NASA Technical Memorandum 85915

Test Experience on an Ultrareliable Computer Communication Network

Larry W. Abbott

December 1984

National Aeronautics and
Space Administration

NASA Technical Memorandum 85915

Test Experience on an Ultrareliable Computer Communication Network

Larry W. Abbott
Ames Research Center, Dryden Flight Research Facility, Edwards, California 93523

1984

National Aeronautics and
Space Administration
Ames Research Center
Dryden Flight Research Facility
Edwards, California 93523

TEST EXPERIENCE ON AN ULTRARELIABLE COMPUTER COMMUNICATION NETWORK

Larry W. Abbott

Flight Systems Engineer
NASA Ames Research Center
Dryden Flight Research Facility
Edwards, California

Abstract

The dispersed sensor processing mesh (DSPM) is an experimental, ultrareliable, fault-tolerant computer communications network that exhibits an organic-like ability to regenerate itself after suffering damage. The regeneration is accomplished by two routines — grow and repair. This paper discusses the DSPM concept for achieving fault tolerance and provides a brief description of the mechanization of both the experiment and the six-node experimental network. The main topic of this paper is the system performance of the growth algorithm contained in the grow routine. The characteristics imbued to DSPM by the growth algorithm are also discussed. Data from an experimental DSPM network and software simulation of larger DSPM-type networks are used to examine the inherent limitation on growth time by the growth algorithm and the relationship of growth time to network size and topology.

Introduction

The dispersed sensor processing mesh (DSPM) is an ultrareliable structure for gathering sensor data and distributing effector data. An ultrareliable system requires an ultrareliable communications structure as a complementary partner to the ultrareliable computational element. The DSPM concept is the forerunner to the ultrareliable input/output (I/O) network specified in Ref. 1 for Charles Stark Draper Laboratory's advanced information processing system (AIPS).

The reliability of the DSPM network is greatly enhanced by the ability of the DSPM communication network to reconfigure (2). Two software algorithms — grow and repair (2,3) — perform the reconfiguration task. This paper describes the growth algorithm resident in the grow routine and the effects of the growth algorithm on the characteristics of the DSPM network. The characteristics of the experimental DSPM system are explained in terms of the experimental results, and the experimental results are used to certify the validity of the DSPM simulation software. The resulting validated DSPM simulation is used to derive the generic characteristics for a broad range of DSPM networks.

Concept

Figure 1 depicts a generic DSPM network taken from several examples given in Ref. 4. The network is formed with nodes (shown as circles), links (shown as lines), and a central bus controller (each channel of the quadraplex bus controller is shown as a rectangle). The growth algorithm and other network creation and maintenance software reside in the bus controller. The links carry the network communications, and the nodes gather and distribute data. A brief description of the growth algorithm (Fig. 2) is given; detailed descriptions can be found in Refs. 2 and 3.

Initially, the growth algorithm grows the network to the nodes surrounding the bus controller by activating the link (shown as a solid line) to each node (nodes 1, 2, 3, and 4) and making the node a member of the network. Activating a link requires that the destination node not be a member of the network and that it responds to the bus controller configuration commands.

In each successive growth cycle, the network is grown from nodes activated in the previous cycle until each node is attached to the network through a tree that has the bus controller as its root. During the growth process, the network is unavailable to process inputs to a node or outputs from a node. If faults exist in the links or in the nodes, the growth algorithm circumvents the fault in an "organic like" regeneration of the network by way of another configuration. The fault tolerance of the DSPM concept is a result of the ability of a DSPM network to reconfigure around failures.

Mechanization

The DSPM system is a complex, experimental communications system. To obtain valuable data on practical implementation issues, the ultrareliable DSPM communication concept was interfaced to a state-of-the-art fault-tolerant system and the tests were run on NASA's F-8 Ironbird simulator. NASA's F-8 digital fly-by-wire (DFBW) Ironbird simulator provides a safe, yet realistic means of testing the complex interaction of a highly reliable communication network with a state-of-the-art, triply redundant, digital flight-control system that contains redundant computers, sensors, actuators, and flight-critical software (5).

The DSPM experimental system (for additional details see Refs. 2 and 3) consists of three major components (see Fig. 3): the F-8 DFBW Ironbird simulator (with triplex flight-control computer and control-law software); a central computer with simulation software; and a six-node (plus a triplex bus controller where each channel is based on a 5 MHz MC68000 microprocessor) version of the DSPM network. The bus controller in the experimental system has relatively low performance; a production system should be expected to run much faster. The details of the DSPM hardware implementation are beyond the scope of this paper, but are presented in Refs. 2 and 6.

Experimental Results

The preliminary tests on the DSPM system measured the actual performance of the growth algorithm under varied fault conditions (for the purposes of this paper, a fault is a static failure of a port to the off condition). Performance data for the growth algorithm used on the Ironbird DSPM network (see Fig. 4) is available for all failure combinations up to three failures per node for both the adjacent and disjoint node pairs. An example of data for both an adjacent and a disjoint node pair is shown in Fig. 5. The data discloses two characteristics of the DSPM growth algorithm: (1) the mean growth time is linearly related to the number of faults, and (2) although the growth algorithm is deterministic, there is a wide deviation in the growth time for a given number of faults.

The deviation in growth time for a given number of failures on a given node pair is caused by the different amounts of time needed to process different types (2) and combinations of faults. The differences in growth time for configurations with the same number of failures is because of the topology of the network, the position of the node within the network, and the preferential order of growth (clockwise in DSPM). These deviations are not generic to the DSPM approach, but are dependent on the topology and the preferential growth algorithm. However, the relationship of growth time to inbound versus outbound ports is a generic DSPM characteristic. Specifically, on the experimental DSPM system, if a failed port is an outbound port (it relays bus controller (BC) messages to other nodes), the growth algorithm spends approximately 4.8 msec processing the failed port. If the failed port is an inbound port (it returns node responses to the BC), the growth algorithm spends about 7.7 msec processing the failed port. The fact that an outbound port must be grown before the inbound port of the next node can be grown accounts for the difference in growth time. Thus a fault on an outbound port can always be detected sooner than a fault on an inbound port.

While the linearity of the growth times are influenced by the implementation of the hardware and software in each DSPM-type system, the linear nature of the growth times is a generic DSPM characteristic. All fully operable DSPM configurations must have exactly the same number of good links as nodes and, because the growth time for good links varies very little, the only difference in growth time is related to the number of faulty links. While the time required to determine if a link is faulty varies, depending on whether it is used as an inbound or an outbound link, the values are roughly the same, and the relationship turns out to be approximately linear.

In Fig. 5, note that for a few faults, the growth times for failure sets on disjoint and adjacent node pairs are appoximately the same; but, as more faults are injected, the disjoint node pair requires more growth time than an adjacent node pair. A single fault in a disjoint node pair appears the same as a single fault in an adjacent node pair. Therefore, it is not surprising that the growth times are similar for a few faults. As the number of faults increase, the failures on the disjoint node pair normally disrupt two or more trees while failures on the adjacent node pairs, which share a common link, normally disrupt a single tree.

Naturally, one might ask what happens if more failures are injected into the network than are shown in Fig. 5. Can we linearly extrapolate the growth time? Unfortunately, the answer is no. There are two reasons for this. First of all, each nonlatent fault (a latent fault is a fault in a link with an existing dominate fault or a fault in a link that will not be activated) causes the link associated with it to fail during an activation attempt. Because valid DSPM configurations must have a good link for each node in the network, the maximum number of failed links (and therefore nonlatent faults) equal the number of links minus the number of nodes (see Eq. (1)).

$$\text{Maximum number of faults} = \text{number of links} - \text{number of nodes} \qquad (1)$$

For the Ironbird DSPM network that is used on the F-8 Ironbird simulation (see Fig. 4), the maximum number of faults is six.

The second reason that the growth time cannot be linearly extrapolated provides much better insight into the behavior of the network. Exceeding the maximum number of nonlatent faults replaces inbound failures with outbound failures. However, to see this behavior, latent faults that occur in links where activation attempts occur must be counted as faults.

As shown in the following maximum growth time scenario, when all the faults previously defined are counted, the growth time still does not exceed the growth time for the maximum number of nonlatent failures, even when the number of faults exceed the maximum number of nonlatent failures. Because the scenario is constructed with a homogenous fault type (that is, worst case faults), the linear relationship of growth time is shown without any deviation. An example of maximum growth time is to grow the DSPM network with the maximum growth time by using the following heuristic rules to choose a worst case fault: (1) failures on an inbound port contribute more to growth time than failures on an outbound port and (2) growth at a node proceeds clockwise with port 0 first and port 3 last.

To achieve the maximum growth time we must grow a network with as many inbound ports failures as possible and with growth occurring at the most counterclockwise port possible. The following growth phases form the maximum growth time scenario for the Ironbird DSPM network.

Phase 1: Grow out of the most counterclockwise port of the BC. To force the growth to BC port 2, the links to BC ports 0 and 1 must be failed at one end or the other (see Fig. 6(a)). To obtain the maximum growth time, the first failure (at an inbound port) is injected at node 1, port 0 (N1P0) and results in a growth time of 24.8 msec. The process is cumulative. For two failures, a failure at N4P0 is added to the failure at N1P0, resulting in a growth time of 31.4 msec.

Number of failures	Fail vector	Growth time
1	N1P0	24.8 msec
2	N1P0, N4P0	31.4 msec

Phase 2: Grow out of the most counterclockwise port of node 6. Force the growth to N6P3 by failing N2P2 and N5P2 (see Fig. 6(b)).

Number of failures	Fail vector	Growth time
3	N1P0, N4P0, N2P2	39.1 msec
4	N1P0, N4P0, N2P2, N5P1	46.8 msec

Phase 3: Grow out of the most counterclockwise port of node 4. Force the growth to port 3 of node 4 by failing N5P0 (see Fig. 6(c)).

Number of failures	Fail vector	Growth time
5	N1P0, N4P0, N2P2, N5P1, N5P0	54.5 msec

Phase 4: Grow out of the most counterclockwise port of node 1. Force the growth to port 3 of node 1 by failing N3P0 (see Fig. 6(d)).

Number of failures	Fail vector	Growth time
6	N1P0, N4P0, N2P2, N5P1, N5P0, N3P0	62.1 msec

When six maximum growth time faults are injected into the Ironbird DSPM network, only six links remain, which is the minimum number of links required for this network to function. Any additional failure must be injected into links that have already failed. Since all the links have failed at the inbound end, any new failure would occur at the outbound end and would reduce the growth time shown in Fig. 7 (each point in Fig. 7 is the accumulation of all the previous failures plus the failure listed as a label for the point). Therefore, the maximum growth time occurs at the maximum number of faults defined by Eq. (1).

Simulation

Research on systems such as the DSPM is expensive. Because of the cost, building large DSPM systems or systems with special attributes solely for research is not feasible. As an alternative, simulation software offers a means of studying larger networks with different attributes while still keeping down the cost.

The simulation software was designed to mimic the detailed growth algorithm flowchart used in the development of the DSPM. Each block or group of blocks in the flowchart was timed experimentally to establish the constituent times for the simulation. Each function in the flowchart was functionally implemented in the simulation, and the appropriate time was added to the total simulation time whenever the function was performed.

The validity of the simulation was established by exhaustively comparing the results of the simulation with the actual data from similar known situations. After many simulation runs the simulated growth time for a fault-free network was a good approximation to the actual time. The simulation then had to be validated for faulty configurations.

The simulation was designed to allow the user to fail links or nodes. Additional tests using the fault injection capability established the validity of the simulation in a faulty environment. Figure 8 demonstrates the accuracy of the simulation.

In Ref. 2, the DSPM6 network (Fig. 9) is determined to be the smallest DSPM-type network acceptable for applications requiring ultrareliability. Accordingly, DSPM6 is used as the lower bounds for DSPM network performance. Because DSPM16 (Fig. 1) appears frequently in literature (2, 3, and 4), it was arbitrarily chosen as the upper bound on DSPM performance. The area between the bounds is filled in with simulated data from DSPM8 (see Fig. 10) and DSPM11 (see Fig. 11). All the simulated networks were grown by using a heuristic algorithm to achieve near-worst case growth times.

The fault-free growth times of all the simulated networks is plotted in Fig. 12. Because of the linear relationship of fault-free growth time to the number of nodes in the network, and the linear relationship of incremental growth time and failure, a simple model of worst case growth time (G(n,f)) is possible for a network of nodes (n) and failures (f). The graphical form for such a model is shown in Fig. 13. The mathematical form is given as Eq. (2).

$$G(n,f) = 3.46n + 7.7f \text{ in msec} \qquad (2)$$

As an example, the worst case growth time for a 32-node network with five failures is

$$G(32,5) = 3.46 \times 32 + 7.7 \times 5 = 149.2 \text{ msec} \qquad (3)$$

By restricting the growth simulation to the maximum growth time scenario, the worst case times are obtained and the growth time relationship is linear with the number of failures. For all simulated networks, there is a 7.7 msec increment between failures. For instance, in DSPM6 the difference in growth time between three failures (40.2 msec actual) and four failures (47.9 msec actual) is 7.7 msec.

Given the fault-free growth time for a specific network, a good approximation of worst case growth times can be obtained with a straight line with a slope of 7.7 msec per failure. Further, the fact that the relationship is true for four networks representing three different topologies (DSPM6, DSPM16, and DSPM8/DSPM11) strongly indicates that the linear relationship is a generic DSPM characteristic.

The growth time of a network should be linearly related to the number of links that the network must grow. Because one link must be grown to every node in the network, the fault-free growth time of a 16-node network should be twice the fault-free growth time of an 8-node network. For DSPM16, the fault-free growth time is 51.6 msec, which is approximately twice the 24.0 msec fault-free growth of DSPM8.

Given 149 msec for G(32,5), or even 127 msec for G(16,10), one could question whether the DSPM

growth process is fast enough. During the growth process, no inputs or outputs can traverse the DSPM network. As a result, the aircraft would be flying open loop, and any departure would continue until the process is over and one control-law update is complete. This break in the control-law update certainly must be considered in any vehicle or system design; however, the maximum growth time is expected to be somewhat shorter given improved computational hardware that would be readily available in any operational application.

Conclusions

Tests show that the generic growth characteristics of DSPM-type systems are independent of the network topology. They also show that growth time is linearly dependent on the number of nodes and the number of failures occurring in the DSPM network. However, tests show that the growth time increases as the number of failures increase and the growth time is bounded. As a result of linear and bounded growth times, the growth time relationship can be modeled by a simple, accurate linear equation.

References

(1) "Advanced Information Processing System (AIPS) System Specification," CSDL-C-5709, Charles Stark Draper Laboratory, May 1984.

(2) Abbott, Larry W., "An Ultra Reliable Computer Communication Network — the Dispersed Sensor Processing Mesh," University of Kansas, Jan. 1984.

(3) Abbott, Larry W., "Operational Characteristics of the Dispersed Sensor Processor Mesh," Proceedings of the IEEE/AIAA 5th Digital Avionics Systems Conference, Seattle, Wash., Oct. 31-Nov. 3, 1983, pp. 9.4.1 to 9.4.8.

(4) Hopkins, Albert L. and Brock, Larry D., "Interim Report on Fault-Tolerant Aircraft Signal and Power Transmission Structures," R-1298, Charles Stark Draper Laboratory, Aug. 1978.

(5) Szalai, Kenneth J., Jarvis, Calvin R., Krier, Gary E., Megna, Vincent A., Brock, Larry D., and O'Donnel, Robert N., "Digital Fly-By-Wire Flight Control Validation Experience," NASA TM-72860, Dec. 1978.

(6) Megna, Vincent A., "Tactical Airborne Distributed Computing and Networks," Dispersed Sensor Processing Mesh Project, AGARD-CCP-303, June 1981, pp. 30-1 to 30-14.

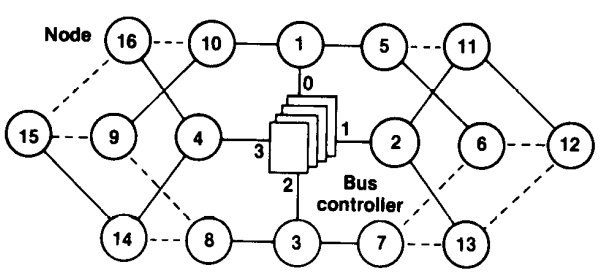

Fig. 1. DSPM16 generic DSPM-type network.

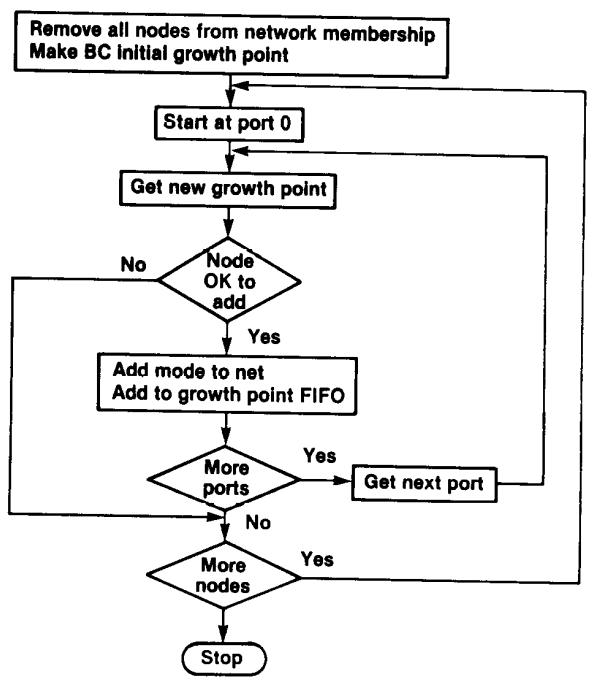

Fig. 2. Simplified growth algorithm.

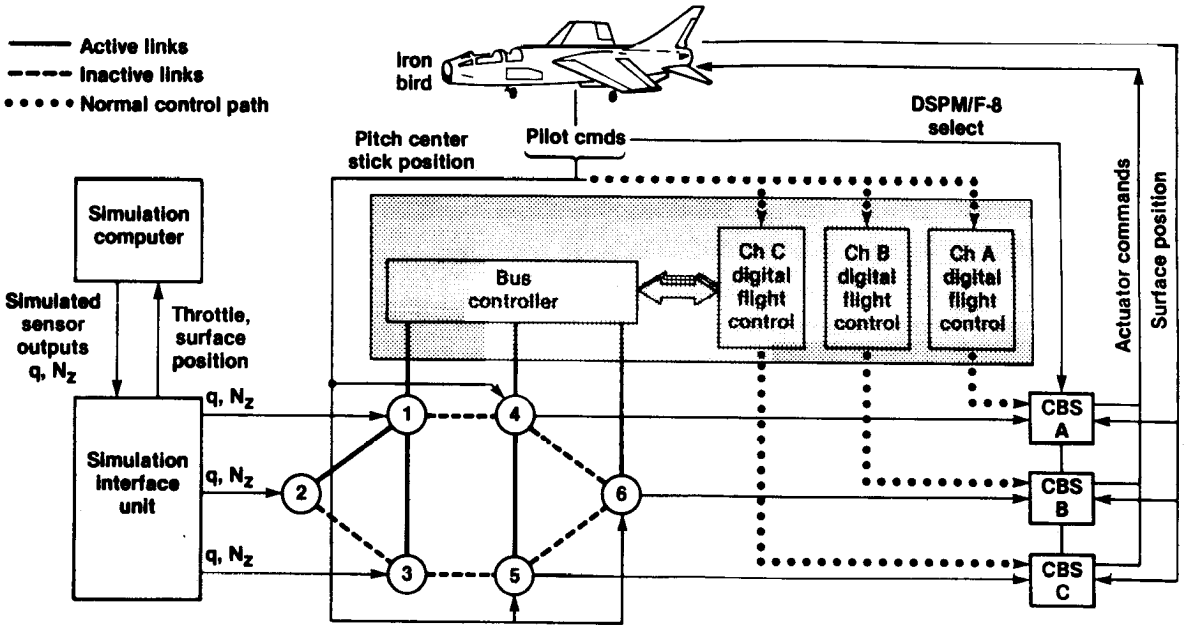

Fig. 3. DSPM experimental system.

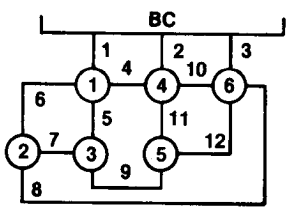

Fig. 4. Ironbird DSPM network used on the experimental system.

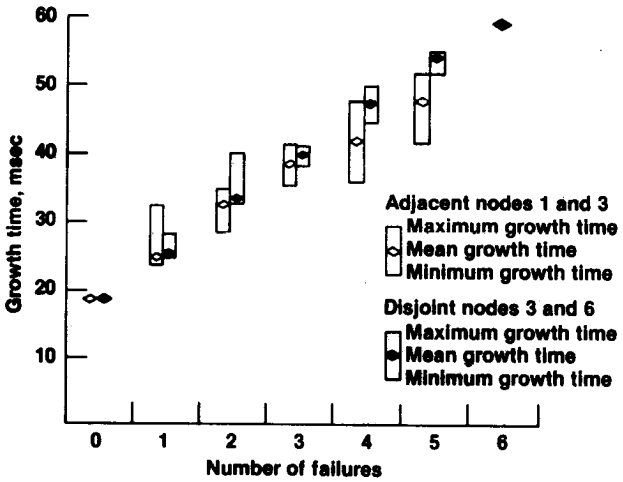

Fig. 5. Number of combined failures occurring on a node pair for two sets of node pairs.

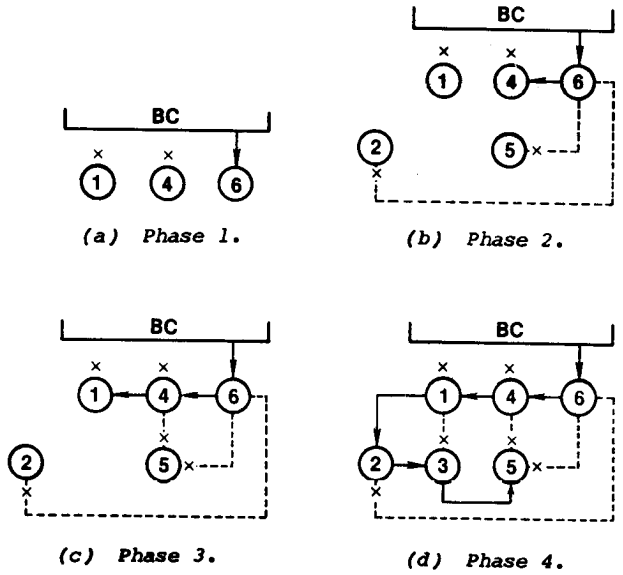

Fig. 6. Maximum growth time scenario.

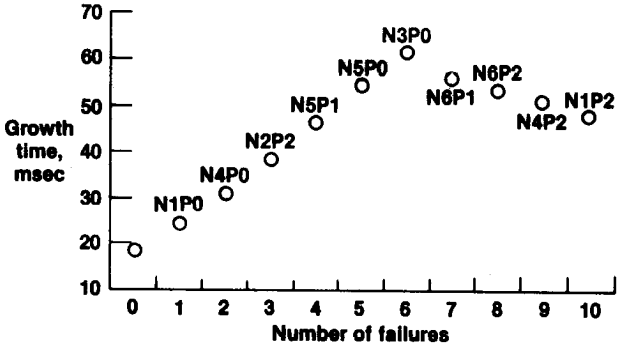

Fig. 7. Maximum growth time as a function of the number of faults.

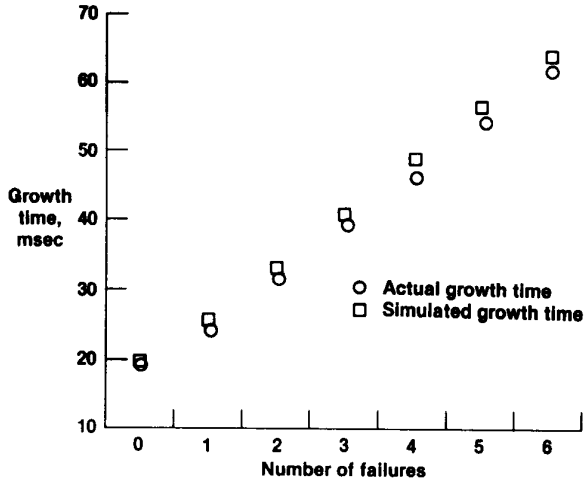

Fig. 8. Simulated versus actual growth time for a maximum growth time example (Ironbird DSPM).

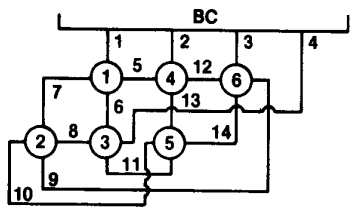

Fig. 9. DSPM6 network.

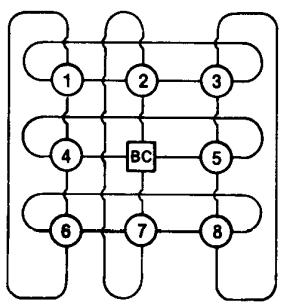

Fig. 10. DSPM8 network.

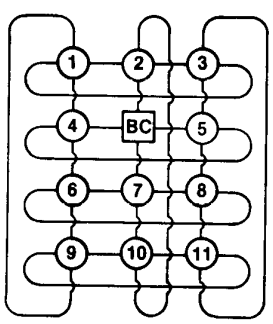

Fig. 11. DSPM11 network.

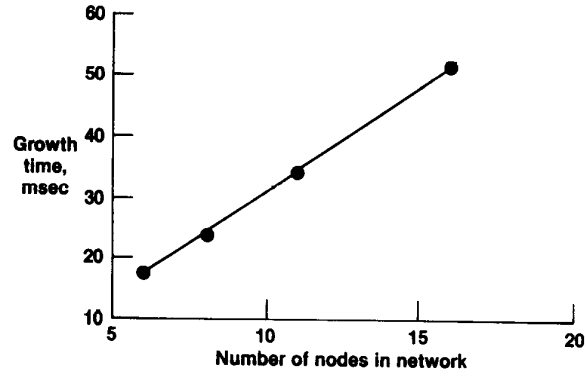

Fig. 12. Fault-free growth time as a function of number of nodes.

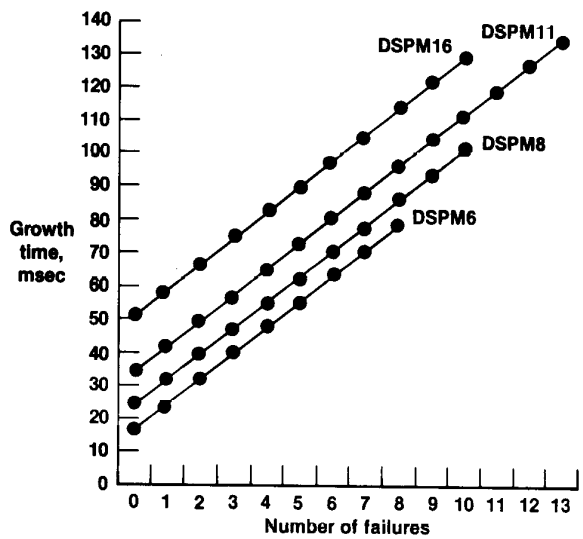

Fig. 13. Worst case growth time G(n,f) for networks with nodes (n) and failures (f).

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
NASA TM-85915		

4. Title and Subtitle	5. Report Date
Test Experience on an Ultrareliable Computer Communication Network	December 1984
	6. Performing Organization Code
7. Author(s) Larry W. Abbott	8. Performing Organization Report No. H-1268
	10. Work Unit No.
9. Performing Organization Name and Address NASA Ames Research Center Dryden Flight Research Facility P.O. Box 273 Edwards, California 93523	11. Contract or Grant No.
	13. Type of Report and Period Covered Technical Memorandum
12. Sponsoring Agency Name and Address National Aeronautics and Space Administration Washington, D.C. 20546	14. Sponsoring Agency Code 505-34-11

15. Supplementary Notes

This report was prepared as Paper 84-2649 for the AIAA/IEEE 6th Digital Avionics Conference, Baltimore, Maryland, Dec. 3-6, 1984.

16. Abstract

The dispersed sensor processing mesh (DSPM) is an experimental, ultra-reliable, fault-tolerant computer communications network that exhibits an organic-like ability to regenerate itself after suffering damage. The regeneration is accomplished by two routines — grow and repair. This paper discusses the DSPM concept for achieving fault tolerance and provides a brief description of the mechanization of both the experiment and the six-node experimental network. The main topic of this paper is the system performance of the growth algorithm contained in the grow routine. The characteristics imbued to DSPM by the growth algorithm are also discussed. Data from an experimental DSPM network and software simulation of larger DSPM-type networks are used to examine the inherent limitation on growth time by the growth algorithm and the relationship of growth time to network size and topology.

17. Key Words (Suggested by Author(s))	18. Distribution Statement
Computer network Ultrareliable Digital communication Local area network	Unclassified-Unlimited STAR category 62

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price*
Unclassified	Unclassified	7	A02

*For sale by the National Technical Information Service, Springfield, Virginia 22161.