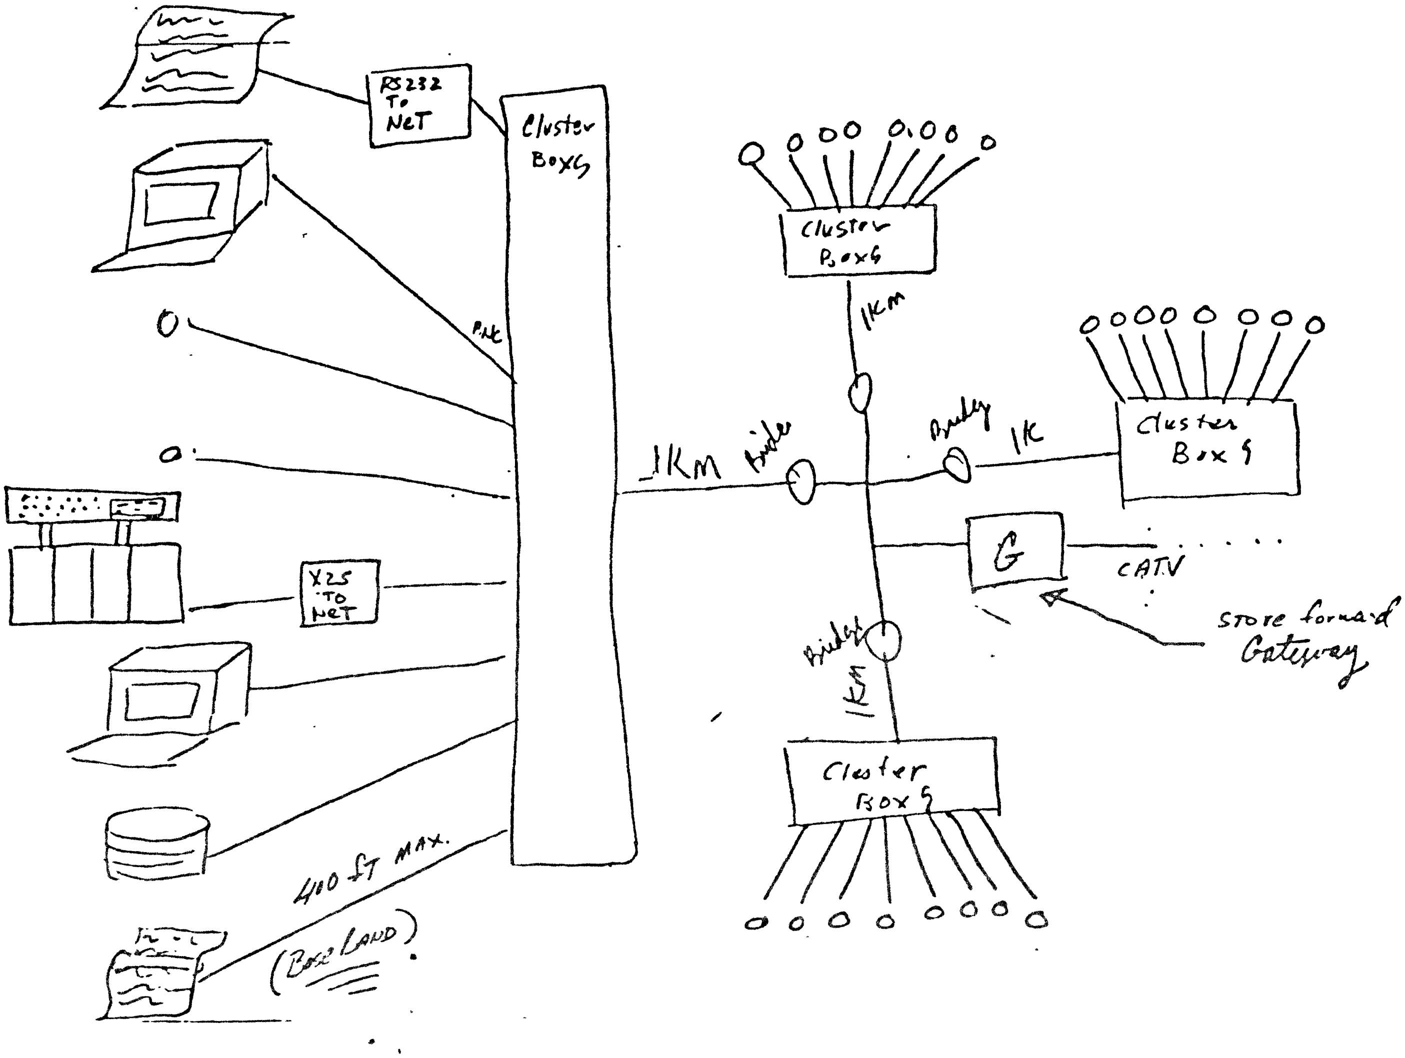

Figure 4-2. Baseband Network With Broadband Access.

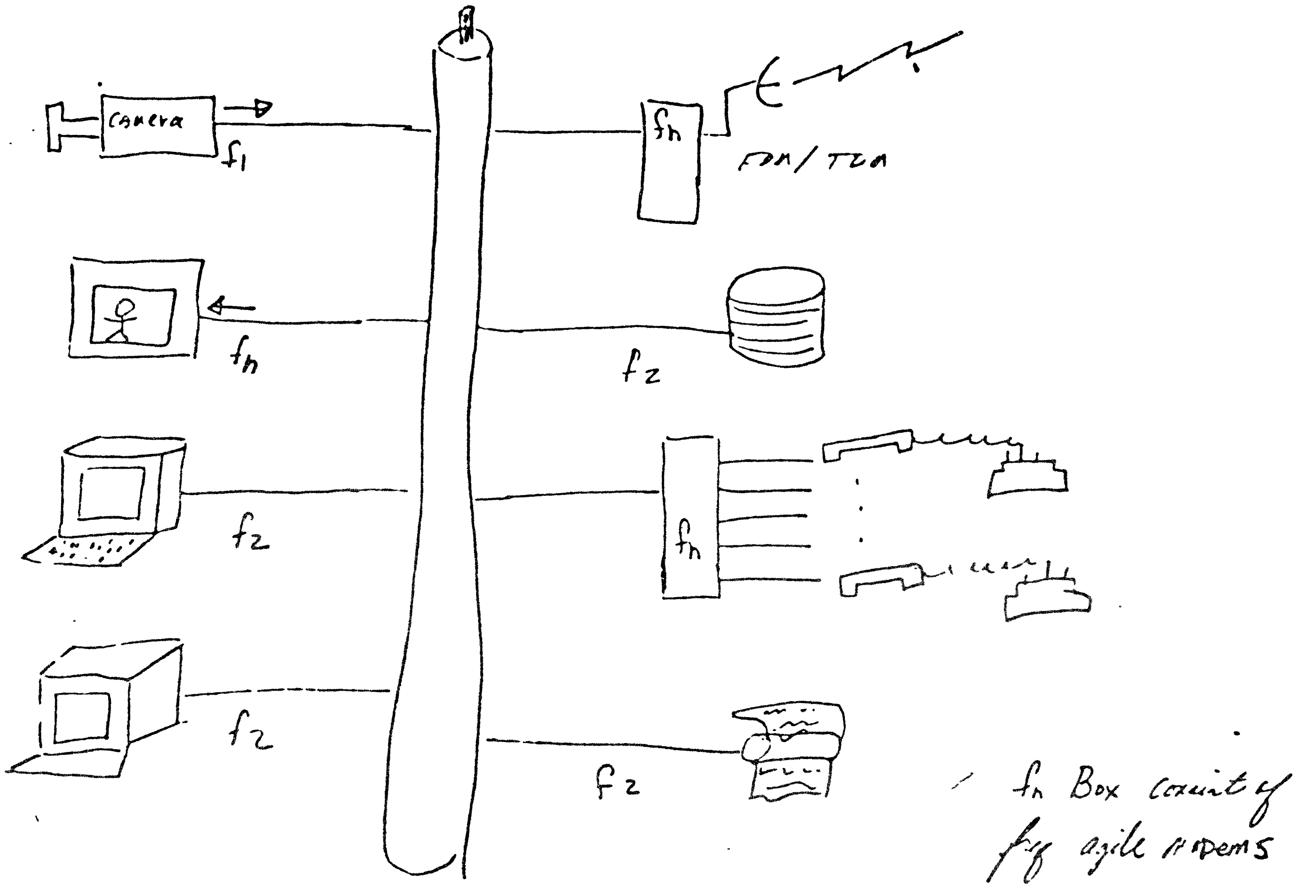

Figure 4-1. Broadband Network.

4. Future Thoughts on Applenet Hardware

Eventually, the market for LISA and Applenet, the office, will want all the following features.

* Database access
* Computer to computer communications
* Teleconferencing
* PABX
* Building security

Each of the above features requires a different bandwidth. Some, like teleconferencing and PABX, need a broadband rather than a baseband network. Currently, broadband networks are very expensive. However, the cost will come down over the next ten years. At that time, we should consider upgrading Applenet to a broadband network (see Figure 4-1). Until then, Applenet can use a gateway to a cable-television-based network to provide broadband services (see Figure 4-2).