ESD-TR-86-196

Semiannual Technical Summary

Distributed Sensor Networks

30 September 1986

Lincoln Laboratory

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

LEXINGTON, MASSACHUSETTS

Prepared for the Defense Advanced Research Projects Agency under Electronic Systems Division Contract F19628-85-C-0002

The work reported in this document was performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology. This work was sponsored by the Defense Advanced Research Projects Agency under Air Force Contract F19628-85-C-0002 (ARPA Order 3345).

This technical report has been reviewed and is approved for publication.

FOR THE COMMANDER

Thomas J. Alpert, Major, USAF
Chief, ESD Lincoln Laboratory Project Office

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LINCOLN LABORATORY

DISTRIBUTED SENSOR NETWORKS

SEMIANNUAL TECHNICAL SUMMARY REPORT
TO THE
DEFENSE ADVANCED RESEARCH PROJECTS AGENCY

1 APRIL — 30 SEPTEMBER 1986

ISSUED 14 MAY 1987

LEXINGTON MASSACHUSETTS

87-7-14 041

ABSTRACT

This report describes the work performed on the DARPA Distributed Sensor Networks Program at Lincoln Laboratory during the period 1 April through 30 September 1986.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

DISTRIBUTED SENSOR NETWORKS

I. INTRODUCTION AND SUMMARY

This is the final Semiannual Technical Summary report of the Distributed Sensor Networks (DSN) program. The program was aimed at developing distributed target surveillance and tracking methods for systems employing multiple spatially distributed sensors and processing resources. Such systems would be made up of sensors, data bases, and processors distributed throughout an area and interconnected by an appropriate digital data-communication system. The working hypothesis of the program was that through netting and distributed processing, the information from many sensors could be combined to yield effective surveillance systems. The overall concept called for a mix of sensor types as well as geographically distributed sensors.

Surveillance and tracking of low-flying aircraft with ground-based acoustic and imaging sensors was used to develop and evaluate DSN concepts in the light of a specific problem. An experimental DSN test bed system was developed and has been used to test and demonstrate DSN techniques. Small arrays of microphones providing directional information were employed as acoustic sensors and visible TV cameras were used as imaging sensors in the test bed system.

The primary accomplishment during this final report period was the demonstration of distributed real-time tracking using both TV and acoustic sensors. Tracking was implemented as a geographically decentralized confederacy of autonomous cooperating nodes. Thus the feasibility of this organization has been established for a DSN system containing multiple sensor types as well as distributed nodes. The demonstrations involved tracking a UH-1 helicopter using six distributed sensor nodes, four acoustic nodes, and two TV nodes. For each of several demonstration runs the helicopter was successfully detected and placed in track by the acoustic nodes operating in surveillance mode. TV subsystems, with narrow fields-of-view but more precise measurement capabilities, were then alerted and used to provide additional azimuth measurements to the tracker. The overall track of the aircraft was obtained by distributed algorithms that integrated all the acoustic and TV measurements. The experiments and results are described in more detail in Section II of this report.

Section III summarizes subsystem development and integration tasks that were completed during this reporting period to support the real-time distributed tracking experiments. Accomplishments include completion of a new real-time acoustic signal-processing subsystem, development of new target-detection and sensor-control algorithms for the TV subsystem, completion of the communication subsystem to allow remote deployment of portions of the test bed, and the integration of all of these elements with the tracker and new situation display capabilities.

Final briefings were given to DARPA in September concerning Lincoln Laboratory DSN accomplishments, lessons learned, and remaining research tasks. Overall project accomplishments cited in those briefings included: (a) validation of the autonomous cooperating nodes architecture;

(b) development of distributed algorithms to perform sensor independent tracking, acoustic tracking, TV tracking, large-area track integration, and system self-location; (c) development of new acoustic direction-finding methods; and (d) implementation and use of a test-bed system to develop and demonstrate distributed tracking methods. Section IV of this report summarizes the material presented in those briefings.

II. TEST-BED DEMONSTRATIONS

A major goal of the DSN project at Lincoln Laboratory has been the demonstration of real-time distributed aircraft tracking with a combination of acoustic and TV sensors. This goal was achieved during this report period.

The experimental test-bed elements used for tracking demonstrations are shown in Figure II-1. The essential elements that were used to perform distributed tracking consisted of five tracking nodes, four acoustic subsystems, and two TV subsystems. These elements were interconnected by means of Ethernets and microwave radio links. The Ethernets and radios were used to implement both broadcast and Point-to-Point (PTP) communications between all elements of the system. This was accomplished by means of special gateway nodes that created a single logical Ethernet using the radios to forward messages between three separate Ethernets.

Three acoustic subsystems, a TV subsystem, and user workstations were all on a single local Ethernet at the main Lincoln Laboratory building complex. These sensor subsystems and the user workstations were located within 500 m of each other. The other sensors, one acoustic subsystem and one TV subsystem, were located approximately 3 km distant at the Lincoln Flight Facility and at a Raytheon site on the other side of Hanscom Field.

Figure II-2 shows the geographic deployment of sensor sites as well as additional information about the demonstration scenario. A UH-1 helicopter was flown through the deployed sensor field as a test target. The demonstration concept called for the helicopter to be acquired and placed in track using microphone data. The initial tracks would then be used to cue the TV subsystems which would acquire the target and provide additional directional measurements to be integrated with acoustic measurements. Error ellipses, which are produced by the tracking algorithms, are indicated in the figure. Initially, the axes of the error ellipses are several hundred meters long. As more data are acquired, especially when the TV data are added, the error dimensions are substantially reduced. The expected large decrease in the error ellipses when TV measurements are added is largely because the TV measurements in the test bed are accurate to a fraction of a degree whereas the acoustic measurements are accurate to only within a few degrees. Later in the track the error ellipses begin to increase again as the target leaves the sensor field and new measurements are no longer provided.

Figure II-3 shows the track obtained for one pass of the UH-1 helicopter during final test-bed demonstrations. Several similar passes were made, some from North to South and others from South to North, all with similar results. The choice of flight paths was very limited by constraints imposed by air traffic control personnel at the Hanscom Field tower. A comparison of Figures II-2 and -3 shows that the expected qualitative performance of the distributed system was achieved. The track was established using data from the acoustic sensors and was subsequently improved by jointly tracking with acoustic sensors and TV sensors.

The display of Figure II-3 is an example of the real-time display provided for users during demonstration runs. Track points are added to the display in real time during the experiment.

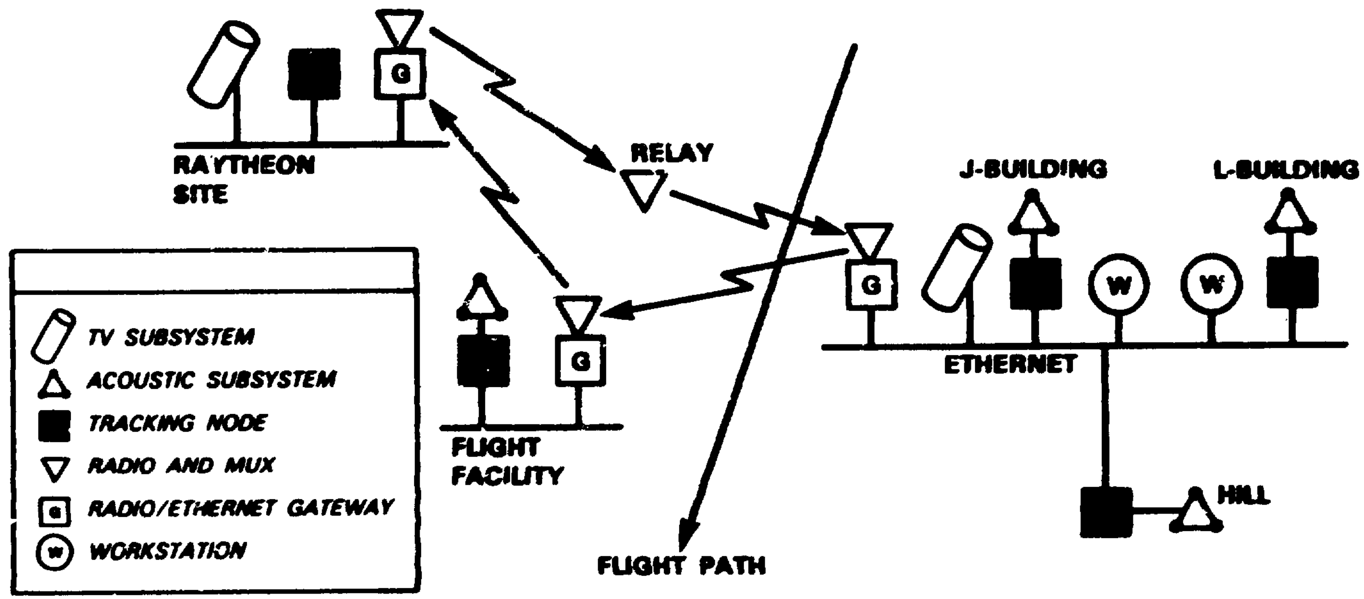

Figure II-1. Experimental test bed for DSN demonstrations.

78433-1

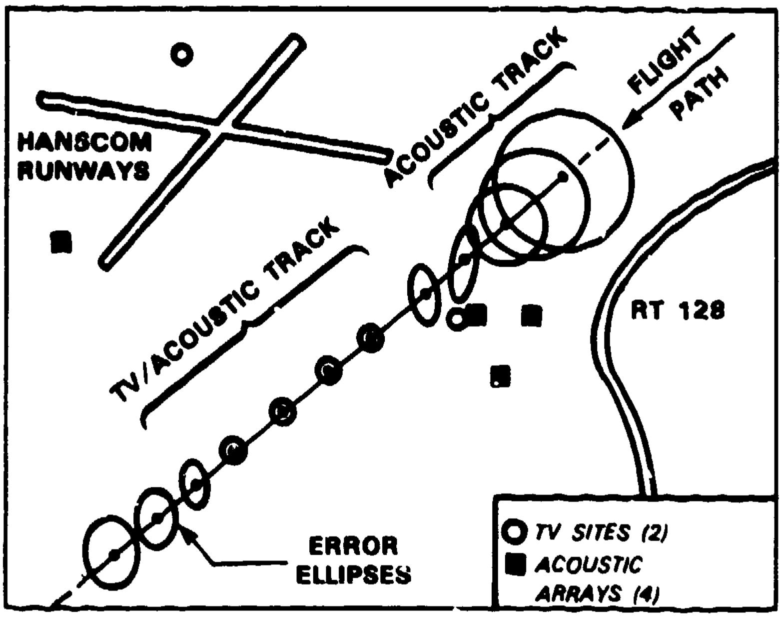

Figure II-2. Demonstration scenario for distributed mixed sensor tracking.

78433-2

Each point shows position, direction, and an error ellipse. A finite history, typically one or two minutes, of each track is retained on the screen. Track data for the situation display are obtained via a PTP communication link between the user workstation and any of the tracking nodes in the test bed. Although these data are obtained from a single node, they represent the overall track obtained using all the test-bed sensors. This results from the nature of the distributed tracking algorithms and the fact that the test bed was operated with all nodes directly interconnected by broadcast communication links. Larger DSN systems, with many more nodes and less communication connectivity, would require additional multisite data collection and integration. Methods to accomplish this, based upon track combining algorithms that are already included in the distributed tracking algorithms, are available.

Additional sensor-specific displays have been developed for the acoustic and TV subsystems in the test bed. The acoustic displays are azimuth vs time histories. Data for these displays are obtained in real time using PTP communication links from any of the acoustic sites to a workstation where they are displayed. The TV subsystem displays are TV images such as that shown in Figure II-4. That figure shows the screen of a TV monitor attached to our J-Building TV subsystem and located so that test bed users can view it. The display shows information from both the local TV subsystem on the roof of J-Building and from the remote TV subsystem. The primary image is the image last obtained by the local TV. If a target was detected in the image, then a box is superimposed upon the screen to indicate the detection; that is box (a) in the figure. If, in addition, the remote site detected a target, a small image around the detection is saved, transmitted across the network, and superimposed upon the local TV screen. This image is a binary image and is image (b) in the figure. This ability to display an image from the remote TV is a recent addition to the system that demonstrates an additional form of remote access to DSN data by means of PTP communication links.

79433-3

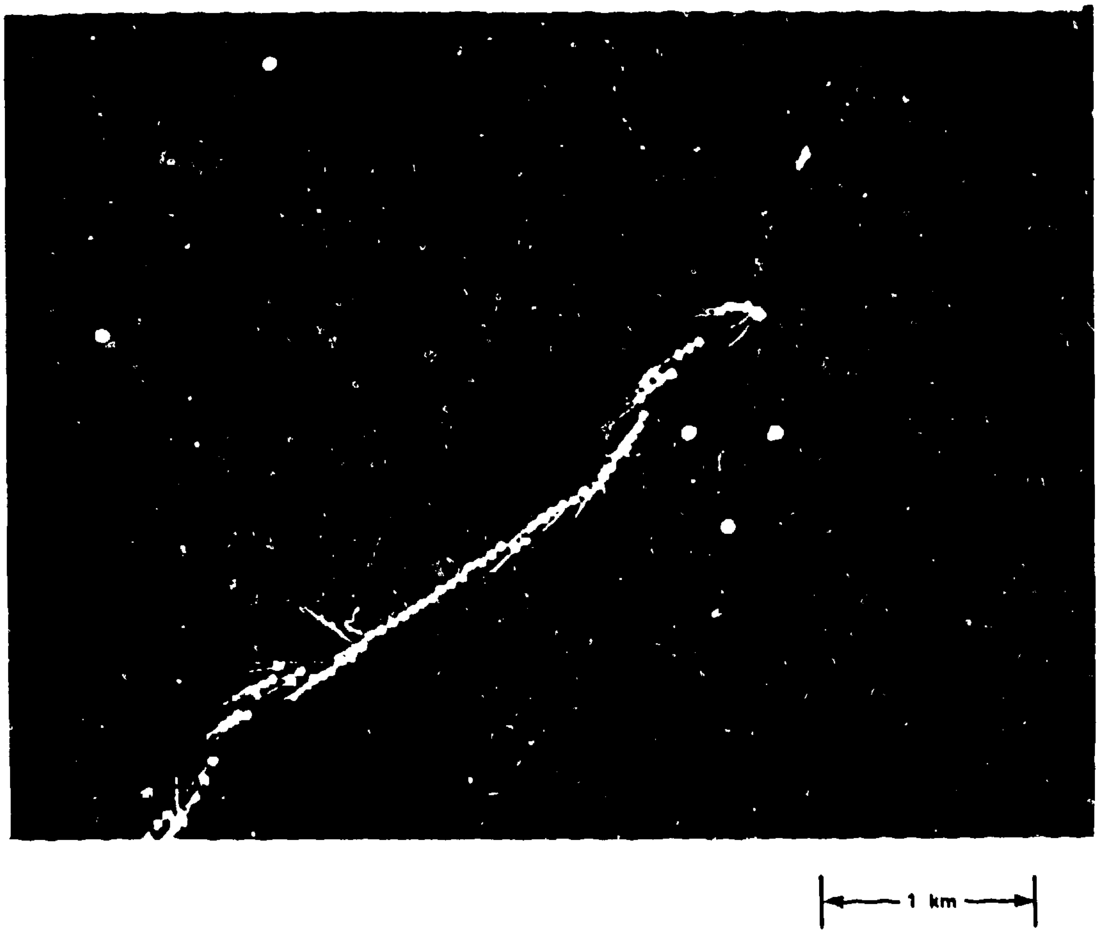

Figure II-3. Representative real-time tracking results for UH-1 helicopter.

79433-4

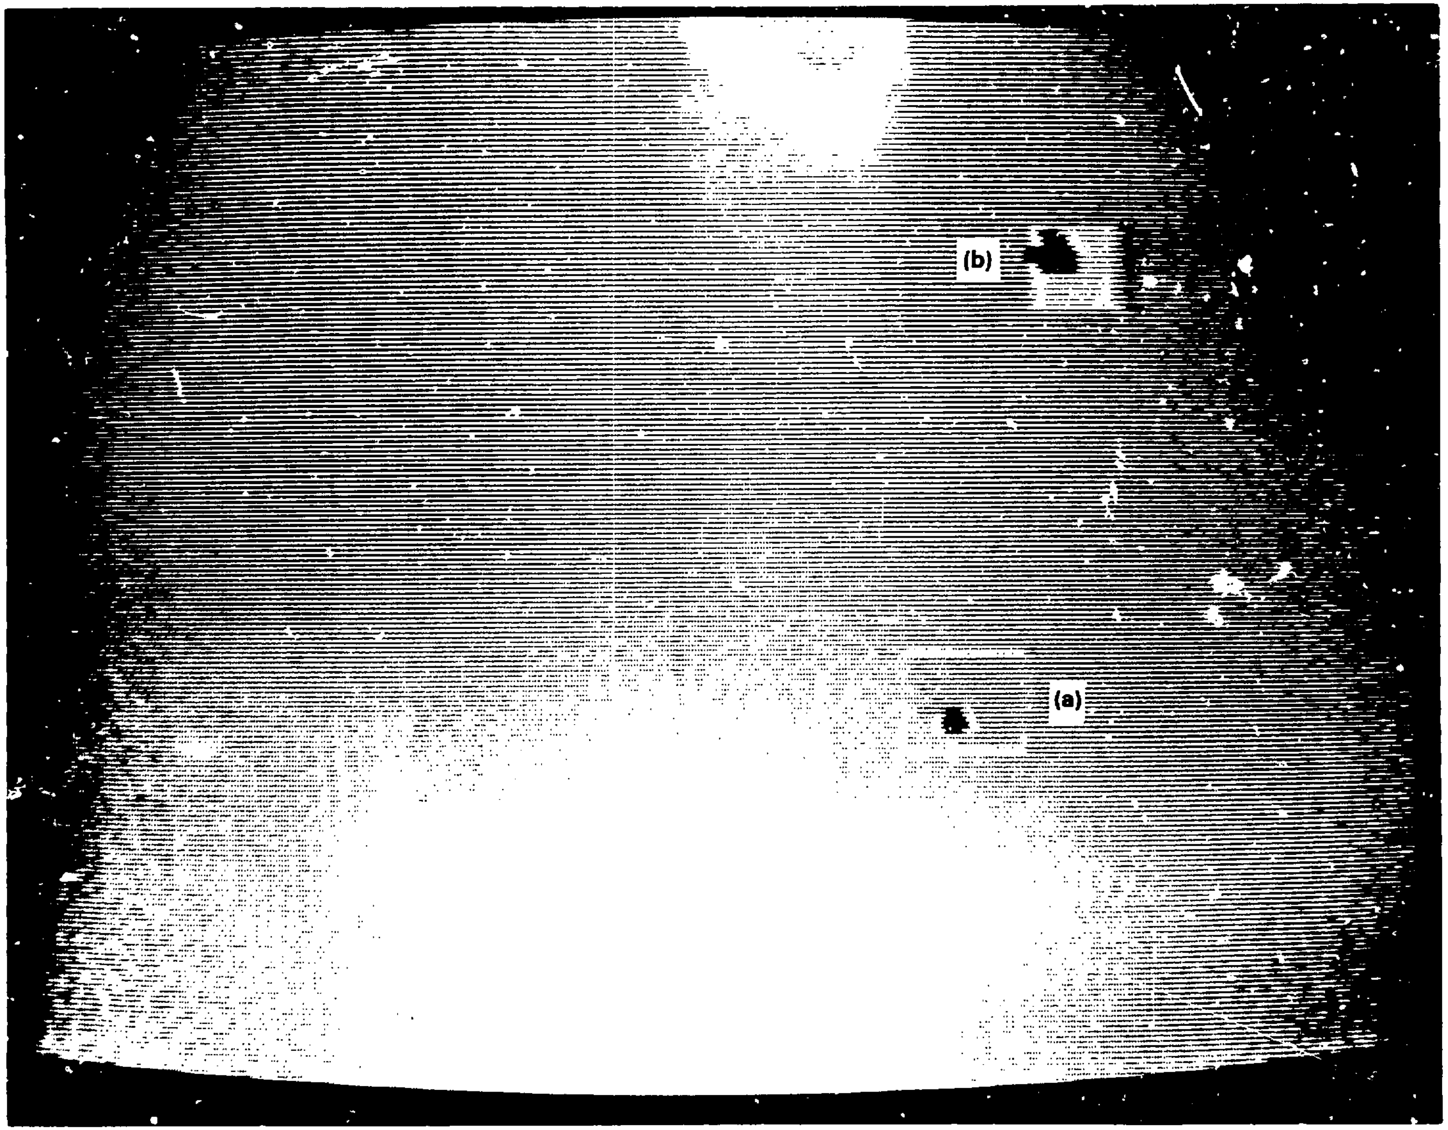

Figure II-4. Photo of user TV display showing (a) helicopter detected by local camera and (b) superimposed helicopter detected by remote camera.

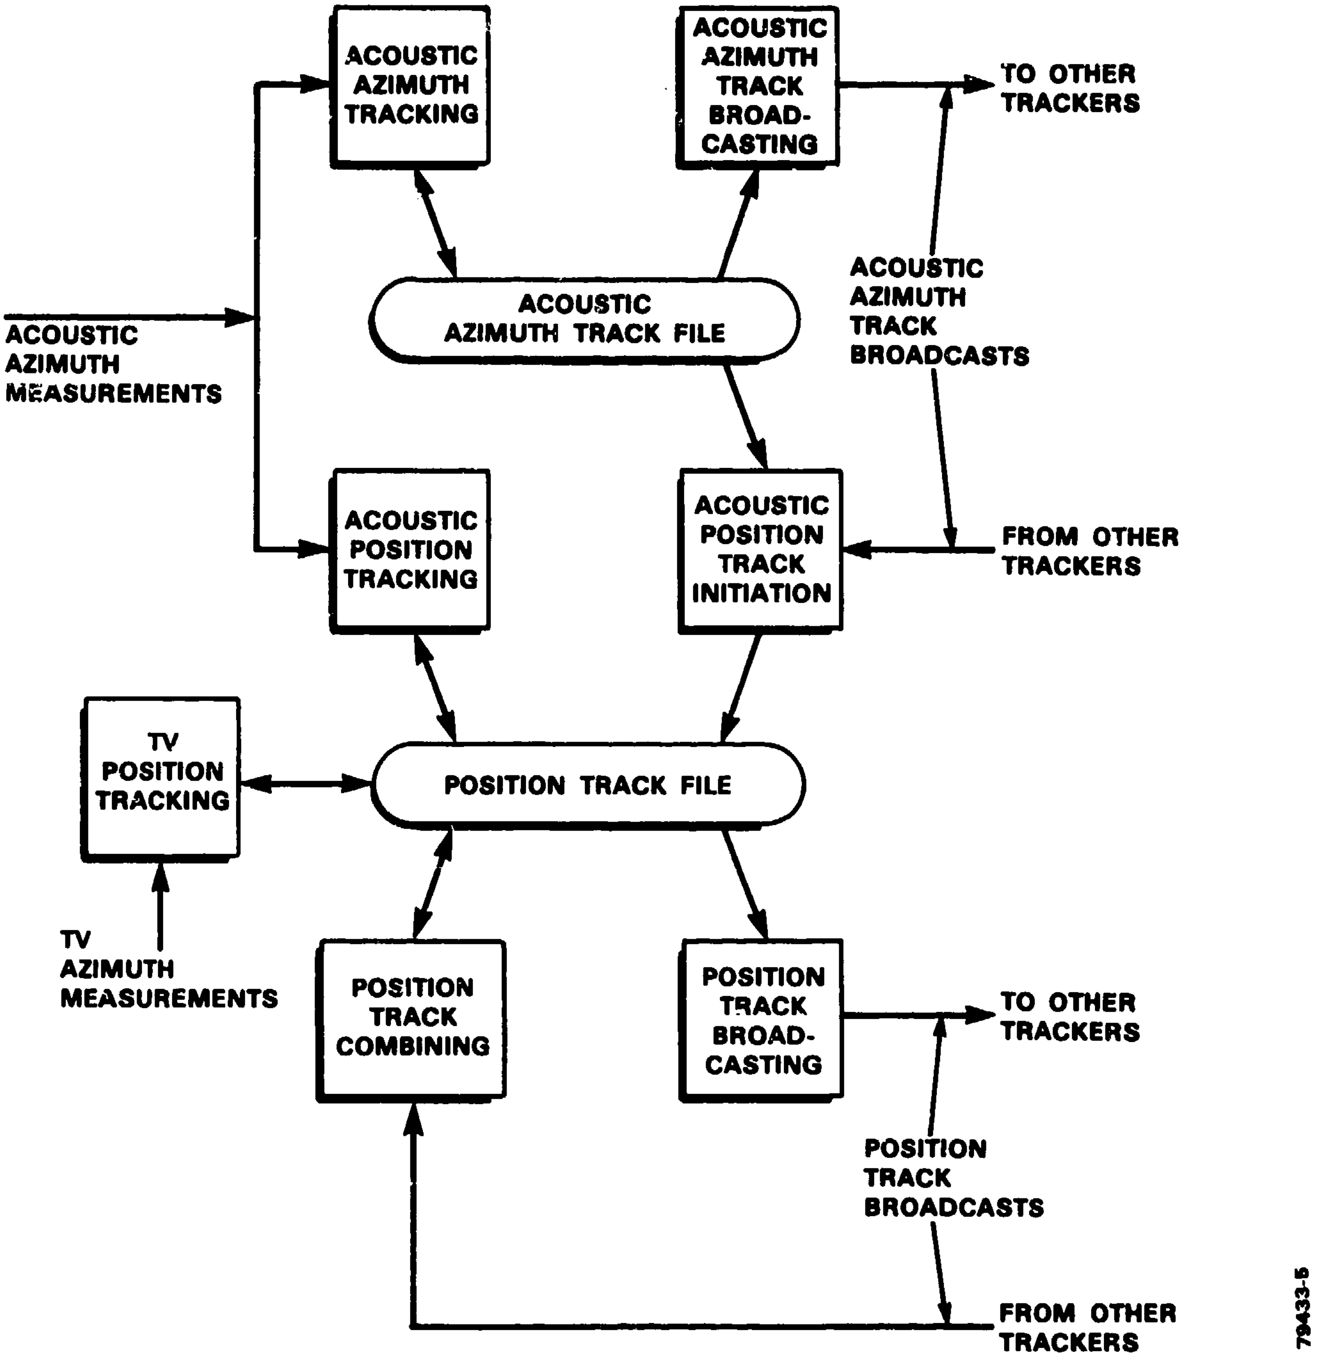

Figure II-5. Distributed mixed sensor tracker organization.

All the tracking functions in the test bed are implemented in the form of distributed cooperative autonomous processes. Figure II-5 shows the overall structure of the tracking algorithms. The elements shown in the figure are duplicated at every tracking node in the system. Each node operates independently and asynchronously, driven by data from local sensors connected to the tracker and by data received from other sensors via broadcast communications. The tracker may be connected to a local acoustic subsystem, a local TV subsystem, both or neither.

Each tracker broadcasts two kinds of tracking formation for the use of other trackers. The essential broadcasts that would be part of any DSN system, independent of sensor types, are the position track broadcasts. Whenever a node updates a track using sensor information available only to that node, it must broadcast the update to other nodes. The broadcast may be immediate or may be deferred to transmit only the cumulative effect of many updates. However, it must be done for the system to operate. The other broadcasts, acoustic azimuth track broadcasts, are specific to a system that employs acoustic measurements for the purpose of position track initiation.

The heart of the tracker is the position track combining function which is completely sensor independent. It combines local target position tracks with those received from other nodes. The algorithm generates an optimal mix of the local and foreign tracks based upon tracking error covariance matrices. Other elements of the tracker are somewhat more sensor specific, but the addition of new sensor types would involve only the addition of new sensor specific modules, not changing any of the existing modules.

The upper portion of Figure II-5 shows the tracking initiation function that is performed using only acoustic sensors since microphone arrays can be used to scan a full 360° in azimuth in a short period of time. Azimuth information from other sites is used to initiate tracks because at least two azimuth measurements from two sites are required to obtain a target position estimate. The track initiation algorithm compensates for acoustic delay and provides a real-time estimate of position and velocity. The acoustic and TV position tracking elements are extended Kalman filters that update position tracks using the sensor measurements. The acoustic tracker compensates for acoustic delays.

Unlike acoustic arrays, TV cameras inherently have a limited field-of-view and are more appropriate for improving tracks than for looking for new targets. Thus, in addition to the Kalman filter algorithms that use TV azimuth measurements, the test bed incorporates a mechanism for cueing the TV subsystems. Figure II-6 shows the interrelationship between the tracker and the TV subsystem. The TV is a user of the tracking system and is provided with target tracks just like any other user. It then selects a specific target, controls the camera to point at the target, processes the image data, and provides the resulting azimuth measurements, if any, to the tracker. The TV subsystem has a closed-loop feedback relationship with the tracking system. The acoustic system operates open loop, providing acoustic measurements without any guidance from the tracker.

Final demonstrations of distributed multisensor tracking in the test bed were preceded by a series of system integration and checkout experiments during this entire report period. Initial experiments involved the testing of the new acoustic signal-processing and TV signal-processing

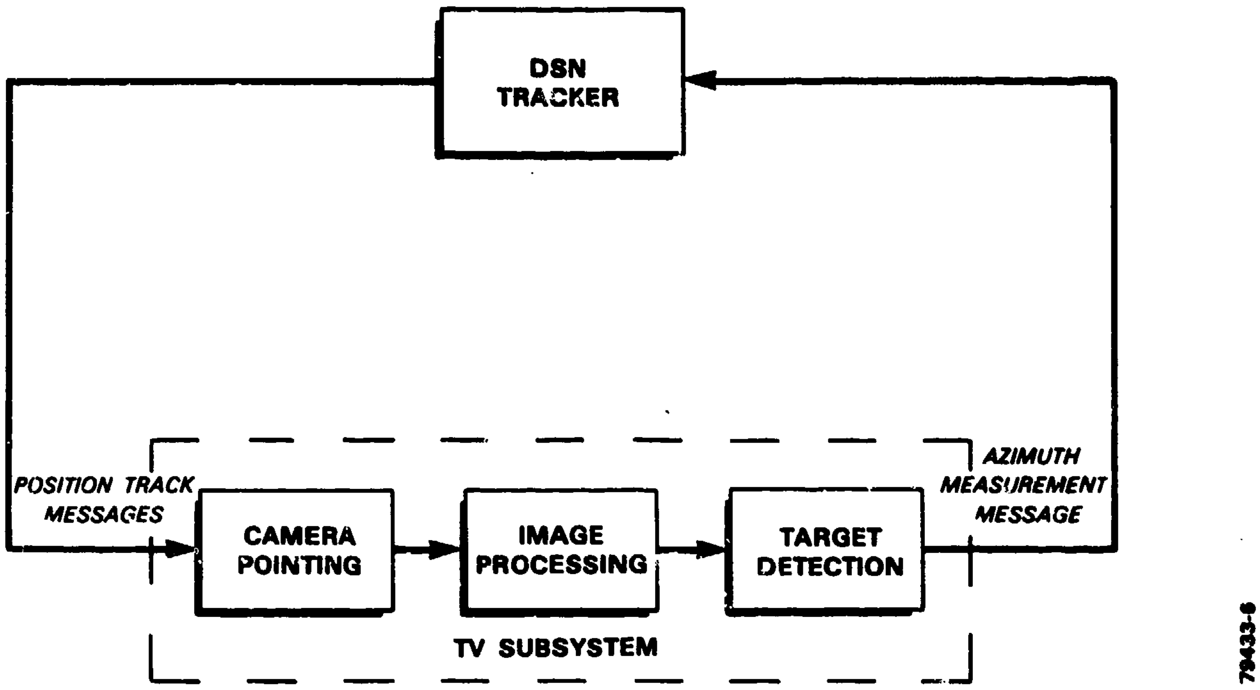

Figure II-6. TV subsystem elements and interactions with DSN tracker.

software with single sensor sites. Subsequent experiments involved up to three acoustic sites and one TV site locally interconnected by means of an Ethernet cable. The last series of experiments involved integrating the microwave radio systems and expanding the experimental baseline to 3 km as well as adding other acoustic sites and TV sites. Experiments utilized targets of opportunity as well as controlled experimental aircraft. Experiments with controlled targets utilized North-South and East-West target paths, depending upon wind conditions that determined which of the Hanscom runways was in use.

Experiments with controlled targets were performed on the average of one each two weeks starting in June. Heavy equipment construction sites within 50 to 100 m of acoustic arrays created severe noise conditions and interfered with experiments in several instances. Heavy winds and bad weather forced the cancellation of some experiments when it was not possible to fly the target aircraft as required by the test plan. Many additional system integration tests were performed using only targets of opportunity.

Algorithms, experimental procedures, and hardware reliability were all significantly improved during this integration and test period preceding final DSN demonstrations.

III. TEST-BED IMPLEMENTATION

Implementation tasks accomplished during this report period to support DSN test-bed experiments and demonstrations are summarized below. Topics covered include TV subsystems, acoustic signal-processing subsystems, communication system software, and the test-bed control and situation display workstations.

A. TV SUBSYSTEMS

The TV subsystem is organized as illustrated in Figure II-6. First, the camera-pointing algorithm selects one of the target tracks in the position message, applies controls to the camera to place the selected target in the camera field-of-view (FOV), and acquires two video frames spaced 1/30 s apart. Second, the image-processing algorithm computes the difference of the two video frames. Third, the target-detection algorithm processes the difference to determine whether a target is present in the FOV and, if it is, to determine its azimuth.

Improved versions of the camera-pointing and target-detection algorithms were developed during this reporting period. Their performance is much superior to the initial algorithms; sufficient to reliably demonstrate cooperative TV/acoustic tracking in the DSN test bed. In addition, a second TV node was integrated into the test bed and software was developed to transmit images from the second TV, over an Ethernet or radio link, to the display monitor of the other TV node. More detailed information concerning the new TV algorithms, their performance, and the remote display capability is given below.

The new capabilities added to the camera-pointing software were adaptive control of elevation angle and zoom setting, and the addition of a false target detector.

Analysis and early test-bed experiments with the TV subsystems revealed that targets were often lost or not acquired because the fixed elevation setting of the camera did not keep the aircraft in the FOV. The elevation-control algorithm solves this problem by adjusting the camera elevation each time a position cue is received. As illustrated in Figure III-1, the algorithm first uses information in the cueing message to determine if the target is a new target. If it is new, the camera elevation angle is selected using a default value for altitude. If the TV subsystem has previously attempted to detect the target, two options exist depending on the outcome of that attempt. First, if a detection was obtained as a result of any of the last three attempts, a 2° change is made in the camera elevation, based on the latest available elevation estimate. The objective is to keep the target near the center of the vertical FOV. Second, if the target was not detected in any of the last three attempts, a search in elevation angle is initiated. This algorithm has significantly increased the probability of target acquisition and the length of time that good measurements are obtained from the TV subsystem.

The probability of detecting a target with a TV camera is a function of camera zoom, which determines the camera FOV, and the target range. It is equal to the product of the probability that the target is in the FOV multiplied by the probability of detection given that the target is in

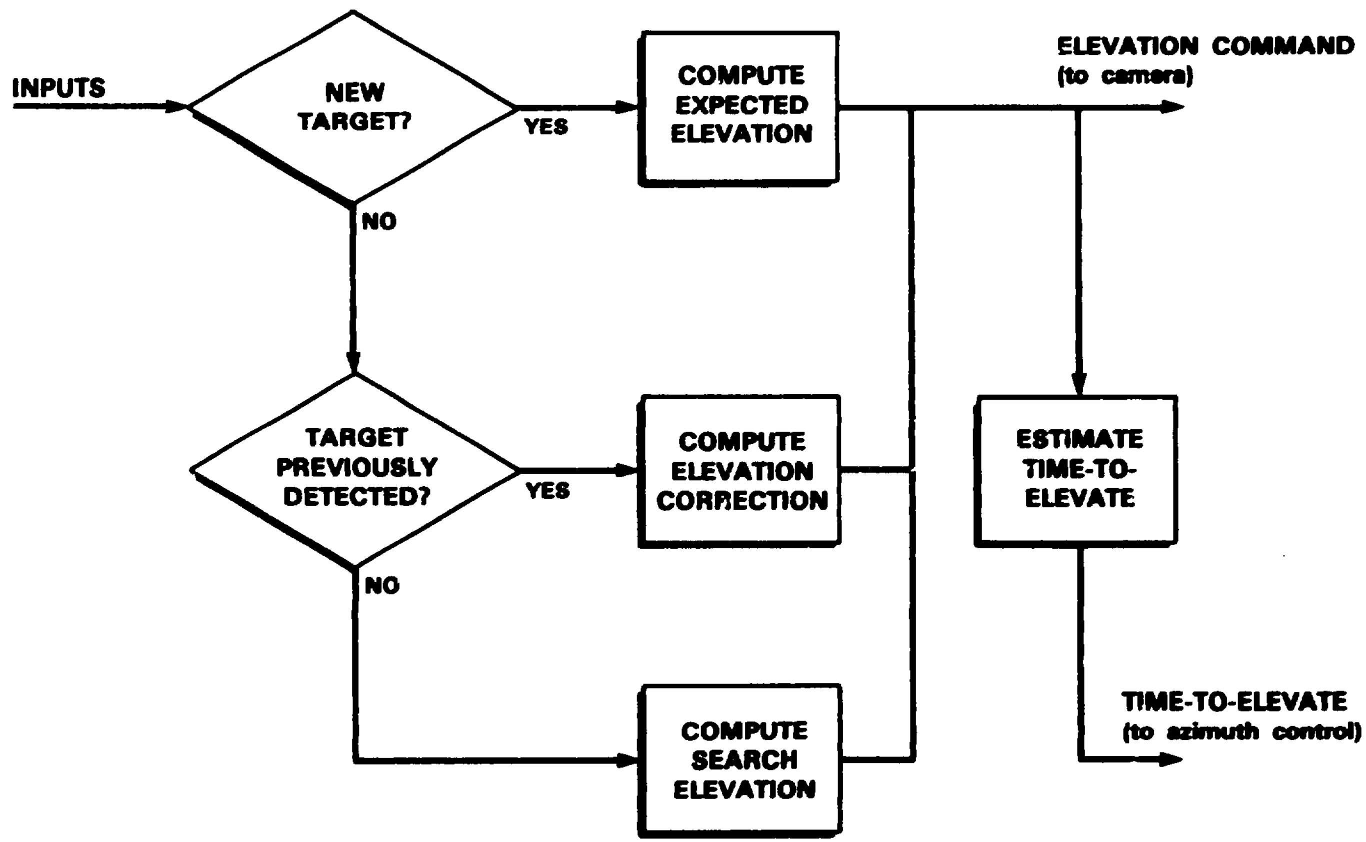

INPUTS

- *TARGET ID*
- *TARGET ESTIMATED STATE* } *FROM DSN TRACKER*
- *TARGET FOUND/NOT FOUND*
- *TARGET ESTIMATED ALTITUDE* } *FROM TARGET DETECTION ALGORITHM*

OUTPUTS

- *ELEVATION COMMAND*
- *TIME-TO-ELEVATE*

Figure III-1. Elevation control algorithm.

79433-7

the FOV. This probability is plotted in Figure III-2 as a function of FOV for a fixed target range and as a function of range for a fixed FOV. First consider the case of a fixed target range. For a small FOV (large magnification), the probability is small due to errors in the position cue and in camera positioning. For a large FOV (small magnification), the probability is small because, even if the target is in the FOV, the size of the target image on the screen is too small to be reliably detected. Similar logic explains the general shape of the probability curve as a function of range for a fixed FOV.

The zoom control algorithm adjusts the FOV in an attempt to maintain a high detection probability independent of target range. It does this by maximizing the FOV while keeping the size of the target image between 1 and 1.5 cm. This size provides a near unity detection probability when the target is in the camera FOV. The calculation of image size is based upon the assumed physical dimensions of targets. The position and velocity provided by the DSN tracker are used to predict future positions and the image size is computed for the time at which image frames will be taken (after positioning the camera in azimuth and elevation). If the estimated size is in the 1- to 1.5-cm range no action is taken; otherwise a table of options is revised with a target size and time-to-zoom calculated for each option. Selection of an option depends on whether enough time is available to complete the zoom change during the camera azimuth and elevation slew time. If a zoom command is issued the option table is updated to correctly represent the state of the zoom control options after the command is executed.

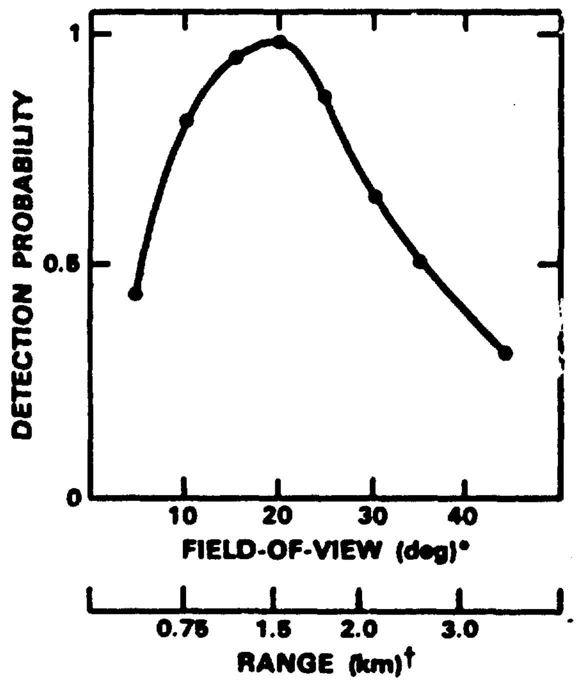

* RANGE FIXED AT ≈ 1.8 km
† FIELD-OF-VIEW FIXED AT ≈ 25°

79433-8

Figure III-2. Probability of target detection as a function of field-of-view and range for a 10-m target and 200-m RMS target position error.

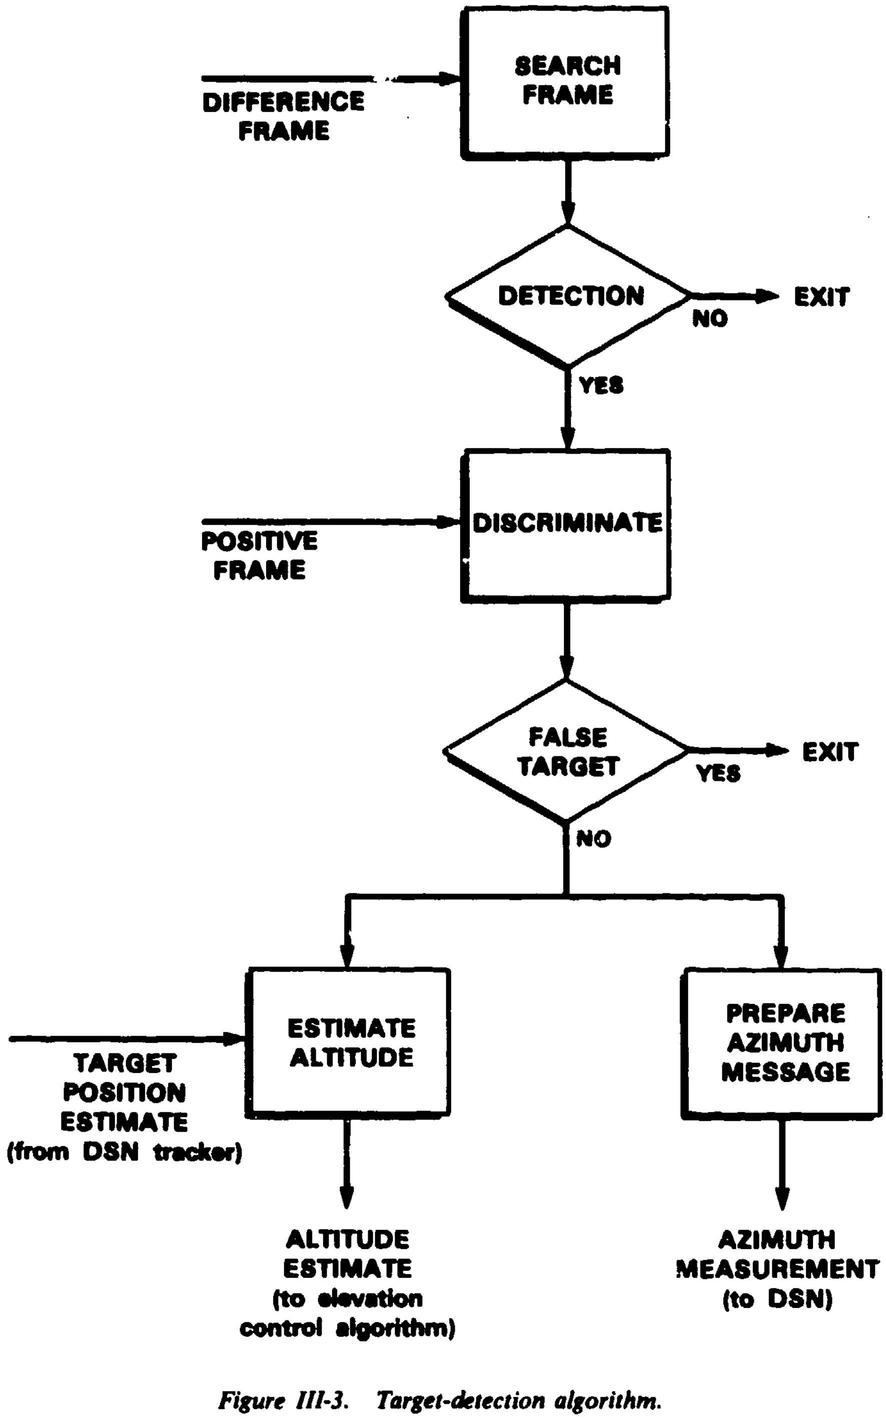

Figure III-3. Target-detection algorithm.

A final feature added to the target-selection and camera-pointing algorithms was a mechanism to eliminate certain kinds of false cues generated by the acoustic system. Often during the course of experiments at Hanscom Air Force Base, position cues were provided that corresponded to noisy ground construction equipment. Two changes were made to avoid wasting TV resources on these cues. First, a false-target list was created. If a target is not detected after a search in elevation, it is entered on this list and all future position cues for this target are disregarded. Second, exclusion regions are defined to specify areas to be excluded from consideration for target selection. This mechanism can be used to exclude known fixed acoustic noise areas from the TV selection process.

The target-detection algorithm in the TV subsystem also has been improved in several ways. It now searches the entire frame (the previous algorithm searched only a fraction of the frame) and locates smaller targets. It measures target elevation (in addition to target azimuth) and estimates target altitude which it provides to the camera-elevation control algorithm. Finally, it discriminates against moving bright clouds and other such false targets.

The improved target-detection algorithm performs the steps illustrated in Figure III-3. First, the difference frame (generated by the image processor by calculating the difference of two positive frames) is scanned in 4 × 6 pixel blocks, which corresponds to image offsets between blocks of about 0.25 cm. This block size is smaller than the target image size which is maintained by the zoom control. This provides the detector more than one chance to detect the target and makes the system more robust in the face of errors in the selection of the zoom. The scan operation consists of summing a subset (10 out of 24 pixels) of intensity values within each block. A subset of pixels is used for each block only to reduce processing time. Next, if one of these sums exceeds a threshold, the brightness of the corresponding block of one of the contributing (non-differenced) frames is tested against another threshold. This discriminates targets from moving bright clouds. Finally, if a detection has been made, the algorithm calculates the corresponding target azimuth and altitude and provides them to the tracking system and the elevation-control algorithm, respectively.

The improved TV algorithms were tested, refined, and demonstrated during a sequence of experiments with live targets conducted during this summer and fall. Typical behavior of the camera-pointing algorithm during a real-time experiment on 15 September is shown in Figure III-4. The figure shows the estimated trajectory of a UH-1 helicopter with annotations that give the camera elevation and field-of-view during the entire period that cues were provided to the TV subsystem. The track is roughly East to West parallel to one of the Hanscom Field runways.

When the initial cue was received, the camera was pointing East with a 3° elevation angle and 7° FOV. The pointing algorithm then predicted a target/camera intercept and, while the camera slewed in azimuth, elevation and zoom settings were changed. Elevation was changed to 11° reflecting the expected target altitude and position. Field-of-view was increased to 17° because the target was close to the camera and very little magnification was needed to obtain a 1-cm target image. When camera settings were completed, two frames were stored and processed,

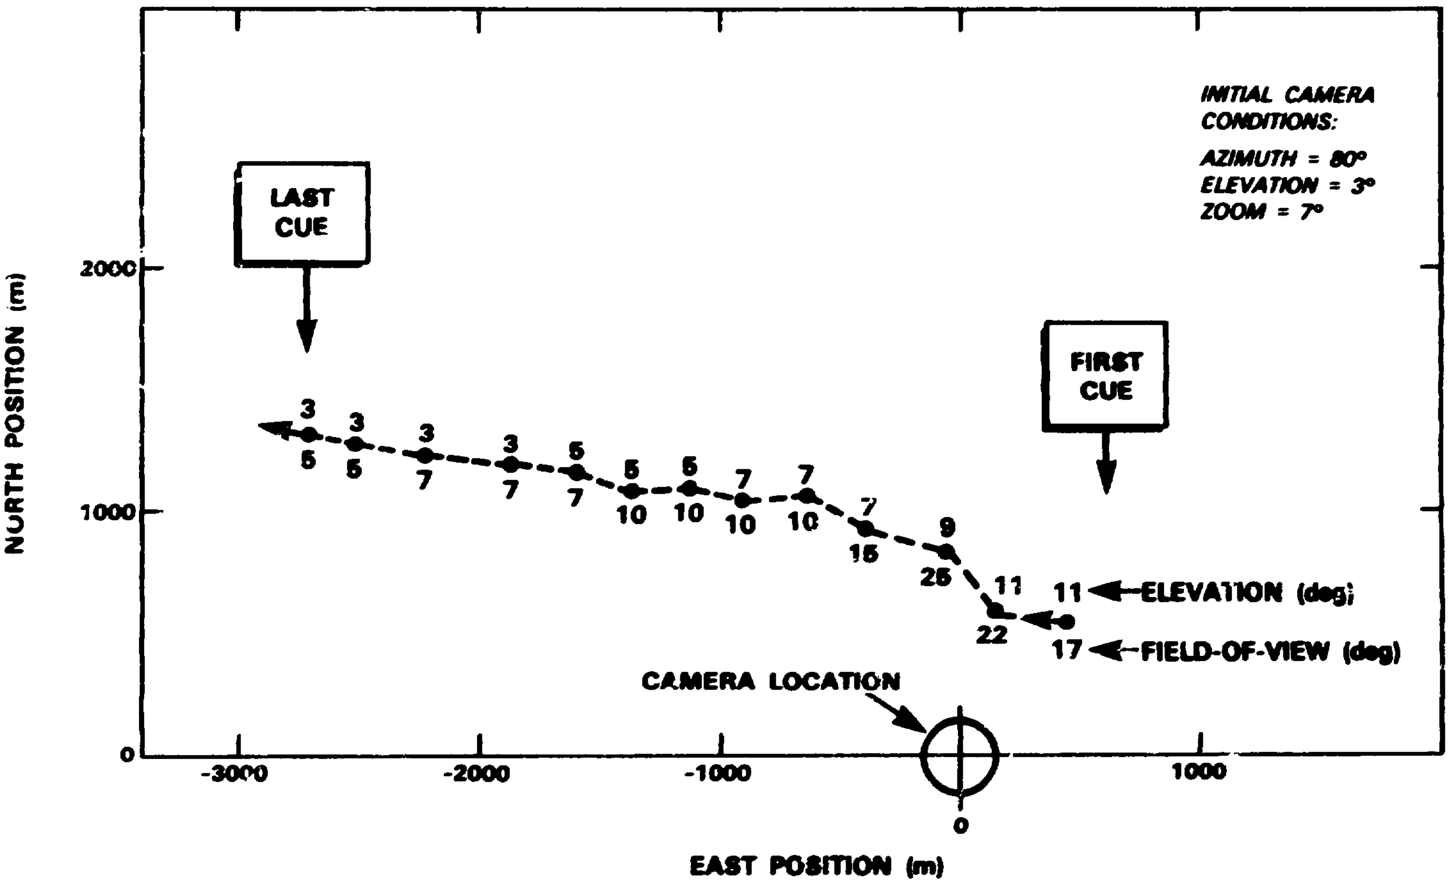

Figure III-4. Variation in zoom and elevation for a live target experiment.

79433-10

a target found, and the azimuth measurement sent to the DSN tracker. Elevation and zoom were adjusted for subsequent cues as shown in the figure. Elevation decreased as the target moved away from the camera because of the effect of geometry and small changes in the target altitude. Field-of-view was increased (magnification decreased) while the target was close to the camera and decreased (magnification increased) as the target moved farther from the camera.

As mentioned previously, a second remote TV subsystem was added to the test bed. Unlike the first TV subsystem, the full TV image from the remote TV node is not available for display to the user located with the first TV because of bandwidth limitations. The zoom-control algorithm, however, regulates the size of the target projection on the monitor screen to, approximately, 1 to 1.5 cm. It is thus possible to specify a 2-cm by 2-cm patch on the remote monitor screen which contains the target's image and to transmit this patch over an Ethernet or radio link in a standard 512-byte DSN message. Software was implemented to do this and to display the patch superimposed upon the user's TV display. The patch is quantized into only two gray scale levels and displayed on the home-base TV monitor at the exact location where the target appeared in the remote TV monitor. Despite the quantization, the helicopter silhouettes were clearly discernible in demonstration tests.

B. ACOUSTIC SIGNAL-PROCESSING SUBSYSTEM

A new version of the acoustic signal-processing system (SPS) software was released during this reporting period. The system utilizes the SPS real-time clock to time-stamp data and thereby provide the synchronization needed for tracking. It also provides more flexibility and convenience of operation in starting and stopping the system and changing algorithm parameters. For example, the system now offers complete flexibility to suppress processing of any subset of microphone channels that may be temporarily inoperative. Previously it was necessary to repair or reconfigure the hardware when a channel was inoperative. System processing bandwidth was increased to enable processing of up to 12 channels at the maximum sampling rate of 2048 samples per second per channel. The new system also provides time-series windowing functions that reduce false alarms.

Address space limitations in the SPS DEC PDP-11 computer complicated the implementation of the new software. In order to add the necessary features, it was necessary to implement an overlay process which time-division multiplexes the virtual address space of the process over a larger physical address space. The overhead resulting from the address space remapping had to be carefully managed to provide real-time operation.

C. SYSTEM AND COMMUNICATION SOFTWARE

The final version of the DSN Nodal Run-Time System (NRTS) for the Standard Nodal Computers (SNC) was completed during this reporting period. It allows the SNCs to be connected to microwave radios as well as to Ethernets. It supports the communications needed to perform DSN test-bed experiments with nodes interconnected by microwave radios as well as Ethernets.

Gateway software that uses the new NRTS was also implemented. It makes the test bed appear to the applications software as though all nodes are interconnected by a single Ethernet, even when there are actually multiple Ethernets interconnected by the microwave radios. A gateway SNC is attached to an Ethernet and to a microwave radio. The gateway functions are invisible to the DSN applications.

The gateway operates as follows. Any message received by an SNC containing a gateway is delivered to the user software in the node, unless it is a Point-to-Point (PTP) message addressed to another node. A message received from the radio link is rebroadcast on the Ethernet. A message received from the Ethernet is rebroadcast over the microwave radio link. The only exception is a PTP message addressed to a local user of the SNC. There is no reason to rebroadcast such a message since it has reached its destination. Only one radio can be attached to a gateway and each Ethernet in the test bed is constrained to have only one node with an attached radio. These constraints made it possible to design a very simple system.

The code to handle the microwave radio and perform gateway function increased the size of the NRTS module in processor P1 in each SNC. These functions also place a greater drain on the computational resources. Therefore the gateway function has been isolated into SNCs with no other DSN functions. The gateway software also provides traffic statistics every few seconds to help identify when there may be problems with the communication system.

D. EXPERIMENT CONTROL AND SITUATION DISPLAY WORKSTATION

During this reporting period a new dynamic situation display was integrated into the experimental test bed. The display, which runs on Silicon Graphics, Inc. workstations, was originally developed for use in multisite integration experiments and has been previously described in detail in the March 1986 semiannual. Modifications, necessary for integration into the test bed, were made to the message receiving routines and Shell scripts to read position track messages from the user interface program in addition to multisite integration messages from a file. Adjustments were made to display messages in real time and to display a local area map and sensor position information. Modifications were made to correct errors in the initial version of the display and to allow a greater number of tracks to be displayed and remain on the screen for a longer period of time.

Integration of the new situation display and related software into the test bed represents an important project milestone. We now have a self-sufficient experimental test bed which can be transported to and used in the field. Experiments can be run from and displayed on a Silicon Graphics, Inc. workstation as conveniently as on our laboratory-based VAX system.

IV. DSN PROGRAM SUMMARY

This is the final Lincoln Laboratory DSN Semiannual Summary Report. The following sections therefore attempt to summarize overall program accomplishments, cite lessons learned concerning the development of DSN systems, and identify remaining research and development areas.

A. ACCOMPLISHMENTS

A primary objective of the Lincoln Laboratory DSN project has been to establish the feasibility of distributed surveillance systems implemented in the form of autonomous cooperative processes. This has been accomplished. The real-time acoustic and TV tracking demonstrations described in Section II of this report constitute a primary demonstration of that feasibility.

Establishing the feasibility of DSN systems has involved many secondary but still important accomplishments. Major items include the development of distributed algorithms, development of acoustic direction-finding methods, and the implementation and use of the DSN test-bed system. These are described in more detail below, followed by a brief summary of other lesser but important project accomplishments.

Distributed algorithms have been developed to: (1) perform sensor independent tracking;[1-5] (2) perform aircraft tracking using small microphone arrays;[1-5] (3) perform tracking with a mix of sensor types, including acoustic and TV sensors;[6,7] (4) integrate track data from many distributed nodes;[6] and (5) determine the configuration of a geographically distributed network by means of internodal range measurements.[2,8] Track integration and network configuration algorithms were demonstrated off-line in the form of multiple cooperating processes implemented on a single computer. Tracking algorithms were demonstrated in real time in the DSN distributed test bed as well as off-line.

The tracking algorithms developed under this program depend upon well-understood estimation methods. They can be easily adapted to accommodate new sensor types by the addition of sensor specific extended Kalman filters for each new type. They provide estimates of present target positions despite acoustic propagation delays, and automatically utilize the data from any number of DSN nodes. These algorithms replace less general acoustic tracking algorithms that were previously developed and evaluated as part of the DSN project.[9-11] The approach for the earlier algorithms was to minimize assumptions concerning target dynamics. This resulted in time-delayed position estimates, ineffective use of the data from the sites nearest to the target, and provided no obvious way to combine acoustic with non-acoustic tracks. The initial algorithms also provided no theoretical model for utilizing data for other than pairs of acoustic sites. The new algorithms solved all of these problems by using models of target dynamics and using them to perform all tracking in terms of present target position, independent of acoustic propagation delays.

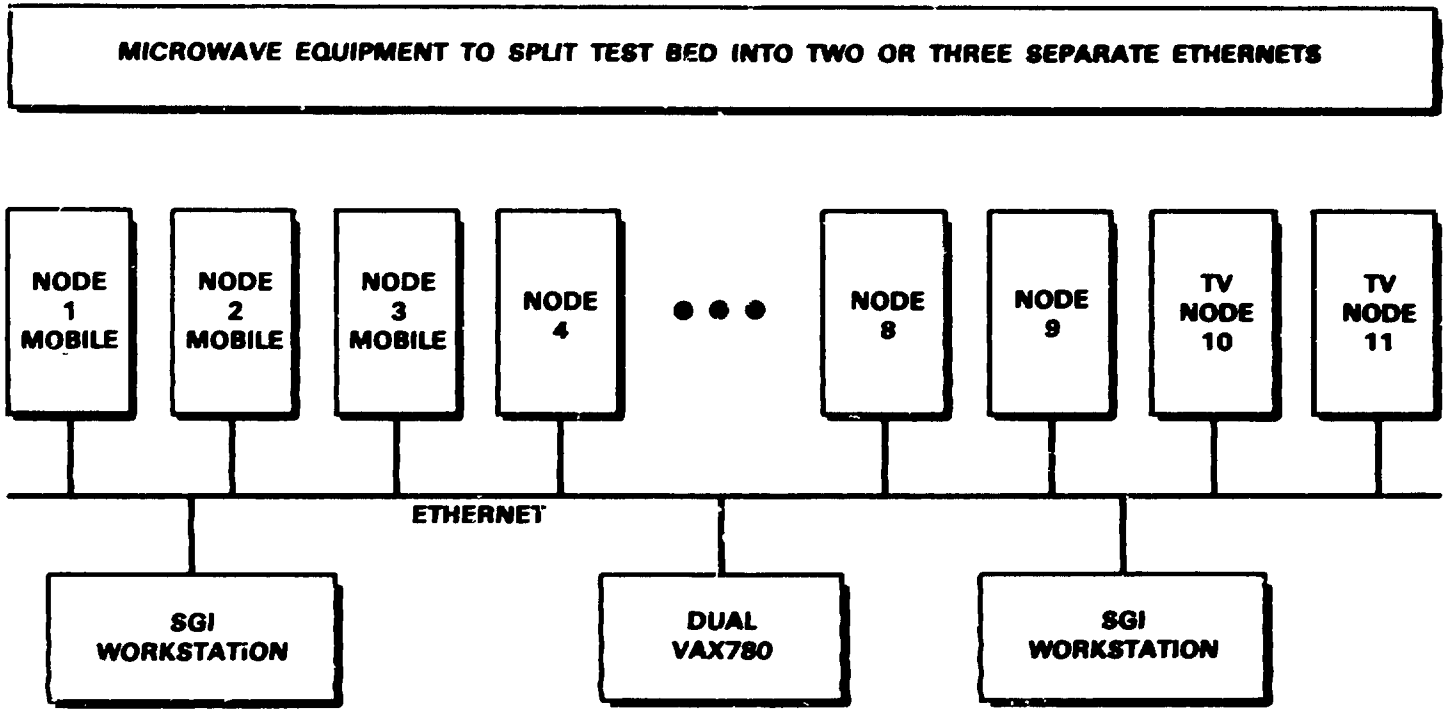

NOTE:

(A) NODES 1 TO 3 IN FIELD DEPLOYABLE TRUCKS
(B) NODES 1 TO 6 HAVE ATTACHED ACOUSTIC ARRAYS

Figure IV-1. DSN test-bed system.

79433-11

The demonstration of real-time aeroacoustic tracking required the development of acoustic signal-processing algorithms to detect targets and determine directions. A wideband array-processing algorithm was developed for this purpose[12,13] and was implemented and used for real-time tracking demonstrations. This algorithm replaced a more conventional one that was developed and used for initial DSN experimentation.[11,14] The new algorithm provided more reliable detections and direction finding while requiring fewer computations than the original.

A major part of the Lincoln Laboratory DSN effort involved the development of the test-bed system that was used for algorithms development and demonstrations. Figure IV-1 shows a block diagram of the test bed. Every node contains from one to three single board microcomputers. These computers normally perform tracking and communication functions. Six of the nodes also contain acoustic subsystems consisting of a small array of microphones and all the electronics required to collect and process acoustic data in real time. Three of the nodes containing the acoustic subsystems are installed in vehicles for field deployment. Two of the nodes are dedicated to TV-specific functions. The test bed also includes microwave communication equipment and software to allow the test bed to be split into up to three elements. This capability is essential for the execution of experiments with mobile nodes separated by more than about a kilometer. Two UNIX workstations and a VAX computer constitute the remainder of the test-bed system. They are employed for software development, test-bed experiment control, and other analysis functions. The test-bed system constitutes a large complex configuration of hardware and software that required substantial effort to develop and use to demonstrate real-time distributed aircraft tracking.

Additional accomplishments included the development of a broadcast protocol for the Communication Network Technology radios developed by Group 86 at Lincoln Laboratory,[6] transfer of technology to the Air Vehicle Survivability Evaluation (AVSE) project at Lincoln Laboratory, and the preliminary investigation of the application of Artificial Intelligence (AI) approaches to some acoustic data interpretation problems.[15,16] The technology transfer entailed the use of DSN hardware and software for the collection of air vehicle data and the use of DSN algorithms to validate estimated detection ranges and demonstrate acoustic aircraft tracking.

B. SYSTEM DEVELOPMENT LESSONS LEARNED

The DSN program at Lincoln Laboratory has established the feasibility of distributed systems for target surveillance and tracking by developing a test bed and using it to demonstrate distributed tracking by means of acoustic and TV sensors. While doing this a number of important expectations were confirmed and lessons learned regarding the development of a distributed DSN system. These concerned: (1) the role of experiments, (2) operating systems and debugging, (3) internodal communication, (4) physical distribution, and (5) system complexity. Following is a summary of observations for each of these areas.

Controlled and repeatable experimental capabilities are essential for the development and test of any complex system or set of algorithms. This is especially important for distributed systems which are more likely to exhibit unexpected behavior than more traditional systems. Simulated data, raw sensor data, and partially processed data recorded by several nodes during live experiments were all used extensively for controlled and repeatable testing. The test bed was designed to support real-time experiments using these data. This allowed well-controlled algorithm and system testing which would have been impossible to achieve with only live experiments.

A related fact is that when live experiments are performed it is essential to analyze and fully understand system behavior, especially unexpected behavior. This is possible only if the system saves sufficient data and provides the capability to reprocess the data to test if the situation has been correctly diagnosed and corrected. Detailed data logging and the reprocessing capability might not be as important in a fielded system but are essential for the development system. This is an example of the general rule that the development system requires more capability than the fielded system. This is especially true for distributed systems.

Neither a distributed operating system nor well-developed remote debugging tools were available at the inception of the DSN effort. It was necessary to develop an operating system and interprocess communication mechanisms for the test bed. The resulting software is necessarily simple and provides limited support for remote debugging. The lack of a fully functional distributed operating system and remote debugging tools substantially increased the difficulty of implementing and testing application level DSN software for the test bed. Although operating systems and tools for distributed systems continue to be areas of active research it is clear that distributed systems, often with sophisticated workstations and file servers interconnected by a local area network, are becoming relatively common. A new implementation of a DSN test bed could make use of these advances and would be a far more flexible and easy use vehicle for research and development of future DSN systems.

The DSN approach is to use broadcast communications for distributed tracking and to use point-to-point communication for other functions such as collection of data from areas and to support system operation and development. As in the case of operating systems, no satisfactory communication system existed for the test bed at the inception of the DSN effort. Experimental Packet Radio systems could probably be modified and used for DSN applications but this was not possible in the DSN time frame. Therefore, part of the DSN effort included the procurement of microwave communication and Ethernet hardware to interconnect test-bed nodes and development of special-purpose software to provide internodal communication services. Although this required considerable effort, it is a minimal system which lacks the general suite of features that one would need for an operational DSN, or even for extensive DSN experimentation in the field.

The communication system is an essential tool for DSN development as well as being an important element of the final system. However, communication system requirements are more stressing during the research and development phase than for the operational case. For example, the development process requires the collection of detailed information for test and evaluation and frequent changes in the network software stress the communication system. As in the case of operating systems, a next generation of DSN experimental systems should make use of advances in local area network and packet radio technology.

Another important point is that the communication system performance is typically limited as much by software and protocol processing as it is by more traditional factors involving the physical link. Throughput in the test-bed system is adequate to support simple distributed tracking experiments. But larger scale experiments or experiments requiring substantially more data to be exchanged would be limited by the present communication system. The communication protocols and message-processing loads are the limiting factors, not the physical communication medium. This is common for computer-to-computer communication systems and should be kept in mind for any future DSN test-bed or system development.

Software development and debugging typically involve an ongoing process of compilation, changing load modules, and testing. This process is often more difficult in a distributed environment than in a centralized one. This was very true for the DSN test bed which was constrained by operating system features, available software tools, and the communication system. Whenever possible the development of all test-bed software was done in the centralized environment of a single computer, including some simulation of the distributed environment. Future system development should probably exploit modern local area network technology to provide a good environment for DSN work. The DSN test bed does include Ethernet hardware, but advanced general-purpose workstations and well-supported network software were not widely available in the earlier stages of the project. That situation has now changed.

Distributed software testing in the test-bed nodes was deferred as long as possible because of the difficulty of working with a distributed system with available tools. During early stages of test-bed development, 9600 baud serial lines were used for communication between nodes and between the software development computer and the nodes. Distributed testing required lengthy downloading to the nodes. The final test-bed version includes floppy disks at each node. Disks are written on the central software development machine and physically loaded on the nodes to make software changes. It would also be possible to download using the test-bed communication system, but the necessary software has not been developed for the test-bed hardware. In addition, the process might still be a bottleneck because a single machine would perform all downloading. This could be somewhat alleviated by the implementation of an incremental downloading system which provided for modular downloading of only small portions of the software that change.

Distributed systems must be designed from the start for remote operation of system elements. This includes remote hardware startup and diagnosis as well as distributed software features. The test bed was implemented using commercial board and higher level products that were not designed for use in a distributed system. The scale and nature of the DSN test-bed effort forced us to take this approach and to minimize attention paid to designing for remote operation and test. As a result, it is clear in retrospect that the test bed is more difficult to operate and use than would otherwise be possible with built-in remote operation and test features.

These are examples of difficulties that arise in the development and testing of distributed systems. A good infrastructure to support the development of distributed systems was essentially nonexistent when the DSN project was started. The situation is better now but there is still a great deal to be done to make the development, testing, and use of DSN systems easier.

And finally, DSN systems are complex systems. This calls for extensive test and maintenance, automated tools for system operation, and automated tools for software development. And most important, hardware and software modularity are essential to keep complexity under control. There are two aspects to this modularity. One relates to how to build DSN systems that are robust and reliable. The other relates to the development process for complicated systems.

One advantage of DSN systems should be their insensitivity to failures and ability to adapt to the addition of new system elements. The implementation in the form of autonomous cooperating processes is essential to achieving this advantage. This is a manifestation of the use of modularity to achieve robust and reliable operation. The Lincoln DSN project has emphasized this approach to tracking. The same approach should be applied to all system elements including the operating system, communications, self-location, etc.

Autonomous processes require modular hardware to be most effective. A DSN system automatically provides some hardware modularity; individual nodes are physically separated. But even a single node can be a complex system performing many different functions. The design should minimize the number of single point failures that can incapacitate a node. A simple problem that could be easily repaired in a centralized system may be more difficult to repair in a remote node. Overall system reliability, because there will be so many parts, requires that failure modes be as small grain as possible. Thus a single powerful nodal computer is less desirable than a multiprocessor design.

Software modularity as an aid to system development and testing is not unique to DSN systems. But the number of different software components in a DSN and the fact that it is a distributed system make it even more important that modular designs and automated system development tools be used. Even the simple experimental DSN test-bed system is complex, containing hundreds of modules. A modular design philosophy and automated tools were employed and we believe that contributed substantially to the success of the project.

C. REMAINING RESEARCH AND DEVELOPMENT ISSUES

Although the DSN program has successfully achieved the goal of demonstrating the feasibility of DSN systems, it is clear there are many areas where additional work could be done or is needed before full-scale DSN development could begin. Following is a summary of such items in three categories: (1) distributed acoustic surveillance, (2) experimental system and subsystem prototypes, and (3) applications of artificial intelligence. Much of the work could be general but some would benefit from emphasizing systems for specific applications.

There has been considerable emphasis upon the use of acoustics during the DSN project. The feasibility of exploiting acoustic emanations for detection and tracking has been established but there are several remaining research questions. First, what are the limits of acoustic tracking for the case of multiple and maneuvering targets? Analysis, algorithm development, and experimentation are required to answer this question. Second, it is clear that some target identification can be achieved using passive acoustics. How well this can be done and how to do it should be determined. Third, acoustic performance will vary considerably with environmental conditions.

This needs to be understood, quantified, and analyzed in the context of specific system deployments and requirements.

There appear to be no fundamental impediments to the development of a DSN system. However, there are a number of engineering problems that need further work to demonstrate solutions before a major commitment is made to DSN system development. A fully functional experimental prototype system should be developed, demonstrating all essential DSN features and subsystems. Major subsystems include modular prototype nodes with remote start, test and operation, and a prototype communication system. Unlike the existing test-bed system, the experimental prototype would be functionally complete. For example, it would maintain time synchronization, provide multi-hop point-to-point communication, and maintain network configuration tables. Each function would be performed in a way intended for a real DSN system. This differs from the existing test-bed system in which, for example, nodes are either manually synchronized or synchronized by a satellite clock, communication services are very specialized and network configuration is by means of manually entered tables that list all nodes in the small test-bed network. The experimental system could be used to demonstrate that no unsolved system problems remain. This would be done by using the experimental system for extended periods of time in such a manner as to be consistent with operational and large-scale use. Development and evaluation of the experimental system would entail the development of improved system development environments to demonstrate that such tools are sufficiently mature to support full-scale system development. This might include the development of new languages and simulation facilities for distributed systems development.

DSN systems offer many opportunities for AI research and for the application of AI approaches to sensor data interpretation and to system diagnosis and control. In the case of system diagnosis and control, automated aids will be essential to the operational utility of DSN systems. In addition to straightforward automation, it may be feasible to exploit knowledge-based methods for the development of diagnosis and control software.

Finally, acoustic surveillance provides at least three interesting data understanding problems. One is to develop acoustic tracking systems with improved performance by explicitly integrating knowledge into the system. This would have broad application not limited to acoustic or distributed systems. Target recognition is another area where AI methods could be applied. The third area is knowledge-based adaptive signal processing which would involve feedback from higher levels of interpretation to the choice of algorithm.

GLOSSARY

AI	Artificial Intelligence
AVSE	Air Vehicle Survivability Evaluation
CNT	Communication Network Technology
DSN	Distributed Sensor Networks
FOV	Field-of-View
NRTS	Nodal Run-Time System
PTP	Point-to-Point
SATS	Semiannual Technical Summary
SGI	Silicon Graphics, Inc.
SNC	Standard Nodal Computer
SPS	Signal-Processing Subsystem

REFERENCES

1. Distributed Sensor Networks Program Semiannual Technical Summary, Lincoln Laboratory, MIT (30 September 1983), DTIC AD-A146209.
2. *Ibid.* (30 September 1984), DTIC AD-A158033.
3. *Ibid.* (31 March 1985), DTIC AD-A160596.
4. R.R. Tenney and J.R. Delaney, "A Distributed Aeroacoustic Tracking Algorithm," American Control Conf., San Diego, California, 6-8 June 1984.
5. J.R. Delaney and R.R. Tenney, "Broadcast Communications Policies for Distributed Aeroacoustic Tracking." American Control Conf., Boston, Massachusetts, 12-14 June 1985.
6. Distributed Sensor Networks Program Semiannual Technical Summary, Lincoln Laboratory, MIT (31 March 1986).
7. *Ibid.* this report, Sections II and III.
8. E.A. Hinzelman, "A Nonlinear Distributed Algorithm for Self-Location," MS Thesis, Department of Electrical Engineering, MIT, June 1985.
9. Distributed Sensor Networks Program Semiannual Technical Summary, Lincoln Laboratory, MIT (30 September 1982), DTIC AD-A131245/3.
10. *Ibid.* (31 March 1980), DTIC AD-A103045/1.
11. *Ibid.* (30 September 1979), DDC AD-A086800.
12. S.H. Nawab, F.U. Dowla, and R.T. Lacoss, "Direction Determination of Wideband Signals," IEEE Trans. Acoustics, Speech, Signal Process. **33**, 1114-1122 (October 1985).
13. S.H. Nawab, F.U. Dowla, and R.T. Lacoss, "A New Method for Wideband Sensor Array Processing," Intl. Conf. Acoustics, Speech, Signal Process., San Diego, California, 19-21 March 1984.
14. Distributed Sensor Networks Program Semiannual Technical Summary, Lincoln Laboratory, MIT (30 September 1978), DDC AD-A070455/1.
15. E.E. Milios, "Signal Processing and Interpretation Using Multilevel Signal Abstractions," PhD Thesis, Department of Electrical Engineering, MIT, May 1986.
16. S.H. Nawab. V. Lesser, and E. Milios, "Conceptual Diagnosis of an Algorithmic Signal Processing System," Intl. Conf. Acoustics, Speech, Signal Process., Tokyo, April 1986.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

A182216

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified	1b. RESTRICTIVE MARKINGS
2a. SECURITY CLASSIFICATION AUTHORITY	3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited.
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE	
4. PERFORMING ORGANIZATION REPORT NUMBER(S) ESD-TR-86-196	5. MONITORING ORGANIZATION REPORT NUMBER(S)
6a. NAME OF PERFORMING ORGANIZATION Lincoln Laboratory, MIT 6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Electronic Systems Division
6c. ADDRESS (City, State, and Zip Code) P.O. Box 73 Lexington, MA 02173-0073	7b. ADDRESS (City, State, and Zip Code) Hanscom AFB, MA 01731
8a. NAME OF FUNDING/SPONSORING ORGANIZATION Defense Advanced Research Projects Agency 8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER
8c. ADDRESS (City, State, and Zip Code) 1400 Wilson Boulevard Arlington, VA 22209	10. SOURCE OF FUNDING NUMBERS PROGRAM ELEMENT NO. 62708E; PROJECT NO. 5D30 5T10; TASK NO.; WORK UNIT ACCESSION NO. AO5345

11. TITLE (Include Security Classification)
Distributed Sensor Networks

12. PERSONAL AUTHOR(S)
Richard T. Lacoss

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT (Year, Month, Day)	15. PAGE COUNT
Semiannual Technical Summary	FROM 1 April TO 30 Sept. 1986	30 September 1986	40

16. SUPPLEMENTARY NOTATION
None

17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	multiple-sensor surveillance system; distributed tracking; target surveillance and tracking; communication network; acoustic sensors; video sensors; low-flying aircraft; acoustic array processing; digital radio; distributed estimation

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

This report describes the work performed on the DARPA Distributed Networks Program at Lincoln Laboratory during the period 1 April through 30 September 1986

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT ■ UNCLASSIFIED/UNLIMITED ■ SAME AS RPT. ☐ DTIC USERS	21. ABSTRACT SECURITY CLASSIFICATION Unclassified
22a. NAME OF RESPONSIBLE INDIVIDUAL Maj. Thomas J. Alpert	22b. TELEPHONE (Include Area Code) / 22c. OFFICE SYMBOL

DD FORM 1473, 84 MAR   83 APR edition may be used until exhausted. All other editions are obsolete.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE