NBSIR 81-2414

A Discussion of Gridnet Simulation Results

U.S. DEPARTMENT OF COMMERCE
National Bureau of Standards
Institute for Computer Sciences and Technology
Center for Computer Systems Engineering
Washington, DC 20234

November 1981

onsored by
efense Nuclear Agency
ashington, DC 20305

NBSIR 81-2414

A DISCUSSION OF GRIDNET SIMULATION RESULTS

R. T. Moore

U.S. DEPARTMENT OF COMMERCE
National Bureau of Standards
Institute for Computer Sciences and Technology
Center for Computer Systems Engineering
Washington, DC 20234

November 1981

Sponsored by
Defense Nuclear Agency
Washington, DC 20305
Task Code RB
Work Unit 00023

U.S. DEPARTMENT OF COMMERCE, Malcolm Baldrige, *Secretary*
NATIONAL BUREAU OF STANDARDS, Ernest Ambler, *Director*

TABLE OF CONTENTS

LIST OF FIGURES

A DISCUSSION OF GRIDNET SIMULATION RESULTS

R. T. Moore

This report is a review and evaluation of the results of computer simulation of GRIDNET conducted under U. S. Department of Commerce contract NB80SBCA0477 during the period from September 30, 1980 to September 4, 1981.

The objectives of the program were to test and modify algorithms which would permit messages to be routed from any source to any destination, in a network having thousands of nodes, and to accomplish this routing in an efficient manner using only limited local knowledge of network operability status. In addition, estimates were developed of the network's delay and throughput characteristics as a function of traffic loading.

Based on the simulation results, GRIDNET could be developed into a highly reliable and highly survivable data communications network that could support traffic between a large number of nodes in an effective and efficient manner.

Key words: Alternate routing; communications networks; distributed control; message delay; network throughput; survivability.

1. INTRODUCTION

GRIDNET is intended to provide a communications capability between a relatively large number of stations by interconnecting a number of CROSSFIRE loops together using "gateway" type stations at their intersections (1, 2). High survivability is achieved by providing many alternate routes and by using distributed control in such a way that continued operation is possible so long as any operable path exists between a message source and its intended destination. Messages are routed on the basis of limited, localized knowledge of the operability status of adjacent loops in order to minimize the bandwidth required for the exchange of routing information between gateway stations.

An initial series of routing algorithms was developed to permit messages to be routed by gateway stations that had knowledge of the operability status of only the up to four immediately adjacent loops. These algorithms were subjected to limited manual testing. The results indicated that optimal route selection occured when there were no breaks in the network, and satisfactory alternate routes were chosen when selected breaks were introduced in loops or nodes. Manual testing was too laborious to use for an exhaustive evaluation of the capabilities of these algorithms, however, so plans were made to issue a contract for a computer simulation.

In the latter part of Fiscal Year 1980, a Request for Proposals was issued. Included in the statement of work were certain requirements for the implementation of a simulation of a GRIDNET in such a way as to test the effectiveness of the original routing algorithms and to develop estimates of the performance for a GRIDNET in terms of information transfer rates and message delays as a function of network size and signalling rates. In addition, the simulation software was to be installed in one of the NBS computers so that it could continue to be used by the Government for other studies after the contract had been completed. Evaluation of the proposals received in response to this solicitation resulted in the selection of General Electric Company, Space Systems Division, Huntsville, Alabama, and contract NB80SBCA0477 was awarded.

2. INITIAL SIMULATION RESULTS

The contractor had proposed to perform the GRIDNET simulation by making appropriate modifications to certain modules of an existing simulation program. This program was the Data System Dynamic Simulator, DSDS, which they had developed under contract to the National Aeronautics and Space Administration. This was an effective approach, and it permitted them to make rapid progress. Within a little more than two months after the contract award, the original GRIDNET routing algorithms had been implemented in the simulation program. These permitted a simulated message to be directed from a source to a destination in the network. A variety of network configurations were used. These ranged from an intact network to networks with increasingly complex groups of breaks or outages to a network that was fragmented so that no possible path existed between source and destination.

2.1 ROUTING RESULTS

The initial series of tests quickly revealed the limitations of the original routing algorithms. Certain minor flaws in these algroithms were disclosed and corrected, but even after these corrections, messages were successfully routed to their destinations in only 61 out of 72 test cases where operable routes existed.

Additional rules and algorithms were then investigated. Some of these were successful in all 72 of the test cases, but the success was obtained at the cost of unacceptably long path lengths. In some of the test cases the message had to go through several hundred loops in order to reach a destination that could have been reached by traversing only a couple of dozen loops using global knowledge of network operability. It was also possible to design configurations of breaks that were more complex than those in the test cases, and that would defeat the algorithms. These would prevent a message from reaching its destination even though there was an operable path to that destination. The performance of these algorithms was clearly unacceptable and led to the conclusion that satisfactory alternate routing around network breaks could not be realized without using operability status information that was derived from an area larger than covered by the next adjacent loops.

2.2 THROUGHPUT AND DELAY PERFORMANCE

Following the demonstration of the limitations of single loop look-ahead routing, simulation efforts were concentrated on developing estimates of the throughput and delay characteristics of a GRIDNET. This decision was based on the assumption that favorable performance in these areas might serve as justification for expending additional resources to develop better routing algorithms. The new routing algorithms would be based on the use of operability information from both the next adjacent and the second adjacent loops. This would require the use of additional bandwidth to disseminate status reports on the operability of second adjacent loops, and would also involve a greater data processing load at the gateway stations to compute alternate routes around network breaks. Unless the overall network throughput and delay characteristics were favorable, this additional load might be unacceptable.

As a first step, the 12-station single loop model shown in Figure 1 was implemented and tested. The model had four gateway stations, and traffic was simulated by creating 8,Ø48 bit messages at each of the gateways and directing these messages randomly to the other three gateways. At a one megabit signalling rate, the loop handled 73 messages during 7ØØ milliseconds of simulation time providing a line utilization of more than 8Ø%. Under this loading, the delay per message was 187.3 milliseconds. At about one third of this loading, the message delay was just over 3Ø milliseconds. These results were judged to be sufficiently promising to justify additional effort on the development of an acceptable routing algorithm, and the contract with General Electric was extended to cover this increased scope of work.

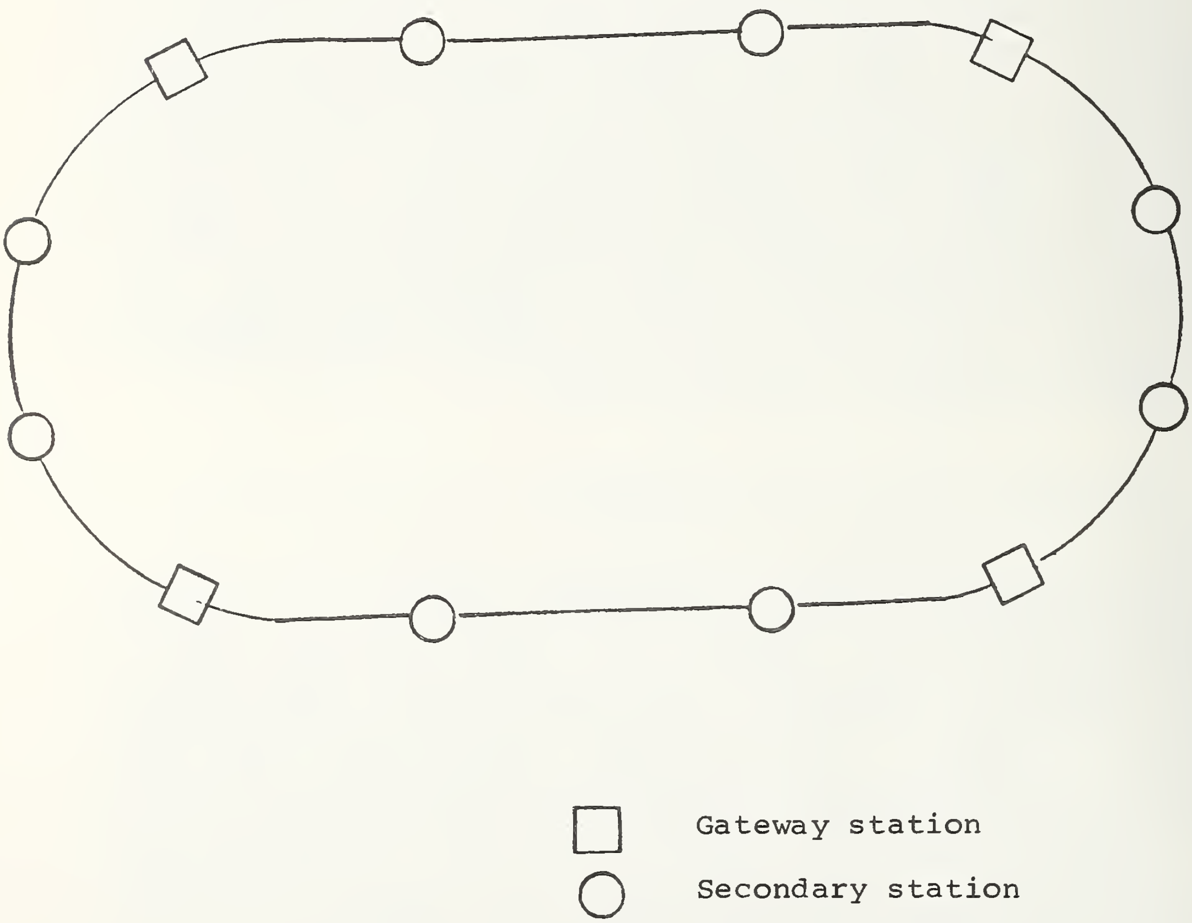

Figure 1. 12-station single loop model.

3. FINAL SIMULATION RESULTS

The estimates of performance, particularly those relating to loading capacity, throughput and delay, that were developed during the first phase of the contractual effort were used as a benchmark against which the results of the second phase were to be judged. In particular, the contractual goals stipulated that these benchmark values should not be degraded in accomplishing the following routing objectives:

i) Every message shall reach its destination if there is an operable path to that destination.

ii) A message that must be routed indirectly to its destination because of cuts introduced in the network by test scenarios that were previously used shall not be required to traverse an average of more than 30% additional loops as compared to those that would be required to establish the shortest path based on global knowledge of network operability.

iii) The routing algorithms shall be capable of handling both individual, group, and global addressing.

These targeted improvements in routing performance were to be sought using operability information derived from look-ahead to the second adjacent loops. This information was made available to each gateway on a loop by increasing the size of the status report that was used in the adjacent loop look-ahead algorithms. The increase in the size of the information field in these status reports is three octets (24 bits) for each of the second adjacent loops.

3.1 ROUTING RESULTS FROM PHASE TWO EFFORTS

Three sets of routing algorithms were developed, tested and evaluated under the Phase Two simulation effort.

The first of these was successful in routing a message to its destination in 71 out of the 72 test cases using an average of 53.5% more loops than would have been required with global operability knowledge. This was not satisfactory.

A second set of algorithms did slightly better. It was successful in all 72 of the test cases, and used an excess of 56.4% of loops, but it was still possible to devise configurations of network breaks (more subtle than the benchmark test cases) that would foil the algorithms.

A third and final set of algorithms were finally developed that met the routing goals. In all 72 test cases, a path was found to the destination that required an average of only 25% excess loops. In the worst, single case of the 72 instances, only 50% excess loops were required. In no configuration that has been devised and tested to date has it been possible to foil the algorithms and cause a message to fail to be delivered to its destination so long as there was an operable path through the network to that destination. These algorithms employ a path predictor routine that calculates a path from the current position to the destination. Then this calculated path is subjected to an optimization routine that examines the predicted path to see if any point along the path returns and approaches closer to the present position. If it does, and there is no obstruction, a new path is calculated directly to that point of "closest approach." This process is repeated until the destination is reached.

3.1.1 Group and Global Addressing

Group and global addressing was accomplished in the model by establishing addressing domains. Essentially, all stations within an upper and lower address bound, expressed in terms of both of its two principal addresss components, can be directed to receive a message. The message is routed to the nearest point within this domain where delivery is made to one station, and one or more replica messages are generated and directed to the next nearest neighbors within the domain. The process of delivery and replication is continued until delivery has been made to all stations within the designated domain. This process is efficient in bandwidth utilization as compared to the process of dispatching an original message to each intended recipient, however, it could not be used if the intended recipients do not occupy a contiguous domain.

The model did not accomodate the use of group and global addressing under circumstances where alternative routing was required to go aroung breaks or outages in the network. There does not appear to be any obvious reason why this capability could not be realized; there was just not sufficient time to address the question under the Phase Two contract.

3.2 PHASE TWO RESULTS ON DELAY AND THROUGHPUT

Two network configurations were employed to develop further estimates of the traffic handling capabilities of a GRIDNET. These are shown in Figure 2.

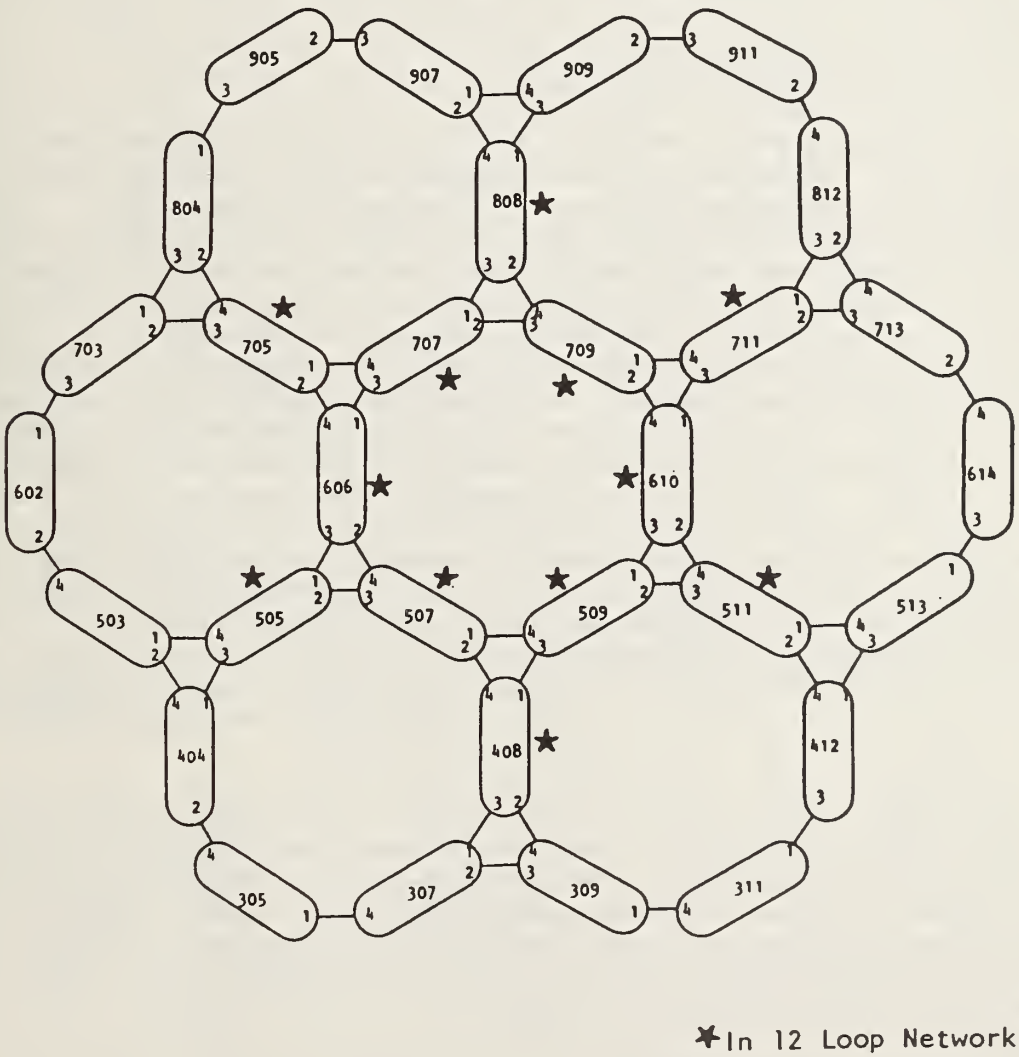

Figure 2. GRIDNET 12 Loop and 3Ø Loop Network

One was a 12 loop model where six loops were arranged in a connected hexagonal shape. Six additional "stub" loops were interconnected, one at the juncture of each pair of loops in the inner hexagonal structure. Messages were generated randomly and transmitted from one stub loop to another via the loops in the inner hexagon. Performance estimates were developed from the traffic handled by the inner loops.

Then, this network was enlarged to 3Ø loops by adding 18 additional loops to the "stubs" in such a way as to form seven interconnected hexagons. Again, traffic in the form of 8,Ø48 bit messages was introduced into the network from random sources addressed to random destinations. Comparison of the throughput and message delay and line utiliztion of the six inner loops in this larger network provided values that were quite comparable to those developed with the 12 loop network.

3.2.1 Loop Throughput Capacity

The throughput capacity of loops in a GRIDNET was established by increasing the rates at which messages were introduced into the network until instability occurred. This instability was evidenced when the number of messages in the network failed to reach some stable maximum value. Load induced instability causes the number of messages within the network to continue to increase with the passage of time. This in turn results in continuously increasing values of message delay, and shortly thereafter, total system collapse. In the simulations that were conducted using both the 12 and the 3Ø loop model, saturation occurred when the inner loops handled approximately 9Ø messages per second when the message length was 8,Ø48 bits with an operating bit rate of one megabit per second. At this throughput the line utilization was calculated to be approximately 84% and the loop delay was approximately 125 milliseconds. These results assume propagation delays associated with a 1ØØ km loop length and a 5ØØ microseconds delay to perform message routing at each gateway. These values represent the approximate upper bounds of network capability.

In both the 12 and the 3Ø loop models, if the loading is reduced to 7Ø% line utilization the loop delay is reduced to about 4Ø millisecond and a throughput of about 75 messages per second can be achieved on the inner loops. The simulation results indicate that the throughput, delay and utilization characteristics of the inner loops are not dependent upon the network size.

3.2.2 Buffer Requirements

In the simulation runs that were used to determine the maximum loading capacity of the 3Ø loop model, estimates were also developed for the number of message buffers that would be required at the gateway stations of the inner loops. At the maximum rate of message input that could be accommodated and still maintain network stability, approximately 5Ø buffers would be required at each gateway station-pair. With a 3Ø% reduction from maximum loading, the number of buffers required is reduced to less than 15 at each gateway pair.

4. EVALUATION OF SIMULATION RESULTS

Based on the simulation results reported from the study (3), it appears that the principal survivability and efficiency goals of the conceptual GRIDNET system can be achieved. Messages can be routed to their destinations by reasonably efficient paths so long as an operable path to the destination exists. This routing is accomplished without the use of a large amount of the network bandwidth for the exchange of operability status information. There is distributed control of network operation and no crucial functions are concentrated in any station or even a small group of stations. Delays and throughput as a function of line utilization are consistent with those that would be anticipated from queuing theory. The parameter values developed by the simulation would be changed with loop sizes or message lengths other than those used, but to a first approximation, they establish a base for future reference.

In evaluating the simulation results, it is perhaps useful to identify a number of the things that the simulation did not not attempt to cover or to develop information about. These are things that will affect network performance, but were not considered appropriate subjects for study until after the more fundamental characteristics of the GRIDNET had been examined.

The simulation did not completely model the communication control protocol. Bit level message simulation with error detection and retransmission was not conducted. The proposed GRIDNET protocol provides for continuous polling on any loop when actual message traffic is not being handled. Early phase one work quickly revealed that simulation at this level of detail would require an unacceptably large amount of computer time. It was necessary to make certain simplifying assumptions regarding status within an assumed polling cycle and provide a time allowance for the completion of that cycle when it became necessary to transfer a message. These assumptions reduced the computer processing time from several thousand times real-time to about 100 times real-time.

The simulation did not accommodate partial outages such as the failure of a single gateway half-station, or a single fiber of the basic CROSSFIRE loop. As a consequence it did not accommodate the logic to detect and localize failures of this type. In the model that was used, the smallest unit of failure was a complete loop including all the stations on that loop.

The routing procedures were not completely integrated with the communication control protocol. Certain information that is necessary to perform message routing must be inserted into and extracted from designated subfields in the message header. What are believed to be reasonable assumptions for processing times and header sizes were substituted for the bit level modeling that would have been required to simulate at this extent of detail. Modest changes to the assumed values were introduced and these caused only small changes in the performance characteristics suggesting that this is not an area of extreme sensitivity.

Routing capabilities were not completely and exhaustively tested by directing a message from every possible source to every possible destination in the face of every possible combination of outages. It is believed that the testing that was done was adequate in scope to have detected any problem or weakness in the final routing algorithms if there had been one present. Also, alternate routing capabilities were not tested using network configurations larger than 150 loops, a size that could accommodate on the order of 3,000 stations.

The group and global addressing simulations were not conducted in conjunction with alternate routing to accommodate outages, nor were they completely integrated with the communications control protocols.

Simulations of loading to develop estimates of delay, throughput, line utilization and buffer requirements were not conducted using network configurations larger than 30 loops, or with message sizes with an information field of other than 8,000 bits. Comparison of 12 loop and 30 loop network results suggests that the number of loops in the network should not be a major factor in affecting performance. On the other hand, message size and loop length can be expected to strongly influence the values of the performance parameters.

5. CONCLUSIONS AND RECOMMENDATIONS

The overall results of the simulation are judged to have met the planned objective of determining whether or not the GRIDNET concept could provide a highly survivable data communication capability between a very large number of stations without devoting an unacceptably large portion of the available network bandwidth to the exchange of routing information. It provides positive evidence that this should be possible together with rough approximations of the performance characteristics that might be expected from such a network. It has also helped to identify areas that will require further work in order to demonstrate the feasibility of the concept by means of a functional, operating model.

It is recommended that a model of GRIDNET be implemented in hardware and software to provide a positive demonstration of the feasibility of the concept and to establish more definitive values for the parameters characterizing its performance.

6. REFERENCES

1) Moore, R. T., et al., Phase II Final Report Computerized Site Security Monitor and Response System, NBSIR 79-1725, September 3Ø, 1978, PB 249 343, NTIS, Springfield, VA 22161.

2) Moore, R. T., GRIDNET, October 3Ø, 198Ø, NBSIR 8Ø-2149, PB81-14437Ø, NTIS, Springfield, VA 22161.

3) Graf, H. A., et al., GRIDNET Simulation Final Report, Volume 1, System Description and Results, General Electric Report No. 81HVØØ8, 7 August 1981, PB , NTIS, Springfield, VA 22161.

7. GLOSSARY OF TERMS

Gateway. A Gateway is the facility that permits a message to be transferred from one loop to another. It consists of two stations that are each able to access a common buffer memory. One of the stations is on one loop and the other station is on an adjacent loop. Each of the stations is capable of functioning both as a primary station and as a secondary station, and each of them are considered to be half of a gateway. When a message is to be transferred from one loop to the other, the sending half-gateway places it in the common buffer memory and notifies the receiving halfgateway. The receiving half-gateway removes the message from the common buffer memory and forwards it along the path toward its destination.

Line Utilization. Line Utilization is the ratio of bits used to convey user information divided by the total of the bits used for user information and overhead functions such as control and supervisory sequences.

Loop. A Loop is a closed, circular, data communications configuration connecting a Primary Station with one or more Secondary stations. In the GRIDNET, dual Loops are used and the same data is transferred around one of the pair in a clockwise direction and around the other in a counter-clockwise direction. Gateway stations interconnect Loops and a message originating at a station on one loop may be routed through several intermediate Loops before reaching its destination station on a distant Loop.

Primary Station. A Primary Station issues commands to the Secondary Stations on a Loop. The Primary Station controls the movement of message traffic. In GRIDNET the role of Primary Station is rotated among the Gateways. These may number from one to four depending upon the configuartion of the network.

Route. Route is the path or sequence of Loops that a message may traverse in going from its source to its destination.

Secondary Station. A Secondary Station responds to commands that are issued by the Primary Station. A Gateway functions as a Secondary Station whenever it is not the acting Primary Station. A Loop may have up to perhaps 20 additional Secondary Stations that are not Gateways.

NBS-114A (REV. 2-80)

U.S. DEPT. OF COMM. BIBLIOGRAPHIC DATA SHEET (See instructions)	1. PUBLICATION OR REPORT NO. NBSIR 81-2414	2. Performing Organ. Report No.	3. Publication Date October 1981

4. TITLE AND SUBTITLE

A Discussion of Gridnet Simulation Results

5. AUTHOR(S)

R. T. Moore

6. PERFORMING ORGANIZATION (If joint or other than NBS, see instructions)

NATIONAL BUREAU OF STANDARDS
DEPARTMENT OF COMMERCE
WASHINGTON, D.C. 20234

7. Contract/Grant No.

NB80SBCA0477

8. Type of Report & Period Covered

9/30/80 - 9/4/81

9. SPONSORING ORGANIZATION NAME AND COMPLETE ADDRESS (Street, City, State, ZIP)

Defense Nuclear Agency
Washington, D.C. 20305

10. SUPPLEMENTARY NOTES

☐ Document describes a computer program; SF-185, FIPS Software Summary, is attached.

11. ABSTRACT (A 200-word or less factual summary of most significant information. If document includes a significant bibliography or literature survey, mention it here)

This report is a review and evaluation of the results of computer simulation of GRIDNET as conducted under U.S. Department of Commerce contract NB80SBCA0477 during the period September 30, 1980, to September 4, 1981.

The objectives of the program were to test and modify algorithms which would permit messages to be routed from any source to any destination, in a network having thousands of nodes, and to accomplish this routing in an efficient manner using only limited local knowledge of network operability status. In addition, estimates were developed of the network's delay and throughput characteristics as a function of traffic loading.

Based on the simulation results, GRIDNET could be developed into a highly reliable and highly survivable data communications network that could support traffic between a large number of nodes in an effective and efficient manner.

12. KEY WORDS (Six to twelve entries; alphabetical order; capitalize only proper names; and separate key words by semicolons)

Alternate routing; communications networks; distributed control; message delay; network throughput; survivability.

13. AVAILABILITY

☒ Unlimited
☐ For Official Distribution. Do Not Release to NTIS
☐ Order From Superintendent of Documents, U.S. Government Printing Office, Washington, D.C. 20402.
☒ Order From National Technical Information Service (NTIS), Springfield, VA. 22161

14. NO. OF PRINTED PAGES

15

15. Price

$5.00

USCOMM-DC 6043-P80