Transient Phenomena in Bridged Local Area Networks[1]

C. Ersoy *S. S. Panwar*

Polytechnic University, Brooklyn, New York

R. Dalias *D. Segal*

SIAC Inc., Brooklyn, New York

Abstract

Local Area Network segments are interconnected by bridges or routers in order to overcome limitations on the number of users, the total traffic and the maximum length of a segment [1]. One of the simplest interconnection methods is the use of spanning tree bridges based on the IEEE 802.1 D standard. This paper describes the transient behavior of the Bridged LAN following failures. In the presence of a network component failure, the bridged LAN may experience high transient traffic periods called **storms**, as well as **virtual circuit disconnects**. The paper summarizes conditions that lead to these problems, describes some experimental tests, and suggests ways of reducing their impact on the bridged LAN.

1. Introduction

Local Area Network segments are interconnected by bridges or routers in order to overcome limitations on the number of users, the total traffic, and the maximum length of a segment. One of the simplest interconnection methods is the use of spanning tree bridges. These bridges have been accepted as the interconnection standard for CSMA/CD LANs by the IEEE 802.3 committee [5].

In this paper, we study transient phenomena in bridged LANs. In the first section, we give a brief description of the components of the system under study. In Sections 2 and 3, we give the results of a detailed study on two transient phenomena: storms and virtual circuit disconnects. We give our conclusions in Section 4.

The system studied consists of Ethernet LANs interconnected by transparent bridges. Each Ethernet LAN consists of different segments connected by repeaters. A description and a performance analysis of Ethernet LANs can be found in [2] and [3], respectively. Spanning Tree Bridges [1] are transparent to end stations. They require no manual operations for configuration. Transparent bridges learn the locations of end stations by observing source addresses. The learned station addresses are stored in Routing Data Base (RDB) tables. The entries in RDB tables are aged out for allowing station movements and topology changes. The period for erasing the inactive station addresses from RDB tables is called the Aging Time. Transparent bridges route traffic by comparing the destination address of frames to the table of learned addresses. Their routing decision can be discarding the packet (filtering), forwarding it to a specific port, or flooding the packet to all ports except the one from which the packet is received. For proper operation, the learning and forwarding processes assume that the logical topology of bridges and LANs is a spanning

---

[1] This paper was published in the Proceedings of IEEE GLOBECOM'90. This work was supported by the Securities Industry Automation Corporation (SIAC); the Center for Advanced Technology in Telecommunications (CATT), Polytechnic University; and by the NSF under Grant NCR-8909719.

tree. From a given physical topology, bridges find a logical spanning tree topology by exchanging hello messages. The spanning tree algorithm used for finding the logical topology and the contents of hello messages are explained in [1], [4], [5]. Bridges adapt to changes in the topology and the movement of stations without network management intervention. However network management can be used to force desired topologies.

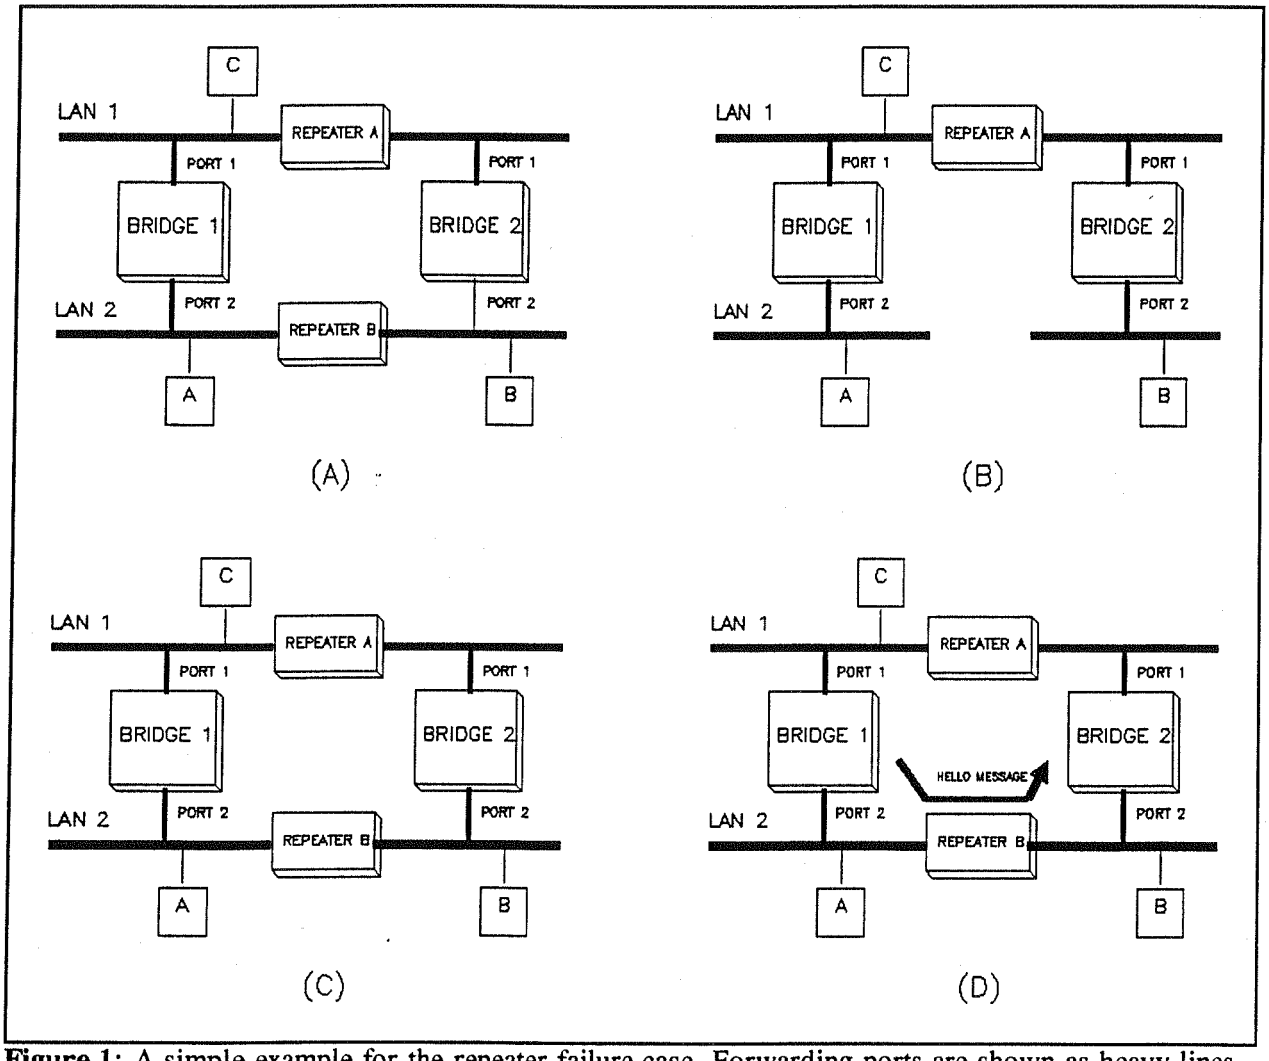

Figure 1: A simple example for the repeater failure case. Forwarding ports are shown as heavy lines.

2. Storms

The saturation of a network or a network segment due to a high level of traffic is called a **storm**. Reasons for storms vary from hardware failures to protocol related failures. We will consider storms caused by intermittent failures and one-directional component failures.

The Spanning Tree Algorithm ensures that there will not be any permanent loops in the network. The addition of a new network element or unexpected device failures may cause temporary loops in the network. In normal circumstances these loops are broken by action taken by the bridges. During the time

when the network has the loop, there can be storms. The sufficient conditions for having a storm of the kind we are considering here are :

1) The existence of a loop in the network, and,
2) A unicast packet with unknown destination or a multicast packet.

Almost all networks have broadcast messages carried by multicast packets. Therefore, presence of a loop will generally result in a storm. In addition, we will show later that unicast messages may also lead to storms. Some of the cases in which the network may have temporary loops are listed below:

1) Failure and recovery of a bridge,
2) Addition of a repeater connecting two previously unconnected LAN segments,
3) Failure and recovery of a repeater,
4) One directional failures, for example, that of a repeater.

Figure 1 shows four phases of a simple Bridged LAN before and after the failure of Repeater B. In Figure 1(A), the logical topology found by the Spanning Tree Algorithm is shown with bold lines. While the system is in the state shown in Figure 1(A), Repeater B fails and disconnects the Port 2 side of the bridges. Bridge 2 stops hearing the Hello Messages generated by Bridge 1 through LAN 2. After waiting for a period specified in the protocol [5], in order to maintain network connectivity, Bridge 2 becomes a forwarding bridge as shown in Figure 1(B).

At a later time, Repeater B recovers and reconnects Port 2 sides of the bridges as shown in Figure 1(C). Just after that, Node C generates a broadcast packet. This packet is received and forwarded by both of the bridges. The forwarded packets on LAN 2 are received and re-forwarded to LAN 1 by each of the bridges. This causes multiple copies of the broadcast packet travelling in the loop created by the two bridges. Bridges forward the packets as fast as they can and this causes a storm on both of the LANs.

Bridges periodically exchange topology information by sending and receiving Hello messages [4], [5]. The period is called the Hello Time. The storm lasts until Bridge 2 hears the hello message sent by Bridge 1 over LAN 2 as shown in Figure 1(D). This Hello message notifies Bridge 2 that a higher priority bridge is active in parallel. After hearing that Hello packet, Bridge 2 terminates the loop by changing the state of its Port 2 to blocking. The system returns to the state shown in Figure 1(A).

In the scenario of Figure 1, if Repeater B fails and recovers in a period of time which is not long enough for Bridge 2 to become a forwarding bridge, then there will not be a temporary loop. The time required for a bridge to become active as a result of a change in the network topology can be changed by the network manager. Therefore, loops caused by intermittent failures can be reduced in frequency by increasing this reaction time. However, this will be at the cost of increasing the time to recovery for the network after a non-intermittent failure. It is important to note that, in the case of one-directional repeater failure, the loop may not be terminated since Bridge 2 will not be able to hear the Hello Message sent by Bridge 1. As a result, the storm will not end until the repeater recovers.

Under the same repeater failure scenario, a packet with a single destination address (a unicast packet) may or may not cause a storm. After the loop is created as shown in Figure 1(C), Node B sends a packet to Node C. This packet is forwarded by both of the bridges. If C has not been transmitting for some time, it will not be on the RDB of the Port 1 side of either bridge. As a result, both copies are forwarded back to LAN 1 by the other bridge. In this case, a storm results. On the other hand, if C has sent packets

recently, the bridges will not cause a storm, because both of them will know that Node C lies on their Port 1 sides, and will filter any packets destined to C.

2.1 Results of the experiments performed on Storms

There are four LANs and five bridges in the experiment set-up as shown in Figure 2. Bold lines represent the logical tree topology. In order to create a storm, Repeater B was disconnected from Repeater C. After L_BRDG_4 became a forwarding bridge, Repeater B was reconnected. This created a loop and a storm around L_BRDG_2 and L_BRDG_4. The duration of the storm corresponds to the lifetime of the active loop. The loop was broken by the first Hello Message generated by L_BRDG_2 towards LAN 4. Upon receiving this Hello Message, L_BRDG_4 put its Port 2 into blocking state. The storm in the system lasted at most a Hello Time. Just before the reconnection of Repeater B, the utilization on LANs was negligible. During the storm, the utilization jumped to approximately 35%. The same experiment was repeated a few times and in each case the utilization on LAN 4 was between 30% and 38%. Similarly, the utilization increased to between 30% and 38% on LAN 3, and between 14% and 20% on LAN 2 and LAN 1.

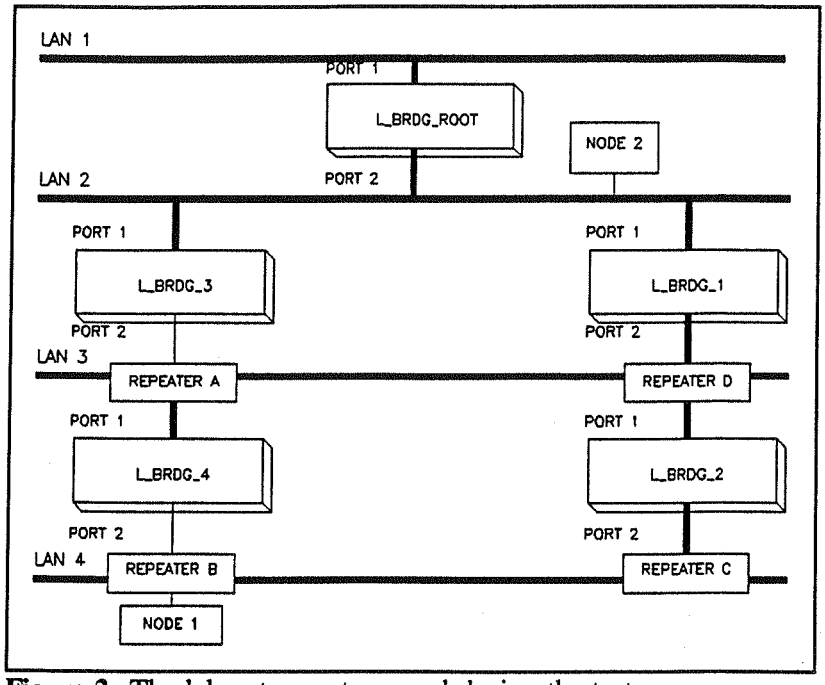

Figure 2: The laboratory set-up used during the tests.

The local bridges used during the experiments can forward at most 2650 64-byte packets per second. This corresponds to approximately 17% utilization on a 10 Mbit/s Ethernet LAN. The storm is caused by packets circulating in both directions in the loop. Since the storm is caused by 64 byte packets, 34% utilization corresponds to the combination of the maximum forwarding rates of two bridges. This shows that the maximum utilization due to the storm is limited by the maximum forwarding rates of the bridges. During the storm both L_BRDG_3 and L_BRDG_4 were forwarding packets on LAN 3 and LAN 4. On the other hand, only L_BRDG_1 forwards these packets on LAN 2 and similarly, only L_BRDG_ROOT forwards packets to LAN 1. This explains why the utilization on LAN 1 and LAN 2 is approximately half of the utilization on LAN 3 and LAN 4. This also shows that bridges may not isolate the storms which are created around an active loop in the topology. In Figure 2, the storm in the lower loop propagated to the upper LANs, because it was caused by a multicast packet or a packet with an unknown destination address. Bridges flood these packets to all of their ports except the one from which they received it. Due to this property of bridges, all the storms created by such packets are carried to the other parts of the network. In other words, they are network-wide storms.

On the other hand, if a storm is created by unicast packets, the filtering of unicast packets by bridges may localize storms in one part of the network. In the example of Figure 2, if L_BRDG_1 knows that the destination address of the storm packet is on its Port 2 side, it will block the storm packets. Therefore,

the storm will stay local to the lower loop and will not be carried to the other parts of the network. However, if the destination is unknown to L_BRDG_1, the storm will propagate to LAN 2. The example above will also be valid for recoveries from other kinds of failures, such as bridge port failures which may also cause temporary loops in bridged LANs.

2.2 Storms due to one-directional failures

So far we have assumed that the repeater failure is two sided. That is to say, if a repeater fails, it does not pass any packets in either direction. Suppose a repeater failure prevents packets from being transmitted from one side to the other, but allows the packets to pass in the other direction. Suppose that for the set up of Figure 1, Repeater B stops passing packets from the Bridge 1 side to the Bridge 2 side, but lets the packets pass in the other direction. As in the previous scenario, Bridge 2 becomes a forwarding bridge a certain amount of time after Repeater B fails. Since this time the failure is one sided, Hello Messages generated by Bridge 1 towards LAN 2 will not reach Bridge 2, but the data packets forwarded by Bridge 2 will reach Bridge 1. Again, as in the previous scenario, any unicast packet with an unknown destination or a multicast packet will cause a storm in the loop. However, this loop is one directional. Bridge 2 will not be aware that it is working in parallel with another active bridge. This time the storm will last until Repeater B recovers or, if the storm is due to a unicast packet, the packet destination becomes known. In the presence of one-sided failure, we know of no simple way of preventing such storms. By closely monitoring the bridge and LAN for abnormal traffic, there may be ways of detecting and flagging such phenomena.

2.3 Suggested means of reducing the impact and frequency of storms

Network management should be aware that a state change of a bridge port may mean a failure. According to the Spanning Tree Algorithm [5], if a loop consist of two bridges, the one with the higher priority will be in the forwarding mode and the other bridge will be in the blocking mode. Under normal circumstances, higher priority bridge should not hear any hello message from the lower priority bridge. Therefore, the higher priority bridge should warn the network manager upon receiving a hello message from the lower priority bridge. An alternative is to wait for two hello messages to prevent an alarm being sent if a bridge has just been connected to the network. If a station address changes port sides within the aging time, the network manager should be alerted to a possible problem. The network manager can set filters on a port by port basis preventing certain types of packets to travel to certain parts of the network. This could be used to limit the propagation of storms.

We described that one-directional component failures may cause storms which are difficult to stop. Thus, components which can have this type of failure should not be included in the network if possible. If they are used, the network management system should be able to monitor one-directional faults in these components.

3. Virtual Circuit Disconnects

We will show how an intermittent fault in a repeater causes a virtual circuit disconnect between a host gateway and a user in Figures 3 to 6. In each figure, the logical topology is shown with bold lines. Dotted, broken or thin lines correspond to the routes taken by different types of packets. At the bottom of each figure, the source and the destination of the packets are indicated in parentheses, i.e., (SOURCE,DESTINATION). The symbols at the lower or upper side of a bridge indicate the elements

in the Routing Data-Base (RDB) of the port closest to that side. Note that bridges in the set-up have separate RDB's for each port. While the system is in the state shown in Figure 3, Repeater_B fails and disconnects the port 1 sides of L_BRDG_3 and L_BRDG_4. After waiting for a period specified in the protocol [5], L_BRDG_4 becomes a forwarding bridge as shown in Figure 4. At a later time Repeater_B reconnects the port 1 sides of L_BRDG_3 and L_BRDG_4. Both L_BRDG_3 and L_BRDG_4 receive and forward (G,MULTICAST) packets. We know from Section 2 that this causes a storm. During the storm, L_BRDG_1, L_BRDG_2, L_BRDG_3, and L_BRDG_4 forward (G,MULTICAST) packets in both directions. These four bridges learn that the gateway is on both sides. In other words, they have the address of the Gateway in both of their RDB tables as shown in Figure 5.

L_BRDG_2 does not forward any packets destined to the Gateway after learning that the Gateway is on both sides. As a result of that, (USER,G) packets cannot reach their destinations as shown in Figure 6. Since there is no packet flow from users to the host, the host gateway breaks its virtual circuit with the user after the expiration of the virtual circuit time-out.

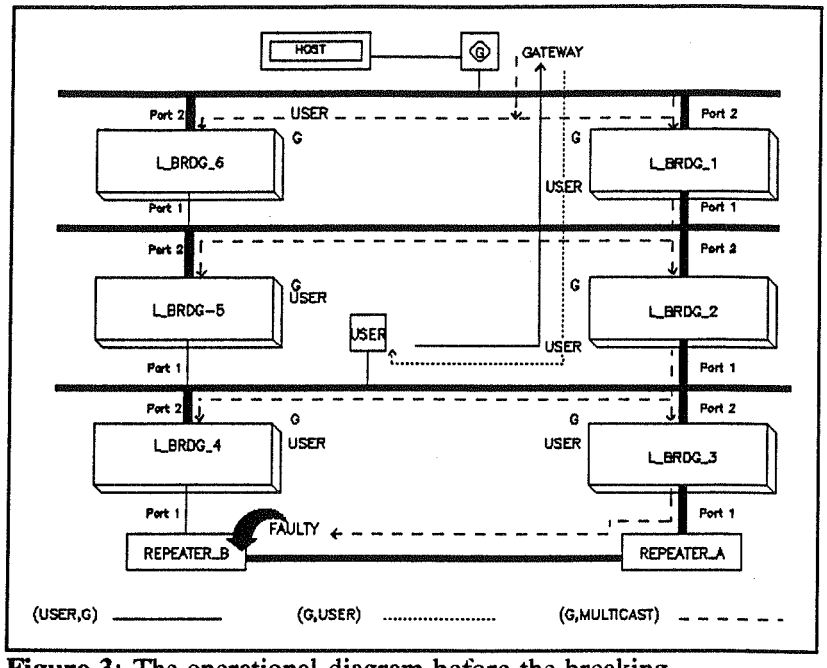

Figure 3: The operational diagram before the breaking.

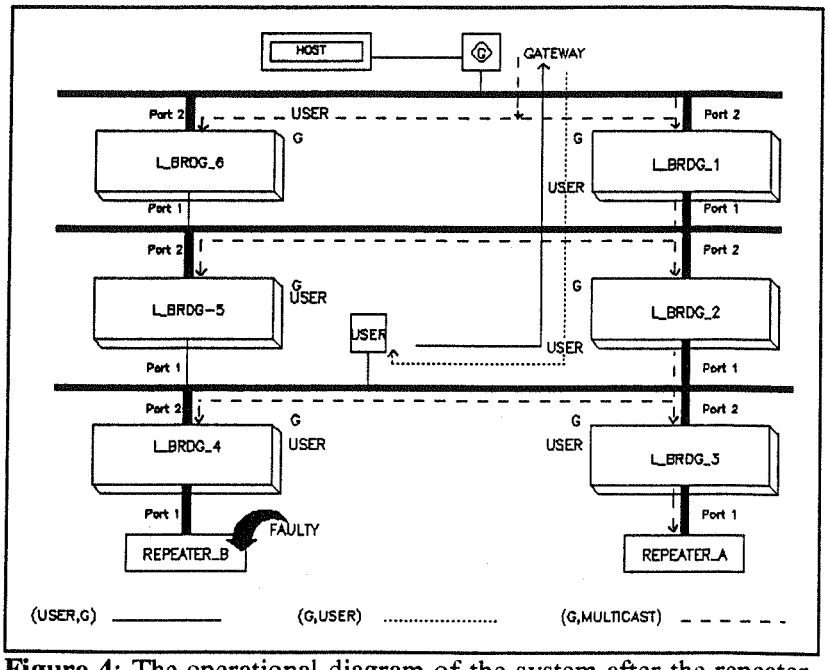

Figure 4: The operational diagram of the system after the repeater failure.

L_BRDG_4 receives the hello message generated by L_BRDG_3 from its Port 1 at most a Hello Time after Repeater_B recovers. L_BRDG_4 becomes inactive after receiving the Hello Message as shown in Figure 6. The storm ends because L_BRDG_4 stops forwarding. The incorrect information about the Gateway in the RDB tables of the bridges are deleted after the Aging Time. Each bridge removes the

address of the Gateway from the RDB table of its Port 1. L_BRDG_2 and L_BRDG_1 start forwarding (USER,G) packets after erasing the address of the gateway from the RDB tables of their Port 1. This allows the host to re-initialize its virtual circuit with the user. The virtual circuit will not break if the timeout is set to a value larger than the sum of Aging Time and Hello Time. However, in this case, repeater failures may be masked from the system operators, decreasing the capability of the network management system to identify faulty components.

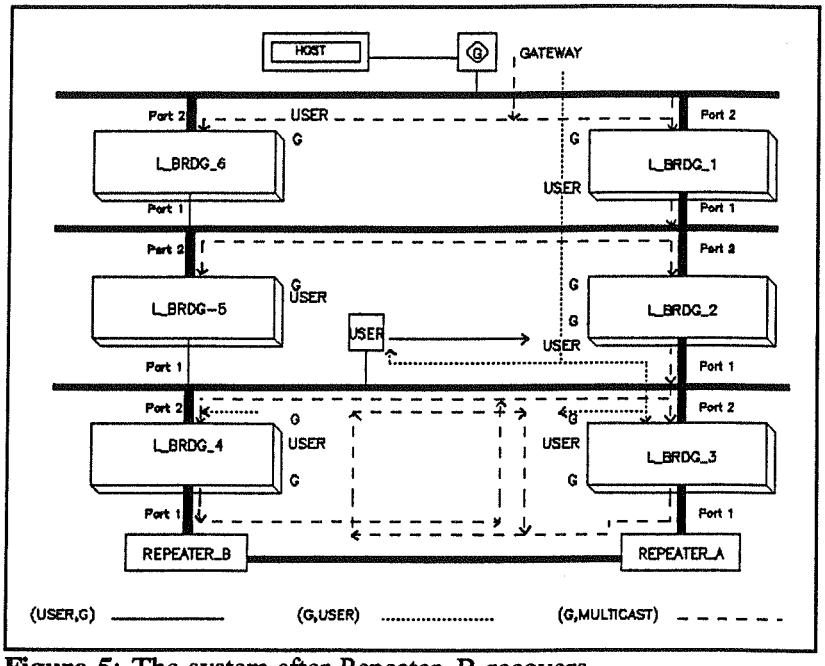

Figure 5: The system after Repeater_B recovers.

In the scenario described above, bridges have separate RDB tables for each port. Now, consider a bridge with a single RDB table which keeps the station addresses and the associated port numbers. This table will not have duplicate entries for the same station address. Thus, there will not be inconsistent routing information in the RDB table. When a loop is formed in the network, packets with the same source address may appear on different sides of the bridge. Therefore, the associated port for a station may change from one port to the other, depending on which port the packet from that station was most recently received. If the RDB table has the correct information, there will be no traffic interruption when the loop is broken by a hello message as shown in Figure 6. If RDB table has the incorrect information, (USER,G) packets will be blocked.

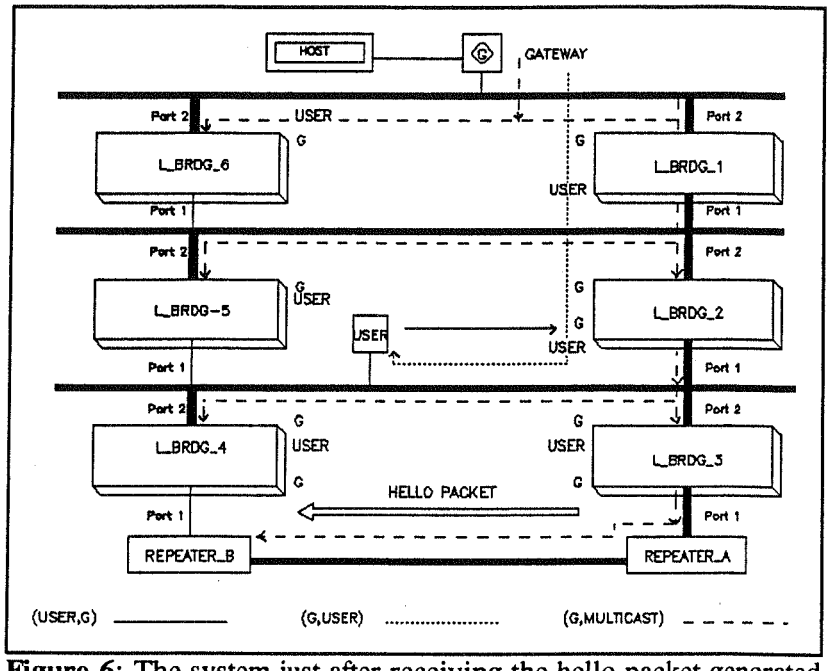

Figure 6: The system just after receiving the hello packet generated by L_BRDG_3.

Any traffic interruption will end when either a packet from the station appears on the "right" side of the bridge, thus correcting the address table, or the incorrect entry is aged out after the Aging Time. Thus traffic interruptions may last for a shorter period of time than when the bridge has an address table for

every port. Since the frequency and duration of traffic interruptions is lower in this case, it is advantageous to have a single address table per bridge. However, the impact of a single table on storms is yet to be determined, and should be a topic of further study.

4. Conclusions

In the presence of an intermittent failure, the bridged LAN may experience high transient traffic periods called storms, as well as virtual circuit disconnects. In most cases observed, storms last less than a Hello Time. The peak traffic on the LAN during the storm is proportional to the peak forwarding rates of the bridges connected to the LAN. Some storms may not remain as local problems and propagate towards other parts of the network. The presence of separate RDB tables for each port increases the probability of virtual circuit disconnects. Some virtual circuit disconnect problems can be prevented by setting certain time-out values appropriately. This solution has the disadvantage of masking component failures from the network manager. We have also suggested other means to detect and reduce the frequency of storms and virtual circuit disconnects. These point to a need for better network management tools for bridged LANs.

References

[1] "Transparent Bridges for Interconnection of IEEE 802 LANs", F. Backes, IEEE Network Magazine, vol 2, no. 1, January 1988.

[2] IEEE CSMA/CD Access Method, American National Standards ANSI/IEEE Std. 802.3, 1985.

[3] "Measured Capacity of an Ethernet: Myths and reality", D. R. Boggs, J. C. Mogul, C. A. Kent, Proc. ACM SIGCOMM'88, pp. 222-234, August 1988.

[4] "An Algorithm for Distributed Computation of a Spanning Tree", R. Perlman, Proceedings, Ninth Data Communications Symposium, 1984.

[5] IEEE Project 802 Local and Metropolitan Area Network Standards, "IEEE Standard 802.1 (D) MAC Bridges".