M.I.T. Laboratory for Computer Science

Network Implementation Note
No. 4
February 5, 1979

A STAR-SHAPED RING NETWORK

by J. H. Saltzer and K. T. Pogran

Summary

The basic concept of a local data communication network consisting of a ring of small-delay repeaters can be realized with a variety of possibilities for the actual routing of the wire cables that link the repeaters. This note proposes that the cables interconnecting the repeaters be arranged in the shape of a star, with every cable passing through a central room known as a wire center. The advantages of this physical arrangement are: reduced risk that the ring is accidentally broken, rapid reconfiguration, simple trouble-shooting, and, because cable lengths can be normalized, use of a less sophisticated transmission subsystem. Disadvantages are: longer wire runs, and a need to be organized to avoid a "rats nest" at the wire center.

The basic ring

Proposers of ring networks usually draw pictures as in figure one, with the implication that the cables interconnecting individual repeaters follow any convenient, reasonably direct route from one repeater to the next. Installing a ring network with that approach could be expected to expose the following problems:

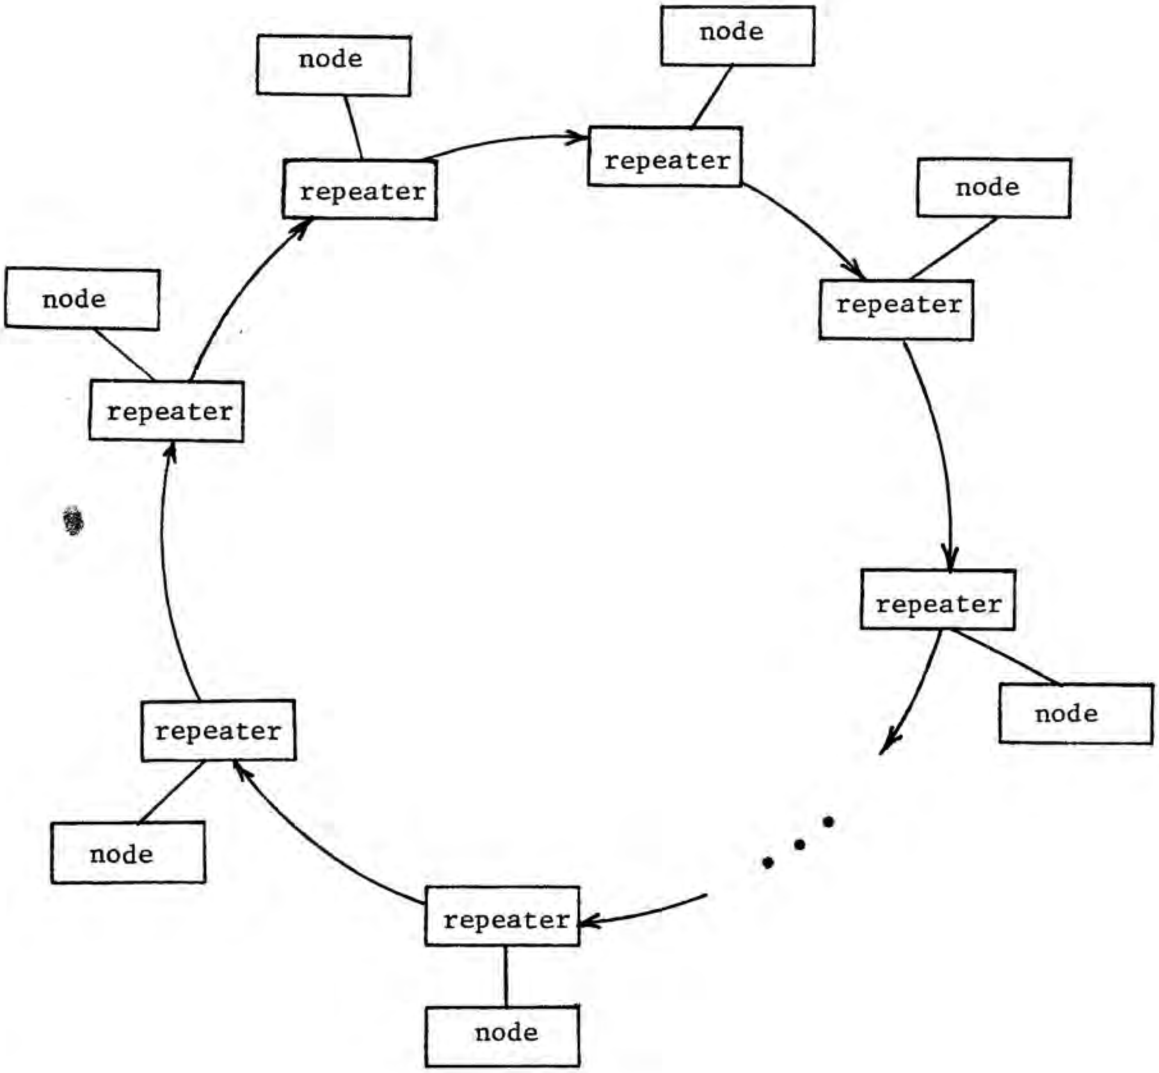

Figure one. Usual ring drawing.

1) Cable Vulnerability. The physical ring trails widely through the building, and is vulnerable at every point to accidents. If a link is accidentally severed or shorted, the entire ring is inoperable until the problem can be isolated and a new cable installed.

2) Repeater Failure. Since the repeaters are all in series, failure of any repeater would make the ring inoperable. For the kind of office environment for which our ring is intended, it will be common for many of the nodes not to be in operation at any time. But the repeaters must always operate.

3) Perambulation. When either a repeater or a cable linking two repeaters fails, locating the failure requires perambulation of the ring, and thus access to all offices containing repeaters and wire runs containing cables. Portable test equipment is also required.

4) Installation headaches. Installation of a new repeater requires selecting two repeaters that are supposed to be directly linked, verifying that they actually are linked (that is, that the documentation of network topology is up-to-date) and installing new cables from each of them to the site of the new repeater. This installation approach has several consequences. The length of cable driven by the source repeater will change, possibly requiring retuning of its transmission circuitry. The old cable, now unused, is likely to be abandoned; one would expect the walls and ceilings to accumulate useless wire. Finally the ring may develop a wildly irregular topology, thereby increasing the effort involved in perambulation.

Bypass relays

The local net interface (LNI) currently in use at L.C.S. is designed to provide a partial solution to the problem of repeater failure, since that is almost certainly the most significant of the four problems. The LNI provides an "I-am-healthy" signal that operates a mechanical relay that would, if not energized, bypass that repeater, as in figure 2. Thus turning off the primary power at a node would kill the "I-am-healthy" signal, the relay would de-energize, and its contacts would cut the repeater out of the ring. This approach replaces the repeater failure problem with four new problems, that are (one hopes) easier to cope with:

2a) Repeater self-check failure. There is still a class of repeater failures that can disrupt the network, namely those in which the repeater's internal self-check circuits insist that the repeater is healthy when it really isn't. The frequency of such failures ought to be substantially smaller than that of all repeater failures, however.

2b) Relay contact failure. The relay contacts may fail, thus disrupting the ring.

2c) Line length variation. The length of transmission cable connecting a healthy repeater to its next healthy neighbor may vary widely, depending on how many intervening repeaters are bypassed. Thus the transmitter/receiver system must be designed to operate over a wide range of transmission distances, perhaps automatically adjusting to the cable length. (This problem is a rapidly varying version of one of the installation headache problems.)

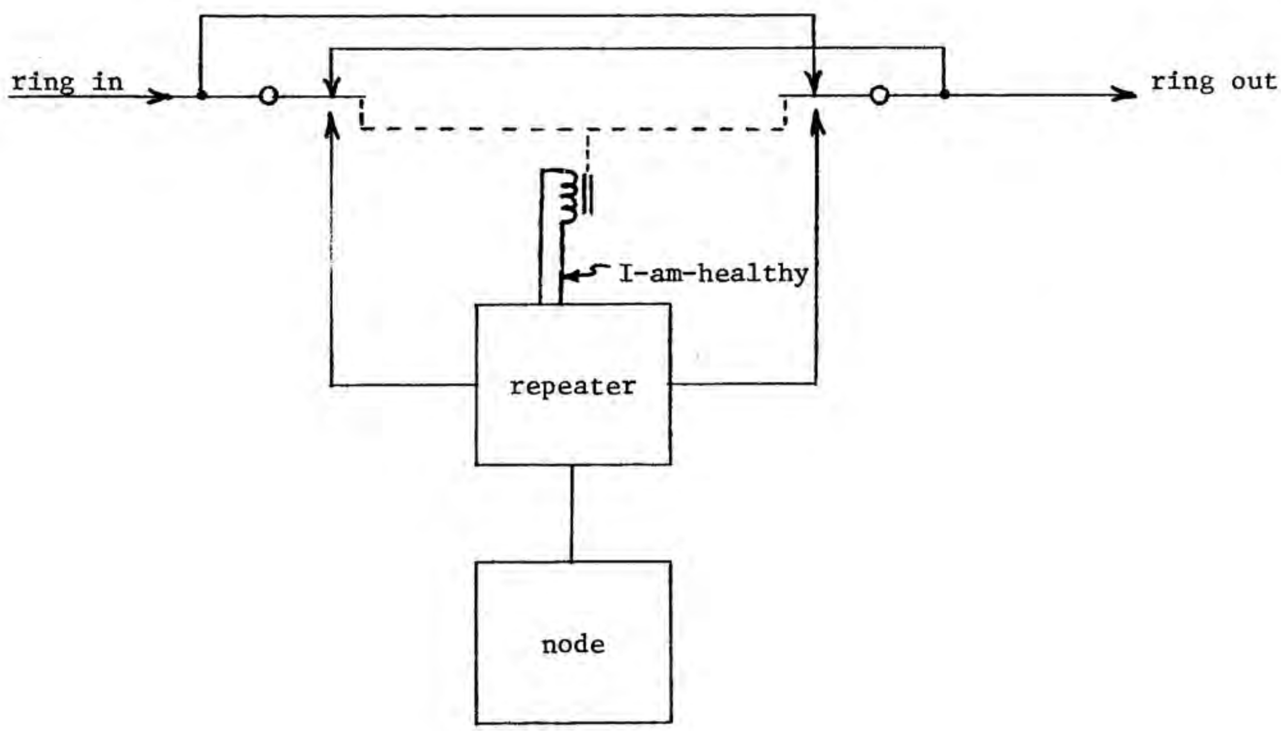

Figure two. Bypass relay energized by "I-am-healthy" line. Only one set of contacts is shown. In practice, three sets would be used, switching a twisted pair and its shield.

2d) Bypass disruption. At any instant a bypass relay may cut in or out, and its contacts may "bounce" for a few tens of milliseconds, thus destroying any message or token currently circulating around the ring.

The software recovery protocols of the ring network are designed to take occasional lost packets, messages, or tokens in stride, so the bypass disruption problem is not of concern, so long as it happens only occasionally.

The relay contact failure problem may be addressed by connecting normally closed relay contacts in parallel as shown in figure two, thus requiring that two relay contacts must fail simultaneously before a de-energized relay can disrupt the ring.

The Star-Shaped ring: step one

The perambulation and installation headache problems can be attacked simultaneously simply by rearranging the inter-repeater cables so that they always loop back through a single room, called the wire center, as in figure three.

With this star-shaped arrangement, it is not necessary for a trouble-shooter to have keys to every room containing a repeater or to carry test equipment from point to point. With access to the signal on every cable as it passes through the wire center, one can launch a message into the ring, observe the signals on successive cables to see how far it gets, and then reconfigure the ring to disconnect temporarily the repeater or the cable that seems not to be working. Then access is needed only to the area containing the troubled repeater or cable.

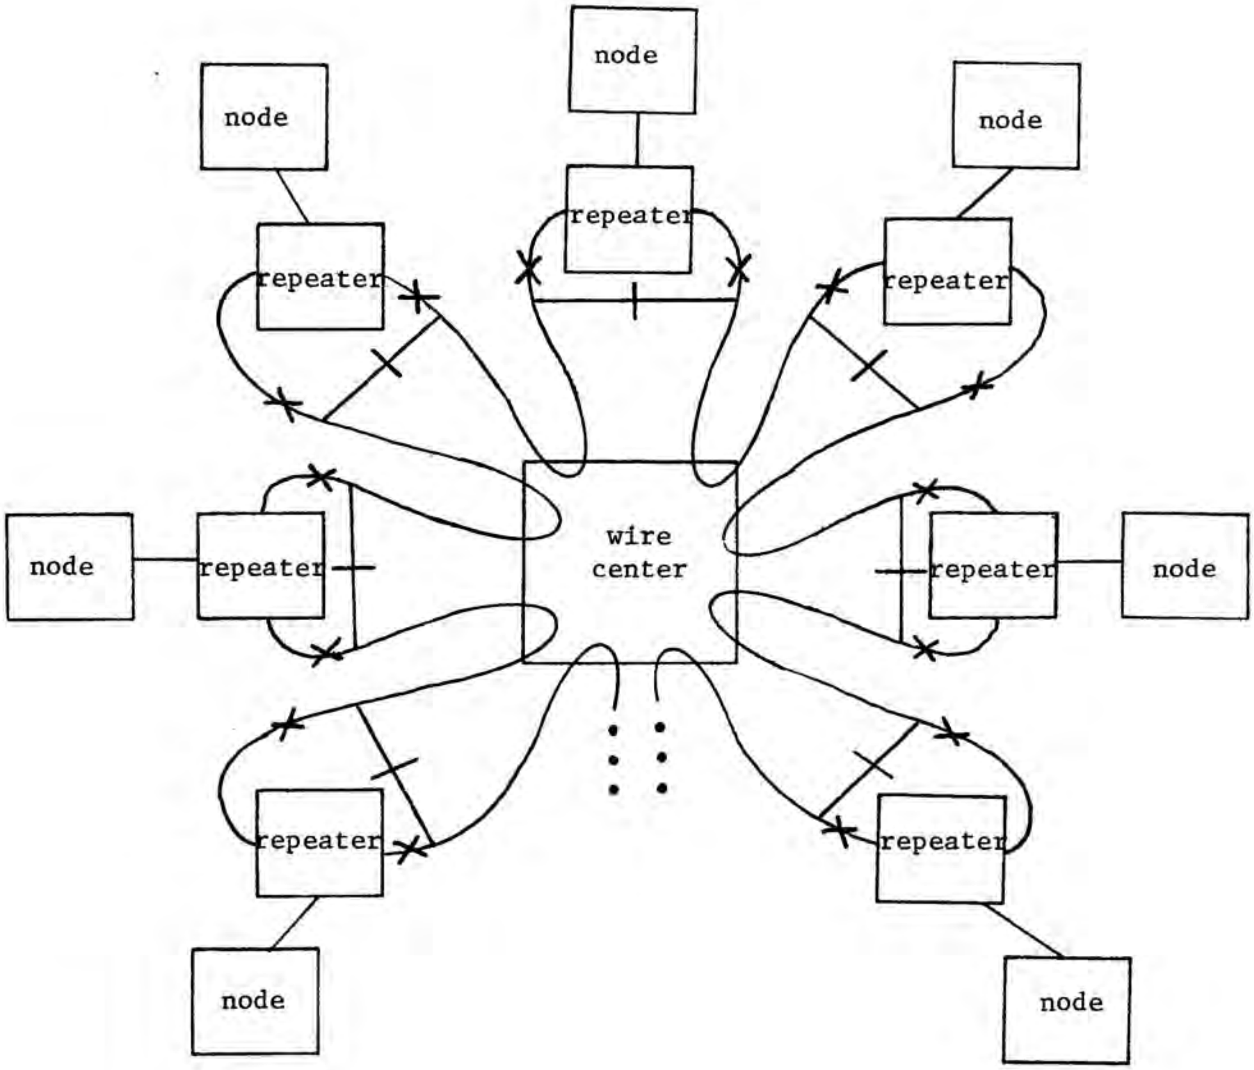

Figure three. Rearrangement of the ring in the shape of a star, to make serviceability easier. Bypass relay contacts are shown thus: normally closed —|—, normally open —X—.

Installation of new repeaters is also regularized. Two cables are installed from the location of a new repeater to the wire center. Then the new repeater is spliced into the ring entirely by rearranging wires at the wire center. At no point is one tracing old cables through walls or ceilings or abandoning them there when making new installations. Unanticipated installations can be handled without creating a hodge podge of wires crisscrossing through the walls and ceilings; physical planning of a network is thus simplified. The order of repeaters on the ring is determined at the wire center and can be rearranged if for some reason reordering seems necessary. (E.g., when trouble-shooting a repeater-repeater transmission failure, quick rearrangement might be a useful technique to isolate the problem to the transmitter or receiver side of the link).

The star-shaped ring: step two

A further refinement of the wire center concept considerably reduces the wire vulnerability problem, simplifies maintenance and installation further, and normalizes transmission line lengths. This refinement involves simply moving the bypass relays from the repeaters to the wire center. The primary impact of centralizing the relays is that it provides for automatic bypass of the vulnerable data transmission lines as well as of the repeaters. To control the relay, an extra pair of wires runs from the repeater to the wire center. In practice, installation of a repeater would be accomplished by pulling a single cable containing two data transmission lines and the relay control pair from the area of the repeater to the wire center. A single plug connects all of these lines to the repeater. Almost any accident, ranging from chopping the cable to kicking out the plug has the effect of bypassing both the repeater and also the data transmission lines to and from that repeater.

An important additional effect of bypassing the data transmission lines that lead to bypassed repeaters is that the path from one working repeater to the next working repeater always consists of the same cable run to the wire room, some number of bypass relays, and exactly one cable run to a receiver. This path varies only in the length of the final cable run; thus the range of signal levels to which the transmission system must automatically adapt is much smaller. Further, at the wire room an appropriate level-setting attenuator can be placed in series with short cables. Then a transmitter will always "see" essentially the same cable length no matter what the configuration of the ring. Thus the line length variation problem can be addressed by this arrangement.

Physical realization of the star-shaped ring

Serviceability of the ring can be further enhanced by a suitable physical realization of the interconnections at the wire center. The design of figure four illustrates. A printed circuit board is constructed with eight bypass relays in a row, and eight connectors to which cables to repeaters can be plugged. Each relay is connected to the next when its coil is de-energized. When energized, the relay cuts into the ring a pair of transmission paths that lead through a socket that can contain an attenuating network to the repeater connector at the edge of the board. Current for the relay coil comes from the repeater. A light-emitting-diode is connected across the relay coil for visual observation of the "I-am-healthy" signal.

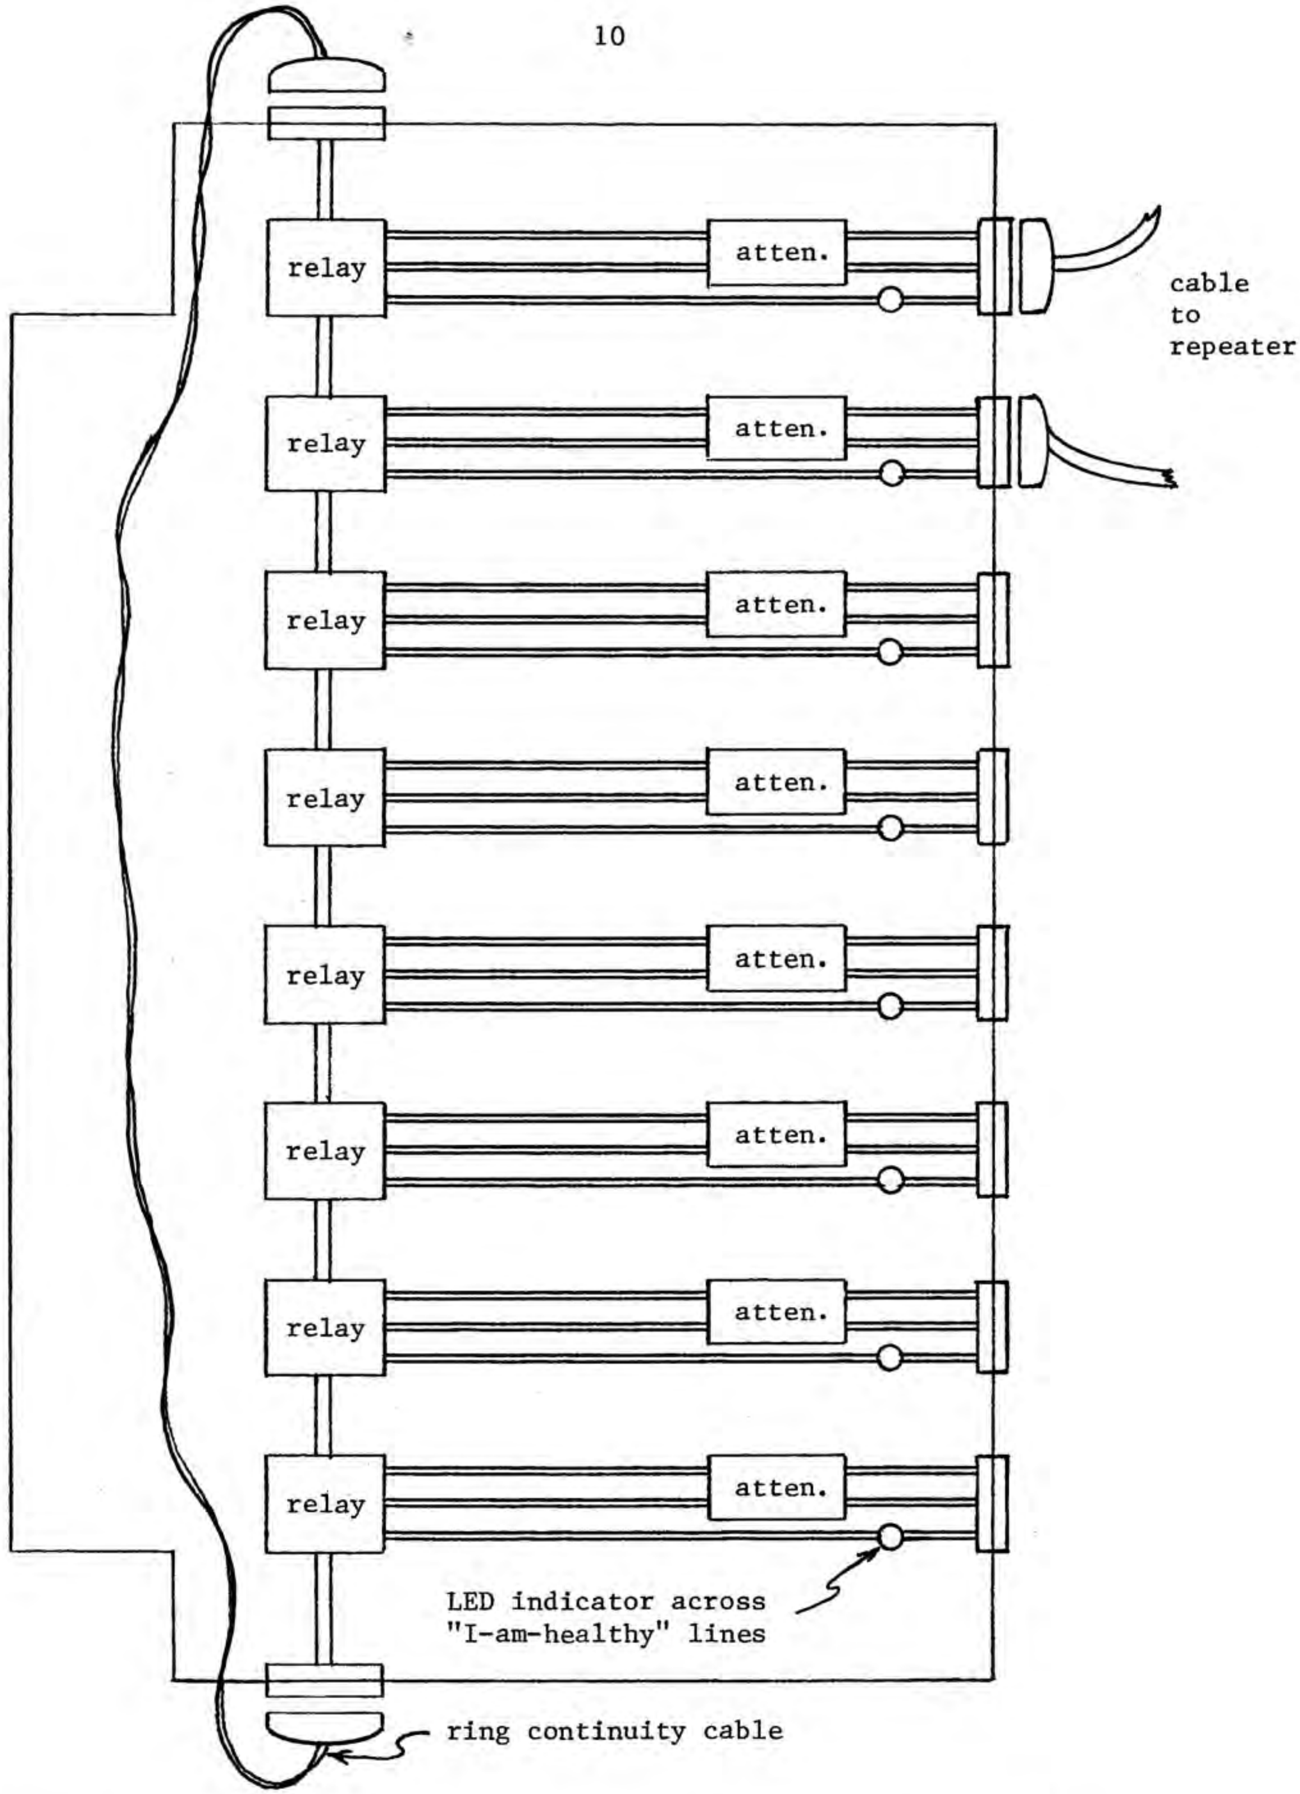

Figure four. Printed circuit board realization of the star-shaped ring.

The sequence of normally closed relay contacts leads to connectors at the top and bottom edges of the board. A "ring continuity cable" runs from the top connector around to the bottom one, completing the ring. In a sense, this ring of normally closed relay contacts on a single board in a controlled environment is the data communication ring, from the point of view of identifying what must be working to allow the ring to operate. If any two repeaters can energize their bypass relays, they can communicate, even though all other potential ring participants may have failed, powered down, or tripped over and disconnected their cables.

Installation of a new ring participant is accomplished by installing a cable from the repeater to the wire center, attaching a connector, and plugging it into an unused bypass relay. No disruption of ring operation occurs. If all of the bypass relays have repeaters attached to them, another printed circuit board containing eight more relays can be installed next to this one, and then in a matter of seconds cabled into the ring by interchanging top ring continuity cable connectors. (This kind of ring expansion may be scheduled at times when a few-second disruption is tolerable. Then, addition of repeaters can be accomplished at any time.)

Note that trouble-shooting with this physical configuration is especially straightforward. If the ring stops working, it is almost certainly because some repeater's "I-am-healthy" line is incorrectly energized. One starts by unplugging all the cable connectors to all the repeaters, checking for ring continuity, and then plugging in the cables to the repeaters one at a time to see which one (or ones) seem to disrupt continuity. Any cable whose reattachment causes trouble can be left unattached for the moment until it can be checked out more carefully. In this way the ring can be brought back into operation for the correctly behaving participants quite rapidly. Further, by rearranging the ring continuity cables that interconnect one printed circuit board to the next it is possible quickly to isolate problems to a group of eight relay/cable combinations. Finally, if a relay or printed circuit board component fails, that board can quickly be replaced or, at worst, bypassed.

<u>Wire center interconnection</u>

One might envision equipping a wire center with up to eight boards, each with eight relays, producing a ring of sixty-four nodes, probably enough to handle a typical building floor or wing. Another floor or wing would have its own wire center. The simplest way to interconnect the wire centers, if they are not too far apart, is to interconnect the ring continuity cables of the two wire centers. However, if the wire centers are very far apart, the center-to-center cables will add to the data transmission path between some repeaters but not between others, and the cable length variability problem will reappear. A better approach would be to build a special wire-center-to-wire-center repeater that would plug into both wire centers in the same way as a node repeater. If traffic is light enough to allow 128 nodes to be arranged in a single ring, this repeater would be a simple non-addressable full duplex repeater. If traffic grows to the point that two 64-node rings would be a better arrangement, the repeater could be replaced by a filtering bridge that forwards to the other ring only messages that are addressed to nodes not on this ring segment. [This same basic filtering bridge design could also be used to divide a 64-node ring at a single wire center into two or more smaller rings, if necessary.]

Automatic ring recovery

A further possibility for automating maintenance is to connect a microprocessor controlled repeater to one of the repeater connectors, and program the microprocessor to occasionally (say, once per second) launch a test packet around the ring. The microprocessor would also control, perhaps through another set of relay coils, the continuity between each repeater's "I-am-healthy" line and its bypass relay. If a test packet fails to make it around the ring, the microprocessor would force de-energization of all of the bypass relays (except its own) at once, check for ring continuity, and then reconnect the "I-am-healthy" lines one at a time, testing ring continuity after each reconnection. When it finds an "I-am-healthy" line that disrupts ring continuity it might broadcast a trouble report around the ring as a way to call for a repairman. Since no other peripheral devices are needed, and the program is fixed, the extra cost of this automatic test system could be very small.

It is not clear how useful this notion of a microprocessor controlled automatic ring maintainer may be in a network consisting only of a single, small wire center. It may be overkill to automate the function because the ring may be very reliable anyway, and there is a danger that the microprocessor will fail in such a way as to disrupt the network, or that its automatic operation will be so effective that it reduces the incentive to repair misbehaving nodes. On the other hand in a campus-wide network with 100 wire centers, the automatic maintenance feature may be very helpful in getting service attention quickly and in rapidly restoring operation.

Conclusion

The star-shaped ring proposal has the decentralized control advantage of a ring network, and at the same time the centralized maintenance advantage of a star network. It thus captures some of the better properties while avoiding the key disadvantages of both. We propose to implement the star-shaped ring (initially without the microprocessor automated maintenance feature) for the L.C.S. ring network.